

PhaseLift is robust to a constant fraction of arbitrary errors

Paul Hand^{a,1}

^a*Rice University*
Department of Computational and Applied Mathematics,
MS-134, 6100 Main Street, Houston, TX 77005

Abstract

Consider the task of recovering an unknown n -vector from phaseless linear measurements. This nonconvex problem may be convexified into a semidefinite rank-one matrix recovery problem, known as PhaseLift. Under a linear number of Gaussian measurements, PhaseLift recovers the unknown vector exactly with high probability. Under noisy measurements, the solution to a variant of PhaseLift has error proportional to the ℓ_1 norm of the noise. In the present paper, we study the robustness of this variant of PhaseLift to gross, arbitrary corruptions. We prove that PhaseLift can tolerate noise and a small, fixed fraction of gross errors, even in the highly underdetermined regime where there are only $O(n)$ measurements. The lifted phase retrieval problem can be viewed as a rank-one robust Principal Component Analysis (PCA) problem under generic rank-one measurements. From this perspective, the proposed convex program is simpler than the semidefinite version of the sparse-plus-low-rank formulation standard in the robust PCA literature. Specifically, the rank penalization through a trace term is unnecessary, and the resulting optimization program has no parameters that need to be chosen.

Keywords. phase retrieval, PhaseLift, matrix completion, robust principal component analysis

1. Introduction

This paper establishes robustness of an algorithm for recovering a vector $x_0 \in \mathbb{R}^n$ from phaseless linear measurements that contain noise and a constant fraction of gross, arbitrary errors. That is, for fixed measurement vectors $a_i \in \mathbb{R}^n$ for $i = 1 \dots m$, our task is to find x_0 satisfying

$$b_i = |\langle x_0, a_i \rangle|^2 + \varepsilon_i + \eta_i \quad (1)$$

for known $b_i \in \mathbb{R}$, known a_i , and unknown η_i and ε_i . Here η_i will represent noise in the measurements, and ε_i will represent gross, arbitrary corruptions. This recovery problem is known as phase retrieval. Measurements of form (1) arise in several applications, such as X-ray crystallography, optics, and microscopy [10, 16, 1]. In such applications, gross corruptions in some measurements may be due to sensor failure, occlusions, or other effects.

Recently, researchers have introduced algorithms for the phase retrieval problem that have provable recovery guarantees [1, 4]. The insight of these methods is that the phase retrieval problem can be convexified by lifting it to the space of matrices. That is, instead of searching for the vector x_0 , one can search for the lifted matrix $x_0 x_0^t$. The quadratic measurements (1) then become linear measurements on this lifted matrix. As the desired matrix is semidefinite and rank-one, one can write a rank minimization problem under the semidefinite and data constraints, which has a convex relaxation known as PhaseLift. In this noiseless case, PhaseLift is the program

$$\min_X \text{tr}(X) \text{ subject to } X \succeq 0, \{a_i^t X a_i = b_i\}_{i=1 \dots m} \quad (2)$$

Here, the trace of X is a convex proxy for the rank of a positive semidefinite X . An estimate for the underlying signal x_0 can be found by computing the leading eigenvector of the optimizer of (2).

¹Email: hand@rice.edu, Phone: 713-348-3702. Fax: 713-348-5318

Similar to [4, 7, 2], we will seek recovery guarantees for independent identically distributed Gaussians

$$a_i \sim \mathcal{N}(0, I_n) \in \mathbb{R}^n$$

Under this data model, [7] and [2] have shown that (2) can be simplified to the semidefinite feasibility problem

$$\text{find } X \succeq 0 \text{ such that } \{a_i^t X a_i = b_i\}_{i=1\dots m}.$$

This feasibility problem succeeds at finding $x_0 x_0^t$ exactly with high probability when $m \geq cn$ for a sufficiently large c [2]. This scaling is quite surprising because there are only $O(n)$ measurements for an $O(n^2)$ dimensional object. As discussed in [7], the semidefinite cone is sufficiently ‘pointy’ that the high-dimensional affine space of data-consistent matrices intersects the semidefinite cone only at exactly one point.

In the noisy case without gross errors, that is for $\varepsilon = 0$, [2] showed that the PhaseLift variant

$$\min \sum_i |a_i^t X a_i - b_i| \text{ subject to } X \succeq 0 \tag{3}$$

successfully recovers a matrix near $x_0 x_0^t$ with high probability. Specifically, they prove that the solution \hat{X} to (3) satisfies $\|\hat{X} - x_0 x_0^t\|_F \leq C_0 \|\eta\|_1 / m$ with high probability. From \hat{X} , an estimate of x_0 can be obtained by

$$\hat{x}_0 = \sqrt{\hat{\lambda}_1} \hat{u}_1$$

where $(\hat{\lambda}_1, \hat{u}_1)$ is the leading eigenvector and eigenvalue pair for \hat{X} . When specialized to the real-valued case, [2] proves that $|\hat{x}_0 - \pm x_0| \leq C_0 \min(\|x_0\|, \|\eta\|_1 / m \|x_0\|)$ for some C_0 .

The contribution of the present paper is to show that the program (3) is additionally robust against a constant fraction of arbitrary errors. For a fixed set of coefficients that contain gross errors, we show that approximate recovery succeeds with high probability for arbitrary signals and arbitrary values of the gross errors.

