

Necessary and Sufficient Conditions for Success of the Nuclear Norm Heuristic for Rank Minimization

Benjamin Recht*, Weiyu Xu† and Babak Hassibi‡

September 7, 2008

Abstract

Minimizing the rank of a matrix subject to constraints is a challenging problem that arises in many applications in control theory, machine learning, and discrete geometry. This class of optimization problems, known as rank minimization, is NP-HARD, and for most practical problems there are no efficient algorithms that yield exact solutions. A popular heuristic algorithm replaces the rank function with the nuclear norm—equal to the sum of the singular values—of the decision variable. In this paper, we provide a necessary and sufficient condition that quantifies when this heuristic successfully finds the minimum rank solution of a linear constraint set. We additionally provide a probability distribution over instances of the affine rank minimization problem such that instances sampled from this distribution satisfy our conditions for success with overwhelming probability provided the number of constraints is appropriately large. Finally, we give empirical evidence that these probabilistic bounds provide accurate predictions of the heuristic’s performance in non-asymptotic scenarios.

AMS (MOC) Subject Classification 90C25; 90C59; 15A52.

Keywords. rank, convex optimization, matrix norms, random matrices, compressed sensing, Gaussian processes.

1 Introduction

Optimization problems involving constraints on the rank of matrices are pervasive in applications. In Control Theory, such problems arise in the context of low-order controller design [9, 19], minimal realization theory [11], and model reduction [4]. In Machine Learning, problems in inference with partial information [23], multi-task learning [1], and manifold learning [28] have been formulated as rank minimization problems. Rank minimization also plays a key role in the study of embeddings of discrete metric spaces in Euclidean space [16]. In certain instances with special structure, rank minimization problems can be solved via the singular value decomposition or can be reduced to the solution of a linear system [19, 20]. In general, however, minimizing the rank of a matrix subject to convex constraints is NP-HARD. The best exact algorithms for this problem involve quantifier

*Center for the Mathematics of Information, California Institute of Technology, 1200 E California Blvd, Pasadena, CA brecht@ist.caltech.edu

†Electrical Engineering, California Institute of Technology, 1200 E California Blvd, Pasadena, CA weiyu@systems.caltech.edu

‡Electrical Engineering, California Institute of Technology, 1200 E California Blvd, Pasadena, CA bhassibi@systems.caltech.edu

elimination and such solution methods require at least exponential time in the dimensions of the matrix variables.

A popular heuristic for solving rank minimization problems in the controls community is the “trace heuristic” where one minimizes the trace of a positive semidefinite decision variable instead of the rank (see, e.g., [4, 19]). A generalization of this heuristic to non-symmetric matrices introduced by Fazel in [10] minimizes the *nuclear norm*, or the sum of the singular values of the matrix, over the constraint set. When the matrix variable is symmetric and positive semidefinite, this heuristic is equivalent to the trace heuristic, as the trace of a positive semidefinite matrix is equal to the sum of its singular values. The nuclear norm is a convex function and can be optimized efficiently via semidefinite programming. Both the trace heuristic and the nuclear norm generalization have been observed to produce very low-rank solutions in practice, but, until very recently, conditions where the heuristic succeeded were only available in cases that could also be solved by elementary linear algebra [20].

The first non-trivial sufficient conditions that guaranteed the success of the nuclear norm heuristic were provided in [21]. Focusing on the special case where one seeks the lowest rank matrix in an affine subspace, the authors provide a “restricted isometry” condition on the linear map defining the affine subspace which guarantees the minimum nuclear norm solution is the minimum rank solution. Moreover, they provide several ensembles of affine constraints where this sufficient condition holds with overwhelming probability. Their work builds on seminal developments in “compressed sensing” that determined conditions for when minimizing the ℓ_1 norm of a vector over an affine space returns the sparsest vector in that space (see, e.g., [6, 5, 3]). There is a strong parallelism between the sparse approximation and rank minimization settings. The rank of a diagonal matrix is equal to the number of non-zeros on the diagonal. Similarly, the sum of the singular values of a diagonal matrix is equal to the ℓ_1 norm of the diagonal. Exploiting the parallels, the authors in [21] were able to extend much of the analysis developed for the ℓ_1 heuristic to provide guarantees for the nuclear norm heuristic.

