

Estimation in the group action channel

Emmanuel Abbe^{1,2}, João M. Pereira¹, and Amit Singer^{1,3}

¹The Program in Applied and Computational Mathematics, Princeton University, Princeton, NJ, USA

²Electrical Engineering Department, Princeton University, Princeton, NJ, USA

³Department of Mathematics, Princeton University, Princeton, NJ, USA

Abstract—We analyze the problem of estimating a signal from multiple measurements on a *group action channel* that linearly transforms a signal by a random group action followed by a fixed projection and additive Gaussian noise. This channel is motivated by applications such as multi-reference alignment and cryo-electron microscopy. We focus on the large noise regime prevalent in these applications. We give a lower bound on the mean square error (MSE) of any asymptotically unbiased estimator of the signal's orbit in terms of the signal's moment tensors, which implies that the MSE is bounded away from 0 when N/σ^{2d} is bounded from above, where N is the number of observations, σ is the noise standard deviation, and d is the so-called *moment order cutoff*. In contrast, the maximum likelihood estimator is shown to be consistent if N/σ^{2d} diverges.

Index Terms—Multi-reference alignment, cryo-EM, Chapman-Robbins bound.

I. INTRODUCTION

In this paper, we consider the problem of estimating $x \in \mathbb{R}^L$ with N measurements from the *group action channel*, defined as

$$Y_j = PG_jx + \sigma Z_j \in \mathbb{R}^K, \quad j = 1, \dots, N, \quad (\text{I.1})$$

where the Z_j are i.i.d and drawn from $\mathcal{N}(0, I_K)$, i.e. $Z_j \in \mathbb{R}^K$ and its entries are i.i.d standard Gaussian variables; $P \in \mathbb{R}^{K \times L}$ is a projection matrix which is known; G_j are i.i.d. matrices drawn from a distribution θ on a compact subgroup Θ of $O(L)$, i.e. the space of orthogonal matrices in \mathbb{R}^L . The distribution θ is not known, however the main goal is to estimate the signal $x \in \mathbb{R}^L$.

The goal of this paper is to understand the sample complexity of (I.1), i.e. the relation between the number of measurements and the noise standard deviation such that an estimator \hat{X} , of x , converges in probability to the true value with N diverging, up to a group action. Allowing for a group action is intrinsic to the problem: if we apply an element g of Θ to x , and its inverse g^{-1} to the right of θ , we will produce exactly the same samples, thus there is no estimator \hat{X} that is able to distinguish the observations that originate from x and the ones from gx .

The model (I.1) is a generalization of multi-reference alignment (MRA), which arises in a variety of engineering and scientific applications, among them structural biology [1], [2], [3], radar [4], [5], robotics [6] and image processing [7], [8], [9]. The one-dimensional MRA problem, where Θ is the group generated by the matrix R that cyclically shifts the elements of the signal, i.e. it maps $(x_1, \dots, x_L) \mapsto (x_L, x_1, \dots, x_{L-1})$, has been recently a topic of active research. In [10], it was shown that the sample complexity is $\omega(\sigma^6)$ when θ is the

uniform distribution and the projection matrix is the identity. In [11] it was presented a provable polynomial time estimator that achieves the sample complexity, while in [12] it was presented a non-convex optimization framework that is more efficient in practice. Note that, when the projection matrix is the identity, we can always enforce a uniform distribution on Θ by applying a random group action, i.i.d. and drawn from the uniform distribution, to the observations. In [13], it was shown that $\omega(\sigma^6)$ is also the sample complexity if θ is unknown beforehand but is uniform or periodic, this is, $\theta = R^\ell \theta$ for some $1 \leq \ell \leq L-1$. However, if θ is aperiodic, the sample complexity is $\omega(\sigma^4)$. It is also presented an efficient estimator that uses the first and second moments of the signal over the group, which can be estimated with order of σ^2 and σ^4 observations, respectively, thus achieving the sample complexity. The main result in this paper is a generalization of the information lower bound presented in [13], however the proof techniques remain the same.

We can also use (I.1) to model the problem of single particle reconstruction in cryo-electron microscopy (cryo-EM), in which a three-dimensional volume is recovered from two-dimensional noisy projections taken at unknown viewing directions [14], [15]. Here x is a linear combination of products of spherical harmonics and radial basis functions, $\Theta \equiv SO(3)$, and its elements act on x by rotating the basis functions. Finally, P is a tomographic projection onto the xy plane. The paper [16] considers the problem (I.1) with θ being known and uniform. It obtains the same result for the sample complexity as this paper, and together with results from computational algebra and invariant theory verifies that in many cases the sample complexity for the considered cryo-EM model is $\omega(\sigma^6)$, and at least $\omega(\sigma^6)$ more generally. They also consider the problem of heterogeneity in cryo-EM.