Theorem 1. *There exist positive numbers $f_{\min}, \gamma, c, C, C'$ such that the following holds. Let $m \geq cn$. Fix a set $S \subset \{1 \dots m\}$ such that $|S|/m \leq f_{\min}$. On an event of probability at least $1 - e^{-\gamma m}$, for any $x_0 \in \mathbb{R}^n$ and for any ε with $\text{supp}(\varepsilon) \subseteq S$, the minimizer \hat{X} to (3) satisfies*

$$\|\hat{X} - x_0 x_0^t\|_F \leq C \frac{\|\eta\|_1}{m}.$$

The resulting estimate for x_0 satisfies

$$\|\hat{x}_0 - \pm x_0\| \leq C' \min\left(\|x_0\|, \frac{\|\eta\|_1}{m \|x_0\|}\right)$$

Note that this high-probability result is universal over signals x_0 and corruptions ε with fixed support. That is, a single realization of the measurement vectors a_i will allow recovery with high probability for an arbitrary and even adversarial x_0 . Further, for a fixed set of corrupted entries, the values of the gross corruptions can be arbitrary and even adversarial.

In the case of no gross errors, in which $\varepsilon = 0$, this theorem reduces to the result in [2] and the rank-one case of [12]. In the noiseless case, in which $\eta = 0$, the theorem guarantees exact recovery of x_0 with high probability under a linear number of measurements, of which a constant fraction are corrupted.

We now explore the optimality of this theorem. The scaling of m versus n is information theoretically optimal and has no unnecessary logarithmic factors. The noise scaling is the same as in [2], and its optimality was established there. For arbitrary errors, the fixed fraction of gross errors can not be extended to a case where $f_{\min} \geq 1/2$ because one could build a problem where half of the measurements are due to an x_0 and the other half are due to some x_1 . In such a case, recovery would be impossible.

As stated, Theorem 1 only applies to real-valued signals. As the applications of phase retrieval typically involve complex measurements, it is an important follow-up question to determine if the corresponding theorem holds for complex measurements. We leave this question for future research. See Section 2.3 for a brief comment on this extension.

1.1. Relation to Robust PCA

Much recent work in matrix completion has studied the recovery of low-rank matrices from arbitrary corruptions to its entries. This is a task known as robust Principal Component Analysis (PCA). Results in this framework typically involve measuring some of the entries of a low rank $n \times n$ matrix X_0 and assuming that some fraction of those measurements are arbitrarily corrupted, giving the data matrix A . The matrix X_0 can then be recovered under certain conditions by a sparse-plus-low-rank convex program:

$$\min \lambda \|X\|_* + \|E\|_1 \text{ such that } \mathcal{P}(X + E) = \mathcal{P}(A) \quad (4)$$

where $\|X\|_*$ is the nuclear norm of X , λ is a constant, $\|E\|_1$ is the ℓ_1 norm of the vectorization of E , and \mathcal{P} is the projection of a matrix onto the observed entries. [3, 5, 11, 20, 8, 6, 14]. Results from this formulation have been quite surprising. For example, under an appropriate choice of λ and under an incoherence assumption, the sparse-plus-low-rank decomposition succeeds for sufficiently low rank X when $O(n^2)$ entries are measured and a small fraction of them have arbitrary errors [3]. Subsequent results have been proved that only require $m \gtrsim rn \text{ polylog}(n)$ measurements, where r is the rank of X [6, 14]. This result is information theoretically optimal except for the polylogarithmic factor.

The present paper can be viewed as a rank-one semidefinite robust PCA problem under generic rank-one measurements. From this perspective, we would naturally formulate the PhaseLift problem under gross errors by

$$\min \lambda \text{tr}(X) + \sum_i |a_i^t X a_i - b_i| \text{ subject to } X \succeq 0. \quad (5)$$

The present paper shows that explicit rank penalization by the trace term is not fundamental for exact recovery in the presence of arbitrary errors. That is, (5) can be simplified by taking $\lambda = 0$. The resulting program has no free parameters that need to be explicitly tuned. As in [7, 2], the positive semidefinite cone provides enough of a constraint to enforce low-rankness. The present work differs from the standard robust PCA literature in that the measurements are generic and are not direct samples of the entries of the unknown matrix. As Gaussian measurements are easier to analyze, we can provably achieve information theoretic scaling without additional logarithmic factors.

1.2. Numerical simulation

We now use numerical simulation to provide empirical evidence for the performance of (3). Let the signal length n vary from 5 to 50, and let the number of measurements m vary from 10 to 250. Let $x_0 = e_1$. For each (n, m) , we consider measurements such that

$$\begin{cases} b_i \sim \text{Uniform}([0, 10^4]) & \text{if } 1 \leq i \leq \lceil 0.05m \rceil, \\ b_i = |\langle x_0, a_i \rangle|^2 & \text{otherwise.} \end{cases} \quad (6)$$

We attempt to recover $x_0 x_0^t$ by solving (3) using the SDPT3 solver [17, 18] and YALMIP [15]. For a given optimizer \hat{X} , define the capped relative error as

$$\min(\|\hat{X} - x_0 x_0^t\|_F / \|x_0 x_0^t\|_F, 1).$$

Figure 1 plots the average capped relative error over 10 independent trials. It provides empirical evidence that the matrix recovery problem (3) succeeds under a linear number of measurements, even when a constant fraction of them contain very large errors.

