

PhaseLift: Exact and Stable Signal Recovery from Magnitude Measurements via Convex Programming

Emmanuel J. Candès*, Thomas Strohmer† and Vladislav Voroninski‡

September 2011

Abstract

Suppose we wish to recover a signal $\mathbf{x} \in \mathbb{C}^n$ from m intensity measurements of the form $|\langle \mathbf{x}, \mathbf{z}_i \rangle|^2$, $i = 1, 2, \dots, m$; that is, from data in which phase information is missing. We prove that if the vectors \mathbf{z}_i are sampled independently and uniformly at random on the unit sphere, then the signal \mathbf{x} can be recovered exactly (up to a global phase factor) by solving a convenient semidefinite program—a trace-norm minimization problem; this holds with large probability provided that m is on the order of $n \log n$, and without any assumption about the signal whatsoever. This novel result demonstrates that in some instances, the combinatorial phase retrieval problem can be solved by convex programming techniques. Finally, we also prove that our methodology is robust vis a vis additive noise.

1 Introduction

In many applications, one would like to acquire information about an object but it is impossible or very difficult to measure and record the phase of the signal. The problem is then to reconstruct the object from intensity measurements only. A problem of this kind that has attracted a considerable amount of attention over the last hundred years or so, is of course that of recovering a signal or image from the intensity measurements of its Fourier transform [15, 16] as in X-ray crystallography. As is well-known, such phase retrieval problems are notoriously difficult to solve numerically.

Formally, suppose $\mathbf{x} \in \mathbb{C}^n$ is a discrete signal and that we are given information about the squared modulus of the inner product between the signal and some vectors \mathbf{z}_i , namely,

$$b_i = |\langle \mathbf{x}, \mathbf{z}_i \rangle|^2, \quad i = 1, \dots, m. \quad (1.1)$$

In truth, we would like to know $\langle \mathbf{x}, \mathbf{z}_i \rangle$ and record both phase and magnitude information but can only record the magnitude; in other words, phase information is lost. In the classical example discussed above, the \mathbf{z}_i 's are complex exponentials at frequency ω_i so that one collects the squared

*Departments of Mathematics and of Statistics, Stanford University, Stanford CA 94305

†Department of Mathematics, University of California at Davis, Davis CA

‡Department of Mathematics, University of California at Berkeley, Berkeley CA

modulus of the Fourier transform of \mathbf{x} . Of course, many other choices for the measurement vectors \mathbf{z}_i are frequently discussed in the literature, see [2, 12] for instance.

We wish to recover \mathbf{x} from the data vector \mathbf{b} , and suppose first that \mathbf{x} is known to be real valued a priori. Then assuming that \mathbf{x} is uniquely determined by \mathbf{b} up to a global sign, the recovery may be cast as a combinatorial optimization problem: find a set of signs σ_i such that the solution to the linear equations $\langle \mathbf{x}, \mathbf{z}_i \rangle = \sigma_i \sqrt{b_i}$, call it $\hat{\mathbf{x}}$, obeys $|\langle \hat{\mathbf{x}}, \mathbf{z}_i \rangle|^2 = b_i$. Clearly, there are 2^m choices for σ_i and only two choices of these signs yield \mathbf{x} up to global phase. The complex case is harder yet, since resolving the phase ambiguities now consists of finding a collection σ_i of complex numbers, each being on the unit circle. Formalizing matters, it has been shown that at least one version of the phase retrieval problem is NP-hard [20]. Thus, one of the major challenges in the field is to find conditions on m and \mathbf{z}_i which guarantee efficient numerical recovery.

A frame-theoretic approach to signal recovery from magnitude measurements has been proposed in [1–3], where the authors derive various necessary and sufficient conditions for the uniqueness of the solution, as well as various polynomial-time numerical algorithms for very specific choices of \mathbf{z}_i . While theoretically quite appealing, the drawbacks are that the methods are (1) either algebraic in nature, thus severely limiting their stability in the presence of noise or slightly inexact data, or (2) the number m of measurements is on the order of n^2 , which is much too large compared to the number of unknowns.

This paper follows a very different route and establishes that if the vectors \mathbf{z}_i are independently and uniformly sampled on the unit sphere, then our signal can be recovered exactly from the magnitude measurements (1.1) by solving a simple convex program we introduce below; this holds with high probability with the proviso that the number of measurements is on the order of $n \log n$. Since there are n complex unknowns, we see that the number of samples is nearly minimal. To the best of our knowledge, this is the first result establishing that under appropriate conditions, the computationally challenging nonconvex problem of reconstructing a signal from magnitude measurements is formally equivalent to a convex program in the sense that they are guaranteed to have the same unique solution.

Finally, our methodology is robust with respect to noise in the measurements. To be sure, when the data are corrupted by a small amount of noise, we also prove that the recovery error is small.

1.1 Methodology

We introduce some notation that shall be used throughout to explain our methodology. Letting \mathcal{A} be the linear transformation

$$\begin{aligned} \mathcal{H}^{n \times n} &\rightarrow \mathbb{R}^m \\ \mathbf{X} &\mapsto \{\mathbf{z}_i^* \mathbf{X} \mathbf{z}_i\}_{1 \leq i \leq m} \end{aligned} \quad (1.2)$$

which maps Hermitian matrices into real-valued vectors, one can express the data collection $b_i = |\langle \mathbf{x}, \mathbf{z}_i \rangle|^2$ as

$$\mathbf{b} = \mathcal{A}(\mathbf{x}\mathbf{x}^*). \quad (1.3)$$

For reference, the adjoint operator \mathcal{A}^* maps real-valued inputs into Hermitian matrices, and is given by

$$\begin{aligned} \mathbb{R}^m &\rightarrow \mathcal{H}^{n \times n} \\ \mathbf{y} &\mapsto \sum_i y_i \mathbf{z}_i \mathbf{z}_i^*. \end{aligned}$$

As observed in [7, 10] (see also [17]), the phase retrieval problem can be cast as the matrix recovery problem

$$\begin{aligned} & \text{minimize} && \text{rank}(\mathbf{X}) \\ & \text{subject to} && \mathcal{A}(\mathbf{X}) = \mathbf{b} \\ & && \mathbf{X} \succeq 0. \end{aligned} \tag{1.4}$$

Indeed, we know that a rank-one solution exists so the optimal \mathbf{X} has rank at most one. We then factorize the solution as $\mathbf{x}\mathbf{x}^*$ in order to obtain solutions to the phase-retrieval problem. This gives \mathbf{x} up to multiplication by a unit-normed scalar. This is all we can hope for since if \mathbf{x} is a solution to the phase retrieval problem, then $c\mathbf{x}$ for any scalar $c \in \mathbb{C}$ obeying $|c| = 1$ is also solution.¹

Rank minimization is in general NP hard, and we propose, instead, solving a trace-norm relaxation. Although this is a fairly standard relaxation in control [4, 18], the idea of casting the phase retrieval problem as a trace-minimization problem over an affine slice of the positive semidefinite cone is very recent [7, 10]. Formally, we suggest solving

$$\begin{aligned} & \text{minimize} && \text{Tr}(\mathbf{X}) \\ & \text{subject to} && \mathcal{A}(\mathbf{X}) = \mathbf{b} \\ & && \mathbf{X} \succeq 0. \end{aligned} \tag{1.5}$$

If the solution has rank one, we factorize it as above to recover our signal. This method which lifts up the problem of vector recovery from quadratic constraints into that of recovering a rank-one matrix from affine constraints via semidefinite programming is known under the name of *PhaseLift* [7].

The program (1.5) is a semidefinite program (SDP) in standard form, and there is a rapidly growing list of algorithms for solving problems of this kind as efficiently as possible. The crucial question is whether and under which conditions the combinatorially hard problem (1.4) and the convex problem (1.5) are formally equivalent.

1.2 Main result

In this paper, we consider the simplest and perhaps most natural model of measurement vectors. In this statistical model, we simply assume that the vectors \mathbf{z}_i are independently and uniformly distributed on the unit sphere of \mathbb{C}^n or \mathbb{R}^n . To be concrete, we distinguish two models.

- *The real-valued model.* Here, the unknown signal \mathbf{x} is real valued and the \mathbf{z}_i 's are independently sampled on the unit sphere of \mathbb{R}^n .
- *The complex-valued model.* The signal \mathbf{x} is now complex valued and the \mathbf{z}_i 's are independently sampled on the unit sphere of \mathbb{C}^n .

Our main result is that the convex program recovers \mathbf{x} exactly (up to global phase) provided the number m of magnitude measurements is on the order of $n \log n$.

Theorem 1.1 *Consider an arbitrary signal \mathbf{x} in \mathbb{R}^n or \mathbb{C}^n and suppose that the number of measurements obeys $m \geq c_0 n \log n$, where c_0 is a sufficiently large constant. Then in both the real and*

¹When the solution is unique up to multiplication by such a scalar, we shall say that unicity holds up to global phase.

complex cases, the solution to the trace-minimization program is exact with high probability in the sense that (1.5) has a unique solution obeying

$$\hat{\mathbf{X}} = \mathbf{x}\mathbf{x}^*. \quad (1.6)$$

This holds with probability at least $1 - 3e^{-\gamma \frac{m}{n}}$, where γ is a positive absolute constant.

Expressed differently, Theorem 1.1 establishes a rigorous equivalence between a class of phase retrieval problems and a class of semidefinite programs. Clearly, any phase retrieval algorithm, no matter how complicated or intractable, would need at least $2n$ quadratic measurements to recover a complex valued object $\mathbf{x} \in \mathbb{C}^n$. In fact recent results, compare Theorem II in [12], show that for complex-valued signals, one needs at least $3n - 2$ intensity measurements to guarantee uniqueness of the solution to (1.4). Further, Balan, Casazza and Edidin have shown that with probability 1, $4n - 2$ generic measurement vectors (which includes the case of random uniform vectors) suffice for uniqueness in the complex case [3]. Hence, Theorem 1.1 shows that the oversampling factor for perfect recovery via convex optimization is rather minimal.

To be absolutely complete, we would like to emphasize that our discrete signals \mathbf{x} may represent 1D, 2D, 3D and higher dimensional objects. For instance, in 2D the vector $\mathbf{x} \in \mathbb{C}^n$ might be a family of samples of the form $x[t_1, t_2]$, $1 \leq t_1 \leq n_1$, $1 \leq t_2 \leq n_2$, and with $n = n_1 n_2$, so that \mathbf{x} is a discrete 2D image. In this case, we would record the squared magnitudes of the dot product

$$\langle \mathbf{x}, \mathbf{z}_i \rangle = \sum_{t_1, t_2} \bar{x}[t_1, t_2] z_i[t_1, t_2].$$

Hence, our framework and theory apply to one- or multi-dimensional signals.

1.3 Geometry

We find it rather remarkable that the only solution to (1.5) is $\hat{\mathbf{X}} = \mathbf{x}\mathbf{x}^*$. To see why this is perhaps unexpected, suppose for simplicity that the trace of the solution were known (we might be given some side information or just have additional measurements giving us this information) and equal to 1, say. In this case, the objective functional is of course constant over the feasible set, and our problem reduces to solving the feasibility problem

$$\begin{aligned} &\text{find} && \mathbf{X} \\ &\text{such that} && \mathcal{A}(\mathbf{X}) = \mathbf{b}, \mathbf{X} \succeq 0 \end{aligned} \quad (1.7)$$

with again the proviso that knowledge of $\mathcal{A}(\mathbf{X})$ determines $\text{Tr}(\mathbf{X})$ (equal to $\text{Tr}(\mathbf{x}\mathbf{x}^*) = \|\mathbf{x}\|_2 = 1$). In this context, our main theorem states that $\mathbf{x}\mathbf{x}^*$ is the unique feasible point. In other words, there is no other positive semidefinite matrix \mathbf{X} in the affine space $\mathcal{A}(\mathbf{X}) = \mathbf{b}$. Naively, we would not expect this affine space of enormous dimension—it is of co-dimension about $n \log n$ and thus of dimension $n^2 - O(n \log n)$ in the complex case—to intersect the positive semidefinite cone in only one point. Indeed, counting degrees of freedom suggests that there are infinitely many candidates in the intersection. The reason why this is not the case, however, is precisely because there is a feasible solution with low rank. Indeed, the slice of the positive semidefinite cone $\{\mathbf{X} : \mathbf{X} \succeq 0\} \cap \{\text{Tr}(\mathbf{X}) = 1\}$ is quite ‘pointy’ at $\mathbf{x}\mathbf{x}^*$ and it is, therefore, possible for the affine space $\{\mathcal{A}(\mathbf{X}) = \mathbf{b}\}$ to be tangent even though it is of very small codimension.




Figure 1: Representation of the affine space $\mathcal{A}(\mathbf{X}) = \mathbf{b}$ (gray) and of the semidefinite cone $\begin{bmatrix} x & y \\ y & z \end{bmatrix} \succeq 0$ (red) which is a subset of \mathbb{R}^3 . These two sets are drawn so that they are tangent to each other at the rank 1 matrix $\frac{1}{2} \begin{bmatrix} 1 & -1 \\ -1 & 1 \end{bmatrix}$ (black dot). Two views of the same 3D figure are provided for convenience.