II. THE MAIN RESULT

Since we can only determine x up to a group action, we define the best alignment of \hat{X} with x by

$$\phi_x(\hat{X}) = \underset{z \in \{g\hat{X}\}_{g \in \Theta}}{\operatorname{argmin}} \|z - x\|. \quad (\text{II.1})$$

and the mean square error (MSE) as

$$\begin{aligned} \text{MSE} &:= \mathbb{E} \left[\min_{g \in \Theta} \|g\hat{X} - x\|^2 \right], \\ &= \mathbb{E} \left[\|\phi_x(\hat{X}) - x\|^2 \right]. \end{aligned} \quad (\text{II.2})$$

The expectation is taken over \widehat{X} , which is a function of the observations with distribution determined by (I.1). Since we are interested in estimators that converge to an orbit of x in probability as N diverges, we only consider estimators which are asymptotically unbiased, i.e., $\mathbb{E}[\phi_x(\widehat{X})] \rightarrow x$ as $N \rightarrow \infty$. However the results presented in this paper can be adapted to biased estimators (see Theorem III.2).

Let us introduce some notation regarding tensors. For a vector $x \in \mathbb{R}^L$, we denote by $x^{\otimes n}$ the $L^{\otimes n}$ dimensional tensor where the entry indexed by $\mathbf{k} = (k_1, \dots, k_n) \in \mathbb{Z}_L^n$ is given by $\prod_{j=1}^n x[k_j]$. The space of n -dimensional tensors forms a vector space, with sum and multiplication defined entry-wise. This vector-space has inner product and norm defined by $\langle A, B \rangle = \sum_{\mathbf{k} \in \mathbb{Z}_L^n} A[\mathbf{k}]B[\mathbf{k}]$ and $\|A\|^2 = \langle A, A \rangle$, respectively.

Definition II.1. The n -th order moment of x over θ , is the tensor of order n and dimension K^n , defined by

$$M_{x,\theta}^n := \mathbb{E}[(PGx)^{\otimes n}], \quad (\text{II.3})$$

where $G \sim \theta$.

In this paper, we provide lower bounds for the MSE in terms of the noise standard deviation and the number of observations. We show that the MSE is bounded below by order of $N/\sigma^{2\bar{d}}$, where \bar{d} is the *moment order cutoff*, defined as the smallest such that the moment tensors up to order \bar{d} define x unequivocally. We also show that if $N \gg \sigma^{2\bar{d}}$, then the marginalized maximum likelihood estimator (MLE) converges in probability to the true signal (up to a group action). We now present the main result of the paper.

Theorem II.2. Consider the estimation problem given by equation (I.1). For any signal $\tilde{x} \in \mathbb{R}^L$ such $\phi_x(\tilde{x}) \neq x$ and for any group distribution θ , let $K_{\tilde{x},\theta}^n = \frac{1}{n!} \|M_{\tilde{x},\theta}^n - M_{x,\theta}^n\|^2$, $d_{\tilde{x},\theta} = \inf \left\{ n : K_{\tilde{x},\theta}^n > 0 \right\}$ and define the **moment order cutoff** as $\bar{d} = \max d_{\tilde{x},\theta}$. Finally let

$$\lambda_N^m = N/\sigma^{2m}.$$

We have

$$\text{MSE} \geq \sup_{\tilde{x}, \theta: d_{\tilde{x},\theta} = \bar{d}} \left\{ \frac{\|\phi_x(\tilde{x}) - x\|^2}{\exp\left(\lambda_N^{\bar{d}} K_{\tilde{x},\theta}^{\bar{d}}\right) - 1 + \mathcal{O}\left(\lambda_N^{\bar{d}} \sigma^{-1}\right)} \right\}, \quad (\text{II.4})$$

thus the MSE is bounded away from zero if $\lambda_N^{\bar{d}}$ is bounded from above. Moreover, if $\lim_{N \rightarrow \infty} \lambda_N^{\bar{d}} = \infty$, then the MLE converges in probability to gx , for some element $g \in \Theta$.

A. Taking the limit $(\tilde{x}, \tilde{\rho}) \rightarrow (x, \rho)$

Theorem II.2 is an application of a modified Chapman-Robbins bound, presented later in Theorem III.2. On the other hand the classical Cramér-Rao bound [17], which gives a lower bound on the variance of an estimator \hat{S} of a parameter $s \in \mathbb{R}$, can be obtained from the Chapman-Robbins bound by taking the limit $\tilde{s} \rightarrow s$. We present an analog version of Theorem II.2 obtained by taking a similar limit.