Building on a different collection of developments in compressed sensing [7, 8, 25], we present a *necessary* and sufficient condition for the solution of the nuclear norm heuristic to coincide with the minimum rank solution in an affine space. The condition characterizes a particular property of the null-space of the linear map which defines the affine space. We show that when the linear map defining the constraint set is generated by sampling its entries independently from a Gaussian distribution, the null-space characterization holds with overwhelming probability provided the dimensions of the equality constraints are of appropriate size. We provide numerical experiments demonstrating that even when matrix dimensions are small, the nuclear norm heuristic does indeed always recover the minimum rank solution when the number of constraints is sufficiently large. Empirically, we observe that our probabilistic bounds accurately predict when the heuristic succeeds.

1.1 Main Results

Let X be an $n_1 \times n_2$ matrix decision variable. Without loss of generality, we will assume throughout that $n_1 \leq n_2$. Let $\mathcal{A} : \mathbb{R}^{n_1 \times n_2} \rightarrow \mathbb{R}^m$ be a linear map, and let $b \in \mathbb{R}^m$. The main optimization problem under study is

$$\begin{aligned} & \text{minimize} && \text{rank}(X) \\ & \text{subject to} && \mathcal{A}(X) = b. \end{aligned} \tag{1.1}$$

This problem is known to be NP-HARD and is also hard to approximate [18]. As mentioned above, a popular heuristic for this problem replaces the rank function with the sum of the singular values of the decision variable. Let $\sigma_i(X)$ denote the i -th largest singular value of X (equal to the square-root of the i -th largest eigenvalue of XX^*). Recall that the rank of X is equal to the number of nonzero singular values. In the case when the singular values are all equal to one, the sum of the singular values is equal to the rank. When the singular values are less than or equal to one, the sum of the singular values is a convex function that is strictly less than the rank. This sum of the singular values is a unitarily invariant matrix norm, called the *nuclear norm*, and is denoted

$$\|X\|_* := \sum_{i=1}^r \sigma_i(X).$$

This norm is alternatively known by several other names including the Schatten 1-norm, the Ky Fan norm, and the trace class norm.

As described in the introduction, our main concern is when the optimal solution of (1.1) coincides with the optimal solution of

$$\begin{aligned} & \text{minimize} && \|X\|_* \\ & \text{subject to} && \mathcal{A}(X) = b. \end{aligned} \tag{1.2}$$

This optimization is convex, and can be efficiently solved via a variety of methods including semidefinite programming (see [21] for a survey).

Whenever $m < n_1 n_2$, the null space of \mathcal{A} , that is the set of Y such that $\mathcal{A}(Y) = 0$, is not empty. Note that X is an optimal solution for (1.2) if and only if for every Y in the null-space of \mathcal{A}

$$\|X + Y\|_* \geq \|X\|_*. \tag{1.3}$$

The following theorem generalizes this null-space criterion to a critical property that guarantees when the nuclear norm heuristic finds the minimum rank solution of $\mathcal{A}(X) = b$ for all values of the vector b . Our main result is the following

Theorem 1.1 *Let X_0 be the optimal solution of (1.1) and assume that X_0 has rank $r < n_1/2$. Then*

1. *If for every Y in the null space of \mathcal{A} and for every decomposition*

$$Y = Y_1 + Y_2,$$

where Y_1 has rank r and Y_2 has rank greater than r , it holds that

$$\|Y_1\|_* < \|Y_2\|_*,$$

then X_0 is the unique minimizer of (1.2).

2. *Conversely, if the condition of part 1 does not hold, then there exists a vector $b \in \mathbb{R}^m$ such that the minimum rank solution of $\mathcal{A}(X) = b$ has rank at most r and is not equal to the minimum nuclear norm solution.*

This result is of interest for multiple reasons. First, as shown in [22], a variety of the rank minimization problems, including those with inequality and semidefinite cone constraints, can be reformulated in the form of (1.1). Secondly, we now present a family of random equality constraints under which the nuclear norm heuristic succeeds with overwhelming probability. We prove both of the following two theorems by showing that \mathcal{A} obeys the null-space criteria of Equation (1.3) and Theorem 1.1 respectively with overwhelming probability.

