Diagonal and Low-Rank Decompositions and Fitting Ellipsoids to Random Points

James Saunderson, Pablo A. Parrilo, Alan S. Willsky

***Abstract*— Identifying a subspace containing signals of interest in additive noise is a basic system identification problem. Under natural assumptions, this problem is known as the Frisch scheme and can be cast as decomposing an $n \times n$ positive definite matrix as the sum of an unknown diagonal matrix (the noise covariance) and an unknown low-rank matrix (the signal covariance). Our focus in this paper is a natural class of random instances, where the low-rank matrix has a uniformly distributed random column space.**

In this setting we analyze the behavior of a well-known convex optimization-based heuristic for diagonal and low-rank decomposition called minimum trace factor analysis (MTFA). Conditions for the success of MTFA have an appealing geometric reformulation as finding a (convex) ellipsoid that exactly interpolates a given set of n points. Under the random model, the points are chosen according to a Gaussian distribution.

Numerical experiments suggest a remarkable threshold phenomenon: if the (random) column space of the $n \times n$ low-rank matrix has codimension as small as $2\sqrt{n}$ then with high probability MTFA successfully performs the decomposition task, otherwise it fails with high probability. In this work we provide numerical evidence and prove partial results in this direction, showing that with high probability MTFA recovers such random low-rank matrices of corank at least cn^{β} for $\beta \in (5/6, 1)$ and some constant c.

I. Introduction

Structured matrix decomposition and recovery problems have received a much attention recently in the context of system identification. Often the system to be identified is represented by a structured matrix (in the linear time-invariant case this may be a Hankel matrix or operator) the rank of which is a natural measure of system complexity (such as Macmillan degree), e.g., [2], [3], [4]. Common problem formulations are to recover a low-rank matrix from a small number of linear measurements [5], [6] or decompose a matrix as the sum of a low-rank part and a structured disturbance (of possibly large magnitude) [7].

A typical approach is to formulate the problem of interest as a rank minimization problem and then convexify the problem by replacing the rank function in the objective with the nuclear norm [5], [8]. A key intuition that has developed from this line of work is that the nuclear norm heuristic works well for 'typical' problem instances in high dimensions. This can be formalized by putting a natural measure on problem instances and showing that the method works well on a set of large measure. Stated another way we can consider a natural random ensemble of problem instances (indexed by dimension) and ask that the heuristic work well with 'high probability'.

This research was supported in part by AFOSR under Grant FA9550-12-1-0287 and in part by AFOSR under Grant FA9550-11-1-0305. A preliminary version of this work appeared in the Master's thesis of the first author [1].

The authors are with the Laboratory for Information and Decision Systems, Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge MA 20139, USA. {jamess,parrilo,willsky}@mit.edu

A. The Frisch scheme and matrix decompositions

In this paper we take this approach to a classical identification problem known as the Frisch scheme [9] or, in the statistics literature, factor analysis [10].

Suppose $\mathcal{U} \subset \mathbb{R}^n$ is a subspace (the signal subspace), u (the signal) is a zero-mean random variable taking values in $\mathcal{U}$, and w (the noise) is a zero-mean random variable independent of u. Consider observing $x = u + w$ and assume we have access to the covariance $X := \mathbb{E}[xx^T]$ of x.

In general we cannot hope to identify $\mathcal{U}$ (the signal subspace) from X without additional assumptions on w and u. In the Frisch scheme [9] we assume w has diagonal covariance $D := \mathbb{E}[ww^T]$ (or that we know a basis in which it is diagonal). Since u takes values in $\mathcal{U}$, $L := \mathbb{E}[uu^T]$ has column space $\mathcal{U}$ and so is rank deficient. Identifying a subspace in noise under the Frisch scheme assumption can be cast as finding a decomposition $X = L + D$ of the covariance X of x as the sum of a diagonal matrix D and a rank-deficient matrix L. In particular, we are often interested in finding such a decomposition with the *minimum rank* of L, corresponding to identifying the lowest dimensional subspace consistent with the data X and our assumptions.

B. Minimum trace factor analysis

A natural tractable convex optimization-based method to recover D and L from their sum X is *minimum trace factor analysis* (MTFA):

$$\min_{D,L} \ \mathrm{tr}(L) \ \text{ s.t. } \ X = D + L, \ L \succeq 0, \ D \text{ diagonal.}$$

This heuristic dates to the work of Ledermann in 1940 [11], and is a natural precursor to the modern nuclear norm-based approach to developing convex optimization-based methods for rank minimization problems [8].

We investigate how well MTFA performs diagonal and low-rank decompositions on problem instances where the column space of L is uniformly distributed (i.e., with respect to Haar measure) on codimension k subspaces of $\mathbb{R}^n$. This leads us to the central problem considered in this paper.

Problem 1: Let $X = D + L$ where the column space of L is uniformly distributed on codimension k subspaces of $\mathbb{R}^n$. For which pairs (n, k) does MTFA correctly decompose X with high probability?

We would like to emphasize that we have stated the problem in terms of the *co*dimension of the column space of L (or equivalently the *co*rank of L). In particular, for a fixed n, we are interested in the *smallest* k such that MTFA correctly decomposes 'most' matrices that are the sum of a diagonal matrix and a corank k positive semidefinite matrix.