Corollary II.3. Under the conditions of Theorem II.2, let $x_h := (1-h)x + h\tilde{x}$, $\theta_h := (1-h)\theta + h\tilde{\theta}$,

$$Q_{\tilde{x},\tilde{\theta}}^n = \lim_{h \rightarrow 0} \frac{1}{n!h^2} \|M_{x_h,\theta_h}^n - M_{x,\theta}^n\|^2, \\ q_{\tilde{x},\tilde{\theta}} = \inf \left\{ n : Q_{\tilde{x},\tilde{\theta}}^n > 0 \right\} \text{ and } \bar{q} = \max q_{\tilde{x},\tilde{\theta}}. \text{ Then}$$

$$\text{MSE} \geq \sup_{\tilde{x}, \tilde{\theta}: q_{\tilde{x},\tilde{\theta}} = \bar{q}} \left\{ \frac{\|\phi_x(\tilde{x}) - x\|^2}{\lambda_N^{\bar{q}} Q_{\tilde{x},\tilde{\theta}}^{\bar{q}} + \mathcal{O}\left(\lambda_N^{\bar{q}} \sigma^{-1}\right)} \right\}. \quad (\text{II.5})$$

We leave the proof of this corollary to [13, Appendix C]. It is interesting to compare this bound with (II.4) when $\lambda_N^{\bar{d}}$ or $\lambda_N^{\bar{q}}$ diverge. If $\bar{q} \geq \bar{d}$, then (II.5) will dominate (II.4), and the lower bound for the MSE will be inversely proportional to $\lambda_N^{\bar{q}}$, which is a behavior typical of estimation problems with continuous model parameters. On the other hand, if $\bar{d} > \bar{q}$, then (II.4) dominates (II.5). the MSE will depend exponentially on $\lambda_N^{\bar{d}}$, which is a behaviour typical of discrete parameter estimation problems [18]. One can show that $\bar{d} > \bar{q}$ only happens when the supremum in (II.4) is attained by some x^* not in the orbit of x . The exponential decay in (II.4) is the same as the probability of error of the hypothesis testing which decides if the observations come from x or x^* .

We conjecture that the lower bounds presented in this paper can be achieved asymptotically by the MLE. In fact when the search space is discrete, the MLE achieves the least probability of error (assuming a uniform prior on the parameters), which behaves like (II.4). Also, when the search space is continuous, the MLE is asymptotically efficient, which means it achieves the Cramér-Rao lower bound. However this bound is obtained from the Chapman-Robbins lower bound (which we use in this paper) by taking a similar limit as in (II.5), and the bound also scales inversely proportional to the number of observations.

B. Prior Knowledge

The result presented can be adapted to improve the bound if we have prior knowledge about the signal and group distribution. If we know beforehand that $(x, \theta) \in \mathcal{A}$ (for instance, x has a zero element or θ is the uniform distribution on Θ), we can instead define $\bar{d} = \max_{(\tilde{x}, \tilde{\theta}) \in \mathcal{A}} d_{\tilde{x},\tilde{\theta}}$ and restrict the supremum in (II.4) to $(\tilde{x}, \tilde{\theta})$ in \mathcal{A} .

C. Examples

1) Let $x = (a, b, c) \in \mathbb{R}^3$; Θ be the group generated by the cyclic shift matrix R that maps $(a, b, c) \mapsto (b, c, a)$; and P projects x into its first two elements, i.e $P(a, b, c) = (a, b)$. Furthermore, we know a-priori that one, and only one, of the elements of x is 0 (let's assume without loss of generality that $a = 0$), the other two elements are distinct and θ is uniform, i.e. $\mathbb{P}(G = I) = \mathbb{P}(G = R) = \mathbb{P}(G = R^2) = \frac{1}{3}$. We have

$$M_{x,\theta}^1 = \mathbb{E}[PGx], \\ = \frac{1}{3}(0, b) + \frac{1}{3}(b, c) + \frac{1}{3}(c, 0), \\ = \frac{b+c}{3}(1, 1).$$

and

$$\begin{aligned} M_{x,\theta}^2 &= \mathbb{E} \left[(PGx)(PGx)^T \right], \\ &= \frac{1}{3} \begin{bmatrix} 0 & 0 \\ 0 & b^2 \end{bmatrix} + \frac{1}{3} \begin{bmatrix} b^2 & bc \\ bc & c^2 \end{bmatrix} + \frac{1}{3} \begin{bmatrix} c^2 & 0 \\ 0 & 0 \end{bmatrix}, \\ &= \frac{1}{3} \begin{bmatrix} b^2 + c^2 & bc \\ bc & b^2 + c^2 \end{bmatrix}. \end{aligned}$$

