

Super-resolution multi-reference alignment

Tamir Bendory¹, Ariel Jaffe², William Leeb³, Nir Sharon⁴, and Amit Singer⁵

¹School of Electrical Engineering, Tel Aviv University

²Applied Mathematics Program, Yale University

³School of Mathematics, University of Minnesota, Twin Cities

⁴School of Mathematical Sciences, Tel Aviv University

⁵Department of Mathematics and Program in Applied and Computational Mathematics, Princeton University

Abstract

We study super-resolution multi-reference alignment, the problem of estimating a signal from many circularly shifted, down-sampled, and noisy observations. We focus on the low SNR regime, and show that a signal in $\mathbb{R}^M$ is uniquely determined when the number L of samples per observation is of the order of the square root of the signal's length ($L = O(\sqrt{M})$). Phrased more informally, one can square the resolution. This result holds if the number of observations is proportional to $1/\text{SNR}^3$. In contrast, with fewer observations recovery is impossible even when the observations are not down-sampled ($L = M$). The analysis combines tools from statistical signal processing and invariant theory. We design an expectation-maximization algorithm and demonstrate that it can super-resolve the signal in challenging SNR regimes.

1 Introduction

Model. We study the problem of estimating a signal from its circularly shifted, sampled, and noisy copies. More precisely, we consider N independent observations sampled from the model

$$y = PR_s x + \varepsilon, \quad s \sim \text{Uniform}[0, \dots, M-1], \quad \varepsilon \sim \mathcal{N}(0, \sigma^2 I), \quad (1.1)$$

where R_s denotes an operator that shifts the target signal $x \in \mathbb{R}^M$ circularly by s entries, that is, $(R_s x)[n] = x[(n-s) \bmod M]$, and P denotes a fixed sampling operator that collects $L \leq M$ equally-spaced samples. We assume that the random variable s is distributed uniformly over $[0, \dots, M-1]$, and the noise $\varepsilon \in \mathbb{R}^L$ is i.i.d. Gaussian. Explicitly, the i -th observation reads:

$$\begin{aligned} y_i[\ell] &= P(R_{s_i} x)[\ell] + \varepsilon_i[\ell] \\ &= x[\ell K - s_i] + \varepsilon_i[\ell], \end{aligned} \quad (1.2)$$

where $\ell = 0, \dots, L-1$, and $K := M/L$ is assumed to be an integer. Importantly, the shifts s_i are all unknown, and thus (1.1) is a special case of the multi-reference alignment (MRA) model, which we review in Section 2. Figure 1 presents an example of two observations with signal-to-noise ratio (SNR) equal to one (namely, the expected squared norm of the noise equals the squared norm of the signal).

Our goal is to estimate x from N observations sampled from (1.1). In contrast to previous works on MRA, the individual observations are down-sampled, and therefore recovering the full signal x is also a special case of the super-resolution problem. Accordingly, we refer to x as the “high-resolution signal,” while $y_1, \dots, y_N$ are the “low-resolution observations.” The parameter K can be thought of as a “super-resolution factor.” The difficulty in estimating x resides chiefly in three factors: the additive noise, the unknown circular shifts (the nuisance variables of the problem), and the sampling operator.

The statistical model (1.1) suffers from an intrinsic symmetry: it is invariant under a global circular shift since $p(y|x) = p(y|R_i x)$ for all $i = 1, \dots, M-1$. In this case, we say that the goal is to recover the signal up to a global circular shift. More formally, the goal is to recover *the orbit of x* :

$$Gx := \{gx \mid g \in G\}, \quad (1.3)$$

where $G := \{R_0, R_1, \dots, R_{M-1}\}$ is the group of cyclic shifts $\mathbb{Z}_M$. However, as will be shown in Section 3, without prior information on the signal, even the orbit Gx is not identifiable from the observations, and thus prior information on x is *necessary* for its identification.

Connection with sampling theory. We think of the discrete signal $x \in \mathbb{R}^M$ as Nyquist-rate samples of a continuous bandlimited signal. Specifically, let us define a real signal with bandlimit B as

$$x_c(t) = \sum_{k=-B}^B \hat{x}[k] e^{2\pi i k t}, \quad t \in [0, 1), \quad (1.4)$$

where $\iota = \sqrt{-1}$, and $\hat{x}$ denotes the Fourier series coefficients of x . Since x_c is real, it follows that $\hat{x}[k] = \overline{\hat{x}[-k]}$. According to the well-known Shannon-Nyquist sampling theorem, the samples

$$\tilde{x}_c[m] := x_c(m/M) = \sum_{k=-B}^B \hat{x}[k] e^{2\pi i k m/M}, \quad m = 0, \dots, M-1, \quad (1.5)$$

characterize x_c uniquely when $M \geq 2B + 1$. Model (1.1) is identical to rotating the discrete signal $\tilde{x}_c$ on the M -point grid $\{m/M\}_{m=0}^{M-1}$, sampling it L times, and adding noise.

With the above interpretation in mind, we identify the length of the signal in (1.1) with twice the signal’s bandwidth, namely, $M = 2B + 1 \approx 2B$. Thus, if $L < M$, we say that each observation is sampled *below the signal’s Nyquist rate*, and thus the recovery process should compensate for an aliasing distortion. To avoid aliasing, the standard signal processing approach in many applications is to remove the high frequency components before sampling [53, 26]—namely, low-passfiltering the signal—and then estimate a down-sampled, smooth approximation of x . While this strategy is generally optimal for a single observation, this is not necessarily true when multiple observations are available. In this work, we show that if sufficiently many observations are acquired (as a function of the noise level), then in principle it suffices to acquire only $L = O(\sqrt{M})$ samples at




Figure 1: Two shifted copies of a signal of length $M = 120$ (the high-resolution signal) are presented in blue. The red squares display $L = 15$ noisy samples with SNR equal to one. The goal is to estimate the high-resolution signal from multiple noisy observations.

each observation to recover high-resolution details, even if the circular shifts are unknown and the noise level might be high.

The analogy between (1.1) and rotating a continuous bandlimited signal (1.4) holds only when the rotations are restricted to the grid $\{m/M\}_{m=0}^{M-1}$. In Section 6 we discuss potential extensions to more intricate models that permit rotations over a continuous interval.

Super-resolution. The model (1.1) is an instance of the *super-resolution from multiple observations* problem: the task of estimating the fine details of a signal from its low-resolution observations. This problem has attracted the attention of numerous researchers in the last couple of decades in a variety of fields, such as computer vision, image processing, and medical imaging; see for instance [46, 28, 32] and references therein. In particular, the statistical model in some of these works is akin to (1.1), see for example [50, 49, 61]. Nevertheless, as far as we know, previous works on super-resolution did not aim to derive and quantify the achievable super-resolution in the low SNR regime. To avoid confusion, we mention that there exists a different thread of research, which is not directly related to this work, that studies super-resolution from a single image based on prior knowledge (such as sparsity [19, 15]), or machine learning techniques [36, 41].

Main contributions. In this paper we provide a detailed analysis of model (1.1) and derive fundamental conditions permitting an accurate estimate of x . In particular, we characterize the interplay between the number of observations N , the noise level σ , the signal's length M , and the number of samples per observation L , in the low SNR regime.

The following theorem summarizes (informally) the theoretical contribution of this paper. Precise formulations and technical details are provided in Section 3.

Theorem 1.1 (informal). *Suppose that N observations from (1.1) are collected and $\sigma \rightarrow \infty$. If $N/\sigma^6 \rightarrow \infty$ and $L \geq C\sqrt{M}$ for some constant C , then the maximum of the likelihood func-*




Figure 2: An example of an accurate estimate when the SNR is equal to 1. In the experiment, $N = 10^4$ observations were generated from a signal of length $M = 120$ (plotted in dashed blue; the same one as in Figure 1). The bandwidth of the signal is $B = 15$ and it was sampled $L = 15$ times at each observation—half of the Nyquist sampling rate. The classical signal processing approach suggests to remove all frequencies beyond $L/2$ and then process the low-resolution data. This low-passed version of the signal is presented in red. Notably, the two peaks in the center of the signal are blurred and merged into one. In contrast, the EM algorithm resolves the two peaks and estimates the high-resolution signal accurately (in green).

tion $p(y_1, \dots, y_N|x)$ is almost surely attained by a finite set of signals that includes the target signal x . If in addition x was drawn from a Gaussian prior, then almost surely there exists a single signal that achieves the maximum of the posterior distribution $p(x|y_1, \dots, y_N)$.