In Problem 1 and throughout the paper, we often refer to statements indexed by n and k (with $n \geq k$) holding 'with high probability' (abbreviated as w.h.p. in the sequel). By this we mean that there are positive constants c_1, c_2 and β such that for large enough n and k the statement holds with probability at least $1 - c_1 e^{-c_2 k^{\beta}}$.

C. Geometric interpretations

The diagonal and low-rank decomposition problem is particularly interesting because it is closely related to an interpolation problem with an appealing geometric interpretation. Given a $k \times k$ symmetric positive semidefinite matrix Y, we call the level sets $\{x \in \mathbb{R}^k : x^T Y x = 1\}$ *ellipsoids*. Note that these ellipsoids are always centered at the origin, and may be 'degenerate'. The geometric interpretation of Problem 1 turns out to be the following problem of fitting an ellipsoid to random Gaussian points.

Problem 2: Let $v_1, v_2, \ldots, v_n$ be independent standard (zero mean, identity covariance) Gaussian vectors in $\mathbb{R}^k$. For which pairs (n, k) is there an ellipsoid passing *exactly* through all of the points v_i w.h.p.?

If we fix the number of points n, it is easier to fit an ellipsoid to n points in a larger dimensional space. As such we are interested in the smallest k such that we can fit an ellipsoid to n standard Gaussian points in $\mathbb{R}^k$ with high probability.

D. Numerical evidence and a conjecture

Numerical simulations (detailed in Section IV) strongly suggest there is a very nice answer to the equivalent Problems 1 and 2. Based on the evidence in Figure 1 we make the following (equivalent) conjectures.

Conjecture 1 (Random diagonal/low-rank decomposition): With high probability MTFA successfully decomposes the sum of a diagonal matrix and a matrix with column space uniformly distributed on codimension k subspaces of $\mathbb{R}^n$ as long as $n \lesssim k^2/4$.

Conjecture 2 (Random ellipsoid fitting): With high probability there is an ellipsoid passing through n standard Gaussian points in $\mathbb{R}^k$ as long as $n \lesssim k^2/4$.

(Here we use the notation $n \lesssim f(k)$ to mean that $n \leq f(k) + g(k)$ where $\lim_{n,k\to\infty} g(k)/n = 0$.)

To gain intuition for how strong these conjectured results are, it is instructive to substitute values for n and k into the statement of the conjectures. Putting $n = 1000$, and $k = 65$ we see that MTFA can correctly recover 'most' matrices of rank up to $n - k = 935$ (hardly just low rank matrices). Furthermore, we can fit an ellipsoid, with high probability to up to 1000 Gaussian points in $\mathbb{R}^{65}$.

An easy dimension counting argument (see Section II-B) shows that the probability of being able to fit any quadratic surface, not necessarily an ellipsoid, to $n > \binom{k+1}{2} \sim k^2/2$ points in $\mathbb{R}^k$ is zero. Similarly if $n > \binom{k+1}{2}$ the diagonal and low-rank decomposition problem is not even generically locally identifiable [12]. Yet if we decrease n, the number of points, from $k^2/2$ to $k^2/4$, we can fit an *ellipsoid* to them with high probability. Similarly, if we increase the corank k of the low-rank matrix from $\sqrt{2n}$ to $\sqrt{4n}$ we can correctly perform diagonal and low-rank matrix decompositions using a tractable convex program with high probability.

E. Our results

At present we are not able to establish Conjectures 1 and 2. Nevertheless in Section V of this paper we outline a proof of the following weaker results (which are equivalent).

Theorem 1: Fix $\alpha \in (0, 1/6)$. There are absolute positive constants $C, \bar{c}, \tilde{c}$ such that for sufficiently large n and k satisfying $n \leq C k^{\frac{6}{5}(1-\alpha)}$ MTFA correctly decomposes the sum of a diagonal matrix and a low-rank matrix with column space uniformly distributed on codimension k subspaces of $\mathbb{R}^n$ with probability at least $1 - \bar{c} n e^{-\tilde{c} k^{3\alpha}}$ (i.e. w.h.p.).

Theorem 2: Fix $\alpha \in (0, 1/6)$. There are absolute positive constants $C, \bar{c}, \tilde{c}$ such that for sufficiently large n and k satisfying $n \leq C k^{\frac{6}{5}(1-\alpha)}$, there is an ellipsoid passing through n standard Gaussian vectors in $\mathbb{R}^k$ with probability at least $1 - \bar{c} n e^{-\tilde{c} k^{3\alpha}}$ (i.e. w.h.p.).

In previous work we have given simple deterministic conditions for the ellipsoid fitting problem [13] based on a notion of the coherence of a subspace (see Section V-A). These conditions can only be satisfied if $n \leq 2k$, i.e., the number of points is at most twice the dimension k of the space in which they live. By comparison, the above results establish much stronger scaling—we can have a number of points in $\mathbb{R}^k$ that grows superlinearly with k and still fit an ellipsoid with high probability.

F. Outline

In Section II we recall the relationships from [13] between the analysis of MTFA and the ellipsoid fitting problem. From that point on we focus on the ellipsoid fitting problem. We discuss different methods for constructing ellipsoids (and hence obtaining sufficient conditions for the problems of interest) in Section III followed by our numerical experiments in Section IV. Section V follows with an outline of our proof of Theorem 2.