From these two moments, we can solve for b and c , however all these equations are symmetric on b and c , thus we can't identify which one of the values obtained is b and which one is c . In other words, both candidate solutions are $x = (0, b, c)$ and $x^* = (0, c, b)$. However $M_{x,\theta}^3$ differs from $M_{x^*,\theta}^3$, if we look for the entry in $M_{x,\theta}^3$, indexed by $(1, 1, 2)$ we note that

$$\begin{aligned} M_{x,\theta}^2[1, 1, 2] &= \mathbb{E} \left[(PGx)_1^2 (PGx)_2 \right], \\ &= \frac{1}{3} 0^2 b + \frac{1}{3} b^2 c + \frac{1}{3} c^2 0 \\ &= \frac{1}{3} b^2 c, \end{aligned}$$

and analogously $M_{x^*,\theta}^3[1, 1, 2] = c^2 b$. From the 8 entries of $M_{x,\theta}^3$, 2 are equal to $M_{x^*,\theta}^3$ and 6 differ by $b^2 c - c^2 b$, in absolute value, so $\|M_{x^*,\theta}^3 - M_{x,\theta}^3\|^2 = 6(b^2 c - c^2 b)^2$. This means $\bar{d} = 3$, $\bar{q} = 2$, thus if λ_N^3 the lower bound (II.4) dominates (II.5), the supremum is attained at x^* and

$$\text{MSE} \geq \frac{\|\phi_x(x^*) - x\|^2}{\exp(\lambda_N^3(b^2 c - c^2 b)^2) - 1 + \mathcal{O}(\lambda_N^3 \sigma^{-1})}.$$

Note that $\|\phi_x(x^*) - x\|^2 = \min(b^2, c^2, 2(b-c)^2)$.

2) Let $x = (a, b) \in \mathbb{R}^2$; Θ be the group generated by the cyclic shift matrix R that maps $(a, b) \mapsto (b, a)$; and P projects x into its first element, i.e $P(a, b) = a$. Furthermore, we know a-priori that θ is uniform, i.e. $\mathbb{P}(G = I) = \mathbb{P}(G = R) = \frac{1}{2}$. We have

$$M_{x,\theta}^1 = \mathbb{E}[PGx] = \frac{1}{2}a + \frac{1}{2}b = \frac{a+b}{2}$$

and

$$M_{x,\theta}^2 = \mathbb{E}[(PGx)^2] = \frac{1}{2}a^2 + \frac{1}{2}b^2 = \frac{a^2 + b^2}{2}.$$

From these two moments we can determine a and b up to an action of the group. Now take $x_h = (a+h, b-h)$, so that $M_{x,\theta}^1 = M_{x_h,\theta}^1$. We have

$$\lim_{h \rightarrow 0} \frac{1}{h} (M_{x_h,\theta}^2 - M_{x,\theta}^2) = a - b.$$

Here $\bar{q} = \bar{d} = 2$, thus if λ_N^2 diverges, (II.5) dominates (II.4), and the lower bound is

$$\text{MSE} \geq \frac{4}{\lambda_N^2(a-b)^2 + \mathcal{O}(\lambda_N^2 \sigma^{-1})}.$$

III. PROOF TECHNIQUES

The outline of the proof is as follows. In Section III-A we use an adaptation of the Chapman-Robbins lower bound [19], to derive a lower bound on the MSE in terms of the χ^2 divergence, this is Theorem III.2. Then, in Section III-B, we express the χ^2 divergence in terms of the Taylor expansion of the posterior probability density and the moment tensors, obtaining Lemma III.3. Finally in section III-C we combine Theorem III.2 and Lemma III.3 to obtain (II.4), use Lemma III.3 to obtain a similar Taylor expansion for the Kullback-Leibler (KL) divergence and use this to show that the MLE is consistent.

Throughout the paper we denote the expectation by \mathbb{E} , use capital letter for random variables and lower case letter for instances of these random variables. Let $Y^N \in \mathbb{R}^{L \times N}$ be the collection of all measurements as columns in a matrix. Let us denote by $f_{x,\theta}^N$ the probability density of the posterior distribution of Y^N ,

$$f_{x,\theta}^N(y^N) = \prod_{j=1}^N f_{x,\theta}(y_j), \quad (\text{III.1})$$

and the expectation of a function g of the measurements under the measure $f_{x,\theta}^N$ by

$$\mathbb{E}_{x,\theta} \left[g(Y^N) \right] := \int_{\mathbb{R}^{L \times N}} g(y^N) f_{x,\theta}^N(y^N) dy^N.$$

For ease of notation, we write $\mathbb{E} \left[g(Y^N) \right]$ when the signal and distribution are implicit. The bias-variance trade-off of the MSE is given by:

$$\text{MSE} = \text{tr}(\text{Cov}[\phi_x(\hat{X})]) + \|\mathbb{E}[\phi_x(\hat{X})] - x\|^2, \quad (\text{III.2})$$

with

$$\text{Cov}[\phi_x(\hat{X})] = \mathbb{E} \left[\phi_x(\hat{X}) \phi_x(\hat{X})^T \right] - \mathbb{E}[\phi_x(\hat{X})] \mathbb{E}[\phi_x(\hat{X})]^T. \quad (\text{III.3})$$

Our last definition is of the χ^2 divergence, which gives a measure of how "far" two probability distributions are.