Expectation-maximization. As a computational scheme, we propose to retrieve the high-resolution signal x from the low-resolution observations $y_1, \dots, y_N$ using an expectation-maximization (EM) algorithm; a detailed description is given in Section 4. Figure 2 shows a numerical example. A high-resolution signal of length $M = 120$ is estimated from $N = 10^4$ observations in a noisy environment, where the SNR is equal to one and each observation is sampled at $L = 15$ points. The bandwidth of the signal is $B = L$, so that the sampling rate is half of the Nyquist rate. If we were to follow the Shannon-Nyquist sampling scheme of filtering out the $L/2$ high frequencies, the two peaks in the center of the signal would have been blurred into one, even with known circular shifts and in the absence of noise. In contrast, the EM algorithm resolves the two adjacent peaks and estimates the signal accurately. A detailed description of this simulation, and additional

numerical experiments, are provided in Section 5. We note, however, that while the theoretical analysis guarantees identifiability in the regime $M \approx L^2$, in our experiments the EM algorithm fails to estimate the high-resolution signal even when $L \approx M^{2/3}$. Following [6, 17, 60], we postulate that this inadequate performance reflects a fundamental statistical-computational gap in the super-resolution problem, rather than a shortcoming of the EM framework.

Remark on terminology and notation. We refer to each realization of the model (1.1) as an observation, and to the entries of each observation as samples. Namely, $y_i[\ell]$ denotes the ℓ -th sample of the i -th observation. In addition, in the sequel all indices should be considered as modulo M or L , depending on the context. When writing $P \gtrsim Q^d$ we mean that Q^d is the leading-order term of the right-hand side expansion in Q . For example, $M \gtrsim L^2/6$ implies that M is greater than $L^2/6$ plus a linear polynomial in L .

2 Background on multi-reference alignment and invariants

The model (1.1) is a special case of the *multi-reference alignment* (MRA) problem. This problem entails estimating a signal from multiple noisy observations; in each observation the signal is acted upon by an unknown element of a known group G . In its most general form, the MRA model reads

$$y = T(g \circ x) + \varepsilon, \quad g \in G, x \in \mathcal{X}, \quad (2.1)$$

where T is a known linear operator, with the group G acting on a vector space $\mathcal{X}$ [8]. Specifically, if $x \in \mathbb{R}^M$, G is identified with the group of circular shifts $\mathbb{Z}_M$, and T is the sampling operator P , then the general MRA model (2.1) reduces to (1.1).

Similarly to many MRA models in the literature [7, 14, 47, 1, 17, 42, 2, 4, 51], this work is inspired by single-particle reconstruction problems using cryo-electron microscopy (cryo-EM) and X-ray free electron lasers (XFEL)—high-resolution structural methods for biological macromolecules [30, 31, 44, 56, 10]. In particular, this work is a first step towards understanding the resolution limits of these modalities; see further discussion in Section 6.

Suppose we collect N observations from (2.1). If the noise level is low, the standard approach is to estimate the group element $g_1, \dots, g_N$. For example, in (1.1) the unknown circular shifts $s_1, \dots, s_N$ can be estimated by simultaneous clustering and synchronization (see Section 3.1). This can be done, for instance, using the Non-Unique Games framework [39]. However, in the low SNR regime—which is the main interest of this work—the group elements cannot be recovered reliably by any method [3, 13]. Therefore, we consider two techniques that circumvent shift determination: estimation based on shift-invariant features, and the EM algorithm. In particular, we formulate EM in detail in Section 4, and present numerical experiments in Section 5.

For the theoretical analysis, we use features that are invariant under circular shifts. Specifically, the q -th order circular-shift invariant feature of a signal $z \in \mathbb{R}^L$ is simply its auto-correlation:

$$M_q(z)[\ell_1, \dots, \ell_{q-1}] = \sum_{i=0}^{L-1} z[i]z[i + \ell_1] \dots z[i + \ell_{q-1}]. \quad (2.2)$$

It is readily seen that this quantity remains unchanged under any circular shift of z , namely, $M_q(z) = M_q(R_{\bar{s}}z)$ for any fixed $\bar{s}$. These invariants can also be presented in Fourier domain.

Specifically, let $\hat{z}[k]$ denote the k -th Fourier coefficient of z . Then the polynomials

$$\hat{M}_q(z)[k_1, \dots, k_{q-1}] = \hat{z}[k_1] \cdots \hat{z}[k_{q-1}] \hat{z}[-k_1 - \cdots - k_{q-1}], \quad (2.3)$$

are also invariant under circular shift. Throughout the paper, we use the terms auto-correlations, invariants, and invariant features interchangeably. Using these invariants, a variety of algorithms were proposed under different MRA setups [14, 47, 1, 17, 21, 42], as well as for cryo-EM and XFEL [34, 37, 40, 12, 45, 58, 55].

In this work, we harness the first three invariants. The first invariant is the zero frequency $\hat{M}_1(z) = \hat{z}[0]$ (equivalently, the mean of the signal). The second invariant is the power spectrum of the signal $\hat{M}_2(z)[k] = |z[k]|^2$ for $k = 0, \dots, L-1$. Unfortunately, the mean and the power spectrum do not determine a general signal uniquely (see for example [11]). Thus, we need the third-order invariant, the *bispectrum*, which determines almost all signals uniquely [57, 52]:

$$\hat{M}_3(z)[k_1, k_2] = \hat{z}[k_1] \hat{z}[k_2] \hat{z}[-k_1 - k_2], \quad k_1, k_2 = 0, \dots, L-1. \quad (2.4)$$

The bispectrum is a useful tool in many data processing applications, such as separating Gaussian and non-Gaussian processes [18], studying the cosmic background radiation, seismic, radar and EEG signals [59, 23, 43], MIMO systems [20], and classification [63].

For large σ , the variance of estimating the q -th order auto-correlation (either M_q or $\hat{M}_q$) is proportional to σ^{2q} since the estimator involves the product of q noisy terms. Thus, reliable estimation requires an order of σ^{2q} observations. For the problem under consideration, it implies that we need to record $N/\sigma^6 \gg 1$ observations to obtain an accurate estimate of the bispectrum. Interestingly, it was shown that for the MRA model (2.1), the invariant features approach is optimal in the following sense. Let $\bar{q}$ be the lowest-order auto-correlation that identifies a generic signal (in our case, $\bar{q} = 3$). Then, in the asymptotic regime where N and σ diverge (while L is fixed), the estimation error of any method is bounded away from zero if $N/\sigma^{2\bar{q}}$ is bounded from above [9, 2]. In other words, $\bar{q}$ determines the minimal number of observations required for an accurate estimate in the low SNR regime. Remarkably, we show that for (1.1) and for a certain range of L , at the same estimation rate (i.e., N scales with σ^6) one can reduce the sampling rate significantly below the Nyquist rate and still achieve an accurate estimate of the signal. In Section 6 we discuss the potential of super-resolution in case higher-order auto-correlations could be computed—that is, if more observations are available.

3 Analysis

The analysis is carried out in the asymptotic regime of $N \rightarrow \infty$, while the dimension M remains fixed. Therefore, we assume, without loss of generality, that $\|x\|^2 = M$ so that $\text{SNR} = 1/\sigma^2$. By the term “accurate recovery” we mean that the recovery error drops to zero almost surely. For example, the condition $N/\sigma^6 \rightarrow \infty$ ensures that we can almost surely estimate the bispectrum accurately.

In Section 3.1, we show that (1.1) can be interpreted as the *heterogeneous multi-reference alignment* (hMRA) model applied to K subsets of x , and formulate the likelihood function of (1.1). This, in turn, immediately implies that the signal is not determined uniquely from the likelihood function (a result implicit in earlier works, such as [61]):

Theorem 3.1. *The likelihood function $p(y_1, \dots, y_N|x)$ does not determine x uniquely, neither its orbit under $\mathbb{Z}_M$ (1.3).*

Nevertheless, the likelihood function allows us to identify a family of signals which can be described as the orbit of x under a parameterized sub-group of the permutation group; we denote this orbit by $G_{\Pi,L}x$ from reasons that will be explained later. Our analysis consists of two stages: identifying the orbit $G_{\Pi,L}x$ from the observations, and finding a unique signal in $G_{\Pi,L}x$ that maximizes the posterior distribution. In particular, in Section 3.2 we use auto-correlation analysis to show that for any $L \leq 192$ satisfying $L \gtrsim \sqrt{6M}$ (more accurately, any pair (L, M) satisfying (3.5) and (3.6)), the orbit $G_{\Pi,L}x$ can be computed from the first three auto-correlations of y ; we conjecture it remains true for any $L > 192$. These auto-correlations can be estimated from the data if $N/\sigma^6 \rightarrow \infty$ in any SNR regime. Finally, in Section 3.3 we show that if x was drawn from almost any Gaussian prior on the signal, then there is a unique signal in $G_{\Pi,L}x$ that maximizes the posterior distribution.