Figure 1 represents this geometry. In this example,

$$\mathbf{x} = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ -1 \end{bmatrix} \implies \mathbf{x}\mathbf{x}^* = \frac{1}{2} \begin{bmatrix} 1 & -1 \\ -1 & 1 \end{bmatrix}$$

and the affine space $\mathcal{A}(\mathbf{X}) = \mathbf{b}$ is tangent to the positive semidefinite cone at the point $\mathbf{x}\mathbf{x}^*$.

Mathematically speaking, phase retrieval is a problem in algebraic geometry since we are trying to find a solution to a set of polynomial equations. The originality in our approach is that we do not use tools from this field. For instance, we prove that there is no other positive semidefinite matrix \mathbf{X} in the affine space $\mathcal{A}(\mathbf{X}) = \mathbf{b}$, or equivalently, that a certain system of polynomial equations (a symmetric matrix is positive semidefinite if and only if the determinants of all the leading principal minors are nonnegative) only has one solution; this is a fact that general techniques from algebraic geometry appear to not detect.

1.4 Stability

In the real world, measurements are contaminated by noise. Using the frameworks developed in [8] and [14], it is possible to extend Theorem 1.1 to accommodate noisy measurements. One could consider a variety of noise models as discussed in [7] but we work here with a simple generic model in which we observe

$$b_i = |\langle \mathbf{x}, \mathbf{z}_i \rangle|^2 + \nu_i, \tag{1.8}$$

where ν_i is a noise term with bounded ℓ_2 norm, $\|\boldsymbol{\nu}\|_2 \leq \epsilon$. This model is nonstandard since the usual statistical linear model posits a relationship of the form $b_i = \langle \mathbf{x}, \mathbf{z}_i \rangle + \nu_i$ in which the mean response is a linear function of the unknown signal, not a quadratic function. Furthermore, we prefer studying (1.8) rather than the related model $b_i = |\langle \mathbf{x}, \mathbf{z}_i \rangle| + \nu_i$ (the modulus is not squared) because in many applications of interest in optics and other areas of physics, one can measure squared magnitudes or intensities—not magnitudes.

We now consider the solution to

$$\begin{aligned} & \text{minimize} && \text{Tr}(\mathbf{X}) \\ & \text{subject to} && \|\mathcal{A}(\mathbf{X}) - \mathbf{b}\|_2 \leq \epsilon \\ & && \mathbf{X} \succeq 0. \end{aligned} \tag{1.9}$$

We do not claim that $\hat{\mathbf{X}}$ has low rank so we suggest estimating \mathbf{x} by extracting the largest rank-1 component. Write $\hat{\mathbf{X}}$ as

$$\hat{\mathbf{X}} = \sum_{k=1}^n \hat{\lambda}_k \hat{\mathbf{u}}_k \hat{\mathbf{u}}_k^*, \quad \hat{\lambda}_1 \geq \dots \geq \hat{\lambda}_n \geq 0,$$

and set

$$\hat{\mathbf{x}} = \sqrt{\hat{\lambda}_1} \hat{\mathbf{u}}_1.$$

We prove the following estimate.

Theorem 1.2 *Fix $\mathbf{x} \in \mathbb{C}^n$ or \mathbb{R}^n and assume the \mathbf{z}_i 's are uniformly sampled on the sphere of radius \sqrt{n} . Under the hypotheses of Theorem 1.1, the solution to (1.9) obeys ($\|\mathbf{X}\|_2$ is the Frobenius norm of \mathbf{X})*

$$\|\hat{\mathbf{X}} - \mathbf{x}\mathbf{x}^*\|_2 \leq C_0 \epsilon \tag{1.10}$$

for some positive numerical constant C_0 . We also have

$$\|\hat{\mathbf{x}} - e^{i\phi} \mathbf{x}\|_2 \leq C_0 \min(\|\mathbf{x}\|_2, \epsilon/\|\mathbf{x}\|_2) \tag{1.11}$$

for some $\phi \in [0, 2\pi]$. Both these estimates hold with nearly the same probability as in the noiseless case.

Thus our approach also provides stable recovery in presence of noise. This important property is not shared by other reconstruction methods, which are of a more algebraic nature and rely on particular properties of the measurement vectors, such as the methods in [2, 3, 12], as well as the methods that appear implicitly in Theorem 3.1 and Theorem 3.3 of [7].

We note that one can further improve the accuracy of the solution $\hat{\mathbf{x}}$ by “debiasing” it. We replace $\hat{\mathbf{x}}$ by its rescaled version $s\hat{\mathbf{x}}$ where $s = \sqrt{\sum_{k=1}^n \hat{\lambda}_k} / \|\hat{\mathbf{x}}\|_2$. This corrects for the energy leakage occurring when $\hat{\mathbf{X}}$ is not exactly a rank-1 solution, which could cause the norm of $\hat{\mathbf{x}}$ to be smaller than that of the actual solution. Other corrections are of course possible.

1.5 Organization of the paper

The remainder of the paper is organized as follows. Subsection 1.6 introduces some notation used throughout the paper. In Section 2 we present the main architecture of the proof of Theorem 1.1, which comprises two key ingredients: approximate ℓ_1 isometries and approximate dual certificates. Section 3 is devoted to establishing approximate ℓ_1 isometries. In Section 4, we construct approximate dual certificates and complete the proof of Theorem 1.1 in the real-valued case. Section 5 shows how the proof for the real-valued case can be adapted to the complex-valued case. Section 6 is concerned with the proof of Theorem 1.2. Numerical simulations, illustrating our theoretical results, are presented in Section 7. We conclude the paper with a short discussion in Section 8.

1.6 Notations

It is useful to introduce notations that shall be used throughout the paper. Matrices and vectors are denoted in boldface (such as \mathbf{X} or \mathbf{x}), while individual entries of a vector or matrix are denoted in normal font; e.g. the i th entry of \mathbf{x} is x_i . For matrices, we define

$$\|\mathbf{X}\|_p = \left[\sum_i \sigma_i^p(\mathbf{X}) \right]^{1/p},$$

(where $\sigma_i(\mathbf{X})$ denotes the i th singular value of \mathbf{X}), so that $\|\mathbf{X}\|_1$ is the nuclear norm, $\|\mathbf{X}\|_2$ is the Frobenius norm and $\|\mathbf{X}\|_\infty$ is the operator norm also denoted by $\|\mathbf{X}\|$. For vectors, $\|\mathbf{x}\|_p$ is the usual ℓ_p norm. We denote the $n - 1$ dimensional sphere by S^{n-1} , i.e. the set $\{\mathbf{x} \in \mathbb{R}^n : \|\mathbf{x}\|_2 = 1\}$.

Next, we define $T_{\mathbf{x}}$ to be the set of symmetric matrices of the form

$$T_{\mathbf{x}} = \{\mathbf{X} = \mathbf{x}\mathbf{y}^* + \mathbf{y}\mathbf{x}^* : \mathbf{y} \in \mathbb{R}^n\} \quad (1.12)$$

and denote $T_{\mathbf{x}}^\perp$ by its orthogonal complement. Note that $\mathbf{X} \in T_{\mathbf{x}}^\perp$ if and only if both the column and row spaces of \mathbf{X} are perpendicular to \mathbf{x} . Further, the operator $\mathcal{P}_{T_{\mathbf{x}}}$ is the orthogonal projector onto $T_{\mathbf{x}}$ and similarly for $\mathcal{P}_{T_{\mathbf{x}}^\perp}$. We shall almost always use $\mathbf{X}_{T_{\mathbf{x}}}$ as a shorthand for $\mathcal{P}_{T_{\mathbf{x}}}(\mathbf{X})$.

Finally, we will abuse language and say that a symmetric matrix \mathbf{H} is feasible if and only if $\mathbf{x}\mathbf{x}^* + \mathbf{H}$ is feasible for our problem (1.5). This means that \mathbf{H} obeys

$$\mathbf{x}\mathbf{x}^* + \mathbf{H} \succeq 0 \quad \text{and} \quad \mathcal{A}(\mathbf{H}) = 0. \quad (1.13)$$

2 Architecture of the Proof

In this section, we introduce the main architecture of the argument and defer the proofs of crucial intermediate results to later sections. We shall prove Theorem 1.1 in the real case first for ease of exposition. Then in Section 5, we shall explain how to modify the argument to the complex and more general case.

Suppose then that $\mathbf{x} \in \mathbb{R}^n$ and that the \mathbf{z}_i 's are sampled on the unit sphere. It is clear that we may assume without loss of generality that \mathbf{x} is unit-normed. Further, since the uniform distribution on the unit sphere is rotationally invariant, it suffices to prove the theorem in the case where $\mathbf{x} = \mathbf{e}_1$. Indeed, we can write any unit vector \mathbf{x} as $\mathbf{x} = \mathbf{U}\mathbf{e}_1$ where \mathbf{U} is orthogonal. Since

$$|\langle \mathbf{x}, \mathbf{z}_i \rangle|^2 = |\langle \mathbf{U}\mathbf{e}_1, \mathbf{z}_i \rangle|^2 = |\langle \mathbf{e}_1, \mathbf{U}^* \mathbf{z}_i \rangle|^2 = |\langle \mathbf{e}_1, \mathbf{z}_i \rangle|^2,$$

the problem is the same as that of finding \mathbf{e}_1 . We henceforth assume that $\mathbf{x} = \mathbf{e}_1$.

Finally, the theorem can be equivalently stated in the case where the \mathbf{z}_i 's are i.i.d. copies of a white noise vector $\mathbf{z} \sim \mathcal{N}(0, I)$ with independent standard normals as components. Indeed, if $\mathbf{z}_i \sim \mathcal{N}(0, I)$,

$$|\langle \mathbf{x}, \mathbf{z}_i \rangle|^2 = b_i \iff |\langle \mathbf{x}, \mathbf{u}_i \rangle|^2 = b_i / \|\mathbf{z}_i\|_2^2,$$

where $\mathbf{u}_i = \mathbf{z}_i / \|\mathbf{z}_i\|_2$ is uniformly sampled on the unit sphere. Since $\|\mathbf{z}_i\|_2$ does not vanish with probability one, establishing the theorem for Gaussian vectors establishes it for uniformly sampled vectors and vice versa. From now on, we assume \mathbf{z}_i i.i.d. $\mathcal{N}(0, I)$.

2.1 Key lemma

The set $T := T_{\mathbf{e}_1}$ defined in (1.12) may be interpreted as the tangent space at $\mathbf{e}_1\mathbf{e}_1^*$ to the manifold of symmetric matrices of rank 1. Now standard duality arguments in semidefinite programming show that a sufficient (and nearly necessary) condition for $\mathbf{x}\mathbf{x}^*$ to be the unique solution to (1.5) is this:

- the restriction of \mathcal{A} to T is injective ($\mathbf{X} \in T$ and $\mathcal{A}(\mathbf{X}) = 0 \Rightarrow \mathbf{X} = 0$),
- and there exists a *dual certificate* \mathbf{Y} in the range of \mathcal{A}^* obeying²

$$\mathbf{Y}_T = \mathbf{e}_1\mathbf{e}_1^* \quad \text{and} \quad \mathbf{Y}_T^\perp \prec I_T^\perp. \quad (2.1)$$

The proof is straightforward and omitted. Our strategy to prove Theorem 1.1 hinges on the fact that a strengthening of the injectivity property allows to relax the properties of the dual certificate, as in the approach pioneered in [13] for matrix completion. We establish the crucial lemma below.

Lemma 2.1 *Suppose that the mapping \mathcal{A} obeys the following two properties: for all positive semidefinite matrices \mathbf{X} ,*

$$m^{-1}\|\mathcal{A}(\mathbf{X})\|_1 < (1 + 1/9)\|\mathbf{X}\|_1; \quad (2.2)$$

and for all matrices $\mathbf{X} \in T$

$$m^{-1}\|\mathcal{A}(\mathbf{X})\|_1 > 0.94(1 - 1/9)\|\mathbf{X}\|. \quad (2.3)$$

Suppose further that there exists \mathbf{Y} in the range of \mathcal{A}^* obeying

$$\|\mathbf{Y}_T - \mathbf{e}_1\mathbf{e}_1^*\|_2 \leq 1/3 \quad \text{and} \quad \|\mathbf{Y}_T^\perp\| \leq 1/2. \quad (2.4)$$

Then $\mathbf{e}_1\mathbf{e}_1^*$ is the unique minimizer to (1.5).

The first property (2.2) is reminiscent of the (one-sided) RIP property in the area of compressed sensing [9]. The difference is that it is expressed in the 1-norm rather than the 2-norm. Having said this, we note that RIP-1 properties have also been used in the compressed sensing literature, see [6] for example. We use this property instead of a property about $\|\mathcal{A}(\mathbf{X})\|_2$, because we actually believe that a RIP property in the 2-norm does not hold here because $\|\mathcal{A}(\mathbf{X})\|_2^2$ involves fourth moments of Gaussian variables. The second property (2.3) is a form of local RIP-1 since it holds only for matrices in T .