Note that for a linear map $\mathcal{A} : \mathbb{R}^{n_1 \times n_2} \rightarrow \mathbb{R}^m$, we can always find an $m \times n_1 n_2$ matrix \mathbf{A} such that

$$\mathcal{A}(X) = \mathbf{A} \operatorname{vec} X. \quad (1.4)$$

In the case where \mathbf{A} has entries sampled independently from a zero-mean, unit-variance Gaussian distribution, then the null space characterization of theorem 1.1 holds with overwhelming probability provided m is large enough. For simplicity of notation in the theorem statements, we consider the case of square matrices. These results can be then translated into rectangular matrices by padding with rows/columns of zeros to make the matrix square. We define the random ensemble of $d_1 \times d_2$ matrices $\mathfrak{G}(d_1, d_2)$ to be the Gaussian ensemble, with each entry sampled i.i.d. from a Gaussian distribution with zero-mean and variance one. We also denote $\mathfrak{G}(d, d)$ by $\mathfrak{G}(d)$.

The first result characterizes when a particular low-rank matrix can be recovered from a random linear system via nuclear norm minimization.

Theorem 1.2 (Weak Bound) *Let X_0 be an $n \times n$ matrix of rank $r = \beta n$. Let $\mathcal{A} : \mathbb{R}^{n \times n} \rightarrow \mathbb{R}^{\mu n^2}$ denote the random linear transformation*

$$\mathcal{A}(X) = \mathbf{A} \operatorname{vec}(X),$$

where \mathbf{A} is sampled from $\mathfrak{G}(\mu n^2, n^2)$. Then whenever

$$\mu \geq 1 - \frac{64}{9\pi^2} \left((1 - \beta)^{3/2} - \beta^{3/2} \right)^2 \quad (1.5)$$

there exists a numerical constant $c_w(\mu, \beta) > 0$ such that with probability exceeding $1 - e^{-c_w(\mu, \beta)n^2}$,

$$X_0 = \arg \min \{ \|Z\|_* : \mathcal{A}(Z) = \mathcal{A}(X_0) \}.$$

In particular, if β and μ satisfy (1.5), then nuclear norm minimization will recover X_0 from a random set of μn^2 constraints drawn from the Gaussian ensemble almost surely as $n \rightarrow \infty$.

The second theorem characterizes when the nuclear norm heuristic succeeds at recovering all low rank matrices.

Theorem 1.3 (Strong Bound) *Let \mathcal{A} be defined as in Theorem 1.2. Define the two functions*

$$f(\beta, \epsilon) = \frac{8}{3\pi} \frac{(1 - \beta)^{3/2} - \beta^{3/2} - 4\epsilon}{1 + 4\epsilon}$$

$$g(\beta, \epsilon) = \sqrt{2\beta(2 - \beta)} \log \left(\frac{3\pi}{2\epsilon} \right).$$

Then there exists a numerical constant $c_s(\mu, \beta) > 0$ such that with probability exceeding $1 - e^{-c_s(\mu, \beta)n^2}$, for all $n \times n$ matrices X_0 of rank $r \leq \beta n$

$$X_0 = \arg \min \{ \|Z\|_* : \mathcal{A}(Z) = \mathcal{A}(X_0) \}$$




Figure 1: The Weak Bound (1.5) versus the Strong Bound (1.6).

whenever

$$\mu \geq 1 - \sup_{\substack{\epsilon > 0 \\ f(\beta, \epsilon) - g(\beta, \epsilon) > 0}} (f(\beta, \epsilon) - g(\beta, \epsilon))^2. \quad (1.6)$$

In particular, if β and μ satisfy (1.5), then nuclear norm minimization will recover all rank r matrices from a random set of μn^2 constraints drawn from the Gaussian ensemble almost surely as $n \rightarrow \infty$.