The following summarizes the main results of this section:

Theorem 3.2. *Suppose that $N \rightarrow \infty$ observations from (1.1) are collected, $N/\sigma^6 \rightarrow \infty$, and that x was drawn from almost any Gaussian prior. Then, for $L \leq 192$ and any K that satisfies (3.5), there exists a single signal that achieves the maximum of the posterior distribution $p(x|y_1, \dots, y_N)$.*

Conjecture 3.3. *Suppose that $N \rightarrow \infty$ observations from (1.1) are collected, $N/\sigma^6 \rightarrow \infty$, and that x was drawn from almost any Gaussian prior. Then, for any fixed M , there exists a single signal that achieves the maximum of the posterior distribution $p(x|y_1, \dots, y_N)$ as long as $M \leq L \cdot \mathcal{P}(L)$, where $\mathcal{P}(L)$ is given in (3.5).*

3.1 Reduction to heterogeneous MRA and the likelihood function

Consider two realizations y_i, y_j generated, respectively, after shifting x by s_i and s_j , and recall that $K = M/L$ is an integer. If $s_i - s_j = cK$ for some integer c , then y_i is equal to a circular shift of y_j , with a different noise realization. It follows that any observation y_i is a noisy and circularly-shifted realization of one of the following K signals,

$$x_k := [x[k], x[k+K], x[k+2K], \dots, x[k+(L-1)K]], \quad k = 0, \dots, K-1. \quad (3.1)$$


Namely, $x_k[\ell] = x[k + K\ell]$ for $\ell = 0, \dots, L-1$. We refer to $x_0, \dots, x_{K-1} \in \mathbb{R}^L$ as sub-signals. Using this notation, the model (1.1) can be written as

$$y = R_\ell x_k + \varepsilon, \quad (3.2)$$


where k is drawn uniformly at random from $\{0, \dots, K-1\}$, R_ℓ is a circular shift on an L -point grid $[0, 1, \dots, L-1]$, and ℓ is distributed uniformly. The model (1.1) is thus equivalent to the hMRA model, recently studied in [47, 6, 17, 42], applied to the sub-signals $x_0, \dots, x_{K-1}$.

Observations from the hMRA model (3.2) enable the recovery of $x_0, \dots, x_{K-1}$ up to a circular shift of each sub-signal and a permutation across signals. This can be seen by considering the marginalized likelihood of a single observation y :


$$p(y|x) = \frac{1}{M} \sum_{\ell=0}^{L-1} \sum_{k=0}^{K-1} p(y|x, \ell, k) = \frac{1}{(2\pi\sigma^2)^{L/2} M} \sum_{\ell=0}^{L-1} \sum_{k=0}^{K-1} e^{-\frac{1}{2\sigma^2} \|y - R_\ell x_k\|_2^2}. \quad (3.3)$$




(a) A signal consisting of 3 sub-signals



(b) A permutation of the sub-signals



(c) Circular shifts of the sub-signals



(d) Permutation + circular shifts of the sub-signals

Figure 3: An illustration of the orbit $G_{\Pi,L}x$; all four signals have the same likelihood function. (a) A signal of length $M = 12$ consists of $K = 3$ sub-signals (drawn in different colors). (b) Permuting the sub-signal $(x_0, x_1, x_2) \mapsto (x_1, x_2, x_0)$. (c) Shifting the sub-signals $(x_0, x_1, x_2) \mapsto (R_{-1}x_0, R_{-2}x_1, x_2)$. (d) Permuting and shifting the sub-signals $(x_0, x_1, x_2) \mapsto (R_{-2}x_1, x_2, R_{-1}x_0)$.

Plainly, $p(y|x)$ is invariant under any permutation π (overall $K!$ permutations)

$$(x_0, \dots, x_{K-1}) \mapsto (x_{\pi(0)}, \dots, x_{\pi(K-1)}),$$

or circular shifts $\ell_0, \dots, \ell_{K-1}$ (overall L^K permutations)

$$(x_0, \dots, x_{K-1}) \mapsto (R_{\ell_0}x_0, \dots, R_{\ell_{K-1}}x_{K-1}).$$

This set of permutations, denoted by $G_{\Pi,L}$, includes $K!L^K$ elements and constitutes a subgroup of the permutation group of M elements. The orbit of x under $G_{\Pi,L}$ is illustrated in Figure 3.

Importantly, previous works on hMRA aimed to retrieve the orbit $G_{\Pi,L}x$. In this work we further wish to recover the high-resolution signal x by ordering the sub-signals $x_0, \dots, x_{K-1}$ properly, which is impossible based solely on the likelihood. Thus, we must impose some additional constraints on the signal. In particular, we show in Section 3.3 that for almost any Gaussian prior there is a single element of $G_{\Pi,L}x$ that achieves the maximum of the posterior distribution.

3.2 Identifying the orbit $G_{\Pi,L}x$

3.2.1 The noiseless case

In the absence of noise, if we have observed each one of the K sub-signals $x_0, \dots, x_{K-1}$, we can determine the orbit $G_{\Pi,L}x$ immediately by considering all of their circular shifts and permutations. Therefore, the only question is how many observations from (1.1) are required to see each sub-signal x_k at least once; this problem is known in the combinatorics literature as the *coupon collector's problem* (see, for instance, [29]). In expectation, it is known that KH_K observations are required to see all K signals, where H_K is the harmonic sum

$$H_K = \frac{1}{1} + \frac{1}{2} + \dots + \frac{1}{K} = \log K + \gamma + \epsilon_K,$$

where $\gamma \approx 0.57721$ is the Euler-Mascheroni constant and $\epsilon_K \sim 1/(2K)$ for large K . If K is large enough, the harmonic sum can be bounded by $H_K \leq C \log K$ for some small constant C . For example, $H_K \leq 2 \log K$ for any $K \geq 3$. Thus, we say that in expectation $N \approx K \log K = M/L \log(M/L)$ observations suffice to characterize the orbit $G_{\Pi,L}x$ from noiseless observations. Yet, even in the absence of noise, Theorem 3.1 suggests that finding x itself from $G_{\Pi,L}x$ is a non-trivial task, requiring additional assumptions; we address this in Section 3.3.

3.2.2 Auto-correlation analysis

In the low SNR regime, we propose to estimate the orbit $G_{\Pi,L}x$ using the first three auto-correlations of the observations, or, equivalently, their Fourier counterparts: the mean, power spectrum, and bispectrum. Assuming $N/\sigma^6 \rightarrow \infty$ and considering (3.2), the invariants of the data

converge to the average of the invariants of the K sub-signals, up to bias terms:

$$\begin{aligned}
\frac{1}{N} \sum_{i=1}^N \hat{M}_1(y_i) &\xrightarrow{N \rightarrow \infty} \frac{1}{K} \sum_{k=0}^{K-1} \hat{M}_1(x_k), \\
\frac{1}{N} \sum_{i=1}^N \hat{M}_2(y_i)[\ell] &\xrightarrow{N \rightarrow \infty} \frac{1}{K} \sum_{k=0}^{K-1} \hat{M}_2(x_k)[\ell] + B_2, \\
\frac{1}{N} \sum_{i=1}^N \hat{M}_3(y_i)[\ell_1, \ell_2] &\xrightarrow{N \rightarrow \infty} \frac{1}{K} \sum_{k=0}^{K-1} \hat{M}_3(x_k)[\ell_1, \ell_2] + B_3[\ell_1, \ell_2],
\end{aligned} \tag{3.4}$$

where $B_2 = \sigma^2 L \mathbf{1}$, $B_3[\ell_1, \ell_2] = \bar{x} \sigma^2 L^2 D[\ell_1, \ell_2]$, $\mathbf{1} \in \mathbb{R}^L$ is a vector of ones, $\bar{x}$ is the average of x , $D \in \mathbb{R}^{L \times L}$ is a zero matrix except $D[0, 0] = 3$ and $D[i, 0] = D[0, i] = D[i, i] = 1$ for $i = 1, \dots, L-1$, and $\ell, \ell_1, \ell_2 = 0, \dots, L-1$. We note that if σ^2 is known, one can easily remove the bias factors from the second- and third-order invariants. As $N \rightarrow \infty$, the left-hand side equals the right-hand side almost surely.