We would like to emphasize that the bound for the dual certificate in (2.4) is loose in the sense that \mathbf{Y}_T and $\mathbf{e}_1\mathbf{e}_1^*$ may not be that close, a fact which will play a crucial role in our proof. This is in stark contrast with the work of David Gross [13], which requires a very tight approximation.

²The notation $A \prec B$ means that $B - A$ is positive definite.

2.2 Proof of Lemma 2.1

We need to show that there is no feasible $\mathbf{x}\mathbf{x}^* + \mathbf{H} \neq \mathbf{x}\mathbf{x}^*$ with $\text{Tr}(\mathbf{x}\mathbf{x}^* + \mathbf{H}) \leq \text{Tr}(\mathbf{x}\mathbf{x}^*)$. Consider then a feasible $\mathbf{H} \neq 0$ obeying $\text{Tr}(\mathbf{H}) \leq 0$, write

$$\mathbf{H} = \mathbf{H}_T + \mathbf{H}_T^\perp,$$

and observe that

$$0 = \|\mathcal{A}(\mathbf{H})\|_1 = \|\mathcal{A}(\mathbf{H}_T)\|_1 - \|\mathcal{A}(\mathbf{H}_T^\perp)\|_1. \quad (2.5)$$

Now it is clear that $\mathbf{x}\mathbf{x}^* + \mathbf{H} \succeq 0 \Rightarrow \mathbf{H}_T^\perp \succeq 0$ and, therefore, (2.2) gives

$$m^{-1} \|\mathcal{A}(\mathbf{H}_T^\perp)\|_1 \leq (1 + \delta) \text{Tr}(\mathbf{H}_T^\perp)$$

for some $\delta < 1/9$. Also, $\text{Tr}(\mathbf{H}_T) \leq -\text{Tr}(\mathbf{H}_T^\perp) \leq 0$, which implies that $|\text{Tr}(\mathbf{H}_T)| \geq \text{Tr}(\mathbf{H}_T^\perp)$. We then show that the operator and Frobenius norms of \mathbf{H}_T must nearly be the same.

Lemma 2.2 *Any feasible matrix \mathbf{H} such that $\text{Tr}(\mathbf{H}) \leq 0$ must obey*

$$\|\mathbf{H}_T\|_2 \leq \sqrt{\frac{17}{16}} \|\mathbf{H}_T\|.$$

Proof Since the matrix \mathbf{H}_T has rank at most 2 and cannot be negative definite, it is of the form

$$-\lambda(\mathbf{u}_1\mathbf{u}_1^* - t\mathbf{u}_2\mathbf{u}_2^*),$$

where \mathbf{u}_1 and \mathbf{u}_2 are orthonormal eigenvectors, $\lambda \geq 0$ and $t \in [0, 1]$. We claim that we cannot have $t \geq 1/4$.³ Suppose the contrary and fix $t \geq 1/4$. By (2.3), we know that

$$m^{-1} \|\mathcal{A}(\mathbf{H}_T)\|_1 \geq 0.94(1 - \delta) \|\mathbf{H}_T\|.$$

Further, since

$$\|\mathbf{H}_T\| = \frac{|\text{Tr}(\mathbf{H}_T)|}{1 - t} \geq \frac{4}{3} |\text{Tr}(\mathbf{H}_T)|$$

for $t \geq 1/4$, it holds that

$$0 \geq \frac{5}{4}(1 - \delta) |\text{Tr}(\mathbf{H}_T)| - (1 + \delta) \text{Tr}(\mathbf{H}_T^\perp).$$

The right-hand side above is positive if $\text{Tr}(\mathbf{H}_T^\perp) < \frac{5(1-\delta)}{4(1+\delta)} |\text{Tr}(\mathbf{H}_T)|$, so that we may assume that

$$\text{Tr}(\mathbf{H}_T^\perp) \geq \frac{5(1-\delta)}{4(1+\delta)} |\text{Tr}(\mathbf{H}_T)|.$$

Since, $|\text{Tr}(\mathbf{H}_T)| \geq \text{Tr}(\mathbf{H}_T^\perp)$, this gives

$$0 \geq \left[\frac{5}{4}(1 - \delta) - (1 + \delta) \right] \text{Tr}(\mathbf{H}_T^\perp).$$

³The choice of 1/4 is somewhat arbitrary here.

If $\delta < 1/9$, the only way this can happen is if $\text{Tr}(\mathbf{H}_T^\perp) = 0 \Rightarrow \mathbf{H}_T^\perp = 0$. So we would have $\mathbf{H} = \mathbf{H}_T$ of rank 2 and $\mathcal{A}(\mathbf{H}_T) = 0$. Clearly, (2.3) implies that $\mathbf{H} = 0$.

Now that it is established that $t \leq 1/4$, the chain of inequalities follow from the relation between the eigenvalues of \mathbf{H}_T . \blacksquare

To conclude the proof of Lemma 2.1, we show that the existence of an inexact dual certificate rules out the existence of matrices obeying the conditions of Lemma 2.2. From

$$0.94(1 - \delta)\|\mathbf{H}_T\| \leq m^{-1}\|\mathcal{A}(\mathbf{H}_T)\|_1 = m^{-1}\|\mathcal{A}(\mathbf{H}_T^\perp)\|_1 \leq (1 + \delta)\text{Tr}(\mathbf{H}_T^\perp),$$

we conclude that

$$\text{Tr}(\mathbf{H}_T^\perp) \geq 0.94\frac{1 - \delta}{1 + \delta}\|\mathbf{H}_T\| \geq 0.94\frac{1 - \delta}{1 + \delta}\sqrt{\frac{16}{17}}\|\mathbf{H}_T\|_2, \quad (2.6)$$

where we used Lemma 2.2. Next,

$$\begin{aligned} 0 &\geq \text{Tr}(\mathbf{H}_T) + \text{Tr}(\mathbf{H}_T^\perp) = \langle \mathbf{H}, \mathbf{e}_1 \mathbf{e}_1^* \rangle + \text{Tr}(\mathbf{H}_T^\perp) \\ &= \langle \mathbf{H}, \mathbf{e}_1 \mathbf{e}_1^* - \mathbf{Y} \rangle + \langle \mathbf{H}, \mathbf{Y} \rangle + \text{Tr}(\mathbf{H}_T^\perp) \\ &= \langle \mathbf{H}_T, \mathbf{e}_1 \mathbf{e}_1^* - \mathbf{Y}_T \rangle - \langle \mathbf{H}_T^\perp, \mathbf{Y}_T^\perp \rangle + \text{Tr}(\mathbf{H}_T^\perp) \\ &\geq \frac{1}{2}\text{Tr}(\mathbf{H}_T^\perp) - \frac{1}{3}\|\mathbf{H}_T\|_2. \end{aligned}$$

The third line above follows from $\langle \mathbf{H}, \mathbf{Y} \rangle = 0$ and the fourth from Cauchy-Schwarz together with $|\langle \mathbf{H}_T^\perp, \mathbf{Y}_T^\perp \rangle| \leq \frac{1}{2}\text{Tr}(\mathbf{H}_T^\perp)$. Hence, it follows from (2.6) that

$$0 \geq \frac{1}{2}\left(0.94\frac{1 - \delta}{1 + \delta}\sqrt{\frac{16}{17}} - \frac{2}{3}\right)\|\mathbf{H}_T\|_2.$$

Since the numerical factor is positive for $\delta < 0.155$, the only way this can happen is if $\mathbf{H}_T = 0$. In turn, $\|\mathcal{A}(\mathbf{H}_T^\perp)\|_1 = 0 \geq (1 - \delta)\text{Tr}(\mathbf{H}_T^\perp)$ which gives $\mathbf{H}_T^\perp = 0$. This concludes the proof.

3 Approximate ℓ_1 Isometries

We have seen that in order to prove our main result, it suffices to show 1) that the measurement operator \mathcal{A} enjoys approximate isometry properties (in an ℓ_1 sense) when acting on low-rank matrices and 2) that an inexact dual certificate exists. This section focuses on the former and establishes that both (2.2) and (2.3) hold with high probability. In fact, we shall prove stronger results than what is strictly required.

Lemma 3.1 *Fix any δ in $(0, 1/2)$ and assume $m \geq 20\delta^{-2}n$. Then for all unit vectors \mathbf{u} ,*

$$(1 - \delta) \leq \frac{1}{m}\|\mathcal{A}(\mathbf{u}\mathbf{u}^*)\|_1 \leq (1 + \delta) \quad (3.1)$$

on an event E_δ of probability at least $1 - 2e^{-m\epsilon^2/2}$, where $\delta/4 = \epsilon^2 + \epsilon$. On the same event,

$$(1 - \delta)\|\mathbf{X}\|_1 \leq \frac{1}{m}\|\mathcal{A}(\mathbf{X})\|_1 \leq (1 + \delta)\|\mathbf{X}\|_1$$

for all positive semidefinite matrices. The right inequality holds for all Hermitian matrices.

Proof This lemma has an easy proof. Let \mathbf{Z} be the $m \times n$ matrix with \mathbf{z}_i 's as rows. Then

$$\|\mathcal{A}(\mathbf{u}\mathbf{u}^*)\|_1 = \sum_i |\langle \mathbf{z}_i, \mathbf{u} \rangle|^2 = \|\mathbf{Z}\mathbf{u}\|^2$$

so that

$$\sigma_{\min}^2(\mathbf{Z}) \leq \|\mathcal{A}(\mathbf{u}\mathbf{u}^*)\|_1 \leq \sigma_{\max}^2(\mathbf{Z}).$$

The claim is a consequence of well-known deviations bounds concerning the singular values of Gaussian random matrices [21], namely,

$$\begin{aligned} \mathbb{P}(\sigma_{\max}(\mathbf{Z}) > \sqrt{m} + \sqrt{n} + t) &\leq e^{-t^2/2} \\ \mathbb{P}(\sigma_{\min}(\mathbf{Z}) < \sqrt{m} - \sqrt{n} - t) &\leq e^{-t^2/2}. \end{aligned}$$

The conclusion follows from taking $m \geq \epsilon^{-2}n$ and $t = \sqrt{m}\epsilon$ (and from $\epsilon^2 \geq \delta^2/20$ for $0 < \delta \leq 1/2$). For the second part of the lemma, observe that $\mathbf{X} = \sum_j \lambda_j \mathbf{u}_j \mathbf{u}_j^*$ with nonnegative eigenvalues λ_j so that

$$\|\mathcal{A}(\mathbf{X})\|_1 = \sum_j \sum_i \lambda_j |\langle \mathbf{u}_j, \mathbf{z}_i \rangle|^2 = \sum_j \lambda_j \|\mathcal{A}(\mathbf{u}_j \mathbf{u}_j^*)\|_1.$$

The claim follows from (3.1). The last claim is a consequence of $\|\mathcal{A}(\mathbf{X})\|_1 \leq \sum_j \sum_i |\lambda_j| |\langle \mathbf{u}_j, \mathbf{z}_i \rangle|^2$ together with $\sum_j |\lambda_j| = \|\mathbf{X}\|_1$. \blacksquare

Our next result is concerned with the mapping of rank-2 matrices.

Lemma 3.2 *Fix $\delta > 0$. Then there are positive numerical constants c_0 and γ_0 such that if $m \geq c_0 [\delta^{-2} \log \delta^{-1}] n$, \mathcal{A} obeys the following property with probability at least $1 - 3e^{-\gamma_0 m \delta^2}$: for any symmetric rank-2 matrix \mathbf{X} ,*

$$\frac{1}{m} \|\mathcal{A}(\mathbf{X})\|_1 \geq 0.94(1 - \delta) \|\mathbf{X}\|. \quad (3.2)$$

Proof By homogeneity, it suffices to consider the case where $\|\mathbf{X}\| = 1$. Consider then a rank-2 matrix \mathbf{X} with eigenvalue decomposition $\mathbf{X} = \mathbf{u}_1 \mathbf{u}_1^* - t \mathbf{u}_2 \mathbf{u}_2^*$ with $t \in [-1, 1]$ and orthonormal \mathbf{u}_i 's. Note that for $t \leq 0$, Lemma 3.1 already claims a tighter lower bound so it only suffices to consider $t \in [0, 1]$. We have

$$\frac{1}{m} \|\mathcal{A}(\mathbf{X})\|_1 = \frac{1}{m} \sum_{i=1}^m \left| |\langle \mathbf{u}_1, \mathbf{z}_i \rangle|^2 - t |\langle \mathbf{u}_2, \mathbf{z}_i \rangle|^2 \right| = \frac{1}{m} \sum_i \xi_i,$$

where the ξ_i 's are independent copies of the random variable

$$\xi = |Z_1^2 - t Z_2^2|$$

in which Z_1 and Z_2 are independent standard normal variables. This comes from the fact that $\langle \mathbf{u}_1, \mathbf{z}_i \rangle$ and $\langle \mathbf{u}_2, \mathbf{z}_i \rangle$ are independent standard normal. We calculate below that

$$\mathbb{E} \xi = f(t) = \frac{2}{\pi} \left(2\sqrt{t} + (1-t)(\pi/2 - 2 \arctan(\sqrt{t})) \right). \quad (3.3)$$

The graph of this function is shown in Figure 2; we check that $f(t) \geq 0.94$ for all $t \in [0, 1]$.