Definition III.1. The χ^2 divergence between two probability densities f_A and f_B is defined by

$$\chi^2(f_A || f_B) := \mathbb{E} \left[\left(\frac{f_A(B)}{f_B(B)} - 1 \right)^2 \right],$$

where $B \sim f_B$.

Due to equation (III.1), the relation between the χ^2 divergence for N and one observations is given by

$$\chi^2(f_{\bar{x},\bar{\theta}}^N || f_{x,\theta}^N) = (1 + \chi^2(f_{\bar{x},\bar{\theta}} || f_{x,\theta}))^N - 1. \quad (\text{III.4})$$

A. Chapman-Robbins lower bound for an orbit

The classical Chapman-Robbins gives a lower bound on an error metric of the form $\mathbb{E}[\|\hat{X} - x\|^2]$, hence we modified it to accommodate to the group invariant metric defined in (II.2). We point out that $\text{Cov}[\phi_x(\hat{X})]$ is related to the MSE by (III.2).

Theorem III.2 (Chapman-Robbins for orbits). *For any $\tilde{x} \in \mathbb{R}^L$ and group distribution $\tilde{\theta}$ in Θ , we have*

$$\text{Cov}[\phi_x(\hat{X})] \succeq \frac{zz^T}{\chi^2(f_{\tilde{x},\tilde{\theta}}^N \| f_{x,\theta}^N)},$$

where $z = \mathbb{E}_{\tilde{x},\tilde{\theta}}[\phi_x(\hat{X})] - \mathbb{E}_{x,\theta}[\phi_x(\hat{X})]$.

Proof. The proof mimics the one of the classical Chapman and Robbins bound, and is also presented in [13, Appendix A]. Define

$$V := \frac{f_{\tilde{x},\tilde{\theta}}^N(Y^N)}{f_{x,\theta}^N(Y^N)}.$$

and note that

- $\mathbb{E}_{x,\theta}[g(Y^N)V] = \mathbb{E}_{\tilde{x},\tilde{\theta}}[g(Y^N)]$,
- $\mathbb{E}_{x,\theta}[V - 1] = 0$,
- $\mathbb{E}_{x,\theta}[(V - 1)^2] = \chi^2(f_{\tilde{x},\tilde{\theta}}^N \| f_{x,\theta}^N)$.

We have

$$\begin{aligned} w^T \left(\mathbb{E}_{\tilde{x},\tilde{\theta}}[\phi_x(\hat{X})] - \mathbb{E}_{x,\theta}[\phi_x(\hat{X})] \right) \\ = \mathbb{E}_{x,\theta} \left[w^T \left(\phi_x(\hat{X}) - \mathbb{E}_{x,\theta}[\phi_x(\hat{X})] \right) (V - 1) \right], \end{aligned}$$

and by Cauchy-Schwarz

$$\begin{aligned} \left[w^T \left(\mathbb{E}_{\tilde{x},\tilde{\theta}}[\phi_x(\hat{X})] - \mathbb{E}_{x,\theta}[\phi_x(\hat{X})] \right) \right]^2 \\ \leq \mathbb{E}_{x,\theta} \left[\left(w^T \left(\phi_x(\hat{X}) - \mathbb{E}_{x,\theta}[\phi_x(\hat{X})] \right) \right)^2 \chi^2(f_{\tilde{x},\tilde{\theta}}^N \| f_{x,\theta}^N) \right]. \end{aligned}$$

□

B. χ^2 divergence and moment tensors

In this subsection we give a characterization of the χ^2 divergence, which appears in the Chapman-Robbins bound, in terms of the moment tensors.

Instead of considering the posterior probability density of Y^N , we will consider its normalized version $\tilde{Y}^N = Y^N/\sigma$. We then have

$$\tilde{Y}_j = \gamma PG_j x + Z_j, \quad (\text{III.5})$$

where $\gamma = 1/\sigma$, $G_j \sim \theta$ and $Z_j \sim \mathcal{N}(0, I)$. While this change of variable does not change the χ^2 divergence, we can now take the Taylor expansion of the probability density around $\gamma = 0$, that is,

$$f_{x,\theta}(y; \gamma) = f_Z(y) \sum_{j=0}^{\infty} \alpha_{x,\theta}^j(y) \frac{\gamma^j}{j!}, \quad (\text{III.6})$$

where $f_Z(y) = f_{x,\theta}(y; 0)$ is the probability density of Z_j (since when $\gamma = 0$, $\tilde{Y}_j = Z_j$) and

$$\alpha_{x,\theta}^j(y) := \frac{1}{f_Z(y)} \frac{\partial^j f_{x,\theta}}{\partial \gamma^j}(y; 0), \quad (\text{III.7})$$

thus $\alpha_{x,\theta}^0(y) = 1$. We note $f_{x,\theta}(y; \gamma)$ is in infinitely differentiable for all $y \in \mathbb{R}^L$, thus $\alpha_{x,\theta}^j(y)$ is always well-defined. We now use (III.6) to give an expression of the χ^2 divergence in terms of the moment tensors.