The reason we require $N/\sigma^6 \rightarrow \infty$ is that the third-order auto-correlation requires taking triple products of three noise terms (thus tripling the effective noise level), and thus for large σ the number of observations needs to scale at least as σ^6 to keep the variance of the estimator under control; more precisely, only when $N/\sigma^6 \rightarrow \infty$ one can estimate the invariants accurately. Therefore, if $N/\sigma^6 \rightarrow \infty$, one can estimate the first three auto-correlations at any SNR levels. If σ is fixed while $N \rightarrow \infty$, then one can estimate all auto-correlations at any SNR level. If $\sigma \rightarrow \infty$ and N does not scale with σ^6 , namely, $N/\sigma^6 < C$ for some finite constant C , then the third-order auto-correlation cannot be consistently estimated from the observations.

3.2.3 Identifiability conditions for the orbit $G_{\Pi, L}x$ from the auto-correlations

As discussed in Section 2, it is well-known that the first three invariants determine a single generic signal uniquely [52, 33, 14]. Using tools from invariant theory and algebraic geometry, this result was recently extended to demixing of $K \geq 1$ invariants as in (3.4) [6]. The framework of [6] is based on checking computationally the rank of the Hessian matrix of the map between a generic¹ signal and the invariants. If the rank is sufficiently high (depending on the algebraic structure of the problem), we say that the Hessian test is passed, implying that the orbit $G_{\Pi, L}x$ can be identified for generic x . The Hessian test is executed on pairs of parameters (K, L) ; if (K, L) passes the test, it implies identifiability for all pairs (K', L) , $K' \leq K$. In particular, it was verified² for all $L \leq 192$ that a set of generic signals is determined uniquely from (3.4), up to the symmetries that form the group $G_{\Pi, L}$, as long as

$$K < \mathcal{P}(L) := \frac{L + 3 + \lfloor L/2 \rfloor + \lceil (L-1)(L-2)/6 \rceil}{L+1} \approx L/6. \tag{3.5}$$

This immediately implies that the number of required samples is bounded by

$$K = \frac{M}{L} \lesssim L/6 \quad \Rightarrow \quad L \gtrsim \sqrt{6M}. \tag{3.6}$$

¹By the notion of generic signals, we mean that all signals that are not recoverable in the regime defined by equation (3.5) satisfy a certain polynomial equation, and thus are of measure zero.

²We extended the range of parameters examined by the authors of [6]. We thank Dr. Joseph Kileel for his assistance to execute this computational verification.

This bound is conjectured to hold true for any pair (L, K) that satisfies (3.5); see [6, 17]. We note that the bound of (3.5) is tight in the sense that it agrees with a simple upper bound based on parameter counting: on the one hand, the bispectrum of a generic signal contains $\mathcal{P}(L) \cdot L \approx L^2/6$ distinct entries (out of L^2 entries in total), and on the other hand K signals consists of KL parameters.

The following proposition and conjecture summarize the result:

Proposition 3.4. *Suppose that we acquire an average of the mean, power spectrum, and bispectrum of K signals as in the right-hand side (3.4). Then, for $L \leq 192$ and any K that satisfies (3.5), one can identify the orbit $G_{\Pi,L}x$ for generic x .*

Conjecture 3.5. *Suppose that we acquire an average of the mean, power spectrum, and bispectrum of K signals as in the right-hand side of (3.4). Then, for any pair (K, L) that satisfies (3.5), one can identify the orbit $G_{\Pi,L}x$ for generic x .*

3.2.4 A note on computational considerations

It is important to note that Proposition 3.4 does not claim that the bound (3.5) can be achieved using a computationally efficient (e.g., polynomial time) algorithm. In the context of the hMRA model, numerical evidence suggests that for i.i.d. standard Gaussian signals, one can estimate the orbit $G_{\Pi,L}x$ from a mix of bispectra using non-convex least squares only in the regime $K \leq \sqrt{L}$ —substantially below the identifiability regime $K \lesssim L/6$ [17]. Recently, it was proven that i.i.d. standard Gaussian signals can be disentangled, with high probability, using a sum-of-squares algorithm as long as $K \leq \sqrt{L}/\text{polylog}(L)$ [60]. In [6, 17], it was conjectured that the $\sqrt{L}$ bound reflects a fundamental statistical-computational gap—namely, while it is statistically possible to recover approximately $L/6$ signals, any efficient (polynomial-time) algorithm can estimate at most $\sqrt{L}$ signals.

If indeed one can recover only up to $\sqrt{L}$ signals efficiently from (3.4), it implies that the orbit $G_{\Pi,L}x$ can be estimated efficiently in the regime $K = \frac{M}{L} \leq \sqrt{L} \Rightarrow M \leq L^{3/2}$ for i.i.d. Gaussian entries. Having said that, in contrast to the hMRA model, the goal of the super-resolution problem is not to recover the orbit $G_{\Pi,L}x$, but rather the signal x itself; the latter task seems to be a significantly more challenging computational problem. In addition, the main interest of this work is in smooth signals (e.g., signals with decaying power spectrum). For such signals, the achievable performance of hMRA deteriorates [9], and even recovering $\sqrt{L}$ signals seem to be unreachable. Indeed, numerical experiments in Section 5 suggest that recovery is not attainable even in the regime $M \approx L^{3/2}$ —at least not with the EM algorithm.

3.3 Identifying a unique high-resolution signal from the orbit $G_{\Pi,L}x$

Until now, we have shown that one can identify the orbit $G_{\Pi,L}x$ for generic x if $L \gtrsim \sqrt{6M}$, $L < 192$, and the first three auto-correlations can be estimated from the observations. Next, we wish to show how to determine a single signal out of $G_{\Pi,L}x$.

Recall that the posterior distribution $p(x|y_1, \dots, y_N)$ is proportional to the likelihood function $p(y_1, \dots, y_N|x)$ times a prior on the signal $p(x)$. According to (3.3), the likelihood is constant over the orbit $G_{\Pi,L}x$. Consequently, we choose the signal in the orbit that best fits the prior. Importantly, this part is independent of the observations.

Many priors can be used. In this section, we focus on Gaussian signals with zero mean and covariance Σ , that is, $p(x) = \frac{1}{\sqrt{(2\pi)^M |\Sigma|}} e^{-\frac{1}{2}x^T \Sigma^{-1}x}$. Priors of this form are ubiquitous in signal processing, and have been considered in previous works, such as [50, 49, 61]. In particular, we wish to show that among all signals in $G_{\Pi,L}x$, there is a single signal that maximizes $p(x)$, or, equivalently, minimizes $x^T \Sigma^{-1}x$ for a positive-definite matrix, Σ^{-1} . The next lemma shows that permuting a signal usually changes this quadratic form. To this end, we define the Σ^{-1} norm by $\|z\|_{\Sigma^{-1}}^2 := z^T \Sigma^{-1}z$ for a positive-definite matrix Σ^{-1} .

Lemma 3.6. *Let $\Omega \subset \mathbb{R}^n$ and denote by R_1 and R_2 two permutation matrices and their ratio by $R = R_2 R_1^T$. Then the set*

$$\mathcal{Z}_{R_1, R_2} = \left\{ z \in \Omega \mid \|R_1 z\|_{\Sigma^{-1}}^2 = \|R_2 z\|_{\Sigma^{-1}}^2 \right\},$$

is a subset of Ω of measure zero if and only if Σ^{-1} and R do not commute.