We now need a deviation bound concerning the fluctuation of $m^{-1} \sum_i \xi_i$ around its mean and this is achieved by classical Chernoff bounds. Note that $\xi \leq Z_1^2 + |t| Z_2^2$ is a sub-exponential variable and thus, $\|\xi\|_{\psi_1} := \sup_{p \geq 1} [\mathbb{E} |\xi|^p]^{1/p}$ is finite.⁴

⁴It would be possible to compute a bound on this quantity but we will not pursue this at the moment.




Figure 2: $f(t) = \mathbb{E} |Z_1^2 - tZ_2^2|$ as a function of t .

Lemma 3.3 (Bernstein-type inequality [21]) *Let X_1, \dots, X_m be i.i.d. sub-exponential random variables. Then*

$$\mathbb{P}\left(\left|\frac{1}{m} \sum_{i=1}^m X_i - \mathbb{E} X_1\right| \geq \epsilon\right) \leq 2 \exp\left[-c_0 m \min\left(\frac{\epsilon^2}{\|X\|_{\psi_1}^2}, \frac{\epsilon}{\|X\|_{\psi_1}}\right)\right]$$

in which c_0 is a positive numerical constant.

We have thus established that for a fixed X ,

$$m^{-1} \|\mathcal{A}(\mathbf{X})\|_1 \geq (0.94 - \epsilon_0) \|\mathbf{X}\|$$

with probability at least $1 - 2e^{-\gamma_0 m \epsilon_0^2}$ (provided $\epsilon_0 \leq \|\xi\|_{\psi_1}$, which we assume).

To complete the argument, let \mathcal{S}_ϵ be an ϵ net of the unit sphere, \mathcal{T}_ϵ be an ϵ net of $[0, 1]$, and set

$$\mathcal{N}_\epsilon = \{\mathbf{X} = \mathbf{u}_1 \mathbf{u}_1^* - t \mathbf{u}_2 \mathbf{u}_2^* : (\mathbf{u}_1, \mathbf{u}_2, t) \in \mathcal{S}_\epsilon \times \mathcal{S}_\epsilon \times \mathcal{T}_\epsilon\}.$$

Since $|\mathcal{S}_\epsilon| \leq (3/\epsilon)^n$, we have

$$|\mathcal{N}_\epsilon| \leq (3/\epsilon)^{2n+1}.$$

Now for any $\mathbf{X} = \mathbf{u} \mathbf{u}^* - t \mathbf{v} \mathbf{v}^*$, consider the approximation $\mathbf{X}_0 = \mathbf{u}_0 \mathbf{u}_0^* - t_0 \mathbf{v}_0 \mathbf{v}_0^* \in \mathcal{N}_\epsilon$, where $\|\mathbf{u}_0 - \mathbf{u}\|_2$, $\|\mathbf{v}_0 - \mathbf{v}\|_2$ and $|t - t_0|$ are each at most ϵ . We claim that

$$\|\mathbf{X} - \mathbf{X}_0\|_1 \leq 9\epsilon, \tag{3.4}$$

and postpone the short proof. On the intersection of $E_1 = \{m^{-1} \|\mathcal{A}(\mathbf{X})\|_1 \leq (1 + \delta_1) \|\mathbf{X}\|_1, \text{ for all } \mathbf{X}\}$ with $E_2 := \{m^{-1} \|\mathcal{A}(\mathbf{X}_0)\|_1 \geq (0.94 - \epsilon) \|\mathbf{X}_0\|, \text{ for all } \mathbf{X}_0 \in \mathcal{N}_\epsilon\}$,

$$\begin{aligned} m^{-1} \|\mathcal{A}(\mathbf{X})\|_1 &\geq \|\mathcal{A}(\mathbf{X}_0)\|_1 - \|\mathcal{A}(\mathbf{X} - \mathbf{X}_0)\|_1 \\ &\geq (0.94 - \epsilon) \|\mathbf{X}_0\| - 9(1 + \delta_1)\epsilon \\ &\geq (0.94 - \epsilon)(\|\mathbf{X}\| - \|\mathbf{X}_0 - \mathbf{X}\|) - 9(1 + \delta_1)\epsilon \\ &\geq (0.94 - \epsilon)(1 - 5\epsilon) - 9(1 + \delta_1)\epsilon \\ &\geq 0.94 - (15 + 9\delta_1)\epsilon, \end{aligned}$$

which is the desired bound by setting $0.94\delta = (15 + 9\delta_1)\epsilon$. In conclusion, set $\delta_1 = 1/2$ and take $\epsilon = 0.94\delta/20$. Then E_1 holds with probability at least $1 - O(e^{-\gamma_1 m \epsilon^2})$ provided m obeys the condition of the theorem. Further, Lemma 3.2 states that E_2 holds with probability at least $1 - 2e^{-\gamma_2 m}$. This concludes the proof provided we check (3.4).

We begin with

$$\|\mathbf{X} - \mathbf{X}_0\|_1 \leq \|\mathbf{u}\mathbf{u}^* - \mathbf{u}_0\mathbf{u}_0^*\|_1 + |t - t_0|\|\mathbf{v}\mathbf{v}^*\|_1 + |t_0|\|\mathbf{v}\mathbf{v}^* - \mathbf{v}_0\mathbf{v}_0^*\|_1.$$

Now

$$\|\mathbf{u}\mathbf{u}^* - \mathbf{u}_0\mathbf{u}_0^*\|_1 \leq 2\|\mathbf{u}\mathbf{u}^* - \mathbf{u}_0\mathbf{u}_0^*\| \leq 4\|\mathbf{u} - \mathbf{u}_0\|_2,$$

where the first inequality follows from the fact that $\mathbf{u}\mathbf{u}^* - \mathbf{u}_0\mathbf{u}_0^*$ is of rank at most 2, and the second follows from

$$\begin{aligned} \|\mathbf{u}\mathbf{u}^* - \mathbf{u}_0\mathbf{u}_0^*\| &= \sup_{\|\mathbf{x}\|_2=1} \left| \langle \mathbf{u}_0, \mathbf{x} \rangle^2 - \langle \mathbf{u}, \mathbf{x} \rangle^2 \right| \\ &= \sup_{\|\mathbf{x}\|_2=1} \left| \langle \mathbf{u} - \mathbf{u}_0, \mathbf{x} \rangle \langle \mathbf{u} + \mathbf{u}_0, \mathbf{x} \rangle \right| \leq \|\mathbf{u} - \mathbf{u}_0\|_2 \|\mathbf{u} + \mathbf{u}_0\|_2 \leq 2\|\mathbf{u} - \mathbf{u}_0\|_2. \end{aligned}$$

Similarly, $\|\mathbf{v}\mathbf{v}^* - \mathbf{v}_0\mathbf{v}_0^*\|_1 \leq 4\epsilon$ and this concludes the proof.⁵ ■

Lemma 3.4 *Let Z_1 and Z_2 be independent $\mathcal{N}(0, 1)$ variables and $t \in [0, 1]$. We have*

$$E|Z_1^2 - tZ_2^2| = f(t),$$

where $f(t)$ is given by (3.3).

Proof Set

$$\rho = \frac{1-t}{1+t} \text{ and } \cos \theta = \rho$$

in which $\theta \in [0, \pi/2]$. By using polar coordinates, we have

$$\begin{aligned} \mathbb{E}|Z_1^2 - tZ_2^2| &= \frac{1}{2\pi} \int_0^\infty r^3 e^{-r^2/2} dr \int_0^{2\pi} |\cos^2 \phi - t \sin^2 \phi| d\phi \\ &= \frac{1}{\pi} \int_0^{2\pi} |\cos^2 \phi - t \sin^2 \phi| d\phi \\ &= \frac{2}{\pi} \int_0^\pi |\cos^2 \phi - t \sin^2 \phi| d\phi \end{aligned}$$

⁵The careful reader will remark that we have also used $\|\mathbf{X} - \mathbf{X}_0\| \leq 5\epsilon$, which also follows from our calculations.

Now using the identities $\cos^2 \phi = (1 + \cos 2\phi)/2$ and $\sin^2 \phi = (1 - \cos 2\phi)/2$, we have

$$\begin{aligned}
\mathbb{E} |Z_1^2 - tZ_2^2| &= \frac{1+t}{\pi} \int_0^\pi |\cos 2\phi + \rho| d\phi \\
&= \frac{1+t}{2\pi} \int_0^{2\pi} |\cos \phi + \rho| d\phi \\
&= \frac{1+t}{\pi} \int_0^\pi |\cos \phi + \rho| d\phi \\
&= \frac{1+t}{\pi} \int_0^\pi |\rho - \cos \phi| d\phi \\
&= \frac{1+t}{\pi} \left[\int_0^\theta \cos \phi - \rho d\phi + \int_\theta^\pi \rho - \cos \phi d\phi \right] \\
&= \frac{2}{\pi} (1+t) [\sin \theta + \rho(\pi/2 - \theta)].
\end{aligned}$$

We recognize (3.3). ■

4 Dual Certificates

To prove our main theorem, it remains to show that one can construct an inexact dual certificate \mathbf{Y} obeying the conditions of Lemma 2.1.

4.1 Preliminaries

The linear mapping $\mathcal{A}^* \mathcal{A}$ is of the form⁶

$$\mathcal{A}^* \mathcal{A} = \sum_{i=1}^m \mathbf{z}_i \mathbf{z}_i^* \otimes \mathbf{z}_i \mathbf{z}_i^*,$$

which is another way to express that $\mathcal{A}^* \mathcal{A}(\mathbf{X}) = \sum_i \langle \mathbf{z}_i \mathbf{z}_i^*, \mathbf{X} \rangle \mathbf{z}_i \mathbf{z}_i^*$. Now observe the simple identity:

$$\mathbb{E}[\mathbf{z}_i \mathbf{z}_i^* \otimes \mathbf{z}_i \mathbf{z}_i^*] = 2\mathcal{I} + \mathbf{I}_n \otimes \mathbf{I}_n := \mathcal{S}, \quad (4.1)$$

where \mathcal{I} is the identity operator and \mathbf{I}_n the n -dimensional identity matrix. Put differently, this means that for all \mathbf{X} ,

$$\mathcal{S}(\mathbf{X}) = 2\mathbf{X} + \text{Tr}(\mathbf{X})\mathbf{I}.$$

The proof is a simple calculation and omitted. It is also not hard to see that the mapping \mathcal{S} is invertible and its inverse is given by

$$\mathcal{S}^{-1} = \frac{1}{2} \left(\mathcal{I} - \frac{1}{n+2} \mathbf{I}_n \otimes \mathbf{I}_n \right) \Leftrightarrow \mathcal{S}^{-1}(\mathbf{X}) = \frac{1}{2} \left(\mathbf{X} - \frac{1}{n+2} \text{Tr}(\mathbf{X})\mathbf{I}_n \right).$$

We will use this object in the definition of our dual certificate.

⁶For symmetric matrices, $\mathbf{A} \otimes \mathbf{B}$ is the linear mapping $\mathbf{H} \mapsto \langle \mathbf{A}, \mathbf{H} \rangle \mathbf{B}$.

4.2 Construction

For pedagogical reasons, we first introduce a possible candidate certificate defined by

$$\bar{\mathbf{Y}} := \frac{1}{m} \mathcal{A}^* \mathcal{A} \mathcal{S}^{-1}(\mathbf{e}_1 \mathbf{e}_1^*). \quad (4.2)$$

Clearly, $\bar{\mathbf{Y}}$ is in the range of \mathcal{A}^* as required. To justify this choice, the law of large numbers gives that in the limit of infinitely many samples,

$$\lim_{m \rightarrow \infty} \frac{1}{m} \sum_i (z_i z_i^* \otimes z_i z_i^*) \mathcal{S}^{-1}(\mathbf{e}_1 \mathbf{e}_1^*) = \mathbb{E}(z_i z_i^* \otimes z_i z_i^*) \mathcal{S}^{-1}(\mathbf{e}_1 \mathbf{e}_1^*) = \mathbf{e}_1 \mathbf{e}_1^*.$$

In other words, in the limit of large samples, we have a perfect certificate since $\bar{\mathbf{Y}}_T = \mathbf{e}_1 \mathbf{e}_1^*$ and $\bar{\mathbf{Y}}_T^\perp = 0$. Our hope is that the sample average is sufficiently close to the population average so that one can check (2.4). In order to show that this is the case, it will be useful to think of $\bar{\mathbf{Y}}$ (4.2) as the random sum

$$\bar{\mathbf{Y}} = \frac{1}{m} \sum_i \mathbf{Y}_i,$$

where each matrix \mathbf{Y}_i is an independent copy of the random matrix

$$\frac{1}{2} \left[z_1^2 - \frac{1}{n+2} \|z\|_2^2 \right] z z^*$$

in which $z = (z_1, \dots, z_n) \sim \mathcal{N}(0, I)$.