Figure 1 plots the bound from Theorems 1.2 and 1.3. We call (1.5) the *Weak Bound* because it is a condition that depends on the optimal solution of (1.1). On the other hand, we call (1.6) the *Strong Bound* as it guarantees the nuclear norm heuristic succeeds no matter what the optimal solution. The Weak Bound is the only bound that can be tested experimentally, and, in Section 4, we will show that it corresponds well to experimental data. Moreover, the Weak Bound provides guaranteed recovery over a far larger region of (β, μ) parameter space. Nonetheless, the mere existence of a Strong Bound is surprising in of itself and results in a much better bound than what was available from previous results (c.f., [21]).

1.2 Notation and Preliminaries

For a rectangular matrix $X \in \mathbb{R}^{n_1 \times n_2}$, X^* denotes the transpose of X . $\text{vec}(X)$ denotes the vector in $\mathbb{R}^{n_1 n_2}$ with the columns of X stacked on top of one another.

For vectors $v \in \mathbb{R}^d$, the only norm we will ever consider is the Euclidean norm

$$\|v\|_{\ell_2} = \left(\sum_{i=1}^d v_i^2 \right)^{1/2}.$$

On the other hand, we will consider a variety of matrix norms. For matrices X and Y of the same dimensions, we define the inner product in $\mathbb{R}^{n_1 \times n_2}$ as $\langle X, Y \rangle := \text{trace}(X^* Y) = \sum_{i=1}^{n_1} \sum_{j=1}^{n_2} X_{ij} Y_{ij}$. The norm associated with this inner product is called the Frobenius (or Hilbert-Schmidt) norm

$\|\cdot\|_F$. The Frobenius norm is also equal to the Euclidean, or ℓ_2 , norm of the vector of singular values, i.e.,

$$\|X\|_F := \left(\sum_{i=1}^r \sigma_i^2 \right)^{\frac{1}{2}} = \sqrt{\langle X, X \rangle} = \left(\sum_{i=1}^{n_1} \sum_{j=1}^{n_2} X_{ij}^2 \right)^{\frac{1}{2}}$$

The operator norm (or induced 2-norm) of a matrix is equal to its largest singular value (i.e., the ℓ_∞ norm of the singular values):

$$\|X\| := \sigma_1(X).$$

The nuclear norm of a matrix is equal to the sum of its singular values, i.e.,

$$\|X\|_* := \sum_{i=1}^r \sigma_i(X).$$

These three norms are related by the following inequalities which hold for any matrix X of rank at most r :

$$\|X\| \leq \|X\|_F \leq \|X\|_* \leq \sqrt{r} \|X\|_F \leq r \|X\|. \quad (1.7)$$

To any norm, we may associate a *dual norm* via the following variational definition

$$\|X\|_d = \sup_{\|Y\|_p=1} \langle Y, X \rangle.$$

One can readily check that the dual norm the Frobenius norm is the Frobenius norm. Less trivially, one can show that the dual norm of the operator norm is the nuclear norm (See, for example, [21]). We will leverage the duality between the operator and nuclear norm several times in our analysis.

2 Necessary and Sufficient Conditions

We first prove our necessary and sufficient condition for success of the nuclear norm heuristic. We will need the following two technical lemmas. The first is an easily verified fact.

Lemma 2.1 *Suppose X and Y are $n_1 \times n_2$ matrices such that $X^*Y = 0$ and $XY^* = 0$. Then $\|X + Y\|_* = \|X\|_* + \|Y\|_*$.*

Indeed, if $X^*Y = 0$ and $XY^* = 0$, we can find a coordinate system in which

$$X = \left\| \left[\begin{array}{cc} A & 0 \\ 0 & 0 \end{array} \right] \right\|_* \quad \text{and} \quad Y = \left\| \left[\begin{array}{cc} 0 & 0 \\ 0 & B \end{array} \right] \right\|_*$$

from which the lemma trivially follows. The next Lemma allows us to exploit Lemma 2.1 in our proof.