2. Proofs

Let $\mathcal{A} : \mathcal{S}^n \rightarrow \mathbb{R}^m$ be defined by the mapping $X \mapsto (a_i^t X a_i)_{i=1 \dots m}$, where \mathcal{S}^n is the space of symmetric real-valued $n \times n$ matrices. Note that $\mathcal{A}^* \lambda = \sum_i \lambda_i a_i a_i^t$. Let \mathcal{A}_S be the restriction of \mathcal{A} onto the coefficients given by the set S . Let e_i be the i th standard basis vector. Let $X_0 = x_0 x_0^t$. We can write the measurements (1) as

$$b = \mathcal{A} X_0 + \varepsilon + \eta.$$




Figure 1: Recovery error for the PhaseLift matrix recovery problem (3) as a function of n and m , when 5% of measurements contain large errors. Black represents an average recovery error of 100%. White represents zero average recovery error. Each block corresponds to the average from 10 independent trials. The solid curve depicts when the number of measurements equals the number of degrees of freedom in a symmetric $n \times n$ matrix. The number of measurements required for successful recovery appears to be linear in n , even with a small fraction of large errors.

Similarly, the optimization program (3) can be written as

$$\min \| \mathcal{A}X - b \|_1 \text{ such that } X \succeq 0.$$

We introduce the following notation. Let $\|X\|_*$ be the nuclear norm of the matrix X . When $X \succeq 0$, $\|X\|_* = \text{tr}(X)$. Denote the Frobenius and spectral norms of X as $\|X\|_F$ and $\|X\|$, respectively. Let $\langle X, Y \rangle = \text{tr}(Y^t X)$ be the Hilbert-Schmidt inner product. Given x_0 , let $T_{x_0} = \{y x_0^t + x_0 y^t \mid y \in \mathbb{R}^n\}$. Note that T_{e_1} is the space of symmetric matrices supported on their first row and column. The orthogonal complement $T_{e_1}^\perp$ is then the space of matrices supported in the lower-right $(n-1) \times (n-1)$ block. When x_0 is clear, we will simply write T instead of T_{x_0} . Let I be the identity matrix, and let $\mathbb{1}(E)$ be the indicator function of the event E .

2.1. Recovery by dual certificates

The proof of Theorem 1 will be based on dual certificates, as in [9, 4, 7, 2]. A dual certificate is an optimal variable for the dual problem to (3). Its existence certifies the correctness of a minimizer to (3).

The first order optimality conditions at X_0 for (3) are given by

$$\tilde{Y} = \mathcal{A}^* \tilde{\lambda} \tag{7}$$

$$\tilde{\lambda} \in \partial \| \cdot \|_1(-\varepsilon) \tag{8}$$

$$\tilde{Y} \succeq 0 \tag{9}$$

$$\langle \tilde{Y}, X_0 \rangle = 0 \tag{10}$$

where $\partial \| \cdot \|_1(-\varepsilon)$ is the subgradient of the ℓ_1 norm evaluated at $-\varepsilon$. Note that (9) and (10) imply $\tilde{Y}_T = 0$. Such a \tilde{Y} would be dual certificate for (3). Unfortunately, constructing such a \tilde{Y} that exactly satisfies these conditions is difficult. As in [9, 4, 7, 2], we seek an inexact dual certificate, which approximately satisfies these conditions. Specifically, we will build a dual certificate $Y = \mathcal{A}^* \lambda$ that satisfies

$$Y_{T^\perp} \succeq I_{T^\perp} \tag{11}$$

$$\|Y_T\|_F \leq 1/2 \tag{12}$$

$$\begin{cases} \lambda_i = -\frac{\varepsilon}{m} \text{sgn}(\varepsilon_i) & \text{if } \varepsilon_i \neq 0 \\ \lambda_i \leq \frac{\varepsilon}{m} & \text{if } \varepsilon_i = 0. \end{cases} \tag{13}$$

To prove that existence of such a Y will guarantee successful recovery of $x_0 x_0^t$ with high probability, we will rely on two technical lemmas. The first technical lemma provides ℓ_1 -isometry bounds on \mathcal{A} and was proven in [4].

Lemma 1 ([4]). *There exist constants c_0, γ_0 such that if $m \geq c_0 n$, then with probability at least $1 - e^{-\gamma_0 m}$,*

$$\frac{1}{m} \|\mathcal{A}(X)\|_1 \leq \left(1 + \frac{1}{16}\right) \|X\|_* \text{ for all } X, \quad (14)$$

$$\frac{1}{m} \|\mathcal{A}(X)\|_1 \geq 0.94 \left(1 - \frac{1}{16}\right) \|X\| \text{ for all symmetric, rank-2 } X \quad (15)$$

We will need simultaneous control of the ℓ_1 -isometry properties over several subsets of measurements.

Lemma 2. *There exists a constant $\tilde{\gamma}_0$ such that the following holds. Let $m \geq 100c_0 n$, and fix a support set S with $|S| = \lceil 0.01m \rceil$. There is an event E_S with probability at least $1 - e^{-\tilde{\gamma}_0 m}$ on which*

$$\frac{1}{|S^c|} \|\mathcal{A}_{S^c} X\|_1 \geq 0.94 \left(1 - \frac{1}{16}\right) \|X\| \text{ for all symmetric rank-2 } X, \quad (16)$$

$$\frac{1}{|S|} \|\mathcal{A}_S X\|_1 \leq \left(1 + \frac{1}{16}\right) \|X\|_* \text{ for all } X, \quad (17)$$

$$\frac{1}{|m|} \|\mathcal{A} X\|_1 \leq \left(1 + \frac{1}{16}\right) \|X\|_* \text{ for all } X. \quad (18)$$

The proof of Lemma 2 is immediate from Lemma 1.

In order to prove that a dual certificate guarantees recovery, we establish a technical result that an optimal solution $X_0 + H$ to (3) lies near the cone $\|H_{T^\perp}\|_* \geq 0.56 \|H_T\|_F$ with high probability.