II. Relating MTFA and Ellipsoid fitting

A. Notions of 'success' for MTFA

We now fix some terminology related to when MTFA is successful in performing diagonal and low-rank decompositions. We are interested in understanding when MTFA correctly decomposes a matrix X made up as the sum of a diagonal matrix D and a low-rank positive semidefinite matrix L, and hence correctly identifies the column space of the low-rank matrix. It is a straightforward consequence of the optimality conditions for MTFA [14] that whether MTFA correctly decomposes $X = D + L$ depends only on the column space of L, motivating the following definition.

Definition 1: A subspace $\mathcal{U}$ is *recoverable by MTFA* if for any diagonal matrix D and any positive semidefinite matrix L with column space $\mathcal{U}$, the optimum of MTFA with input $X = D + L$ is (D, L).

When considering random problem instances, we are interested in whether 'most' codimension k subspaces of $\mathbb{R}^n$ are recoverable by MTFA, a notion formalized as follows.

Definition 2: A pair (n, k) is *recoverable by MTFA with high probability* if a uniformly distributed codimension k subspace of $\mathbb{R}^n$ is recoverable by MTFA w.h.p.

B. Ellipsoid fitting

Consider the following elementary geometric problem.

Problem 3 (Ellipsoid fitting): Given a collection of points $v_1, v_2, \ldots, v_n \in \mathbb{R}^k$, does there exist an ellipsoid passing *exactly* through them?

Note that the property of a collection of points having an interpolating ellipsoid is invariant under invertible linear transformations of $\mathbb{R}^k$. As such, whether we can fit an ellipsoid to $v_1, v_2, \ldots, v_n \in \mathbb{R}^k$ depends only on a subspace associated with the points. This subspace is the column space of the $n \times k$ matrix V with rows given by the v_i. This motivates the following definition (from [13]).

Definition 3: A subspace $\mathcal{V}$ of $\mathbb{R}^n$ has the *ellipsoid fitting property* if whenever V is an $n \times k$ matrix with columns that are a basis for $\mathcal{V}$ there is an ellipsoid in $\mathbb{R}^k$ passing through the rows of V (thought of as points in $\mathbb{R}^k$).

In the random setting, we are interested in the pairs (n, k) such that 'most' dimension k subspaces of $\mathbb{R}^n$ have the ellipsoid fitting property, for which we use the following terminology.

Definition 4: A pair (n, k) has the *ellipsoid fitting property with high probability* if a uniformly distributed dimension k subspace of $\mathbb{R}^n$ has the ellipsoid fitting property w.h.p.

Because the subspace corresponding to n Gaussian points in $\mathbb{R}^k$ is uniformly distributed, the random ellipsoid fitting problem (Problem 2) exactly corresponds to determining which pairs (n, k) have the ellipsoid fitting property w.h.p.

There is a simple upper bound on the number n of random points $v_1, \ldots, v_n$ to which we can fit an ellipsoid in $\mathbb{R}^k$. In order to fit an ellipsoid it is necessary that there is a symmetric matrix Y (not necessary positive semidefinite) satisfying $v_i^T Y v_i = 1$ for all i. This imposes n independent (for random, hence suitably generic, points) linear equations on Y which can only be satisfied if $n \leq \binom{k+1}{2}$.

C. Relating ellipsoid fitting and MTFA

In previous work [13] we established the following link between the recovery properties of MTFA and the ellipsoid fitting problem.

Proposition 1: A subspace $\mathcal{U}$ is recoverable by MTFA if and only if $\mathcal{U}^\perp$ has the ellipsoid fitting property. A pair (n, k) is recoverable by MTFA w.h.p. if and only if it has the ellipsoid fitting property w.h.p.

The proof of the first statement follows from the optimality conditions for MTFA. The second statement follows from the definitions and the observation that if $\mathcal{U}$ is uniformly distributed on k codimensional subspaces of $\mathbb{R}^n$ then its orthogonal complement $\mathcal{U}^\perp$ is uniformly distributed on k dimensional subspaces of $\mathbb{R}^n$.

From now on, rather than discussing both MTFA and ellipsoid fitting in parallel, we focus only on the ellipsoid fitting formulation of the problem.

III. Constructing ellipsoids

We can determine if there is an ellipsoid passing through the points $v_1, v_2, \ldots, v_n \in \mathbb{R}^k$ by solving a semidefinite feasibility problem. In this section we present that semidefinite program and a simple projection-based method to solve it. We also explain some simpler least-squares based techniques for fitting ellipsoids that give sufficient conditions for a collection of points to have an interpolating ellipsoid.

A. The semidefinite feasibility problem

Deciding whether there is an ellipsoid passing through the points $v_1, v_2, \ldots, v_n \in \mathbb{R}^k$ is clearly equivalent to deciding whether there exists

$$Y \succeq 0 \text{ such that } v_i^T Y v_i = 1 \text{ for } i = 1, 2, \ldots, n. \quad (1)$$

Geometrically, this is equivalent to deciding whether the positive semidefinite cone and an affine subspace $\mathcal{L}$ of $k \times k$ symmetric matrices have non-empty intersection.

A simple way to solve such decision problems is by alternating projection. Starting from an initial symmetric $k \times k$ matrix Y_0, define a sequence Y_t by alternately performing the Euclidean projection onto $\mathcal{L}$ and the Euclidean projection onto the positive semidefinite cone. It is well known that this sequence converges if and only if the two sets intersect and, moreover, converges to a point in their intersection.