Proof. Let $A := R_1^T \Sigma^{-1} R_1 - R_2^T \Sigma^{-1} R_2$. Observe that $A = 0$ if and only if R and Σ^{-1} commute. The condition $\|R_1 z\|_{\Sigma^{-1}}^2 = \|R_2 z\|_{\Sigma^{-1}}^2$ is equivalent to $z^T A z = 0$. Thus, if $A = 0$ then $\mathcal{Z}_{R_1, R_2} = \Omega$. Otherwise, $A \neq 0$ and symmetric, i.e., there exists a basis Q of orthogonal eigenvectors such that $(Qy)^T A Qy = \sum_{i=1}^n \lambda_i y_i^2$, with $\lambda_1 \neq 0$. Therefore, the condition $(Qy)^T A Qy = 0$ means that $y_1^2 = \sum_{i=2}^n \frac{\lambda_i}{-\lambda_1} y_i^2$. Denote the indicator functions χ_Z and χ_Y for the two sets $\{z \in \Omega \mid z^T A z = 0\}$ and $\{y \in \mathbb{R}^n \mid y^T (Q^T A Q)y = 0\}$, respectively. Since orthogonal transformations preserve integrals and $\Omega \subset \mathbb{R}^n$, we have

$$|\mathcal{Z}_{R_1, R_2}| = \int_{\Omega} \chi_Z(z) dz \leq \int_{\mathbb{R}^n} \chi_Y(y) dy = \int_{\mathbb{R}^{n-1}} \left(\int_{\mathbb{R}} \chi_Y(y_1, y_2, \dots, y_n) dy_1 \right) dy_2 \cdots dy_n.$$

Recall that $(Qy)^T A Qy = 0$ implies that $y_1 = \pm \left(\sum_{i=2}^n \frac{\lambda_i}{-\lambda_1} y_i^2 \right)^{\frac{1}{2}}$. Thus, for any fixed $y_2, \dots, y_n$, the indicator function χ_Y is nonzero on only two points, implying $\int_{\mathbb{R}} \chi_Y(y_1, y_2, \dots, y_n) dy_1 = 0$. Consequently, $|\mathcal{Z}_{R_1, R_2}| = 0$. $\square$

Any group of permutations over a finite set is finite, and thus there are finitely many pairs R_1 and R_2 of permutation matrices. Consequently, the set of signals for which $\|R_1 z\|_{\Sigma^{-1}}^2 = \|R_2 z\|_{\Sigma^{-1}}^2$ for some pairs of permutations is also of measure zero.

Corollary 3.7. *Assume the conditions of Lemma 3.6 are met, that is, Σ^{-1} and $R_2 R_1^T$ do not commute for all pairs $R_1 \neq R_2$ in the permutation group. Then for almost every x in Ω the minimum of the quadratic form $y^T \Sigma^{-1} y$ is unique among all signals y in $G_{\Pi,L}x$.*

The case when Σ^{-1} is a circulant matrix is of particular importance. Such a prior, reflecting a prior on the signal's power spectrum, is popular in many signal processing tasks, as well as in cryo-EM; see for instance [54]. In this case, both the prior and the likelihood function (and thus the posterior) are shift-invariant, and therefore any signal is indistinguishable from its cyclic shifts. To account for this symmetry, we derive a distinct uniqueness result—up to a circular shift—for circulant matrices.

Proposition 3.8. *Assume Σ^{-1} is a circulant, positive-definite matrix of size $n > 2$. Then:*

1. The Σ^{-1} norm $\|\cdot\|_{\Sigma^{-1}}$ is invariant under cyclic shifts. Consequently, we may consider the quadratic form $y^T \Sigma^{-1} y$ as a function over equivalence classes in $G_{\Pi, L} x$, where two vectors are equivalent if one is a cyclic shift of the other.
2. If all eigenvalues of Σ^{-1} are distinct, then for almost every signal x the minimum of the quadratic form $y^T \Sigma^{-1} y$ is unique over the equivalence classes of $G_{\Pi, L} x$.

The proposition is proved in Appendix A.

4 An expectation-maximization algorithm

Our theoretical study is based on invariant features. Conceptually, it suggests a two-stage procedure: it begins by identifying the orbit $G_{\Pi, L} x$, and then choosing a unique signal according to the prior. While identifying the orbit $G_{\Pi, L} x$ efficiently using bispectrum demixing is possible [17, 60], it is unclear how to devise a tractable algorithm for the second step. As an alternative, we formulate an EM algorithm, described below, which aims to achieve the maximum of the posterior distribution by maximizing the likelihood function and the prior simultaneously [25]. EM is known to work quite well in many practical scenarios; see for instance its application to cryo-EM experimental datasets [54, 48]. In what follows, we formulate EM for MRA with a general linear operator

$$y = TR_s x + \varepsilon, \quad (4.1)$$

where s is drawn from a uniform distribution on a discrete grid S_M with M points, and T is a general linear operator (not necessarily a sampling matrix). For the special case of super-resolution, a similar algorithm was already derived by [61].

EM is a common framework to compute the maximum a posteriori estimator (MAP). Hereafter, we formulate EM for the general model. Given a set of N independent observations $y_1, \dots, y_N$, the log-posterior distribution $\log p(x|y_1, \dots, y_N)$ is proportional to $\log p(y_1, \dots, y_N|x)p(x)$, where

$$\log p(y_1, \dots, y_N|x) = \sum_{i=1}^N \log \sum_{s_\ell \in S_M} e^{-\frac{1}{2\sigma^2} \|y_i - TR_{s_\ell} x\|_2^2} + \text{constant}, \quad (4.2)$$

is the log-likelihood function, and $p(x)$ is a prior on the signal. We assume that the signal is drawn from a Gaussian prior with zero mean and known covariance Σ so that $p(x) \sim \mathcal{N}(0, \Sigma)$. EM aims to maximize the posterior iteratively, where each iteration consists of two steps. The first step, called E-step, computes the expected value of the likelihood x (note, not the marginalized likelihood) with respect to the circular shifts (i.e., the nuisance variables), given the current estimate of the signal x_t and the data $y_1, \dots, y_N$:

$$\begin{aligned} Q(x|x_t) &= \mathbb{E}_{s_1, \dots, s_N | y_1, \dots, y_N, x_t} \{ \log p(y_1, \dots, y_N, s_1, \dots, s_N|x) + \log p(x) \} \\ &= -\frac{1}{2\sigma^2} \mathbb{E}_{s_1, \dots, s_N | y_1, \dots, y_N, x_t} \left\{ \sum_{i=1}^N \|y_i - TR_{s_i} x\| \right\} - \frac{1}{2} x^T \Sigma^{-1} x + \text{constant} \\ &= -\frac{1}{2\sigma^2} \sum_{i=1}^N \sum_{s_\ell \in S_M} w_{i,\ell} \|y_i - TR_{s_\ell} x\| - \frac{1}{2} x^T \Sigma^{-1} x + \text{constant}, \end{aligned} \quad (4.3)$$

where

$$w_{i,\ell} = \frac{e^{\frac{-1}{2\sigma^2}\|y_i - TR_{s_\ell}x_t\|^2}}{\sum_{s_\ell \in S_M} e^{\frac{-1}{2\sigma^2}\|y_i - TR_{s_\ell}x_t\|^2}}. \quad (4.4)$$

The second step, called M-step, maximizes Q with respect to x . The solution is obtained by solving the linear system of equations

$$Ax = b, \quad (4.5)$$

where

$$A := \Sigma^{-1} + \frac{1}{\sigma^2} \sum_{i,\ell} \omega_{i,\ell} (R_\ell^{-1} T^T T R_\ell), \quad (4.6)$$

$$b := \frac{1}{\sigma^2} \sum_{i,\ell} \omega_{i,\ell} R_\ell^{-1} T^T y_i. \quad (4.7)$$

The EM algorithm iterates between computing the weights (4.4) and solving the linear system (4.5) until convergences. In our implementation, the algorithm halts when the relative difference between the posterior of two consecutive iterations falls below 10^{-5} . In general, EM does not converge to the global maximum of the posterior distribution; however, each iteration is guaranteed not to decrease the posterior [25]. In addition, several recent works derived intriguing theoretical results for EM under specific statistical models. See for example [5, 24], and in particular [27] that analyzes EM for the MRA model.

In the special case in which the linear operator T is the sampling operator (1.1), computing the weights and constructing A and b reduces to computing a set of correlations; this can be executed efficiently using FFT. For EM implementation when a blurring kernel is included, see [61].

On the connection between maximum likelihood estimation and the invariant approach. A recent paper by Katsevich and Bandeira [35] studies Gaussian mixture models, for which heterogeneous MRA is a special case, in the parametric setup considered in this work: $N \rightarrow \infty$, $\text{SNR} \rightarrow 0$, and fixed M . In particular, they show that log-likelihood maximization is equivalent to an asymptotic series of successively higher moment matching problems. In this sense, a method based on the bispectrum (a third-order moment), as we use for the analysis, can be thought of as a third-order approximation of the likelihood function.

5 Numerical results

We conducted three experiments to examine the performance of the EM algorithm. The first experiment demonstrates resolving two adjacent peaks from low-resolution observations. The next two experiments study the performance of the algorithm as a function of noise level and the number of samples. The code for all experiments and the EM algorithm is publicly available at <https://github.com/TamirBendory/MRA-SR>.