Lemma 3. *Fix a support set S with $|S| = \lceil 0.01m \rceil$ and let $m \geq 100c_0 n$. On the event E_S from Lemma 2, for all x and for all ε with $\text{supp}(\varepsilon) \subseteq S$, any optimal $X_0 + H$ satisfies $\|H_{T^\perp}\|_* \geq 0.56 \|H_T\|_F - \frac{2}{m} \|\eta\|_1$.*

Proof. By assumption, $\|\mathcal{A}H - \varepsilon - \eta\|_1 \leq \|\varepsilon + \eta\|_1$. By the additivity of the ℓ_1 norm over vectors with disjoint supports,

$$\|\mathcal{A}_{S^c} H - \eta_{S^c}\|_1 + \|\mathcal{A}_S H - \varepsilon - \eta_S\|_1 \leq \|\varepsilon + \eta_S\|_1 + \|\eta_{S^c}\|_1. \quad (19)$$

By the triangle inequality, we have

$$\|\mathcal{A}_{S^c} H - \eta_{S^c}\|_1 + \|\varepsilon + \eta_S\|_1 - \|\mathcal{A}_S H\|_1 \leq \|\varepsilon + \eta_S\|_1 + \|\eta_{S^c}\|_1, \quad (20)$$

which implies

$$\|\mathcal{A}_{S^c} H\|_1 \leq 2\|\eta_{S^c}\|_1 + \|\mathcal{A}_S H\|_1 \quad (21)$$

Breaking (21) into its components on T and T^\perp and applying the triangle inequality, we have

$$\|\mathcal{A}_{S^c} H_T\|_1 - \|\mathcal{A}_{S^c} H_{T^\perp}\|_1 \leq 2\|\eta_{S^c}\|_1 + \|\mathcal{A}_S H_T\|_1 + \|\mathcal{A}_S H_{T^\perp}\|_1 \quad (22)$$

$$\Rightarrow \|\mathcal{A}_{S^c} H_T\|_1 \leq 2\|\eta_{S^c}\|_1 + \|\mathcal{A}_S H_T\|_1 + \|\mathcal{A}_S H_{T^\perp}\|_1 \quad (23)$$

We now apply the ℓ_1 isometry bounds from Lemma 1 on each term of (23). On the event E_S ,

$$\|\mathcal{A}_{S^c} H_T\|_1 \geq 0.94 \left(1 - \frac{1}{16}\right) |S^c| \|H_T\| \quad (24)$$

$$\geq 0.94 \left(1 - \frac{1}{16}\right) \frac{|S^c|}{\sqrt{2}} \|H_T\|_F, \quad (25)$$

where the second inequality follows because H_T has rank at most 2. On the event E_S ,

$$\|\mathcal{A}_S H_T\|_1 \leq |S| \left(1 + \frac{1}{16}\right) \|H_T\|_* \quad (26)$$

$$\leq |S| \left(1 + \frac{1}{16}\right) \sqrt{2} \|H_T\|_F \quad (27)$$

On the event E_S ,

$$\|\mathcal{A}H_{T^\perp}\|_1 \leq m \left(1 + \frac{1}{16}\right) \|H_{T^\perp}\|_* \quad (28)$$

Combining (23)–(28), we have

$$\left(\frac{0.94}{\sqrt{2}} \frac{(1 - \frac{1}{16}) |S^c|}{(1 + \frac{1}{16}) m} - \sqrt{2} \frac{|S|}{m}\right) \|H_T\|_F \leq \frac{2}{m} \|\eta_{S^c}\|_1 + \|H_{T^\perp}\|_* \quad (29)$$

Thus, $0.56\|H_T\|_F \leq \frac{2}{m}\|\eta\|_1 + \|H_{T^\perp}\|_*$ on the event E_S . \square

We may now prove that existence of an inexact dual certificate (11)–(13) will guarantee successful recovery of a matrix near $X_0 = x_0 x_0^t$ with high probability, provided that there are few enough arbitrary errors.

Lemma 4. *There exists a C such that the following holds. Fix S such that $|S| = \lceil 0.01m \rceil$. Let $m \geq 100c_0n$. Fix $x_0 \in \mathbb{R}^n$. Fix $\varepsilon \in \mathbb{R}^m$ such that $\text{supp}(\varepsilon) \subseteq S$. Suppose that there exists $Y = \mathcal{A}^* \lambda$ satisfying (11)–(13). Then, on the event E_S from Lemma 2, a minimizer \hat{X} of (3) satisfies $\|\hat{X} - X_0\|_F \leq C \frac{\|\eta\|_1}{m}$.*

Proof. Let $\hat{X} = X_0 + H$ be a minimizer for (3), which implies that $\|\mathcal{A}(X_0 + H) - b\|_1 \leq \|\mathcal{A}X_0 - b\|_1$. That is,

$$\|\mathcal{A}H - \varepsilon - \eta\|_1 \leq \|\varepsilon + \eta\|_1.$$

Letting $\alpha = 7/m$, condition (13) gives

$$\lambda/\alpha \in \partial \|\cdot\|_1(-\varepsilon) \quad (30)$$

Hence,

$$\|-\varepsilon\|_1 + \langle \lambda/\alpha, \mathcal{A}H - \eta \rangle \leq \|\varepsilon\|_1 + \|\eta\|_1 \quad (31)$$

$$\Rightarrow \langle \lambda, \mathcal{A}H - \eta \rangle \leq \alpha \|\eta\|_1 \quad (32)$$

$$\Rightarrow \langle Y, H \rangle \leq \langle \lambda, \eta \rangle + \alpha \|\eta\|_1 \quad (33)$$

$$\Rightarrow \langle Y, H \rangle \leq 2\alpha \|\eta\|_1 \quad (34)$$

Decomposing (34) into T and T^\perp , we have

$$\langle Y_T, H_T \rangle \leq -\langle Y_{T^\perp}, H_{T^\perp} \rangle + 2\alpha \|\eta\|_1 \quad (35)$$