The projection onto the positive semidefinite cone is $\Pi_{\text{psd}}(Y) = \sum_{i=1}^n (\lambda_i)_+ v_i v_i^T$ where $Y = \sum_{i=1}^n \lambda_i u_i u_i^T$ is an eigendecomposition of Y and $(t)_+ = \max\{t, 0\}$.

To concisely describe the projection onto $\mathcal{L}$, observe that $\mathcal{L} = \{Y \in \mathcal{S}^k : \mathcal{A}(Y) = \mathbf{1}\}$ where we define the linear map

$$\mathcal{A} : \mathcal{S}^k \to \mathbb{R}^n \quad \text{by} \quad [\mathcal{A}(Y)]_i = v_i^T Y v_i. \quad (2)$$

The adjoint $\mathcal{A}^* : \mathbb{R}^n \to \mathcal{S}^k$ of the map $\mathcal{A}$ is $\mathcal{A}^*(x) = \sum_{i=1}^n x_i v_i v_i^T$, and its pseudoinverse is $\mathcal{A}^\dagger = \mathcal{A}^*(\mathcal{A}\mathcal{A}^*)^{-1}$. The Euclidean projection onto $\mathcal{L}$ is then

$$\Pi_{\mathcal{L}}(Y) = \mathcal{A}^\dagger(\mathbf{1}) + (Y - \mathcal{A}^\dagger \mathcal{A}(Y)). \quad (3)$$

B. Least-squares based constructions

If we seek only sufficient conditions on a set of points that ensure there is an ellipsoid passing through them, we need not directly solve the semidefinite feasibility problem (1). A simpler alternative is to start with a symmetric matrix Y_0 which is certainly positive definite (and satisfies $v_i^T Y_0 v_i \approx 1$ for all i) and project it onto the subspace $\mathcal{L}$ of symmetric matrices Y that do satisfy $v_i^T Y v_i = 1$ exactly for all i. If we are lucky and the projection does not change Y_0 too much, the resulting symmetric matrix is still positive semidefinite, and we have constructed an ellipsoid passing through the points. If this method fails, we do not know whether there is an ellipsoid passing through the points. This idea is behind many of the dual certificate construction methods in the sparse recovery literature.

We now formalize this construction. Given a set of points, $v_1, v_2, \ldots, v_n \in \mathbb{R}^k$, let $\mathcal{A} : \mathcal{S}^k \to \mathbb{R}^n$ be defined as in (2). If Y_0 is our initial 'guess' then the corresponding least squares construction is simply $Y = \Pi_{\mathcal{L}}(Y_0)$. If $Y \succeq 0$ then we have found an ellipsoid passing through the points.

We typically take $Y_0 = \mathcal{A}^*(\eta \mathbf{1}) = \eta \sum_{i=1}^n v_i v_i^T$ for some positive scalar η (chosen so that Y_0 approximately fits the points). This is clearly positive definite and is well adapted to the position of the points. Then from (3) the least squares construction is

$$Y = \mathcal{A}^\dagger(\mathbf{1}).$$

This construction is *not* invariant in the sense that it depends on the choice of points, not just on the subspace corresponding to the points. To obtain an invariant construction, we can first put the points in *isotropic* position, i.e. change coordinates so that $\sum_{i=1}^n v_i v_i^T = I$. We call the least squares construction based on this choice of points the *isotropic* construction.

IV. Numerical experiments

In this section we describe our numerical experiments. These have two primary aims. The first is to determine, experimentally, which pairs (n, k) have the ellipsoid fitting property w.h.p., giving support for Conjecture 2. The second is to determine for which pairs (n, k) we can find an ellipsoid using the simple least squares-based methods described in Section III-B, providing evidence that the least squares construction is a useful proof technique.

For each $n = 100, 101, \ldots, 1000$ and each $k = \lfloor 2\sqrt{n} \rfloor, \ldots, n$ we repeated the following procedure 10 times:

1) Sample an $n \times k$ matrix V with i.i.d. standard Gaussian entries and let Q be an $n \times k$ matrix with columns that are an orthonormal basis for the column space of V.
2) Try to fit an ellipsoid in three ways
   - by solving the semidefinite feasibility problem (1) using alternating projections (see Section III-A)
   - by the least squares construction applied to
     - the rows of Q (the isotropic construction)
     - the rows of V.

We record the proportion of trials, for each pair (n, k), where each method successfully fits an ellipsoid. Figure 1 shows the results. We see a clear phase transition between being able and unable to fit an ellipsoid to random Gaussian points. Furthermore, superimposed on the phase transition plot for the feasibility problem is the curve $n = k^2/4$, providing supporting evidence for the random ellipsoid fitting conjecture (Conjecture 2).

As expected, directly solving the feasibility problem outperforms both of the least squares constructions, with the isotropic least squares construction outperforming the version applied directly to the Gaussian points. Both seem to successfully fit an ellipsoid to a number of points growing superlinearly, and perhaps even quadratically, with the dimension.