In all experiments, we set a prior on the signal's power spectrum, and therefore, to account for the circular shift symmetry, the relative recovery error is defined as

$$\text{relative error} = \min_{\ell \in \mathbb{Z}^M} \frac{\|R_\ell x_{\text{est}} - x\|}{\|x\|}, \quad (5.1)$$

where $x \in \mathbb{R}^M$ is the underlying signal, and x_{est} is the output of the EM algorithm. The SNR is defined as $\text{SNR} = \|x\|^2 / (M\sigma^2)$. The EM iterations terminate when either a maximal number of iterations is reached, or the relative absolute difference of the posterior between two consecutive iterations drop below a tolerance parameter. In all experiments, the maximal number of EM iterations was set to be 100 iterations, and the tolerance parameters was 10^{-5} . The EM may be initialized from multiple random points and thus produce different estimators. Among those estimators, we choose the one with the largest posterior. The posterior is computed at each EM iteration.

Experiment 1. The signal in this experiment is of length $M = 120$ with bandwidth (the largest non-zero frequency, see (1.5)) $B = 15$. We generated $N = 10^4$ observations by shifting the signal and sampling it at $L = 15$ points, corresponding to half of the Nyquist sampling rate. Then, an i.i.d. Gaussian noise was added, corresponding to $\text{SNR} = 1$. We ran the EM algorithm from five random initial points; each trial required 13 to 17 iterations to converge. Since the signal (only in this experiment) is bandlimited, in each iteration the current estimate is projected onto the low $B = 15$ frequencies. The target and estimated signals are presented in Figure 2; the relative recovery error is 0.0614. Figure 4 displays the relative error per frequency. As can be seen, the relative error of frequencies above the Nyquist rate is still quite low, indicating that the EM algorithm resolves high frequencies accurately.




Figure 4: Recovery error per frequency of the experiment presented in Figure 2. The figure indicates that the EM algorithm succeeds to resolve frequencies beyond the largest frequency determined by the Nyquist sampling rate $L/2$ (vertical red line).

Experiment 2. Figure 5 presents the error curve as a function of the SNR, in the high and low SNR regimes. For each SNR value, 50 trials were conducted and we present the median error. A signal of length $M = 64$ was drawn from a Gaussian distribution with zero mean, and a circulant




Figure 5: Relative estimation error as a function of the SNR. In the high SNR regime, the relative error scales as $\text{SNR}^{-1/2}$, which is the same estimation rate as if there were no shifts (namely, the estimation rate of averaging independent Gaussian variables). In the low SNR regime, the error decays faster than SNR^{-1} , demonstrating a sharp transition from the high SNR regime.

covariance matrix, corresponding to power spectrum decaying linearly, that is, as $1/f$. Following the circular shift, each observation was sampled at $L = 32$ equally-spaced points—corresponding to half of the Nyquist sampling rate.

Figure 5b presents the relative error as a function of the SNR, for 30 SNR values sampled uniformly on a logarithmic scale between $10^{0.2}$ to 10^2 , and $N = 10^2$ observations; this reflects the high SNR regime. Unfortunately, the EM seems to suffer from a flaw: it should be initialized from many points (among them we choose the one with the largest posterior value) in order to result in a consistent recovery. In this experiment, we initialized the algorithm from 1000 points; the computational load is still quite cheap since each trial requires only a few iterations (around 5). Yet, it suggests that EM may not be the optimal computational scheme in the high SNR regime. The slope of the error curve is approximately $-1/2$. Since the SNR is proportional to $1/\sigma^2$, this indicates that the error scales as σ —the optimal estimation rate even if the circular shifts were known.

Figure 5a shows a similar experiment for SNR values ranging between $10^{-0.6}$ to 1 and $N = 10^5$ observations. In this low SNR regime, the EM algorithm seems to be more consistent, and thus we initialized it from merely 20 random points. In this regime, the slope of the error curve becomes steeper and the error slope is smaller than -1 , implying that the error scales faster than σ^4 . This indicates, in line with previous works on MRA (with finite M), that the estimation rate in the high and low SNR regimes is drastically different [9, 47, 1, 2]. In fact, our analysis predicts that as $\text{SNR} \rightarrow 0$, the slope of the error curve would tend to $\text{SNR}^{-3/2}$; see Section 3.

Experiment 3. Figure 6 examines the recovery error for different values of M and L . For each M , we chose values of L so that M/L is an integer. The signals were generated as in Experiment 2 with $\text{SNR} = 5$ and $N = 1000$, and the EM was initialized from 50 random locations in each trial. For each pair (M, L) , the mean error over 50 trials was recorded. The red vertical dashed




Figure 6: Relative recovery error as a function of L for different values of M . The red dashed line indicates $L = M^{2/3}$. The results suggest that the super-resolution problem is significantly harder than hMRA.

line indicates $L = M^{2/3}$; this is the conjectured computational recovery limit for hMRA, namely, for recovering the orbit $G_{\Pi,L}x$ (see discussion in Section 3.2.4). Notwithstanding, we get relatively small recovery error only for much larger values of L , suggesting that the super-resolution problem is computationally more challenging than hMRA. In particular, our theoretical analysis is split into two stages: recovering the orbit $G_{\Pi,L}x$, and recovering x from the orbit; the latter depends only on the prior, and not on the data. In contrary, the EM algorithm aims to implicitly carry out both stages simultaneously. We believe that the second stage, together with the smoothness of the signals (see Section 3.2.4), is the reason the performance of EM for super-resolution is inferior to what was demonstrated in previous MRA setups [14, 1, 42].

6 Discussion

Super-resolution limits. This work analyzes the super-resolution from multiple observations problem in a noisy environment using the third-order auto-correlation. To use higher-order auto-correlations, more observations should be collected: the number of observations needs to scale as σ^{2q} to estimate the q -th order auto-correlation accurately. The q -th auto-correlation provides $O(L^{q-1})$ polynomial equations of the sought signals. Based on our analysis and the reduction of the super-resolution problem to the hMRA model (3.2), we expect that the $q \geq 3$ auto-correlation would identify $M = O(L^{q-1})$ grid points. Such a result will follow directly from a generalization of [6] to higher-order auto-correlations. This leads us to the following conjecture:

Conjecture 6.1. *Suppose that N observations from (1.1) are collected and each observation is sampled at L equally-spaced locations. Then, in the low SNR regime $\sigma \rightarrow \infty$, if $N/\sigma^{2q} \rightarrow \infty$ for some $q \geq 3$, one can identify up to $M = O(L^{q-1})$ grid points. In other words, only $L = O(M^{1/(q-1)})$ samples per observation suffice for signal identification. In particular, for $N \rightarrow \infty$ and any fixed noise level (that might be arbitrarily high), there is no theoretical limit on the achievable resolution.*

Continuous super-resolution. A natural generalization of the model considered in this work is the following. Let $x : S^1 \rightarrow \mathbb{R}$ be a band-limited signal on the circle (1.4), and let R_θ denote

a rotation, that is, $(R_\theta x)(t) = x(t - \theta)$, where θ is distributed uniformly on the circle. Together with i.i.d. Gaussian noise $\varepsilon \in \mathbb{R}^L$, the data generative model reads

$$y = P(R_\theta x) + \varepsilon, \quad \theta \sim \text{Uniform}[0, 1), \quad \varepsilon \sim \mathcal{N}(0, \sigma^2 I), \quad (6.1)$$

where P denotes a sampling operator that collects L equally-spaced point-wise samples. The goal is to estimate x from N observations sampled from (6.1). This setup is interesting in the sub-Nyquist regime, where P samples x below its Nyquist sampling rate. While this model shares many similarities with (1.1), it poses some additional challenges that are beyond the scope of this work; we intend to address them in a follow-up work.

Super-resolution of images. While this paper deals with 1-D signals, the methodology can be extended to higher-dimensions. For example, an interesting MRA setup that was studied in [42] considers rotating 2-D “bandlimited” images. Specifically, suppose that an image X belongs to the vector space of images that can be expanded by finitely many coefficients in a steerable basis (such as Fourier-Bessel [62] or prolate spheroidal wave functions [38]):

$$X(r, \phi) = \sum_{k=1}^{K_{\max}} \sum_{q=1}^{Q_{\max}} a_{k,q} u_{k,q}(r, \phi), \quad (6.2)$$

where (r, ϕ) are polar coordinates, $a_{k,q}$ are the expansion coefficients, and $u_{k,q}(r, \phi)$ are the basis functions of the steerable basis. The images are acted upon by unknown elements of the group of in-plane rotations $SO(2)$. The steerability property implies that rotating an image by an angle α amounts to multiplying the expansion coefficients by $e^{ik\alpha}$:

$$X(r, \phi - \alpha) = \sum_{k=1}^{K_{\max}} \sum_{q=1}^{Q_{\max}} a_{k,q} e^{-ik\alpha} u_{k,q}(r, \phi). \quad (6.3)$$

Accordingly, it is easy to see that the triple products $a_{k_1,q_1} a_{k_2,q_2} a_{-k_1-k_2,q_3}$ are invariant under rotations—these products form the bispectrum [62, 42]. This in turn implies that for an image expanded by M coefficients, there are $O(M^{5/2})$ bispectrum entries. In this case, our framework suggests that, perhaps, one can identify an image from sufficiently many observations with merely $L = O(M^{-2/5})$ samples per observation.