As $Y_{T^\perp} \succeq 0$ and $H_{T^\perp} \succeq 0$, we have

$$\langle Y_{T^\perp}, H_{T^\perp} \rangle \leq |\langle Y_T, H_T \rangle| + 2\alpha \|\eta\|_1 \quad (36)$$

By conditions (11)–(12)

$$\|H_{T^\perp}\|_* \leq |\langle Y_{T^\perp}, H_{T^\perp} \rangle| \leq |\langle Y_T, H_T \rangle| + 2\alpha \|\eta\|_1 \leq \frac{1}{2} \|H_T\|_F + 2\alpha \|\eta\|_1 \quad (37)$$

By Lemma 3, on the event E_S ,

$$\|H_{T^\perp}\|_* \geq 0.56 \|H_T\|_F - \frac{2}{m} \|\eta\|_1. \quad (38)$$

Combining (37) and (38), and using $\alpha = 7/m$, we get

$$0.56 \|H_T\|_F \leq \frac{2}{m} \|\eta\|_1 + 0.5 \|H_T\|_F + \frac{14}{m} \|\eta\|_1 \quad (39)$$

So,

$$\|H_T\|_F \leq \frac{C'}{m} \|\eta\|_1 \text{ and thus } \|H_{T^\perp}\|_F \leq \frac{C''}{m} \|\eta\|_1 \quad (40)$$

We conclude $\|\hat{X} - X_0\|_F = \|H\|_F \leq \frac{C}{m} \|\eta\|_1$ for some C . \square

2.2. Construction of the dual certificate

We now construct the dual certificate for arbitrary x_0 . Our construction will be a modification to the dual certificate in [2]. Also similar to [2], we will build dual certificates with high probability on a net of x_0 . We will then use a continuity argument to get a dual certificate for a arbitrary x_0 .

Let S^+ and S^- be disjoint supersets of the indices over which ε is positive or negative, respectively. Let $S = S^+ \cup S^-$. For pedagogical purposes, S^+ and S^- should be thought of as exactly the indices over which ε is positive or negative. For technical reasons, we let them be supersets of cardinality linear in n , in order to use standard probability bounds. For a fixed choice of S^+ and S^- , let the inexact dual certificate Y be defined by

$$Y = \frac{1}{m} \left[\sum_{i \in S^+} -7a_i a_i^t + \sum_{i \in S^-} 7a_i a_i^t + \sum_{i \in S^c} \left[\beta_0 - \left| \langle a_i, \frac{x_0}{\|x_0\|} \rangle \right|^2 \mathbb{1} \left(\left| \langle a_i, \frac{x_0}{\|x_0\|} \rangle \right| \leq 3 \right) \right] a_i a_i^t \right] \quad (41)$$

where $\beta_0 = \mathbb{E} z^4 \mathbb{1}(|z| \leq 3) \approx 2.6728$ and z is a standard normal random variable. We will refer to each of the terms of in the right hand side of (41) as Y_+ , Y_- , and Y_0 , respectively.

The form of Y_0 is due to [2], and the intuition behind it is as follows. Note that $\mathbb{E}(a_i a_i^t) = I_n$ and $\mathbb{E}(|\langle a_i, e_1 \rangle|^2 a_i a_i^t) = \begin{pmatrix} \tilde{\beta}_0 & 0 \\ 0 & I_{n-1} \end{pmatrix}$, where $\tilde{\beta}_0 = \mathbb{E} z^4$ for a standard normal z . The construction $\frac{1}{m} \sum_{i=1}^m [\tilde{\beta}_0 - |\langle a_i, e_1 \rangle|^2] a_i a_i^t$ thus has expected value $(\tilde{\beta}_0 - 1)I_{T^\perp}$, which would provide the exact dual certificate conditions (9)–(10). As shown in [2], a satisfactory inexact dual certificate can be built with $m = O(n)$ and coefficients that are truncated to be no larger than $7/m$. In the present formulation, the terms Y_+ and Y_- are then set to have coefficients $\mp 7/m$ in order to satisfy (13).

We now show that for a fixed signal and a fixed pattern of signs of ε , that a dual certificate exists with high probability.

Lemma 5. *There exists constants c, γ^* such that the following holds. Let $m \geq cn$. Fix $x_0 \in \mathbb{R}^n$ and $\varepsilon \in \mathbb{R}^m$. Let S^+ and S^- be fixed disjoint sets of cardinality $\lceil 0.001m \rceil$. Then the dual certificate Y from (41) satisfies (13), $\|Y_T\|_F \leq 1/4$, and $\|Y_{T^\perp} - \frac{17}{10}I_{T^\perp}\| \leq \frac{3}{10}$ with probability at least $1 - e^{-\gamma^* m}$.*