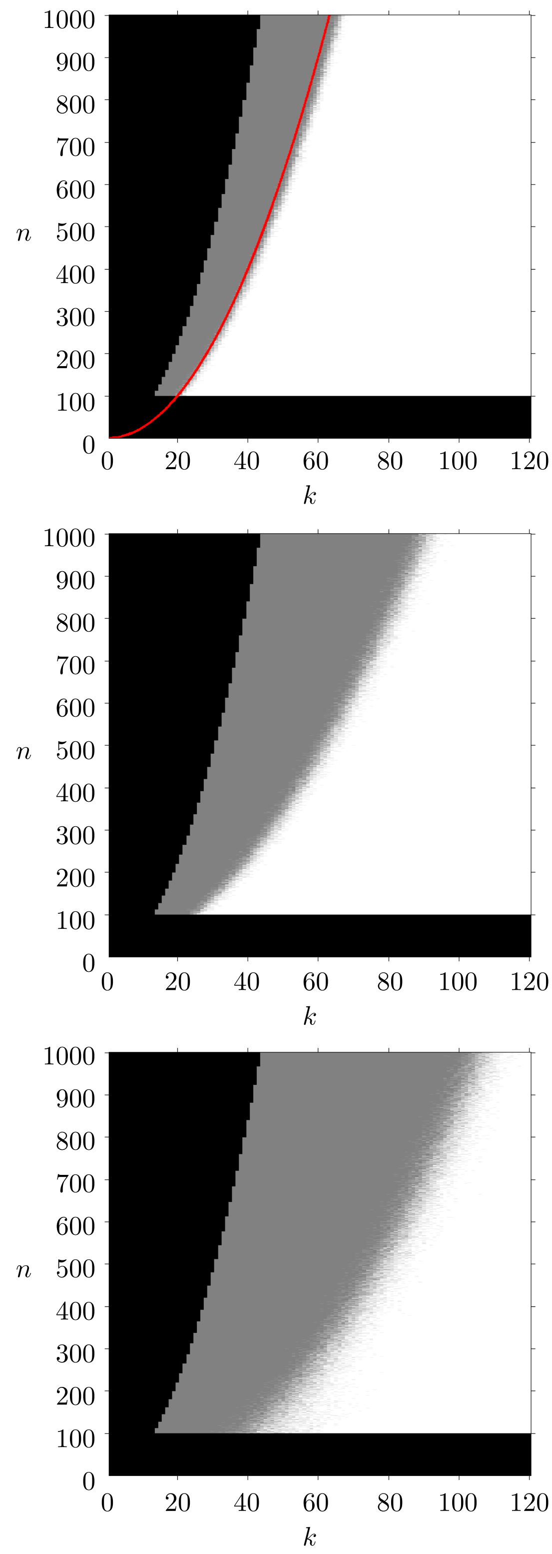

Fig. 1. Phase transition plots for three methods of fitting an ellipsoid to n standard Gaussian vectors in $\mathbb{R}^k$. The top plot corresponds to solving the semidefinite feasibility problem, showing the pairs (n, k) that have the ellipsoid fitting property w.h.p. For the middle plot and bottom plots we use the isotropic and Gaussian least squares constructions, respectively. In each case the cell corresponding to (n, k) is white or dark gray respectively if we found an ellipsoid for all or no trials respectively with intermediate shading proportional to the number of successful trials (lighter is better). A cell is black if we did not run the experiment for the corresponding (n, k) pair. The red line on the top plot is the graph of function $n = k^2/4$.

V. Analysis of random ellipsoid fitting

In this section we analyze the random ellipsoid fitting problem (Problem 2). We first recall a known sufficient condition for ellipsoid fitting, and show how it is not powerful enough to establish results as strong as our main ellipsoid fitting result (Theorem 2), let alone our stronger conjecture (Conjecture 2).

A. Limitations of coherence-based results

Define the *coherence* of a subspace $\mathcal{U}$ of $\mathbb{R}^n$ as

$$\mu(\mathcal{U}) = \max_i \|\Pi_{\mathcal{U}} e_i\|_2^2.$$

Coherence measures how well a subspace aligns with the coordinate axes and satisfies the basic inequality

$$\dim(\mathcal{U})/n \leq \mu(\mathcal{U}) \leq 1. \tag{4}$$

The main result of [13] gives the following coherence-threshold type sufficient condition for ellipsoid fitting.

Theorem 3: If $\mu(\mathcal{V}^\perp) < 1/2$ then $\mathcal{V}$ has the ellipsoid fitting property. Furthermore for any $\epsilon > 0$ there is a subspace $\mathcal{V}$ with $\mu(\mathcal{V}^\perp) = 1/2 + \epsilon$ that does not have the ellipsoid fitting property.

This sufficient condition is never satisfied for subspaces $\mathcal{V}$ of $\mathbb{R}^n$ with $\dim(\mathcal{V}) < n/2$, because, by (4) such a subspace has $\mu(\mathcal{V}^\perp) > 1/2$. Yet our experiments, and our conjecture indicate that many subspaces of dimension much smaller than half the ambient dimension have the ellipsoid fitting property. The basic problem is that coherence can only 'see' the *ratio* of the dimension of the subspace and the ambient dimension, whereas our experiments, our main result, and our conjecture indicate that the scaling depends on more than this ratio. To establish our main result we require a new approach, which we outline in the sequel.

B. Analysis of random ellipsoid fitting

In this section we give an outline of the proof of Theorem 2. We first give a deterministic condition on a set of points that ensures they have the ellipsoid fitting property, and then show that it is satisfied w.h.p. under the assumptions of Theorem 2.