Super-resolution in high dimensions. This work studies the finite-dimensional regime (finite M) in which invariant features achieve the optimal estimation rate as $\text{SNR} \rightarrow 0$. A recent work [51] uncovered that this is not the case in the high-dimensional regime $M \rightarrow \infty$. In particular, it was shown that the parameter that controls the “hardness” of the model is $\alpha = M/(\log M \sigma^2)$: when $\alpha < 2$ the samples complexity of the problem rapidly increases, whereas for $\alpha > 2$ the effect of the unknown shifts is minor. Nevertheless, the authors of [51] only investigated random sub-sampling operators that do not include the super-resolution setup (1.1). The high-dimensional regime is of particular interest since it seems as a good model for modern cryo-EM datasets, where the dimensionality and the number of samples are of the same order of a few millions. In fact, high-dimensional statistical analysis has been already proven effective for cryo-EM data processing. For example, a covariance estimation technique based on high-dimensional analysis (the so-called spiked model) has significantly improved image denoising [16].

Cryo-EM and XFEL. The main motivation of this work arises from cryo-EM and XFEL. The measurements in these applications (under simplifying assumptions, see for example [10]) agree with the general MRA model (1.1), where $g \in SO(3)$ (the group of 3-D rotations), the 3-D Fourier transform of the signal x is assumed to be bounded in a ball (“bandlimited” volume), and the linear operator T collects samples of the 2-D tomographic projection of the rotated volume. The question then would be whether the maximal resolution of a 3-D reconstruction algorithm can surpass the resolution dictated by the detectors acquiring the data—that is, the resolution of the 2-D tomographic projection images. A recent proof of concept (on simulated data) promises an affirmative answer [22]. Extending our analysis to this case requires sophisticated tools and we leave it for a future research.

Acknowledgment

The authors are grateful to Joseph Kileel and Dan Edidin for insightful discussions about algebraic geometry. The authors also thank the anonymous reviewers for their valuable comments and suggestions. A.S. and W.L. were partially supported by NSF BIGDATA award IIS-1837992. T.B., N.S., and A.S. were partially supported by BSF grant no. 2019752, and NSF grant no. 2009753. W.L. and N.S. were partially supported by BSF grant no. 2018230. A.S. was also partially supported by NIH/NIGMS award 1R01GM136780-01, award FA9550-17-1-0291 from AFOSR, award FA9550-20-1-0266 from AFOSR, the Simons Foundation Math+X Investigator Award, and the Moore Foundation Data-Driven Discovery Investigator Award. T.B. was also partially supported by the Zimin Institute for Engineering Solutions Advancing Better Lives.

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. *arXiv preprint arXiv:1905.12442*, 2019.
- [5] Sivaraman Balakrishnan, Martin J Wainwright, and Bin Yu. Statistical guarantees for the EM algorithm: From population to sample-based analysis. *The Annals of Statistics*, 45(1):77–120, 2017.

- [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, Moses Charikar, Amit Singer, and Andy Zhu. Multireference alignment using semidefinite programming. In *Proceedings of the 5th conference on Innovations in theoretical computer science*, pages 459–470. ACM, 2014.
- [8] 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.
- [9] Afonso S Bandeira, Jonathan Niles-Weed, and Philippe Rigollet. Optimal rates of estimation for multi-reference alignment. *Mathematical Statistics and Learning*, 2(1):25–75, 2020.
- [10] 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.
- [11] 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.
- [12] 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.
- [13] 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.
- [14] 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.
- [15] Tamir Bendory, Shai Dekel, and Arie Feuer. Robust recovery of stream of pulses using convex optimization. *Journal of mathematical analysis and applications*, 442(2):511–536, 2016.
- [16] Tejal Bhamre, Teng Zhang, and Amit Singer. Denoising and covariance estimation of single particle cryo-EM images. *Journal of structural biology*, 195(1):72–81, 2016.
- [17] Nicolas Boumal, Tamir Bendory, Roy R Lederman, and Amit Singer. Heterogeneous multireference alignment: A single pass approach. In *Information Sciences and Systems (CISS), 2018 52nd Annual Conference on*, pages 1–6. IEEE, 2018.
- [18] 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.
- [19] Emmanuel J Candès and Carlos Fernandez-Granda. Towards a mathematical theory of super-resolution. *Communications on pure and applied Mathematics*, 67(6):906–956, 2014.

- [20] Binning Chen and Athina P Petropulu. Frequency domain blind MIMO system identification based on second-and higher order statistics. *IEEE Transactions on Signal Processing*, 49(8):1677–1688, 2001.
- [21] Hua Chen, Mona Zehni, and Zhizhen Zhao. A spectral method for stable bispectrum inversion with application to multireference alignment. *IEEE Signal Processing Letters*, 25(7):911–915, 2018.
- [22] James Z Chen. Single-particle 3D reconstruction beyond the Nyquist frequency. In *2018 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*, pages 2442–2445. IEEE, 2018.
- [23] 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.
- [24] Constantinos Daskalakis, Christos Tzamos, and Manolis Zampetakis. Ten steps of EM suffice for mixtures of two gaussians. In *Conference on Learning Theory*, pages 704–710, 2017.
- [25] Arthur P Dempster, Nan M Laird, and Donald B Rubin. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the royal statistical . Series B (methodological)*, pages 1–38, 1977.
- [26] Yonina C Eldar. *Sampling theory: Beyond bandlimited systems*. Cambridge University Press, 2015.
- [27] Zhou Fan, Yi Sun, Tianhao Wang, and Yihong Wu. Likelihood landscape and maximum likelihood estimation for the discrete orbit recovery model. *arXiv preprint arXiv:2004.00041*, 2020.
- [28] Sina Farsiu, Dirk Robinson, Michael Elad, and Peyman Milanfar. Advances and challenges in super-resolution. *International Journal of Imaging Systems and Technology*, 14(2):47–57, 2004.
- [29] William Feller. *An Introduction to Probability Theory and Its Applications*, volume 1. John Wiley & Sons, Inc., 1968.
- [30] Joachim Frank. *Three-dimensional electron microscopy of macromolecular assemblies: visualization of biological molecules in their native state*. Oxford University Press, 2006.
- [31] KJ Gaffney and HN Chapman. Imaging atomic structure and dynamics with ultrafast X-ray scattering. *Science*, 316(5830):1444–1448, 2007.
- [32] Hayit Greenspan. Super-resolution in medical imaging. *The Computer Journal*, 52(1):43–63, 2008.
- [33] Ramakrishna Kakarala. The bispectrum as a source of phase-sensitive invariants for Fourier descriptors: a group-theoretic approach. *Journal of Mathematical Imaging and Vision*, 44(3):341–353, 2012.

- [34] Zvi Kam. The reconstruction of structure from electron micrographs of randomly oriented particles. *Journal of Theoretical Biology*, 82(1):15–39, 1980.
- [35] Anya Katsevich and Afonso Bandeira. Likelihood maximization and moment matching in low SNR gaussian mixture models. *arXiv preprint arXiv:2006.15202*, 2020.
- [36] Jiwon Kim, Jung Kwon Lee, and Kyoung Mu Lee. Accurate image super-resolution using very deep convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1646–1654, 2016.
- [37] 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.
- [38] Boris Landa and Yoel Shkolnisky. Approximation scheme for essentially bandlimited and space-concentrated functions on a disk. *Applied and Computational Harmonic Analysis*, 43(3):381–403, 2017.
- [39] Roy R Lederman and Amit Singer. A representation theory perspective on simultaneous alignment and classification. *Applied and Computational Harmonic Analysis*, 49(3):1001–1024, 2020.
- [40] Eitan Levin, Tamir Bendory, Nicolas Boumal, Joe Kileel, and Amit Singer. 3D ab initio modeling in cryo-EM by autocorrelation analysis. In *Biomedical Imaging (ISBI 2018), 2018 IEEE 15th International Symposium on*, pages 1569–1573. IEEE, 2018.
- [41] Bee Lim, Sanghyun Son, Heewon Kim, Seungjun Nah, and Kyoung Mu Lee. Enhanced deep residual networks for single image super-resolution. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 136–144, 2017.
- [42] 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.
- [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] Kanupriya Pande, Jeffrey J Donatelli, Erik Malmerberg, Lutz Foucar, Christoph Bostedt, Ilme Schlichting, and Petrus H Zwart. Ab initio structure determination from experimental fluctuation X-ray scattering data. *Proceedings of the National Academy of Sciences*, 115(46):11772–11777, 2018.
- [46] Sung Cheol Park, Min Kyu Park, and Moon Gi Kang. Super-resolution image reconstruction: a technical overview. *IEEE signal processing magazine*, 20(3):21–36, 2003.