Proof. Without loss of generality, it suffices to assume $x_0 = e_1$. It suffices to show that with high probability

$$\left\| Y_{0, T^\perp} - \frac{17}{10}I_{T^\perp} \right\| \leq 0.15, \quad (42)$$

$$\|Y_{0, T}\|_F \leq \frac{3}{20}, \quad (43)$$

$$\|Y_{\pm, T^\perp}\| \leq 0.015, \quad (44)$$

$$\|Y_{\pm, T}\|_F \leq 0.035. \quad (45)$$

First, we establish (42)–(43). By Lemma 2.3 in [2], there exist $\tilde{c}, \tilde{\gamma}$ such that if $|S^c| \geq \tilde{c}n$, then with probability at least $1 - e^{-\tilde{\gamma}|S^c|}$,

$$\left\| \frac{m}{|S^c|} Y_{0, T^\perp} - \frac{17}{10}I_{T^\perp} \right\| \leq 1/10, \text{ and} \quad (46)$$

$$\left\| \frac{m}{|S^c|} Y_{0, T} \right\| \leq 3/20. \quad (47)$$

Thus,

$$\left\| Y_{0, T^\perp} - \frac{17}{10}I_{T^\perp} \right\| \leq \left\| Y_{0, T^\perp} - \frac{|S^c|}{m} \frac{17}{10}I_{T^\perp} \right\| + \left(1 - \frac{|S^c|}{m} \right) \frac{17}{10} \leq 0.15$$

which establishes (42). By (47), we get (43) immediately.

Next, we establish (44). Let a' be the vector formed by the last $n - 1$ components of a . Observe that $\sum_{i \in S^+} a'_i a_i'^*$ is a Wishart random matrix. Standard estimates for singular values of random matrices with

Gaussian i.i.d. entries, such as Corollary 5.35 in [19], apply. If $|S^+| = \lceil 0.001m \rceil \geq \tilde{c}_0 n$, with probability at least $1 - e^{-\tilde{\gamma}_1 m}$ for some $\tilde{\gamma}_1$,

$$\left\| \frac{1}{|S^+|} \sum_{i \in S^+} a'_i a_i'^* - I_{T^\perp} \right\| \leq 1/2$$

Hence,

$$\left\| \frac{1}{|S^+|} \sum_{i \in S^+} a'_i a_i'^* \right\| \leq 3/2$$

Thus, $\|Y_{+,T^\perp}\| \leq \frac{3}{2} \frac{|S^+|}{m}$, and we arrive at (44). The bound for the Y_{-,T^\perp} term is identical.

Next, we establish (45). Note that $Y_+ = -7 \frac{|S^+|}{m} \cdot \frac{1}{|S^+|} \sum_{i \in S^+} a_i a_i^t$. As per Lemma 6, if $|S^+| = \lceil 0.001m \rceil \geq \tilde{c}_1 n$, then $\|Y_{+,T}\|_{\mathbb{F}} \leq 7 \frac{|S^+|}{m} \cdot 5 \leq 7 \cdot 0.001 \cdot 5$, with probability at least $1 - e^{-\tilde{\gamma}_1 m}$.

Thus (42)–(45) hold simultaneously with probability at least $1 - e^{-\gamma^* m}$ for some γ^* provided that $c \geq \max(2\tilde{c}, 1000\tilde{c}_0, 1000\tilde{c}_1)$. \square

The behavior of $Y_{\pm,T}$ relies on the following probability estimate for the behavior of a Gaussian Wishart matrix on T .

Lemma 6. *Let $x_0 = e_1$. Let $A = \frac{1}{m} \sum_{i=1}^m a_i a_i^t$. There exists $\tilde{c}_1, \tilde{\gamma}_1$ such that if $m \geq \tilde{c}_1 n$ then $\|A_T\|_F \leq 5$ with probability at least $1 - e^{-\tilde{\gamma}_1 m}$.*

Proof. Let y be the $(1, 1)$ entry of A , and let y' be the rest of the first column of A . Then $\|A_T\|_{\mathbb{F}}^2 = y^2 + \|y'\|_2^2$. So, $y = \frac{1}{m} \sum_{i=1}^m a_i^2(1)$. Hence, $my \sim \chi_m^2$. Standard results on the concentration of chi-squared variables, such as Lemma 1 in [13], give

$$\mathbb{P}(my \geq 4m) \leq e^{-\tilde{\gamma}_1 \cdot 2m}.$$

for some $\tilde{\gamma}_1$. Hence $\mathbb{P}(y^2 \geq 16) \leq e^{-\tilde{\gamma}_1 \cdot 2m}$.

Now, it remains to bound $\|y'\|_2^2$. We can write $y' = \frac{1}{m} Z' c$, where $Z' = [a'_1, \dots, a'_m]$ and $c_i = a_i(i)$, where Z' and c are independent. Note that $\|c\|_2^2$ is a Chi-squared random variable with m degrees of freedom. Hence, with probability at least $1 - e^{-\tilde{\gamma}_1 \cdot 2m}$,

$$\|c\|_2^2 \leq 4m$$

For fixed $\|x\|_2 = 1$, $\|Z'x\|_2^2 \sim \chi_{n-1}^2$, and hence with probability at least $1 - e^{-\gamma_1 \cdot 3m}$

$$\|Z'x\|_2^2 \leq m$$

when $m \geq \tilde{c}_1 n$. Hence, $m^2 \|y'\|_2^2 \leq 4m \cdot m$ with probability at least $1 - 2e^{-\tilde{\gamma}_1 \cdot 4m}$. So, $\|y'\|_2 < 2$ with probability at least $1 - 2e^{-\tilde{\gamma}_1 \cdot 4m}$.

So $\|A_T\|_{\mathbb{F}}^2 \leq 25$, and hence $\|A_T\|_{\mathbb{F}} \leq 5$ with probability at least $1 - e^{-\tilde{\gamma}_1 m}$ for some $\tilde{\gamma}_1$. \square

We now show that for a fixed signal and support set of gross errors, there is a high probability that a dual certificate exists simultaneously for all gross errors.