We briefly review some notation used in this section. For a symmetric matrix the spectral norm $\|\cdot\|_{\text{sp}}$ is the largest singular value. For an element of $\mathbb{R}^n$, $\|\cdot\|_2$ is the Euclidean norm, and $\|\cdot\|_\infty$ the maximum absolute value of the entries. For a map $\mathcal{B}$ we denote by $\|\mathcal{B}\|_{a\to b}$ for the induced norm $\sup_{\|x\|_a \leq 1} \|\mathcal{B}(x)\|_b$. To simplify notation, instead of writing $\|\mathcal{B}\|_{2\to 2}$ we simply write $\|\mathcal{B}\|$.

1) A deterministic condition: Let $\mathcal{A}$ be defined as in (2) with respect to a collection of points $v_1, v_2, \ldots, v_n \in \mathbb{R}^k$. Our aim is to establish conditions under which the following procedure yields an ellipsoid passing through these points.

Take the 'nominal ellipsoid' $Y_0 = \mathcal{A}^*(\eta\mathbf{1})$ where η is a positive constant to be chosen later so that $v_i^T Y_0 v_i \approx 1$ for all i. Project it onto the subspace $\mathcal{L}$ to obtain the least squares construction $\mathcal{A}^\dagger(\mathbf{1}) = \mathcal{A}^*(\mathcal{A}\mathcal{A}^*)^{-1}(\mathbf{1})$. If $\mathcal{A}^\dagger(\mathbf{1}) \succ 0$ we have succeeded (see Section III-B).

It is useful to think of $\mathcal{A}^\dagger(\mathbf{1})$ as a perturbation of the nominal ellipsoid Y_0. Then to establish that $\mathcal{A}^\dagger(\mathbf{1}) \succ 0$ it suffices to show that Y_0 is sufficiently positive definite and control the following difference in *spectral norm*:

$$\mathcal{A}^\dagger(\mathbf{1}) - Y_0 = \mathcal{A}^\dagger(\mathbf{1} - \mathcal{A}\mathcal{A}^*(\eta\mathbf{1})). \tag{5}$$

Note that the ith entry of $\mathbf{1} - \mathcal{A}\mathcal{A}^*(\eta\mathbf{1})$ is just $1 - v_i^T Y_0 v_i$, i.e., how close our 'nominal' ellipsoid is to fitting the ith point. Equation (5) then suggests that we need to control $\|\mathcal{A}^\dagger\|_{\infty\to\text{sp}}$. This controls the sensitivity of the least squares system with respect to the norms relevant to us.

Let us now summarize the above discussion.

Proposition 2: Let $v_1, v_2, \ldots, v_n \in \mathbb{R}^k$ and let $\mathcal{A} : \mathcal{S}^k \to \mathbb{R}^n$ be defined by $[\mathcal{A}(X)]_i = v_i^T X v_i$. If $\eta > 0$ satisfies

$$\lambda_{\text{min}}(\mathcal{A}^*(\eta\mathbf{1})) > \|\mathcal{A}^\dagger\|_{\infty\to\text{sp}} \|\mathbf{1} - \mathcal{A}\mathcal{A}^*(\eta\mathbf{1})\|_\infty \tag{6}$$

then $\mathcal{A}^\dagger(\mathbf{1}) \succ 0$ and so there is an ellipsoid passing through $v_1, v_2, \ldots, v_n$.

Proof: From (5), $\mathcal{A}^\dagger(\mathbf{1}) = \mathcal{A}^*(\eta\mathbf{1}) + \mathcal{A}^\dagger(\mathbf{1} - \mathcal{A}\mathcal{A}^*(\eta\mathbf{1}))$. Hence under the stated assumption

$$\mathcal{A}^\dagger(\mathbf{1}) \succeq (\lambda_{\text{min}}(\mathcal{A}^*(\eta\mathbf{1})) - \|\mathcal{A}^\dagger\|_{\infty\to\text{sp}} \|\mathbf{1} - \mathcal{A}\mathcal{A}^*(\eta\mathbf{1})\|_\infty) I$$

and the right-hand side is positive definite. ∎

2) Outline of the proof of Theorem 2: We now turn to the random setting. Let $v_1, v_2, \ldots, v_n$ be standard Gaussian vectors in $\mathbb{R}^k$. We define $\mathcal{A}$ (and hence the least squares construction) directly using the Gaussian vectors, rather than using the corresponding isotropic version, because the additional dependencies introduced by passing to the isotropic construction seem to complicate some things considerably.

To prove Theorem 2 our task is to choose η in Proposition 2, and show that the three relevant quantities $\lambda_{\text{min}}(\mathcal{A}^*(\eta\mathbf{1}))$, $\|\mathcal{A}^\dagger\|_{\infty\to\text{sp}}$ and $\|\mathbf{1} - \mathcal{A}\mathcal{A}^*(\eta\mathbf{1})\|_\infty$ satisfy (6) under the assumptions of Theorem 2. From here on we make the choice $\eta = 1/k(n + k - 1)$ so that the expectation of $\mathbf{1} - \mathcal{A}\mathcal{A}^*(\eta\mathbf{1})$ is zero.

Controlling the smallest eigenvalue of the nominal ellipsoid, and how well it fits the points, are both fairly straightforward. The relevant results are summarized in the following two lemmas. The first is standard [15, Lemma 36 and Corollary 35], the second is proved in [1, Section 3.9.3].