We would like to make an important point before continuing. We have seen that all we need from $\bar{\mathbf{Y}}$ is

$$\|\bar{\mathbf{Y}}_T - \mathbf{e}_1 \mathbf{e}_1^*\|_2 \leq 1/3$$

(and $\|\bar{\mathbf{Y}}_T^\perp\| \leq 1/2$). This is in stark contrast with David Gross' approach [13] which requires a very small misfit, i.e. an error of at most $1/n^2$. In turn, this loose bound has an enormous implication: it eliminates the need for the golfing scheme and allows for the simple certificate candidate (4.2). In fact, our certificate can be seen as the first iteration of Gross' golfing scheme.

4.3 Truncation

For technical reasons, it is easier to work with a truncated version of $\bar{\mathbf{Y}}$ and our dual certificate is taken to be

$$\mathbf{Y} = \frac{1}{m} \sum_i \mathbf{Y}_i 1_{E_i}, \quad (4.3)$$

where the \mathbf{Y}_i 's are as before and 1_{E_i} are independent copies of 1_E with

$$E = \{|z_1| \leq \sqrt{2\beta \log n}\} \cap \{\|z\|_2 \leq \sqrt{3n}\}.$$

We shall work with $\beta = 3$ so that $|z_1| \leq \sqrt{6 \log n}$.

Lemma 4.1 *Let \mathbf{Y} be as in (4.3). Then*

$$\mathbb{P}\left(\|\mathbf{Y}_T - \mathbf{e}_1 \mathbf{e}_1^*\|_2 \geq \frac{1}{3}\right) \leq 2 \exp\left(-\gamma \frac{m}{n}\right), \quad (4.4)$$

where $\gamma > 0$ is an absolute constant. This holds with the proviso that $m \geq c_1 n$ for some numerical constant $c_1 > 0$, and that n is sufficiently large.

Lemma 4.2 *Let \mathbf{Y} be as in (4.3). Then*

$$\mathbb{P}\left(\|\mathbf{Y}_T^\perp\| \geq \frac{1}{2}\right) \leq 4 \exp\left(-\gamma \frac{m}{\log n}\right). \quad (4.5)$$

where $\gamma > 0$ is an absolute constant. This holds with the proviso that $m \geq c_1 n \log n$ for some numerical constant $c_1 > 0$, and that n is sufficiently large.

4.4 \mathbf{Y} on T and proof of Lemma 4.1

It is obvious that for any symmetric matrix $\mathbf{X} \in T$,

$$\|\mathbf{X}\|_2 \leq \sqrt{2} \|\mathbf{X} \mathbf{e}_1\|_2$$

since only the first row and column are nonzero. We have

$$\mathbf{Y}_T \mathbf{e}_1 - \mathbf{e}_1 = \frac{1}{m} \sum_{i=1}^m \mathbf{y}_i 1_{E_i} - \frac{1}{m} \sum_{i=1}^m \mathbf{e}_1 1_{E_i^c}, \quad (4.6)$$

where the \mathbf{y}_i 's are independent copies of the random vector

$$\mathbf{y} = \frac{1}{2} \left[z_1^2 - \frac{1}{n+2} \|\mathbf{z}\|_2^2 \right] z_1 \mathbf{z} - \mathbf{e}_1 := (\xi z_1) \mathbf{z} - \mathbf{e}_1. \quad (4.7)$$

We claim that

$$\left\| \frac{1}{m} \sum_{i=1}^m \mathbf{e}_1 1_{E_i^c} \right\|_2 \leq 1/9,$$

with probability at least $1 - 2e^{-\gamma m}$ for some $\gamma > 0$. This is a simple application of Bernstein's inequality. Set $\pi(\beta) = \mathbb{P}(E_i^c)$ and observe that

$$\pi(\beta) = \mathbb{P}(|z_1| \geq \sqrt{2\beta \log n}) + \mathbb{P}(\|\mathbf{z}\|_2^2 \geq 3n) \leq n^{-\beta} + e^{-\frac{n}{3}}. \quad (4.8)$$

The right-hand side follows from $\mathbb{P}(|z_1| \geq t) \leq e^{-t^2/2}$ which holds for $t \geq 1$ and from $\mathbb{P}(\|\mathbf{z}\|_2^2 \geq 3n) \leq e^{-n/3}$. In turn, this last bound follows from

$$\mathbb{P}(\|\mathbf{z}\|_2^2 - n \geq \sqrt{2nt} + t^2) \leq e^{-t^2/2}.$$

Returning to Bernstein, this gives

$$\mathbb{P}\left(\left| \frac{1}{m} \sum_{i=1}^m 1_{E_i^c} - \pi(\beta) \right| \geq t\right) \leq 2 \exp\left(-\frac{mt^2}{2\pi(\beta) + 2t/3}\right).$$

Setting $t = 1/18$, $\beta = 3$ and taking n large enough so that $\pi(3) \leq 1/18$ proves the claim.

The main task is to bound the 2-norm of the sum $\sum_{i=1}^m \mathbf{y}_i 1_{E_i}$ and a convenient way to do this is via the vector Bernstein inequality.

Theorem 4.3 (Vector Bernstein inequality) *Let \mathbf{x}_i be a sequence of independent random vectors and set $V \geq \sum_i \mathbb{E} \|\mathbf{x}_i\|_2^2$. Then for all $t \leq V/\max\|\mathbf{x}_i\|_2$, we have*

$$\mathbb{P}(\|\sum_i (\mathbf{x}_i - \mathbb{E} \mathbf{x}_i)\|_2 \geq \sqrt{V} + t) \leq e^{-t^2/4V}.$$

It is because this inequality requires bounded random vectors that we work with the truncation $\sum_{i=1}^m \mathbf{y}_i 1_{E_i}$.

Put $\bar{\mathbf{y}} = \mathbf{y} 1_E$. Since $\|\bar{\mathbf{y}}\|_2^2 \leq \|\mathbf{y}\|_2^2$, we first compute $\mathbb{E} \|\mathbf{y}\|_2^2$. We have

$$\|\mathbf{y}\|_2^2 = \|\mathbf{z}\|_2^2 z_1^2 \xi^2 - 2z_1^2 \xi + 1, \quad \xi = \frac{1}{2} \left[z_1^2 - \frac{1}{n+2} \|\mathbf{z}\|_2^2 \right],$$

and a little bit of algebra yields

$$\|\mathbf{y}\|_2^2 = \frac{1}{4} z_1^6 \|\mathbf{z}\|_2^2 - \frac{1}{2(n+2)} z_1^4 \|\mathbf{z}\|_2^4 + \frac{1}{4(n+2)^2} z_1^2 \|\mathbf{z}\|_2^6 - z_1^4 + \frac{1}{n+2} z_1^2 \|\mathbf{z}\|_2^2 + 1.$$

Thus,

$$\begin{aligned} \mathbb{E} [\|\mathbf{y}\|_2^2] &= \frac{1}{4}(15n+90) - \frac{1}{2(n+2)}(3n^2+30n+72) + \frac{1}{4(n+2)^2}(n+4)(n+6) - 1 \\ &\leq 4(n+4), \end{aligned} \tag{4.9}$$

where we have used the following identities

$$\begin{aligned} \mathbb{E} [z_1^2 \|\mathbf{z}\|_2^2] &= n+2, \\ \mathbb{E} [z_1^2 \|\mathbf{z}\|_2^6] &= (n+2)(n+4)(n+6), \\ \mathbb{E} [z_1^4 \|\mathbf{z}\|_2^4] &= 3n^2+30n+72, \\ \mathbb{E} [z_1^6 \|\mathbf{z}\|_2^2] &= 15n+90. \end{aligned}$$

Second, on the event of interest we have $|\xi| \leq \beta \log n$ (assuming $2\beta \log n \geq 3$), $|z_1| \leq \sqrt{2\beta \log n}$ and $\|\mathbf{z}\|_2 \leq \sqrt{3n}$ and, therefore,

$$\|\bar{\mathbf{y}}\|_2 \leq \sqrt{6n} (\beta \log n)^{3/2} + 1 \leq \sqrt{7n} (\beta \log n)^{3/2}$$

provided n is large enough.

Third, observe that by symmetry, all the entries of $\bar{\mathbf{y}}$ but the first have mean zero. Hence,

$$\|\mathbb{E} \bar{\mathbf{y}}\|_2 = |\mathbb{E} y_1 - \bar{y}_1| = |\mathbb{E} 1_{E^c} y_1| \leq \sqrt{\mathbb{P}(E^c)} \sqrt{\mathbb{E} y_1^2}.$$

We have

$$y_1^2 = (\xi z_1^2 - 1)^2 = \frac{1}{4} z_1^8 - z_1^4 + \frac{1}{n+2} \|\mathbf{z}\|_2^2 z_1^2 - \frac{1}{2(n+2)} \|\mathbf{z}\|_2^2 z_1^6 + \frac{1}{4(n+2)^2} \|\mathbf{z}\|_2^4 z_1^4 + 1$$

and using the identities above

$$\mathbb{E} y_1^2 = \frac{101}{4} - \frac{27n^2 + 210n + 288}{4(n+2)^2} \leq 22,$$

which gives

$$\|\mathbb{E} \bar{\mathbf{y}}\|_2 \leq \sqrt{22(n^{-\beta} + e^{-\frac{n}{3}})}.$$

Finally, with $V = 4m(n+4)$, Bernstein's inequality gives that for each $t \leq 4(n+4)/[\sqrt{7n}(\beta \log n)^{3/2}]$,

$$\|m^{-1} \sum_i (\bar{\mathbf{y}}_i - \mathbb{E} \bar{\mathbf{y}}_i)\|_2 \geq 2\sqrt{\frac{n+4}{m}} + t$$

with probability at most $\exp(-\frac{mt^2}{16(n+4)})$. It follows that

$$\|m^{-1} \sum_i \bar{\mathbf{y}}_i\|_2 \geq \sqrt{22(n^{-\beta} + e^{-\frac{n}{3}})} + 2\sqrt{\frac{n+4}{m}} + t$$

with at most the same probability. Our result follows by taking $t = 1/6$, $\beta = 3$, $m \geq c_1 n$ where n and c_1 are sufficiently large such that

$$\sqrt{22(n^{-\beta} + e^{-\frac{n}{3}})} + 2\sqrt{\frac{n+4}{m}} + \frac{1}{6} \leq \frac{2}{9}.$$

4.5 \mathbf{Y} on T^\perp and proof of Lemma 4.2

We have

$$\mathbf{Y}_T^\perp = \frac{1}{m} \sum_i \mathbf{X}_i 1_{E_i},$$

where the \mathbf{X}_i 's are independent copies of the random matrix

$$\mathbf{X} = \frac{1}{2} \left[z_1^2 - \frac{1}{n+2} \|\mathbf{z}\|_2^2 \right] \mathcal{P}_{T^\perp}(\mathbf{z}\mathbf{z}^T). \quad (4.10)$$

One natural way to bound the norm of this random sum is via the operator Bernstein's inequality. We develop a more customized approach, which gives sharper results.

Decompose \mathbf{X} as

$$\mathbf{X} = \frac{1}{2} \left[z_1^2 - 1 \right] \mathcal{P}_{T^\perp}(\mathbf{z}\mathbf{z}^T) + \frac{1}{2} \left[1 - \frac{1}{n+2} \|\mathbf{z}\|_2^2 \right] \mathcal{P}_{T^\perp}(\mathbf{z}\mathbf{z}^T) := \mathbf{X}^{(0)} + \mathbf{X}^{(1)}.$$

Note that since z_1 and $\mathcal{P}_{T^\perp}(\mathbf{z}\mathbf{z}^T)$ are independent, we have $\mathbb{E} \mathbf{X}^{(0)} = 0$ and thus, $\mathbb{E} \mathbf{X}^{(1)} = 0$ since $\mathbb{E} \mathbf{X} = 0$. With $\bar{\mathbf{X}}_i^{(0)} = \mathbf{X}_i^{(0)} 1_{E_i}$ and similarly for $\bar{\mathbf{X}}_i^{(1)}$, it then suffices to show that

$$\left\| \sum_i \bar{\mathbf{X}}_i^{(0)} \right\| \leq m/4 \quad \text{and} \quad \left\| \sum_i \bar{\mathbf{X}}_i^{(1)} \right\| \leq m/4 \quad (4.11)$$

with large probability. Write the norm as

$$\left\| \sum_i \bar{\mathbf{X}}_i^{(0)} \right\| = \sup_{\mathbf{u}} \left| \sum_i \langle \mathbf{u}, \bar{\mathbf{X}}_i^{(0)} \mathbf{u} \rangle \right|,$$

where the supremum is over all unit vectors \mathbf{u} that are orthogonal to \mathbf{e}_1 . The strategy is now to find a bound on the right-hand side for each fixed \mathbf{u} and apply a covering argument to control the supremum over the whole unit sphere. In order to do this, we shall make use of a classical large deviation result.