Lemma 7. *Fix x_0 and a support set S . If $m \geq cn$ and $|S|/m \leq \min(0.001, \gamma^*/2 \log 2)$, then there is an event \tilde{E}_{S,x_0} on which for all ε with $\text{supp}(\varepsilon) \subseteq S$, there exists a Y satisfying (13), $\|Y_T\|_F \leq 1/4$, and $\|Y_{T^\perp} - \frac{17}{10} I_{T^\perp}\| \leq 3/10$. The probability of \tilde{E}_{S,x_0} is at least $1 - e^{-\gamma^* m/2}$.*

Proof. Consider all of the $2^{|S|}$ possible assignments of sign to the entries of ε on S . For each, choose an S^+ and S^- that are disjoint, have cardinality $\lceil 0.001m \rceil$, and are supersets of the indices assigned a positive or negative sign, respectively. Let \tilde{E}_{S,x_0} be the event on which all sign assignments yield a Y satisfying (13), $\|Y_T\|_{\mathbb{F}} \leq 1/4$, and $\|Y_{T^\perp} - \frac{17}{10} I_{T^\perp}\| \leq 3/10$. By Lemma 5, this event has probability at least

$$1 - 2^{|S|} e^{-\gamma^* m} \geq 1 - e^{-\gamma^* m/2}.$$

\square

We now show that for a fixed support set of gross errors, there is a high probability that a dual certificate exists simultaneously for all signals and for all gross errors.

Lemma 8. *Fix a support set S . If $m \geq \max(c, 4 \log(201)/\gamma^*)n$ and $|S|/m \leq \min(0.001, \gamma^*/2 \log 2)$, then on an event of probability at least $1 - e^{-\gamma^* m/4}$, for all x_0 and for all ε with $\text{supp}(\varepsilon) \subseteq S$, there exists $Y = \mathcal{A}^* \lambda$ satisfying (13) with $\alpha = 7/m$ and $\|Y_T\|_F \leq 0.44$ and $\|Y_{T^\perp} - \frac{17}{10}I_{T^\perp}\| \leq 4/10$.*

Proof. By Lemma 7, for any fixed all x_0 such that $\|x_0\| = 1$, for all ε with $\text{supp}(\varepsilon) \subseteq S$, there exists a $Y = \mathcal{A}^* \lambda$ such that

$$\|\lambda\|_\infty \leq \frac{7}{m} \quad (48)$$

$$\lambda_S = \frac{7}{m} \text{sgn}(\varepsilon_S) \quad (49)$$

$$\|Y_{T^\perp} + 1.7I_{T^\perp}\| \leq 0.3 \quad (50)$$

$$\|Y_T\|_F \leq 0.25 \quad (51)$$

on the event \tilde{E}_{S, x_0} , which has probability at least $1 - e^{-\gamma^* m/2}$. By Lemma 5.2 in [19], there exists a net \mathcal{N}_ε such that $|\mathcal{N}_\varepsilon| \leq (1 + 2/\varepsilon)^n$. Hence, such a Y exists simultaneously for all $x_0 \in \mathcal{N}_\varepsilon$ on an event of probability at least $1 - (1 + 2/\varepsilon)^n e^{-\gamma^* m/2}$. If $m > 4n \log(1 + 2/\varepsilon)/\gamma^*$, then such a Y exists simultaneously for all $x_0 \in \mathcal{N}_\varepsilon$ with probability at least $1 - e^{-\gamma^* cn/4} \geq 1 - e^{-\gamma^* m/4}$.

We now appeal to a continuity argument to show that a dual certificate exists for points not on the net \mathcal{N}_ε . For an arbitrary x such that $\|x\|_2 = 1$, we consider the Y corresponding to the nearest $x_0 \in \mathcal{N}_\varepsilon$. Note that $\|x - x_0\| \leq \varepsilon$ by definition of the net \mathcal{N}_ε . We now closely follow the proof and notation of Corollary 2.4 in [2] to show that Y is a satisfactory approximate dual certificate for x .

Note that $\|Y\| \leq 2.5$. Let $\Delta = xx^t - x_0x_0^t$ and note that $\|\Delta\|_F \leq 2\varepsilon$. Let $T = T_x$. Now we have

$$Y_{T^\perp} + 1.7I_{T^\perp} = Y_{T_{x_0}^\perp} + 1.7I_{T_{x_0}^\perp} - R_1 \quad (52)$$

$$Y_T = Y_{T_{x_0}} + R_2 \quad (53)$$

where

$$R_1 = \Delta Y(I - x_0x_0^t) + (I - x_0x_0^t)Y\Delta - \Delta Y\Delta + 1.7\Delta \quad (54)$$

$$R_2 = \Delta Y(I - x_0x_0^t) + (I - x_0x_0^t)Y\Delta - \Delta Y\Delta \quad (55)$$

We observe that

$$\|R_1\| \leq 2\|Y\|\|\Delta\|\|I - x_0x_0^t\| + \|Y\|\|\Delta\|^2 + 1.7\|\Delta\| \leq 13.4\varepsilon + 10\varepsilon^2 \quad (56)$$

$$\|R_2\| \leq \sqrt{2}\|R_2\| \leq \sqrt{2}(2\|Y\|\|\Delta\|\|I - x_0x_0^t\| + \|Y\|\|\Delta\|^2) \leq 10\sqrt{2}(\varepsilon + \varepsilon^2) \quad (57)$$

If we choose $\varepsilon = 0.01$, $\|R_1\| \leq 0.135$ and $\|R_2\|_F \leq 0.143$. and

$$\|Y_{T^\perp} + 1.7I_{T^\perp}\| \leq 0.3 + 0.135 = 0.435 \quad (58)$$

$$\|Y_T\|_F \leq 0.25 + 0.143 = 0.393 \quad (59)$$

□

We can now prove Theorem 1.