Lemma 1 (Smallest eigenvalue of the nominal ellipsoid): If $k = o(n)$ then there are positive constants $\tilde{c}_1$ and c_1 such that with probability at least $1 - 2e^{-\tilde{c}_1 n}$,

$$\lambda_{\text{min}}(\mathcal{A}^*(\mathbf{1})) = \lambda_{\text{min}}\left(\textstyle\sum_{i=1}^n v_i v_i^T\right) \geq n c_1.$$

Lemma 2 (How well the nominal ellipsoid fits the points): If $0 < \alpha < 1/6$ there are positive constants $\tilde{c}_2, c_2$ and $\bar{c}_2$ such that if $\eta = 1/k(n + k - 1)$ then with probability at least $1 - \bar{c}_2 n e^{-\tilde{c}_2 k^{3\alpha}}$

$$\|\mathbf{1} - \mathcal{A}\mathcal{A}^*(\eta\mathbf{1})\|_\infty \leq c_2 n k^{1/2+3\alpha/2}.$$

Controlling the sensitivity of the least squares system, i.e., $\|\mathcal{A}^\dagger\|_{\infty\to\text{sp}}$ is more involved. To do this we expand $\mathcal{A}^\dagger$ as a series about a natural nominal point $M := \mathbb{E}[\mathcal{A}\mathcal{A}^*]$, establish validity of the expansion, and bound the first term and tail of the series separately. We expand $\mathcal{A}^\dagger$ as

$$\mathcal{A}^\dagger = \mathcal{A}^*\left(M^{-1} + \textstyle\sum_{i=1}^\infty (M^{-1}\Delta)^i M^{-1}\right) \tag{7}$$

where $\Delta = M - \mathcal{A}\mathcal{A}^*$. The series expansion is clearly valid if $\|M^{-1}\Delta\| < 1$. Straightforward calculations show that $M = (k^2 - k)I + k\mathbf{1}\mathbf{1}^T$ and that $\|M^{-1}\| \leq 2/k^2$ and $\|M^{-1}\|_{\infty\to\infty} \leq 6/k^2$. The next lemma, the proof of which involves computing the first few moments of Δ, controls $\|\Delta\|$ and hence convergence of the series.

Lemma 3: If $k \geq \sqrt{n}$ there are positive constants $\tilde{c}_4, c_4$ and $\bar{c}_4$ such that with probability at least $1 - \bar{c}_4 n e^{-\tilde{c}_4\sqrt{n}}$

$$\|\Delta\| \leq c_4 k n^{3/4}.$$

We omit the proof, referring the reader to [1, Section 3.9.4]. The series expansion is valid if $\|M^{-1}\|\|\Delta\| < 1$, which occurs w.h.p. as long as $n^{3/4}/k = o(1)$. Taking (7), applying the triangle inequality and $\|X\|_{\infty\to 2} \leq \sqrt{n}\|X\|$ (valid for $n \times n$ matrices) gives

$$\begin{aligned}\|\mathcal{A}^\dagger\|_{\infty\to\text{sp}} &\leq \|\mathcal{A}^*\|_{\infty\to\text{sp}}\|M^{-1}\|_{\infty\to\infty} + \\ &\qquad \sqrt{n}\|\mathcal{A}^*\|_{2\to\text{sp}}\frac{\|M^{-1}\|\|\Delta\|}{1-\|M^{-1}\|\|\Delta\|}\|M^{-1}\| \\ &\leq \frac{6}{k^2}\|\mathcal{A}^*\|_{\infty\to\text{sp}} + \frac{2\sqrt{n}}{k^2}\|\mathcal{A}^*\|_{2\to\text{sp}}\frac{2k^{-2}\|\Delta\|}{1-2k^{-2}\|\Delta\|}. \qquad (8)\end{aligned}$$

Note that we used the infinity norm for the first term in the series, and the Euclidean norm for the tail, (paying a price of $\sqrt{n}$ for doing so), as $\|\Delta\|_{\infty\to\infty}$ is not small enough to control the tail of the series.

Finally to bound $\|\mathcal{A}^\dagger\|_{\infty\to\text{sp}}$ we need to bound $\|\mathcal{A}^*\|_{\infty\to\text{sp}}$ and $\|\mathcal{A}^*\|_{2\to\text{sp}}$. The following lemmas give such bounds. The first is quite straightforward, essentially using the fact that $\mathcal{A}^*$ maps the positive orthant into the positive semidefinite cone. The second is more involved. Proofs are in [1, Sections 3.9.1 and 3.9.2].

Lemma 4: If $k = o(n)$, there are positive constants $\tilde{c}_1$ and c_1' such that $\|\mathcal{A}^*\|_{\infty\to\text{sp}} \leq c_1' n$ with probability at least $1 - 2e^{-\tilde{c}_1 n}$.

Lemma 5: With probability at least $1 - 2e^{-\frac{1}{2}(k+\sqrt{n})}$, $\|\mathcal{A}^*\|_{2\to\text{sp}} \leq 8(k + \sqrt{n})$.

By combining (8) with Lemmas 4 and 5, and keeping track of the dominant terms, we finally obtain the desired control on $\|\mathcal{A}\|_{\infty\to\text{sp}}$.

Proposition 3: If $k = o(n)$ and $n^{3/4}/k = o(1)$, there are positive constants $c_3, \bar{c}_3$ and $\tilde{c}_3$ such that $\|\mathcal{A}^\dagger\|_{\infty\to\text{sp}} \leq c_3 n^{5/4} k^{-2}$ with probability at least $1 - \bar{c}_3 n e^{-\tilde{c}_3\sqrt{n}}$.