Theorem 4.4 (Bernstein inequality) *Let $\{X_i\}$ be a finite sequence of independent random variables. Suppose that there exists V and c such that for all X_i and all $k \geq 3$,*

$$\mathbb{E} |X_i|^k \leq \frac{1}{2} k! V c_0^{k-2}.$$

Then for all $t \geq 0$,

$$\mathbb{P}\left(\left|\sum_i X_i - \mathbb{E} X_i\right| \geq t\right) \leq 2 \exp\left(-\frac{t^2}{2V + 2c_0 t}\right). \quad (4.12)$$

For the first sum in (4.11), we write

$$\sum_i \langle \mathbf{u}, \bar{\mathbf{X}}_i^{(0)} \mathbf{u} \rangle = \sum_i \eta_i 1_{E_i},$$

where the η_i 's are independent copies of

$$\eta = \frac{1}{2} \left[z_1^2 - 1 \right] \langle \mathbf{z}, \mathbf{u} \rangle^2.$$

The point of the decomposition $\mathbf{X}^{(0)} + \mathbf{X}^{(1)}$ is that z_1 and $\langle \mathbf{z}, \mathbf{u} \rangle$ are independent since \mathbf{u} is orthogonal to \mathbf{e}_1 . We have $\mathbb{E} \eta = 0$ and for $k \geq 2$,

$$\mathbb{E} |\eta 1_E|^k \leq 2^{-k} \mathbb{E} |(z_1^2 - 1) 1_{\{z_1^2 \leq 2\beta \log n\}}|^k \mathbb{E} |\langle \mathbf{z}, \mathbf{u} \rangle|^{2k}.$$

First,

$$\mathbb{E} |(z_1^2 - 1) 1_{\{z_1^2 \leq 2\beta \log n\}}|^k \leq (2\beta \log n)^{k-2} \mathbb{E} (z_1^2 - 1)^2 = 2(2\beta \log n)^{k-2}.$$

Second, the moments of a chi-square variable with one degree of freedom are well known:

$$\mathbb{E} |\langle \mathbf{z}, \mathbf{u} \rangle|^{2k} = 1 \times 3 \times \dots \times (2k - 1) \leq 2^k k!$$

Hence we can apply Bernstein inequality with $V = 4m$ and $c_0 = 2\beta \log n$ and, obtain

$$\mathbb{P}\left(\left|\sum_i \eta_i 1_{E_i} - \mathbb{E}[\eta_i 1_{E_i}]\right| \geq mt\right) \leq 2 \exp\left(-\frac{m}{4} \frac{t^2}{2 + \beta t \log n}\right).$$

We now note that

$$|\mathbb{E} \eta_i 1_{E_i}| = |\mathbb{E} \eta_i 1_{E_i^c}| \leq \sqrt{\mathbb{P}(E_i^c)} \sqrt{\mathbb{E} \eta_i^2} = \sqrt{\frac{3\pi(\beta)}{2}}$$

which gives

$$\mathbb{P}\left(m^{-1} \left|\sum_i \eta_i 1_{E_i}\right| \geq t + \sqrt{\frac{3\pi(\beta)}{2}}\right) \leq 2 \exp\left(-\frac{m}{4} \frac{t^2}{2 + \beta t \log n}\right).$$

For instance, take $t = 1/12$, $\beta = 3$, $m \geq c_1 n$ and n large enough to get

$$\mathbb{P}\left(m^{-1} \left|\sum_i \eta_i 1_{E_i}\right| \geq 1/8\right) \leq 2 \exp\left(-\gamma \frac{m}{\log n}\right).$$

To derive a bound about $\|\bar{\mathbf{X}}^{(0)}\|$, we use (see Lemma 4 in [21])

$$\sup_{\mathbf{u}} \left| \langle \mathbf{u}, \bar{\mathbf{X}}^{(0)} \mathbf{u} \rangle \right| \leq 2 \sup_{\mathbf{u} \in \mathcal{N}_{1/4}} \left| \langle \mathbf{u}, \bar{\mathbf{X}}^{(0)} \mathbf{u} \rangle \right|,$$

where $\mathcal{N}_{1/4}$ is a $1/4$ -net of the unit sphere $\{\mathbf{u} : \|\mathbf{u}\|_2 = 1, \mathbf{u} \perp \mathbf{e}_1\}$. Since $|\mathcal{N}_{1/4}| \leq 9^n$,

$$\mathbb{P}(m^{-1} \|\bar{\mathbf{X}}^{(0)}\| > 1/4) \leq \mathbb{P}\left(m^{-1} \sup_{\mathbf{u} \in \mathcal{N}_{1/4}} \left| \langle \mathbf{u}, \bar{\mathbf{X}}^{(0)} \mathbf{u} \rangle \right| > 1/8\right) \leq 9^n \times 2 \exp\left(-\gamma \frac{m}{\log n}\right).$$

We deal with the second term in a similar way, and write

$$\sum_i \langle \mathbf{u}, \bar{\mathbf{X}}_i^{(1)} \mathbf{u} \rangle = \sum_i \eta_i \mathbf{1}_{E_i},$$

where the η_i 's are now independent copies of

$$\eta = \frac{1}{2} \left[1 - \frac{\|\mathbf{z}\|_2^2}{n+2} \right] \langle \mathbf{z}, \mathbf{u} \rangle^2.$$

On E , $\|\mathbf{z}\|_2^2 \leq 3n$ and, therefore, $\mathbb{E} |\eta \mathbf{1}_{E_i}|^k \leq 2^k k!$. We can apply Bernstein's inequality with $c_0 = 2$ and $V = 8m$, which gives

$$\mathbb{P}\left(\left| \sum_i \eta_i \mathbf{1}_{E_i} - \mathbb{E}[\eta_i \mathbf{1}_{E_i}] \right| \geq mt\right) \leq 2 \exp\left(-\frac{m}{4} \frac{t^2}{4+t}\right).$$

The remainder of the proof is identical to that above and is therefore omitted.

4.6 Proof of Theorem 1.1

We now assemble the various intermediate results to establish Theorem 1.1. As pointed out, Theorem 1.1 follows immediately from Lemma 2.1, which in turn hinges on the validity of the conditions stated in (2.2), (2.3), and (2.4).

Lemma 3.1 asserts that condition (2.2) holds with probability of failure at most p_1 , where $p_1 = 2e^{-\gamma_1 m}$ and here and below, $\gamma_1, \dots, \gamma_4$ are positive numerical constants. Similarly, Lemma 3.2 shows that condition (2.3) holds with probability of failure at most p_2 , where $p_2 = 3e^{-\gamma_2 m}$. In both cases we need that $m > cn$ for an absolute constant $c > 0$.

Proceeding to the dual certificate in (2.4), we note that Lemma 4.1 establishes the first part of the dual certificate with a probability of failure at most p_3 , where $p_3 = 3e^{-\gamma_3 m/n}$. The second part of the dual certificate in (2.4) is shown in Lemma 4.2 to hold with probability of failure at most p_4 , where $p_4 = 4e^{-\gamma_4 \frac{m}{\log n}}$. In the former case we need $m > cn$ for an absolute constant $c > 0$ and in the latter $m > c'n \log n$.

Finally, the union bound gives that under the hypotheses of Theorem 1.1, exact recovery holds with probability at least $1 - 3e^{-\gamma m/n}$ for some $\gamma > 0$, as claimed.

5 The Complex Model

This section proves that Theorem 1.1 holds for the complex model as well. Not surprisingly, the main steps of the proof are the same as in the real case, but there are here and there some noteworthy differences. Instead of deriving the whole proof, we will carefully indicate the nontrivial changes that need to be carried out.

First, we can work with $\mathbf{x} = \mathbf{e}_1$ because of rotational invariance, and with independent complex valued Gaussian sequences $z_i \sim \mathcal{CN}(0, I, 0)$. This means that the real and imaginary parts of z_i are independent white noise sequences with variance $1/2$.

The key Lemma 2.1 only requires a slight adjustment in the numerical constants. The reason for this is that while Lemma 3.1 does not require any modification, Lemma 3.2 changes slightly; in particular, the numerical constants are somewhat different. Here is the properly adjusted complex version.

Lemma 5.1 *Fix $\delta > 0$. Then there are positive numerical constants c_0 and γ_0 such that if $m \geq c_0 [\delta^{-2} \log \delta^{-1}] n$, \mathcal{A} has the following property with probability at least $1 - 3e^{-\gamma_0 m \delta^2}$: for any Hermitian rank-2 matrix X ,*

$$\frac{1}{m} \|\mathcal{A}(\mathbf{X})\|_1 \geq 2(\sqrt{2} - 1)(1 - \delta) \|\mathbf{X}\| \geq 0.828(1 - \delta) \|\mathbf{X}\|. \quad (5.1)$$

The proof of this lemma follows essentially the proof of Lemma 3.2. The function $f(t)$ (cf. equation (3.3)) now takes the form

$$\mathbb{E} \xi = f(t) = \frac{1 + t^2}{1 + t}, \quad (5.2)$$

where $\xi = ||Z_1|^2 - t|Z_2|^2|$, with Z_1 and Z_2 independent $\mathcal{CN}(0, 1, 0)$, as demonstrated in the following lemma.

Lemma 5.2 *Let Z_1 and Z_2 be independent $\mathcal{CN}(0, 1, 0)$ variables and $t \in [0, 1]$. We have*

$$E||Z_1|^2 - t|Z_2|^2| = f(t),$$

where $f(t)$ is given by (5.2).

Proof Set

$$\rho = \frac{1 - t}{1 + t} \text{ and } \cos \theta = \rho$$

in which $\theta \in [0, \pi/2]$. By using polar coordinates for the variables (x_1, y_1) associated with Z_1 and (x_2, y_2) , associated with Z_2 we have

$$\begin{aligned} \mathbb{E} ||Z_1|^2 - t|Z_2|^2| &= \frac{1}{2} \int_0^\infty \int_0^\infty |r_1^2 - tr_2^2| r_1 r_2 e^{-r_1^2/2} e^{-r_2^2/2} dr_1 dr_2 \\ &= \frac{1}{8} \int_0^\infty r^5 e^{-r^2/2} dr \int_0^{2\pi} |\sin \phi \cos \phi| |\cos^2 \phi - t \sin^2 \phi| d\phi, \end{aligned}$$




Figure 3: The function $f(t)$ in (5.2) as a function of t .

where we used polar coordinates again in variables (r_1, r_2) . Now using the identities $\cos^2 \phi = (1 + \cos 2\phi)/2$, $\sin^2 \phi = (1 - \cos 2\phi)/2$ and $2 \sin \phi \cos \phi = \sin 2\phi$ we have

$$\begin{aligned}
 \mathbb{E} |Z_1^2 - tZ_2^2| &= \frac{1}{2} \int_0^\pi |\sin 2\phi| |\cos 2\phi + \rho| d\phi \\
 &= \frac{1}{2} \left[\int_0^\theta \sin \phi (\cos \phi - \rho) d\phi + \int_\theta^\pi \sin \phi (\rho - \cos \phi) d\phi \right] \\
 &= \frac{1}{2} (1+t) \left[-\frac{1}{2} \cos 2\theta + 2\rho \cos \theta + \frac{1}{2} \right] \\
 &= \frac{1}{2} (1+t) [\rho^2 + 1] \\
 &= \frac{1+t^2}{1+t}
 \end{aligned}$$

as claimed. ■

The graph of $f(t)$ is shown in Figure 3. The minimum of this function on $[0, 1]$ is $2(\sqrt{2}-1) > 0.828$. Furthermore, the covering argument in that proof has to be adapted; for example, unit spheres need to be replaced by complex unit spheres.

A consequence of this change in numerical values is that the numerical factors in Lemma 2.2 need to be adjusted.