Proof of Theorem 1. Assume that $|S|/m \leq \min(0.001, \gamma^*/2 \log 2)$ and $m \geq \max(c, 4 \log(201)/\gamma^*)$. By Lemma 8, there is an event of probability at least $1 - e^{-\gamma^* m/4}$ such that for all x_0 and for all ε with $\text{supp}(\varepsilon) \subseteq S$, there exists a $Y = \mathcal{A}^* \lambda$ satisfying (11)–(13). Choose a superset \bar{S} such that $|\bar{S}| = 0.01m$. On the intersection of this event with $E_{\bar{S}}$, Lemma 4 guarantees that $\|\hat{X} - X_0\|_F \leq C\|\eta\|_1/m$. The intersection of these events has probability at least $1 - e^{-\gamma m}$ for some γ .

The proof of the error estimate $\|\hat{x} - x_0\| \leq C' \min\left(\|x_0\|, \frac{\|\eta\|_1}{m\|x_0\|}\right)$ can be found in [4]. □

2.3. Comment on the complex-valued case

Extending Theorem 1 to the complex-valued case may be possible. In the proofs above, the most central use of the real-valued assumption is in the concentration estimates of the dual certificate. The certificate is broken into pieces where ε is positive, negative, and zero. In the complex case, there are an infinite number of directions that the values of ε can take. Hence, the complex case will require a more involved analysis.

- [1] E. Candès, Y. Eldar, T. Strohmer, and V. Voroninski. Phase retrieval via matrix completion. *SIAM J. Imaging Sci.*, 6(1):199–225, 2013.
- [2] Emmanuel J Candès and Xiaodong Li. Solving quadratic equations via phaselift when there are about as many equations as unknowns. *Found. of Comput. Math.*, pages 1–10, 2012.
- [3] Emmanuel J Candès, Xiaodong Li, Yi Ma, and John Wright. Robust principal component analysis? *J. ACM*, 58(3):11, 2011.
- [4] Emmanuel J Candès, Thomas Strohmer, and Vladislav Voroninski. Phaselift: Exact and stable signal recovery from magnitude measurements via convex programming. *Comm. Pure Appl. Math.*, 66(8):1241–1274, 2013.
- [5] Venkat Chandrasekaran, Sujay Sanghavi, Pablo A Parrilo, and Alan S Willsky. Rank-sparsity incoherence for matrix decomposition. *SIAM Journal on Optimization*, 21(2):572–596, 2011.
- [6] Yudong Chen, Ali Jalali, Sujay Sanghavi, and Constantine Caramanis. Low-rank matrix recovery from errors and erasures. *Information Theory, IEEE Transactions on*, 59(7):4324–4337, 2013.
- [7] Laurent Demanet and Paul Hand. Stable optimizationless recovery from phaseless linear measurements. *Journal of Fourier Analysis and Applications*, 20(1):199–221, 2014.
- [8] Arvind Ganesh, John Wright, Xiaodong Li, Emmanuel J Candes, and Yi Ma. Dense error correction for low-rank matrices via principal component pursuit. In *Information Theory Proceedings (ISIT), 2010 IEEE International Symposium on*, pages 1513–1517. IEEE, 2010.
- [9] David Gross. Recovering low-rank matrices from few coefficients in any basis. *Information Theory, IEEE Transactions on*, 57(3):1548–1566, 2011.
- [10] Robert W Harrison. Phase problem in crystallography. *J. Opt. Soc. Am. A*, 10(5):1046–1055, 1993.
- [11] Daniel Hsu, Sham M Kakade, and Tong Zhang. Robust matrix decomposition with sparse corruptions. *Information Theory, IEEE Transactions on*, 57(11):7221–7234, 2011.
- [12] Richard Kueng, Holger Rauhut, and Ulrich Terstiege. Low rank matrix recovery from rank one measurements. *CoRR*, abs/1410.6913, 2014.
- [13] Beatrice Laurent and Pascal Massart. Adaptive estimation of a quadratic functional by model selection. *Annals of Statistics*, pages 1302–1338, 2000.
- [14] Xiaodong Li. Compressed sensing and matrix completion with constant proportion of corruptions. *Constructive Approximation*, 37(1):73–99, 2013.
- [15] J. Löfberg. Yalmip : A toolbox for modeling and optimization in MATLAB. In *Proceedings of the CACSD Conference*, Taipei, Taiwan, 2004.
- [16] RP Millane. Phase retrieval in crystallography and optics. *J. Opt. Soc. Am. A*, 7(3):394–411, 1990.
- [17] K. C. Toh, M.J. Todd, and R. H. Tutuncu. Sdpt3 - a matlab software package for semidefinite programming. *Optimization Methods and Software*, 11:545–581, 1998.
- [18] R.H. Tutuncu, K.C. Toh, and M.J. Todd. Solving semidefinite-quadratic-linear programs using sdpt3. *Mathematical Programming Ser. B*, 95:189–217, 2003.

- [19] 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, 2012.
- [20] John Wright, Arvind Ganesh, Shankar Rao, Yigang Peng, and Yi Ma. Robust principal component analysis: Exact recovery of corrupted low-rank matrices via convex optimization. In *Advances in neural information processing systems*, pages 2080–2088, 2009.