- [47] 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.
- [48] Ali Punjani, John L Rubinstein, David J Fleet, and Marcus A Brubaker. cryoSPARC: algorithms for rapid unsupervised cryo-EM structure determination. *Nature methods*, 14(3):290, 2017.
- [49] Dirk Robinson, Sina Farsiu, and Peyman Milanfar. Optimal registration of aliased images using variable projection with applications to super-resolution. *The Computer Journal*, 52(1):31–42, 2009.
- [50] Dirk Robinson and Peyman Milanfar. Statistical performance analysis of super-resolution. *IEEE Transactions on Image Processing*, 15(6):1413–1428, 2006.
- [51] Elad Romanov, Tamir Bendory, and Or Ordentlich. Multi-reference alignment in high dimensions: sample complexity and phase transition. *arXiv preprint arXiv:2007.11482*, 2020.
- [52] Brian M Sadler and Georgios B Giannakis. Shift-and rotation-invariant object reconstruction using the bispectrum. *JOSA A*, 9(1):57–69, 1992.
- [53] Ronald William Schafer and Alan V Oppenheim. *Discrete-time signal processing*. Prentice Hall Englewood Cliffs, NJ, 1989.
- [54] Sjors HW Scheres. RELION: implementation of a Bayesian approach to cryo-EM structure determination. *Journal of structural biology*, 180(3):519–530, 2012.
- [55] 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.
- [56] Amit Singer. Mathematics for cryo-electron microscopy. *Proceedings of the International Congress of Mathematicians*, 2018.
- [57] 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.
- [58] Benjamin von Ardenne, Martin Mechelke, and Helmut Grubmüller. Structure determination from single molecule X-ray scattering with three photons per image. *Nature communications*, 9(1):2375, 2018.
- [59] Limin Wang and Marc Kamionkowski. Cosmic microwave background bispectrum and inflation. *Physical Review D*, 61(6):063504, 2000.
- [60] Alexander Wein. *Statistical Estimation in the Presence of Group Actions*. PhD thesis, Massachusetts Institute of Technology, 2018.

- [61] Nathan A Woods, Nikolas P Galatsanos, and Aggelos K Katsaggelos. Stochastic methods for joint registration, restoration, and interpolation of multiple undersampled images. *IEEE Transactions on Image Processing*, 15(1):201–213, 2005.
- [62] Zhizhen Zhao and Amit Singer. Fourier–Bessel rotational invariant eigenimages. *JOSA A*, 30(5):871–877, 2013.
- [63] 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.

A Proof of Proposition 3.8

The first lemma refines the condition of Lemma 3.6 in terms of invariant subspaces. Recall that a subspace V is R -invariant if $R(V) \subset V$.

Lemma A.1. *Let S be a symmetric matrix. Then, R and S commute if and only if every eigenspace is R -invariant.*

Proof. Since S is symmetric, we can decompose the space into real eigenspaces. First direction: if R and S commute then for any eigenvector v of S with eigenvalue λ , we have

$$S(Rv) = RSv = \lambda(Rv).$$

Namely, Rv is also in the eigenspace. Second direction: if any eigenspace V_λ is R -invariant then for any $v \in V_\lambda$ we can write Rv in terms of the basis of V_λ , and to have

$$Rv = \sum_i \alpha_i v_i \implies S(Rv) = S\left(\sum_i \alpha_i v_i\right) = \sum_i \alpha_i \lambda v_i = \lambda Rv = R(Sv). \quad \square$$

Lemma A.2. *Let Σ^{-1} be a circulant matrix. Then*

1. Σ^{-1} commutes with any cyclic permutation matrix R .
2. Conversely, if each eigenvalue of Σ^{-1} has multiplicity 1 and Σ^{-1} commutes with a permutation matrix R , then R must be a cyclic permutation.

Proof. Let W denote the DFT matrix, with entries $W_{j,\ell} = \frac{1}{\sqrt{n}} \omega^{(\ell-1)(j-1)}$, $\omega = e^{-2\pi i/n}$, $1 \leq j, \ell \leq n$. Then the columns $w_0, \dots, w_{n-1}$ of W are the eigenvectors of Σ^{-1} . Furthermore, if R is the cyclic permutation matrix corresponding to the permutation $j \mapsto j + k \pmod n$, then $Rw_\ell = \omega^{(\ell-1)k} w_\ell$. Consequently, R preserves each eigenspace of Σ^{-1} , and so by Lemma A.1 R and Σ^{-1} commute. This proves the first statement.

For the converse, suppose each eigenvalue of Σ^{-1} has multiplicity 1, and take any permutation matrix R that commutes with Σ^{-1} . From Lemma A.1, R must leave each eigenspace of Σ^{-1} fixed; consequently, for each eigenvector w_ℓ of Σ^{-1} there is a scalar α_ℓ so that $Rw_\ell = \alpha_\ell w_\ell$. Suppose the permutation corresponding to R sends index 1 to index $k + 1$; then

$$\frac{1}{\sqrt{n}} = W_{1,\ell} = (Rw_\ell)_{k+1} = \alpha_\ell W_{k+1,\ell} = \alpha_\ell \frac{1}{\sqrt{n}} \omega^{(\ell-1)k}, \quad (\text{A.1})$$

and so $\alpha_\ell = \omega^{-(\ell-1)k}$. Since $Rw_\ell = \alpha_\ell w_\ell$, for any index j ,

$$(Rw_\ell)_j = \omega^{-(\ell-1)k} W_{j,\ell} = \omega^{-(\ell-1)k} \frac{1}{\sqrt{n}} \omega^{(\ell-1)(j-1)} = \frac{1}{\sqrt{n}} \omega^{(\ell-1)(j-k-1)} = W_{j-k,\ell}, \quad (\text{A.2})$$

meaning that R cyclically shifts the entries of w_ℓ by k . Since this holds for all basis vectors w_ℓ , R is a cyclic shift. $\square$

We may now prove Proposition 3.8. The first statement of Proposition 3.8 is identical to the first statement of Lemma A.2. For the second statement, Lemma 3.6 tells us that for almost every signal x , the quadratic form $y^T \Sigma^{-1} y$ takes on distinct values on each equivalence class of vectors in the orbit $G_{\Pi,L}x$ (where two vectors are equivalent if one is a cyclic shift of the other). Indeed, for any two permutation matrices R_1 and R_2 , the set $\{x \mid x^T R_1^T \Sigma^{-1} R_1 x = x^T R_2^T \Sigma^{-1} R_2 x\}$ has measure zero if and only if $R_2 R_1^T$ does not commute with Σ^{-1} ; from Lemma A.2, this latter condition is equivalent to $R_2 R_1^T$ not being a cyclic permutation. Since there are only finitely many permutation matrices, the set

$$\left\{x \mid \exists R_1, R_2 \text{ s.t. } R_2 R_1^T \text{ not cyclic and } x^T R_1^T \Sigma^{-1} R_1 x = x^T R_2^T \Sigma^{-1} R_2 x\right\} \quad (\text{A.3})$$

also has measure 0. Consequently, for almost every signal x , the equality $x^T R_1^T \Sigma^{-1} R_1 x = x^T R_2^T \Sigma^{-1} R_2 x$ can hold only when $R_2 R_1^T$ is cyclic, i.e. when $R_1 x$ and $R_2 x$ are in the same equivalence class of $G_{\Pi,L}x$.

Because the quadratic form $y^T \Sigma^{-1} y$ takes on distinct values on each equivalence class in $G_{\Pi,L}x$ for almost every x , it immediately follows that for almost every x the minimum of $y^T \Sigma^{-1} y$ over equivalence classes in $G_{\Pi,L}x$ is unique. This is the desired result.