To obtain a proof of Theorem 2 we combine Proposition 3 with Lemmas 1 and 2, to show that the deterministic condition of Proposition 2 is satisfied w.h.p. under the assumptions of Theorem 2. Indeed assume that $0 < \alpha < 1/6$ is fixed and $n = Ck^{\frac{6}{5}(1-\alpha)}$ where $C < (c_1/(c_2c_3))^{4/5}$ (here c_1, c_2, c_3 are as defined in Lemmas 1 and 2 and Proposition 3). Observe that this choice satisfies $n^{3/4}/k = o(1)$ and $k = o(n)$ as required by Proposition 3 and Lemma 1. Combining the previous probability estimates, there are positive constants $\bar{c}$ and $\tilde{c}$ such that with probability at least $1 - \bar{c}ne^{-\tilde{c}k^{3\alpha}}$

$$\begin{aligned}\lambda_{\min}(\mathcal{A}^*(\eta\mathbf{1})) - \|\mathcal{A}^\dagger\|_{\infty\to\text{sp}}&\|\mathbf{1} - \mathcal{A}\mathcal{A}^*(\eta\mathbf{1})\|_\infty \\ &> c_1 n - c_3 n^{5/4}k^{-2} \cdot c_2 n k^{1/2+3\alpha/2} \\ &= nc_1\left(1 - \frac{c_2c_3}{c_1}n^{5/4}k^{-3/2(1-\alpha)}\right).\end{aligned}$$

The right hand side is positive if $n^{5/4} < \left(\frac{c_1}{c_2c_3}\right)k^{3/2(1-\alpha)}$ which holds given our choice of n above.

For fixed k, the probability of fitting an ellipsoid to n Gaussian points in $\mathbb{R}^k$ is monotonically decreasing in n, so establishing the result for $n = Ck^{\frac{6}{5}(1-\alpha)}$ is sufficient to establish it for all $n \leq Ck^{\frac{6}{5}(1-\alpha)}$.

VI. Conclusion

In this paper we considered two equivalent problems: analyzing the behaviour of a convex optimization method (MTFA) for random instances of diagonal and low-rank matrix decompositions, and determining when it is possible w.h.p. to fit an ellipsoid to random Gaussian points. We established that it is possible to fit w.h.p. an ellipsoid to superlinearly (in k) many Gaussian points in $\mathbb{R}^k$, and presented linear evidence of much stronger scaling. We are hopeful that our main conjecture has a conceptual proof, and that the techniques required to establish it will prove useful for the analysis of recovery problems with structured but not isotropic measurement ensembles.

References

[1] J. Saunderson, "Subspace identification via convex optimization," Master's thesis, Massachusetts Institute of Technology, 2011.

[2] M. Fazel, H. Hindi, and S. P. Boyd, "A rank minimization heuristic with application to minimum order system approximation," in *Proc. 2001 American Control Conference*, vol. 6, 2001, pp. 4734–4739.

[3] Z. Liu and L. Vandenberghe, "Interior-point method for nuclear norm approximation with application to system identification," *SIAM J. Matrix Anal. Appl.*, vol. 31, no. 3, pp. 1235–1256, 2009.

[4] P. Shah, B. N. Bhaskar, G. Tang, and B. Recht, "Linear system identification via atomic norm regularization," in *51st IEEE Conf. on Decision and Control (CDC)*, 2012, pp. 6265–6270.

[5] B. Recht, M. Fazel, and P. Parrilo, "Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization," *SIAM Rev.*, vol. 52, no. 3, pp. 471–501, 2010.

[6] E. Candès and B. Recht, "Exact matrix completion via convex optimization," *Found. Comput. Math.*, vol. 9, no. 6, pp. 717–772, 2009.

[7] V. Chandrasekaran, S. Sanghavi, P. Parrilo, and A. Willsky, "Rank-sparsity incoherence for matrix decomposition," *SIAM J. Optim.*, vol. 21, pp. 572–596, 2011.

[8] M. Fazel, "Matrix rank minimization with applications," Ph.D. dissertation, Stanford University, 2002.

[9] R. Kalman, "Identification of noisy systems," *Russian Math. Surveys*, vol. 40, no. 4, pp. 25–42, 1985.

[10] C. Spearman, "'General intelligence,' objectively determined and measured," *American J. Psychol.*, pp. 201–292, 1904.

[11] W. Ledermann, "On a problem concerning matrices with variable diagonal elements," *Proc. Roy. Soc. Edinburgh*, vol. 60, pp. 1–17, 1940.

[12] A. Shapiro, "Identifiability of factor analysis: some results and open problems," *Linear Algebra Appl.*, vol. 70, pp. 1–7, 1985.

[13] J. Saunderson, V. Chandrasekaran, P. Parrilo, and A. Willsky, "Diagonal and low-rank matrix decompositions, correlation matrices, and ellipsoid fitting," *SIAM J. Matrix Anal. Appl.*, vol. 33, no. 4, pp. 1395–1416, 2012.

[14] G. Della Riccia and A. Shapiro, "Minimum rank and minimum trace of covariance matrices," *Psychometrika*, vol. 47, no. 4, pp. 443–448, 1982.

[15] R. Vershynin, "Introduction to the non-asymptotic analysis of random matrices," *arXiv preprint arXiv:1011.3027*, 2010.