Lemma 5.3 *Any feasible matrix \mathbf{H} such that $\text{Tr}(\mathbf{H}) \leq 0$ must obey*

$$\|\mathbf{H}_T\|_2 \leq \sqrt{\frac{5}{4}} \|\mathbf{H}_T\|.$$

Finally, with all of this in place, Lemma 2.1 becomes this:

Lemma 5.4 *Suppose that the mapping \mathcal{A} obeys the following two properties: for some $\delta \leq 3/13$:*

1) *for all positive semidefinite matrices \mathbf{X} ,*

$$m^{-1}\|\mathcal{A}(\mathbf{X})\|_1 \leq (1 + \delta)\|\mathbf{X}\|_1; \quad (5.3)$$

2) *for all matrices $\mathbf{X} \in T$*

$$m^{-1}\|\mathcal{A}(\mathbf{X})\|_1 \geq 2(\sqrt{2} - 1)(1 - \delta)\|\mathbf{X}\| \geq 0.828(1 - \delta)\|\mathbf{X}\|. \quad (5.4)$$

Suppose further that there exists Y in the range of \mathcal{A}^ obeying*

$$\|\mathbf{Y}_T - \mathbf{e}_1 \mathbf{e}_1^*\|_2 \leq 1/5 \quad \text{and} \quad \|\mathbf{Y}_T^\perp\| \leq 1/2. \quad (5.5)$$

Then $\mathbf{e}_1 \mathbf{e}_1^$ is the unique minimizer to (1.5).*

We now turn our attention to the properties of the dual certificate we studied in Section 4. The first difference is that the expectation of $\mathcal{A}^* \mathcal{A}$ in (4.1) is different in the complex case. A simple calculation yields

$$\mathbb{E} \frac{1}{m} \mathcal{A}^* \mathcal{A} = \mathcal{I} + I_n \otimes I_n := \mathcal{S}.$$

This means that for all \mathbf{X} ,

$$\mathcal{S}(\mathbf{X}) = \mathbf{X} + \text{Tr}(\mathbf{X})\mathbf{I}. \quad (5.6)$$

We note that in this case

$$\mathcal{S}^{-1} = \mathcal{I} - \frac{1}{n+1} I_n \otimes I_n \quad \Leftrightarrow \quad \mathcal{S}^{-1}(\mathbf{X}) = \mathbf{X} - \frac{1}{n+1} \text{Tr}(\mathbf{X})\mathbf{I}_n. \quad (5.7)$$

We of course use this new \mathcal{S}^{-1} in the complex analog of the candidate certificate (4.3). A consequence is that in the proof of Lemma 4.1, for instance, (4.7) now takes the form

$$\mathbf{X} = \left[|z_1|^2 - \frac{1}{n+1} \|\mathbf{z}\|_2^2 \right] \bar{z}_1 \mathbf{z} - \mathbf{e}_1 := (\xi \bar{z}_1) \mathbf{z} - \mathbf{e}_1. \quad (5.8)$$

To bound the 2-norm of a sum of i.i.d. such random variables (as in Lemma 4.1), we employ the same Bernstein inequality for real vectors, using the fact that $\|\mathbf{z}\|_2 = \|(\Re(\mathbf{z}), \Im(\mathbf{z}))\|_2$ for any complex vector \mathbf{z} . Similarly (4.10) becomes

$$\mathbf{X} = \left[|z_1|^2 - \frac{1}{n+1} \|\mathbf{z}\|_2^2 \right] \mathcal{P}_{T^\perp}(\mathbf{z}\mathbf{z}^*). \quad (5.9)$$

To bound the operator norm of a sum of i.i.d. such random matrices (as in Lemma 4.2), we again use a covering argument, this time working with chi-square variables with two degrees of freedom, since $|\langle \mathbf{z}, \mathbf{u} \rangle|^2$ is distributed as $\frac{1}{2} \chi^2(2)$. Since $|\langle \mathbf{z}, \mathbf{u} \rangle|^2$ are real random variables, we use the same version of the Bernstein inequality as in the real-valued case. The only difference is that the moments are now

$$\mathbb{E} |\langle \mathbf{z}, \mathbf{u} \rangle|^{2k} = 2^{-k} \times (2+0) \times (2+2) \times (2+4) \times \dots \times (2+2k-2) = k!$$

6 Stability

This section proves the stability of our approach, namely, Theorem 1.2. Our proof parallels the argument of Candès and Plan for showing the stability of matrix completion [8] as well as that of Gross et al. in [14].

Just as before, we prove the theorem in the real case since the complex case is essentially the same. Further, we may still take $\mathbf{x} = \mathbf{e}_1$ without loss of generality. We shall prove stability when the \mathbf{z}_i 's are i.i.d. $\mathcal{N}(0, \mathbf{I}_n)$ and later explain how one can easily transfer a result for Gaussian vectors to a result for vectors sampled on the sphere. Under the assumptions of the theorem, the RIP-1-like properties, namely, Lemmas 3.1 and 3.2 hold with a numerical constant δ_1 we shall specify later. Under the same hypotheses, the dual certificate \mathbf{Y} (4.2) obeys

$$\|\mathcal{P}_T(\mathbf{Y} - \mathbf{e}_1 \mathbf{e}_1^*)\|_2 \leq \gamma, \quad \|\mathbf{Y}_{T^\perp}\| \leq \frac{1}{2},$$

in which γ is a numerical constant also specified later.

Set $\mathbf{X} = \mathbf{x}\mathbf{x}^* = \mathbf{e}_1 \mathbf{e}_1^*$ and write $\hat{\mathbf{X}} = \mathbf{X} + \mathbf{H}$. We begin by recording two useful properties. First, since \mathbf{X} is feasible for our optimization problem, we have

$$\text{Tr}(\mathbf{X} + \mathbf{H}) \leq \text{Tr}(\mathbf{X}) \iff \text{Tr}(\mathbf{H}) \leq 0. \quad (6.1)$$

Second, the triangle inequality gives

$$\|\mathcal{A}(\mathbf{H})\|_2 = \|\mathcal{A}(\hat{\mathbf{X}} - \mathbf{X})\|_2 \leq \|\mathcal{A}(\hat{\mathbf{X}}) - \mathbf{b}\|_2 + \|\mathbf{b} - \mathcal{A}(\mathbf{X})\|_2 \leq 2\epsilon. \quad (6.2)$$

In the noiseless case, $\mathcal{A}(\mathbf{H}) = 0 \implies \langle \mathbf{H}, \mathbf{Y} \rangle = 0$, by construction. In the noisy case, a third property is that $|\langle \mathbf{H}, \mathbf{Y} \rangle|$ is at most on the order of ϵ . Indeed,

$$m|\langle \mathbf{H}, \mathbf{Y} \rangle| = |\langle \mathcal{A}(\mathbf{H}), \mathcal{A}\mathcal{S}^{-1}(\mathbf{X}) \rangle| \leq \|\mathcal{A}(\mathbf{H})\|_\infty \|\mathcal{A}\mathcal{S}^{-1}(\mathbf{X})\|_1.$$

Since, $\|\mathcal{A}(\mathbf{H})\|_\infty \leq \|\mathcal{A}(\mathbf{H})\|_2$ and

$$\|\mathcal{A}\mathcal{S}^{-1}(\mathbf{X})\|_1 \leq m(1 + \delta_1) \|\mathcal{S}^{-1}(\mathbf{X})\|_1 \leq m(1 + \delta_1),$$

we obtain

$$|\langle \mathbf{H}, \mathbf{Y} \rangle| \leq 2\epsilon(1 + \delta_1). \quad (6.3)$$

We now reproduce the steps of the proof of Lemma 2.1, and obtain

$$0 \geq \text{Tr}(\mathbf{H}_T) + \text{Tr}(\mathbf{H}_T^\perp) \geq \frac{1}{2} \text{Tr}(\mathbf{H}_T^\perp) - \gamma \|\mathbf{H}_T\|_2 - |\langle \mathbf{H}, \mathbf{Y} \rangle|,$$

which gives

$$\text{Tr}(\mathbf{H}_T^\perp) \leq 4\epsilon(1 + \delta_1) + 2\gamma \|\mathbf{H}_T\|_2 \leq 4\epsilon(1 + \delta_1) + 2\sqrt{2}\gamma \|\mathbf{H}_T\|, \quad (6.4)$$

where we recall that \mathbf{H}_T has rank at most 2. We also have

$$\begin{aligned} 0.94(1 - \delta_1) \|\mathbf{H}_T\| &\leq m^{-1} \|\mathcal{A}(\mathbf{H}_T)\|_1 \leq m^{-1} \|\mathcal{A}(\mathbf{H})\|_1 + m^{-1} \|\mathcal{A}(\mathbf{H}_T^\perp)\|_1 \\ &\leq m^{-1/2} \|\mathcal{A}(\mathbf{H})\|_2 + (1 + \delta_1) \text{Tr}(\mathbf{H}_T^\perp) \end{aligned} \quad (6.5)$$

$$\leq 2m^{-1/2}\epsilon + (1 + \delta_1) \text{Tr}(\mathbf{H}_{T^\perp}), \quad (6.6)$$

where the second inequality follows from the RIP-1 property together with the Cauchy-Schwarz inequality. Plugging this last bound into (6.4) gives

$$\mathrm{Tr}(\mathbf{H}_T^\perp) \leq 4\epsilon(1 + \delta_1 + \gamma\alpha m^{-1/2}) + \beta\gamma \mathrm{Tr}(\mathbf{H}_T^\perp),$$

where

$$\alpha = \frac{\sqrt{2}}{0.94(1 - \delta_1)}, \quad \beta = 2\alpha(1 + \delta_1).$$

Hence, when $\beta\gamma < 1$, we have

$$\mathrm{Tr}(\mathbf{H}_T^\perp) = \|\mathbf{H}_T^\perp\|_1 \leq \frac{4(1 + \delta_1 + \gamma\alpha m^{-1/2})}{1 - \beta\gamma} \epsilon = c_1 \epsilon.$$

In addition, (6.6) then gives

$$\|\mathbf{H}_T\| \leq \frac{2m^{-1/2} + (1 + \delta_1)c_1}{0.94(1 - \delta_1)} \epsilon = c_2 \epsilon.$$

In conclusion,

$$\|\mathbf{H}\|_2 \leq \|\mathbf{H}_T\|_2 + \|\mathbf{H}_T^\perp\|_2 \leq \sqrt{2}\|\mathbf{H}_T\| + \|\mathbf{H}_T^\perp\|_1 \leq (\sqrt{2}c_2 + c_1)\epsilon = c_0 \epsilon,$$

and we also have $\|\mathbf{H}\| \leq (c_2 + c_1)\epsilon$.

It remains to show why the fact that $\hat{\mathbf{X}}$ is close to $\mathbf{X} = \mathbf{x}\mathbf{x}^*$ in the Frobenius or operator norm produces a good estimate of \mathbf{x} (recall that $\mathbf{x} = \mathbf{e}_1$). Set $\epsilon_0 := \|\hat{\mathbf{X}} - \mathbf{X}\| \leq c_0 \epsilon$. Below, $\hat{\lambda}_1 \geq 0$ is the largest eigenvalue of $\hat{\mathbf{X}} \succeq 0$, and $\hat{\mathbf{u}}_1$ the first eigenvector. Likewise, $\lambda_1 = 1$ is the top eigenvalue of $\mathbf{X} = \mathbf{e}_1\mathbf{e}_1^*$. Since $\mathrm{Tr}(\hat{\mathbf{X}}) \leq \mathrm{Tr}(\mathbf{X})$,

$$\hat{\lambda}_1 \leq \lambda_1.$$

In the other direction, we know from perturbation theory that

$$|\lambda_1 - \hat{\lambda}_1| \leq \|\hat{\mathbf{X}} - \mathbf{X}\| = \epsilon_0.$$

Assuming that $\epsilon_0 < 1$, this gives $\hat{\lambda}_1 \in [1 - \epsilon_0, 1]$. The sin- θ -Theorem [11] implies that

$$|\sin \theta| \leq \frac{\|\hat{\mathbf{X}} - \mathbf{X}\|}{|\hat{\lambda}_1|} \leq \frac{\epsilon_0}{1 - \epsilon_0},$$

where $0 \leq \theta \leq \pi/2$ is the angle between the spaces spanned by $\hat{\mathbf{u}}_1$ and \mathbf{e}_1 . Writing

$$\hat{\mathbf{u}}_1 = \cos \theta \mathbf{e}_1 + \sin \theta \mathbf{e}_1^\perp$$

in which \mathbf{e}_1^\perp is a unit vector orthogonal to \mathbf{e}_1 , Pythagoras' relationship gives

$$\|\mathbf{e}_1 - \sqrt{\hat{\lambda}_1} \hat{\mathbf{u}}_1\|_2^2 = (1 - \sqrt{\hat{\lambda}_1} \cos \theta)^2 + \hat{\lambda}_1 \sin^2 \theta.$$

Since $\cos \theta = \sqrt{1 - \sin^2 \theta}$, we have

$$1 \geq \sqrt{\hat{\lambda}_1} \cos \theta \geq \sqrt{1 - \epsilon_0 - \frac{\epsilon_0^2}{1 - \epsilon_0}} \geq 1 - \epsilon_0$$

for $\epsilon_0 < 1/3$. Hence,

$$\|\mathbf{e}_1 - \sqrt{\hat{\lambda}_1} \hat{\mathbf{u}}_1\|_2^2 \leq \epsilon_0^2 + \frac{\epsilon_0^2}{(1 - \epsilon_0)^2} \leq \frac{13}{4} \epsilon_0^2$$

provided $\epsilon_0 < 1/3$. Since we always have

$$\|\mathbf{e}_1 - \sqrt{\hat{\lambda}_1} \hat{\mathbf{u}}_1\|_2 \leq \|\mathbf{e}_1\|_2 + \sqrt{\hat{\lambda}_1} \|\hat{\mathbf{u}}_1\|_2 \leq 2,$$

we have established

$$\|\mathbf{e}_1 - \sqrt{\hat{\lambda}_1} \hat{\mathbf{u}}_1\|_2 \leq C_0 \min(\epsilon, 1).$$

This holds for all values of ϵ_0 and proves the claim in the case where $\|\mathbf{x}\|_2 = 1$. The general case is obtained via a simple rescaling.