Lemma III.3. *The divergence $\chi^2(f_{\tilde{x},\tilde{\theta}} \| f_{x,\theta})$ is expressed in terms of the moment tensors as:*

$$\begin{aligned} \chi^2(f_{\tilde{x},\tilde{\theta}} \| f_{x,\theta}) \\ = \frac{\sigma^{-2d}}{(d!)^2} \mathbb{E} \left[\left(\alpha_{\tilde{x},\tilde{\theta}}^d(Z) - \alpha_{x,\theta}^d(Z) \right)^2 \right] + \mathcal{O}(\sigma^{-2d-1}), \end{aligned} \quad (\text{III.8})$$

$$= \frac{\sigma^{-2d}}{d!} \|M_{\tilde{x},\tilde{\theta}}^d - M_{x,\theta}^d\|^2 + \mathcal{O}(\sigma^{-2d-1}), \quad (\text{III.9})$$

where $d = \inf \{ n : \|M_{\tilde{x},\tilde{\theta}}^n - M_{x,\theta}^n\|^2 > 0 \}$.

Proof. This proof is presented in more detail in [13, Appendix B]. Equation (III.8) is obtained by Taylor expanding the χ^2 divergence around $\gamma = 0$, using (III.6) and the fact that $\alpha_{\tilde{x},\tilde{\theta}}^n(z) = \alpha_{x,\theta}^n(z)$ almost surely for all $n < d$, which follows from the definition of d and equation (III.9). Now to prove (III.9), it is enough to show that

$$\mathbb{E} \left[\alpha_{\tilde{x},\tilde{\theta}}^d(Z) \alpha_{x,\theta}^d(Z) \right] = d! \left\langle M_{\tilde{x},\tilde{\theta}}^d, M_{x,\theta}^d \right\rangle, \quad (\text{III.10})$$

Let G and \tilde{G} be two independent random variables such that $G \sim \theta$ and $\tilde{G} \sim \tilde{\theta}$. On one hand we have

$$\left\langle M_{\tilde{x},\tilde{\theta}}^d, M_{x,\theta}^d \right\rangle = \mathbb{E} \left[\left\langle P\tilde{G}\tilde{x}, PGx \right\rangle^d \right]. \quad (\text{III.11})$$

On the other hand, we can write $f_{x,\theta}$ explicitly by

$$f_{x,\theta}(y) = \mathbb{E}_G[f_Z(y - \gamma PGx)], \quad (\text{III.12})$$

where $G \sim \theta$, and using equation (III.7) we can write

$$\begin{aligned} \mathbb{E} \left[\alpha_{\tilde{x},\tilde{\theta}}^d(Z) \alpha_{x,\theta}^d(Z) \right] \\ = \frac{\partial^{2d}}{\partial \tilde{\gamma}^d \partial \gamma^d} \mathbb{E} \left[\frac{f_Z(Z - \tilde{\gamma} P\tilde{G}\tilde{x})}{f_Z(Z)} \frac{f_Z(Z - \gamma PGx)}{f_Z(Z)} \right]_{\tilde{\gamma},\gamma=0} \\ = \mathbb{E} \left[\frac{\partial^{2d}}{\partial \tilde{\gamma}^d \partial \gamma^d} \exp \left(\gamma \tilde{\gamma} \left\langle P\tilde{G}\tilde{x}, PGx \right\rangle \right) \right]_{\tilde{\gamma},\gamma=0} \\ = d! \mathbb{E} \left[\left\langle P\tilde{G}\tilde{x}, PGx \right\rangle^d \right], \end{aligned}$$

where G and \tilde{G} are defined as in (III.11), and (III.10) finally follows from equation (III.11). □

C. Final details of the proof of Theorem II.2

By Theorem III.2, Lemma III.3, equations (III.3) and (III.4) we obtain

$$\text{MSE} \geq \frac{\|\phi_x(\tilde{x}) - x\|^2}{\left(1 + \sigma^{-2d} K_d + \mathcal{O}(\sigma^{-2d-1})\right)^N - 1}. \quad (\text{III.13})$$

Equation (II.4) now follows from

$$\begin{aligned} \left(1 + \sigma^{-2d} K_d + \mathcal{O}(\sigma^{-2d-1})\right)^N = \\ \exp \left(\lambda_N^d K_{\tilde{x},\tilde{\theta}}^d \right) + \mathcal{O} \left(\lambda_N^d \sigma^{-1} \right) \end{aligned}$$

and taking the supremum over \tilde{x} and $\tilde{\theta}$.