As mentioned above, we proved the theorem for Gaussian \mathbf{z}_i 's but it is clear that our results hold true for vectors sampled uniformly at random on the sphere of radius \sqrt{n} . The reason is that of course, $\|\mathbf{z}_i\|_2$ deviates very little from \sqrt{n} . Formally, set $\tilde{\mathbf{z}}_i = \lceil \sqrt{n}/\|\mathbf{z}_i\|_2 \rceil \mathbf{z}_i$ so that these new vectors are independently and uniformly distributed on the sphere of radius \sqrt{n} . Then

$$\langle \mathbf{X}, \tilde{\mathbf{z}}_i \tilde{\mathbf{z}}_i^* \rangle = \frac{n}{\|\mathbf{z}_i\|_2^2} \langle \mathbf{X}, \mathbf{z}_i \mathbf{z}_i^* \rangle,$$

and thus $\langle \mathbf{X}, \mathbf{z}_i \mathbf{z}_i^* \rangle$ is between $(1 - \delta_2) \langle \mathbf{X}, \tilde{\mathbf{z}}_i \tilde{\mathbf{z}}_i^* \rangle$ and $(1 + \delta_2) \langle \mathbf{X}, \tilde{\mathbf{z}}_i \tilde{\mathbf{z}}_i^* \rangle$ with very high probability. This holds uniformly over all Hermitian matrices. Thus if $\tilde{\mathcal{A}}(\mathbf{X}) = \{\tilde{\mathbf{z}}_i^* \mathbf{X} \tilde{\mathbf{z}}_i\}_{1 \leq i \leq m}$,

$$(1 - \delta_2) \|\tilde{\mathcal{A}}(\mathbf{X})\|_q \leq \|\mathcal{A}(\mathbf{X})\|_q \leq (1 + \delta_2) \|\tilde{\mathcal{A}}(\mathbf{X})\|_q$$

for any $1 \leq q \leq \infty$.

Now take $b_i = |\langle \mathbf{x}, \tilde{\mathbf{z}}_i \rangle|^2 + \nu_i$ and solve (1.9) to get $\tilde{\mathbf{X}} = \mathbf{X} + \tilde{\mathbf{H}}$. Going through the same steps as above by using the relationships between \mathcal{A} and $\tilde{\mathcal{A}}$ throughout, and by using the dual certificate \mathbf{Y} associated with \mathcal{A} , we obtain

$$\|\tilde{\mathcal{A}}(\tilde{\mathbf{H}})\|_2 \leq 2\epsilon, \quad |\langle \tilde{\mathbf{H}}, \mathbf{Y} \rangle| \leq 2\epsilon(1 + \delta_1)(1 + \delta_2),$$

and

$$\text{Tr}(\tilde{\mathbf{H}}_{T^\perp}) \leq (1 + \delta_2)c_1\epsilon, \quad \|\tilde{\mathbf{H}}_T\| \leq (1 + \delta_2)c_2\epsilon.$$

Therefore,

$$\|\tilde{\mathbf{H}}\|_2 \leq (1 + \delta_2)(\sqrt{2}c_2 + c_1)\epsilon.$$

The rest of the proof goes through just the same.

7 Numerical Simulations

In this section we illustrate our theoretical results with numerical simulations. In particular, we will demonstrate PhaseLift's robustness vis a vis additive noise.

We consider the setup in Section 1.4, where the measurements are contaminated with additive noise. The solution to (1.9) is computed using the following regularized nuclear-norm minimization problem:

$$\begin{aligned} & \text{minimize} && \frac{1}{2} \|\mathcal{A}(\mathbf{X}) - \mathbf{b}\|_2^2 + \lambda \text{Tr}(\mathbf{X}) \\ & \text{subject to} && \mathbf{X} \succeq 0. \end{aligned} \tag{7.1}$$

It follows from standard optimization theory [19] that (7.1) is equivalent to (1.9) for some value of λ . Hence, we use (7.1) to compute the solution of (1.9) by determining via a simple and efficient bisection search the largest value $\lambda(\epsilon)$ such that $\|\mathcal{A}(\mathbf{X}) - \mathbf{b}\|_2 \leq \epsilon$. The numerical algorithm to solve (7.1) was implemented in Matlab using TFOCS [5]. We then extract the largest rank-1 component as described in Section 1.4 to obtain an approximation $\hat{\mathbf{x}}$.

We will use the relative mean squared error (MSE) and the relative root mean squared error (RMS) to measure performance. However, since a solution is only unique up to global phase, it does not make sense to compute the distance between \mathbf{x} and its approximation $\hat{\mathbf{x}}$. Instead we compute the distance modulo a global phase term and define the relative MSE between \mathbf{x} and $\hat{\mathbf{x}}$ as

$$\min_{c:|c|=1} \frac{\|c\mathbf{x} - \hat{\mathbf{x}}\|_2^2}{\|\mathbf{x}\|_2^2}.$$

The (relative) RMS is just the square root of the (relative) MSE.


In the first set of experiments, we investigate how the reconstruction algorithm performs as the noise level increases. The test signal is a complex-valued signal of length $n = 128$ with independent Gaussian complex entries (each entry is of the form $a + ib$ where a and b are independent $\mathcal{N}(0, 1)$ variables) so that the real and imaginary parts are independent white noise sequences. Obviously, the signal is arbitrary. We use $m = 6n$ measurement vectors sampled independently on the unit sphere \mathbb{C}^n .

We generate noisy data from both a Gaussian model and a Poisson model. In the Gaussian model, $b_i \sim \mathcal{N}(\mu_i, \sigma^2)$ where $\mu_i = |\langle \mathbf{x}, \mathbf{z}_i \rangle|^2$ and σ is adjusted so that the total noise power is bounded by ϵ^2 . In the Poisson model, $b_i \sim \text{Poi}(\mu_i)$ and the noise $b_i - \mu_i$ is rescaled to achieve a desired total power as above (we might do without this rescaling as well but have decided to work with a prescribed signal-to-noise ratio SNR for simplicity of exposition). We do this for five different SNR levels,⁷ ranging from 5dB to 100dB. However, we point out that we do not make use of the noise statistics in our reconstruction algorithm⁸, since our purpose is only to assume an upper bound on the total noise power, as in Theorem 1.2.


For each SNR level, we repeat the experiment ten times with different noise terms, different signals, and different random measurement vectors; we then record the average relative RMS over these ten experiments. Figure 4(a) shows the average relative MSE in dB (the values of $10 \log_{10}(\text{rel. MSE})$ are plotted) versus the SNR for Poisson noise. In each case, the performance degrades very gracefully with decreasing SNR, as predicted by Theorem 1.2. Debiasing as described at the end of Section 1.4 leads to a further improvement in the reconstruction for low SNR, as illustrated in Figure 4(b). The results for Gaussian noise are comparable, see Figure 5.

⁷The SNR of two signals $\mathbf{x}, \hat{\mathbf{x}}$ with respect to \mathbf{x} is defined as $10 \log_{10} \|\mathbf{x}\|_2^2 / \|\mathbf{x} - \hat{\mathbf{x}}\|_2^2$. So we say that the SNR is 10dB if $10 \log_{10} \|\mathbf{x}\|_2^2 / \|\mathbf{x} - \hat{\mathbf{x}}\|_2^2 = 10$.

⁸We refer to [7] for efficient ways to incorporate statistical noise models into the reconstruction algorithm.




(a)




(b)

Figure 4: Performance of PhaseLift for Poisson noise. The stability of the algorithm is apparent as its performance degrades gracefully with decreasing SNR. (a) Relative MSE on a log-scale for the non-debiased recovery. (b) Relative RMS for the original and debiased recovery.



(a)



(b)

Figure 5: Performance of PhaseLift for Gaussian noise. (a) Relative MSE on a log-scale for the non-debiased recovery. (b) Relative RMS for the original and the debiased recovery.




Figure 6: Oversampling rate versus relative RMS.

In the next experiment, we collect Poisson data about a complex-valued random signal just as above, and work with a fixed SNR set to 15dB. The number of measurements varies so that the oversampling rate m/n is between 5 and 22 (m is thus between $n \log n$ and $4.5n \log n$). We repeat the experiment ten times with different noise terms and different random measurement vectors for each oversampling rate; we then record the average relative RMS. Figure 6 shows the average relative RMS of the solution to (1.5) versus the oversampling rate. We observe that the decrease in the RMS is inversely proportional to the number of measurements. For instance, the error reduces by a factor of two when we double the number of measurements. If instead we hold the standard deviation of the errors at a constant level, the mean squared error (MSE) reduces by a factor of about two when we double the number of measurements.

8 Discussion

In this paper, we have shown that it is possible to recover a signal exactly (up to a global phase factor) from the knowledge of the magnitude of its inner products with a family of sensing vectors $\{z_i\}$. The fact that on the order of $n \log n$ magnitude measurements $|\langle x, z_i \rangle|^2$ uniquely determine x is not surprising. The part we find unexpected, however, is that what appears to be a combinatorial problem is solved exactly by a convex program. Further, we have established the existence of a noise-aware recovery procedure—also based on a tractable convex program—which is robust vis a vis additive noise. To the best of our knowledge, there are no other results—about the recovery of an arbitrary signal from noisy quadratic data—of this kind.

An appealing research direction is to study the recovery of a signal from other types of intensity measurements, and consider other families of sensing vectors. In particular, *structured* random families would be of great interest. It also seems plausible that assuming stochastic errors in

Theorem 1.2 would allow to derive sharper error bounds; it would be of interest to know if this is indeed the case. We leave this to future work.

Acknowledgements

E. C. is partially supported by NSF via grant CCF-0963835 and the 2006 Waterman Award, and by AFOSR under grant FA9550-09-1-0643. T. S. acknowledges partial support from the NSF via grants DMS-0811169 and DTRA-DMS 1042939, and from DARPA via grant N66001-11-1-4090. V. V. is supported by the Department of Defense (DoD) through the National Defense Science and Engineering Graduate Fellowship (NDSEG) Program. E. C. thanks Mahdi Soltanolkotabi for help with Figure 1. V. V. acknowledges fruitful conversations with Bernd Sturmfels.

References

- [1] R. Balan, B. Bodmann, P. Casazza, and D. Edidin. Fast algorithms for signal reconstruction without phase. In *Wavelets XII*, volume 6701 of *Proc. SPIE*, pages 670111920–670111932, 2007.
- [2] R. Balan, B. Bodmann, P. Casazza, and D. Edidin. Painless reconstruction from magnitudes of frame coefficients. *J. Four. Anal. Appl.*, 15:488–501, 2009.
- [3] R. Balan, P. Casazza, and D. Edidin. On signal reconstruction without noisy phase. *Appl. Comp. Harm. Anal.*, 20:345–356, 2006.
- [4] C. Beck and R. D’Andrea. Computational study and comparisons of LFT reducibility methods. In *Proceedings of the American Control Conference*, pages 1013–1017, 1998.
- [5] S. Becker, E. Candès, and M. Grant. Templates for convex cone problems with applications to sparse signal recovery. Technical report, Department of Statistics, Stanford University, 2010. Preprint available at <http://tfocs.stanford.edu/tfocs/paper.shtml>.
- [6] R. Berinde, A. Gilbert, P. Indyk, H. Karloff, and M. Strauss. Combining geometry and combinatorics: A unified approach to sparse signal recovery. *CoRR*, pages –1–1, 2008.
- [7] E. Candès, Y. Eldar, T. Strohmer, and V. Voroninski. Phase retrieval via matrix completion. *Preprint*, 2011.
- [8] E. Candès and Y. Plan. Matrix completion with noise. *Proceedings of the IEEE*, 98(6):925–936, 2010.
- [9] E. J. Candès and T. Tao. Decoding by linear programming. *IEEE Trans. on Information Theory*, 51(2):4203–4215, 2005.
- [10] A. Chai, M. Moscoso, and G. Papanicolaou. Array imaging using intensity-only measurements. Technical report, Stanford University, 2010.
- [11] C. Davis and W. M. Kahan. The rotation of eigenvectors by a perturbation. III. *SIAM J. Numer. Anal.*, 7:1–46, 1970.
- [12] J. Finkelstein. Pure-state informationally complete and “really” complete measurements. *Phys. Rev. A*, 70:052107, 2004.
- [13] D. Gross. Recovering low-rank matrices from few coefficients in any basis, 2009. Available at <http://arxiv.org/abs/0910.1879>.
- [14] D. Gross, Y. Liu, S. T. Flammia, S. Becker, and J. Eisert. Quantum-state tomography via compressed sensing. *Physical Review Letters*, 105(15), 2010.
- [15] N. Hurt. *Phase Retrieval and Zero Crossings*. Kluwer Academic Publishers, Norwell, MA, 1989.

- [16] M. Klivanov, P. Sacks, and A. Tikhonravov. The phase retrieval problem. *Inverse problems*, 11:1995, 1–28.
- [17] Y. Lu and M. Vetterli. Sparse spectral factorization: Unicity and reconstruction algorithms. In *The 36th International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Prague, Czech Republic, 2011.
- [18] M. Mesbahi and G. P. Papavassilopoulos. On the rank minimization problem over a positive semidefinite linear matrix inequality. *IEEE Transactions on Automatic Control*, 42(2):239–243, 1997.
- [19] R. Rockafellar. *Convex analysis*. Princeton Landmarks in Mathematics and Physics. Princeton University Press, 1970.
- [20] H. Sahinoglou and S. Cabrera. On phase retrieval of finite-length sequences using the initial time sample. *IEEE Transactions on Circuits and Systems*, 38(5):954–958, 1991.
- [21] R. Vershynin. Introduction to the non-asymptotic analysis of random matrices. In Y. C. Eldar and G. Kutyniok, editors, *Compressed Sensing: Theory and Applications*. Cambridge University Press, 2010. To Appear. Preprint available at <http://www-personal.umich.edu/~romanv/papers/papers.html>.