Finally we prove that the MLE is consistent, i.e. it converges to the true signal in probability, when $\rho = \infty$. Let

$$\mathcal{L}_N(\tilde{x}, \tilde{\theta}) := \frac{\sigma^{2\bar{d}}}{N} \sum_{i=1}^N \log \frac{f_{\tilde{x}, \tilde{\theta}}(\tilde{Y}_i)}{f_{x, \theta}(\tilde{Y}_i)}. \quad (\text{III.14})$$

The MLE is given by

$$\hat{X}_{\text{MLE}} = \underset{\tilde{x}}{\operatorname{argmax}} \underset{\tilde{\theta}}{\operatorname{max}} \mathcal{L}_N(\tilde{x}, \tilde{\theta}).$$

Fix \tilde{x} and $\tilde{\theta}$, and for ease of notation let $d = d_{\tilde{x}, \tilde{\theta}}$. We can write

$$\mathcal{L}_N(\tilde{x}, \tilde{\theta}) = \frac{\sigma^{2(\bar{d}-d)}}{N/\sigma^{2d}} \sum_{i=1}^{N/\sigma^{2d}} \sum_{j=1}^{\sigma^{2d}} \log \frac{f_{\tilde{x}, \tilde{\theta}}(\tilde{Y}_{\sigma^{2d}(i-1)+j})}{f_{x, \theta}(\tilde{Y}_{\sigma^{2d}(i-1)+j})}$$

We have

$$\begin{aligned} \mathbb{E} \left[\sum_{j=1}^{\sigma^{2d}} \log \frac{f_{\tilde{x}, \tilde{\theta}}(\tilde{Y}_j)}{f_{x, \theta}(\tilde{Y}_j)} \right] &= \sigma^{2d} \mathbb{E} \left[\log \frac{f_{\tilde{x}, \tilde{\theta}}(\tilde{Y})}{f_{x, \theta}(\tilde{Y})} \right] \\ &= -\sigma^{2d} D(f_{\tilde{x}, \tilde{\theta}} \| f_{x, \theta}) \end{aligned}$$

where D denotes the KL divergence, defined for two probability densities f_A and f_B as

$$D(f_A \| f_B) := \mathbb{E} \left[\log \left(\frac{f_A(A)}{f_B(A)} \right) \right],$$

where $A \sim f_A$.

Using (III.6), with $\gamma = 1/\sigma$, we have $f_{\tilde{x}, \tilde{\theta}} \rightarrow f_{x, \theta}$ as $\gamma \rightarrow 0$, which implies by [20, Section F, Theorem 9] that

$$\lim_{\gamma \rightarrow 0} \frac{D(f_{\tilde{x}, \tilde{\theta}} \| f_{x, \theta})}{\chi^2(f_{\tilde{x}, \tilde{\theta}} \| f_{x, \theta})} = \frac{1}{2},$$

and

$$D(f_{\tilde{x}, \tilde{\theta}} \| f_{x, \theta}) = \frac{\sigma^{-2d}}{2d!} \|M_{\tilde{x}, \tilde{\theta}}^d - M_{x, \theta}^d\|^2 + \mathcal{O}(\sigma^{-2d-1}),$$

thus by the law of large numbers, since N/σ^{2d} diverges,

$$\mathcal{L}_N(\tilde{x}, \tilde{\theta}) \rightarrow \begin{cases} -\infty & \text{if } d_{\tilde{x}, \tilde{\theta}} < \bar{d} \\ -\frac{1}{2d!} \|M_{\tilde{x}, \tilde{\theta}}^{\bar{d}} - M_{x, \theta}^{\bar{d}}\|^2 & \text{otherwise} \end{cases}$$

As $N \rightarrow \infty$, the maximum of $\mathcal{L}_N(\tilde{x}, \tilde{\theta})$ tends to 0 in probability, and is achieved when $\phi_x(\tilde{x}) = x$, thus the MLE must converge in probability to gx for some $g \in \Theta$.

ACKNOWLEDGMENTS

EA was partly supported by the NSF CAREER Award CCF-1552131, ARO grant W911NF-16-1-0051 and NSF Center for the Science of Information CCF-0939370. JP and AS were partially supported by Award Number R01GM090200 from the NIGMS, the Simons Foundation Investigator Award and Simons Collaborations on Algorithms and Geometry, the Moore Foundation Data-Driven Discovery Investigator Award, and AFOSR FA9550-17-1-0291.

We would like to thank Afonso Bandeira, Tamir Bendory, Joseph Kileel, William Leeb and Nir Sharon for many insightful discussions.

REFERENCES

- [1] W. Park, C. R. Midgett, D. R. Madden, and G. S. Chirikjian, "A stochastic kinematic model of class averaging in single-particle electron microscopy," *The International journal of robotics research*, vol. 30, no. 6, pp. 730–754, 2011.
- [2] W. Park and G. S. Chirikjian, "An assembly automation approach to alignment of noncircular projections in electron microscopy," *IEEE Transactions on Automation Science and Engineering*, vol. 11, no. 3, pp. 668–679, 2014.
- [3] S. H. Scheres, M. Valle, R. Nuñez, C. O. Sorzano, R. Marabini, G. T. Herman, and J.-M. Carazo, "Maximum-likelihood multi-reference refinement for electron microscopy images," *Journal of molecular biology*, vol. 348, no. 1, pp. 139–149, 2005.
- [4] J. P. Zwart, R. van der Heiden, S. Gelsema, and F. Groen, "Fast translation invariant classification of HRR range profiles in a zero phase representation," *IEE Proceedings-Radar, Sonar and Navigation*, vol. 150, no. 6, pp. 411–418, 2003.
- [5] R. Gil-Pita, M. Rosa-Zurera, P. Jarabo-Amores, and F. López-Ferreras, "Using multilayer perceptrons to align high range resolution radar signals," in *International Conference on Artificial Neural Networks*, pp. 911–916, Springer, 2005.
- [6] D. M. Rosen, L. Carlone, A. S. Bandeira, and J. J. Leonard, "A certifiably correct algorithm for synchronization over the special euclidean group," *arXiv:1611.00128*, 2016.
- [7] I. L. Dryden and K. V. Mardia, *Statistical shape analysis*, vol. 4. J. Wiley Chichester, 1998.
- [8] H. Foroosh, J. B. Zerubia, and M. Berthod, "Extension of phase correlation to subpixel registration," *IEEE transactions on image processing*, vol. 11, no. 3, pp. 188–200, 2002.
- [9] D. Robinson, S. Farsiu, and P. Milanfar, "Optimal registration of aliased images using variable projection with applications to super-resolution," *The Computer Journal*, vol. 52, no. 1, pp. 31–42, 2009.
- [10] A. Bandeira, P. Rigollet, and J. Weed, "Optimal rates of estimation for multi-reference alignment," *arXiv preprint 1702.08546*, 2017.
- [11] A. Perry, J. Weed, A. Bandeira, P. Rigollet, and A. Singer, "The sample complexity of multi-reference alignment," *arXiv preprint 1707.00943*, 2017.
- [12] T. Bendory, N. Boumal, C. Ma, Z. Zhao, and A. Singer, "Bispectrum inversion with application to multireference alignment," *IEEE Transactions on Signal Processing*. To appear.
- [13] E. Abbe, T. Bendory, W. Leeb, J. M. Pereira, N. Sharon, and A. Singer, "Multireference Alignment is Easier with an Aperiodic Translation Distribution," *arXiv:1710.02793*, 2017.
- [14] A. Bartesaghi, A. Merk, S. Banerjee, D. Matthies, X. Wu, J. L. Milne, and S. Subramaniam, "2.2 Å resolution cryo-EM structure of β -galactosidase in complex with a cell-permeant inhibitor," *Science*, vol. 348, no. 6239, pp. 1147–1151, 2015.
- [15] D. Sirohi, Z. Chen, L. Sun, T. Klose, T. C. Pierson, M. G. Rossmann, and R. J. Kuhn, "The 3.8 Å resolution cryo-EM structure of Zika virus," *Science*, vol. 352, no. 6284, pp. 467–470, 2016.
- [16] A. S. Bandeira, B. Blum-Smith, A. Perry, J. Weed, and A. S. Wein, "Estimation under group actions: recovering orbits from invariants," *arXiv:1712.10163*, 2017.
- [17] H. Cramér, *Mathematical Methods of Statistics (PMS-9)*, vol. 9. Princeton university press, 2016.
- [18] E. Abbe, J. M. Pereira, and A. Singer, "Sample complexity of the Boolean multireference alignment problem," in *2017 IEEE International Symposium on Information Theory (ISIT)*, pp. 1316–1320, June 2017.
- [19] D. G. Chapman and H. Robbins, "Minimum variance estimation without regularity assumptions," *Ann. Math. Statist.*, vol. 22, pp. 581–586, 12 1951.
- [20] I. Sason and S. Verdú, " f -divergence inequalities," *IEEE Transactions on Information Theory*, vol. 62, no. 11, pp. 5973–6006, 2016.