Lemma 2.2 *Let X be an $n_1 \times n_2$ matrix with rank $r < \frac{n_1}{2}$ and Y be an arbitrary $n_1 \times n_2$ matrix. Let P_X^c and P_X^r be the matrices that project onto the column and row spaces of X respectively. Then if $P_X^c Y P_X^r$ has full rank, Y can be decomposed as*

$$Y = Y_1 + Y_2,$$

where Y_1 has rank r , and

$$\|X + Y_2\|_* = \|X\|_* + \|Y_2\|_*.$$

Proof Without loss of generality, we can write X as

$$X = \begin{bmatrix} X_{11} & 0 \\ 0 & 0 \end{bmatrix},$$

where X_{11} is $r \times r$ and full rank. Accordingly, Y becomes

$$Y = \begin{bmatrix} Y_{11} & Y_{12} \\ Y_{21} & Y_{22} \end{bmatrix},$$

where Y_{11} is full rank since $P_X^r Y P_X^c$ is. The decomposition is now clearly

$$Y = \underbrace{\begin{bmatrix} Y_{11} & Y_{12} \\ Y_{21} & Y_{21}Y_{11}^{-1}Y_{12} \end{bmatrix}}_{Y_1} + \underbrace{\begin{bmatrix} 0 & 0 \\ 0 & Y_{22} - Y_{21}Y_{11}^{-1}Y_{12} \end{bmatrix}}_{Y_2}.$$

That Y_1 has rank r follows from the fact that the rank of a block matrix is equal to the rank of a diagonal block plus the rank of its Schur complement (see, e.g., [14, §2.2]). That $\|X_1 + Y_2\|_* = \|X_1\|_* + \|Y_2\|_*$ follows from Lemma 2.1. \blacksquare

We can now provide a proof of Theorem 1.1.

Proof We begin by proving the converse. Assume the condition of part 1 is violated, i.e., there exists some Y , such that $\mathcal{A}(Y) = 0$, $Y = Y_1 + Y_2$, $\text{rank}(Y_2) > \text{rank}(Y_1) = r$, yet $\|Y_1\|_* > \|Y_2\|_*$. Now take $X_0 = Y_1$ and $b = \mathcal{A}(X_0)$. Clearly, $\mathcal{A}(-Y_2) = b$ (since Y is in the null space) and so we have found a matrix of higher rank, but lower nuclear norm.

For the other direction, assume the condition of part 1 holds. Now use Lemma 2.2 with $X = X_0$ and $Y = X_* - X_0$. That is, let P_X^c and P_X^r be the matrices that project onto the column and row spaces of X_0 respectively and assume that $P_{X_0}^c(X_* - X_0)P_{X_0}^r$ has full rank. Write $X_* - X_0 = Y_1 + Y_2$ where Y_1 has rank r and $\|X_0 + Y_2\|_* = \|X_0\|_* + \|Y_2\|_*$. Assume further that Y_2 has rank larger than r (recall $r < n/2$). We will consider the case where $P_{X_0}^c(X_* - X_0)P_{X_0}^r$ does not have full rank and/or Y_2 has rank less than or equal to r in the appendix. We now have:

$$\begin{aligned} \|X_*\|_* &= \|X_0 + X_* - X_0\|_* \\ &= \|X_0 + Y_1 + Y_2\|_* \\ &\geq \|X_0 + Y_2\|_* - \|Y_1\|_* \\ &= \|X_0\|_* + \|Y_2\|_* - \|Y_1\|_* \quad \text{by Lemma 2.2.} \end{aligned}$$

But $\mathcal{A}(Y_1 + Y_2) = 0$, so $\|Y_2\|_* - \|Y_1\|_*$ non-negative and therefore $\|X_*\|_* \geq \|X_0\|_*$. Since X_* is the minimum nuclear norm solution, implies that $X_0 = X_*$. \blacksquare

For the interested reader, the argument for the case where $P_{X_0}^r(X_* - X_0)P_{X_0}^c$ does not have full rank or Y_2 has rank less than or equal to r can be found in the appendix.

3 Proofs of the Probabilistic Bounds

We now turn to the proofs of the probabilistic bounds 1.5 and 1.6. We first provide a sufficient condition which implies the necessary and sufficient null-space conditions. Then, noting that the null space of \mathcal{A} is spanned by Gaussian vectors, we use bounds from probability on Banach Spaces to show that the sufficient conditions are met. This will require the introduction of two useful auxiliary functions whose actions on Gaussian processes are explored in Section 3.4.

3.1 Sufficient Condition for Null-space Characterizations

The following theorem gives us a new condition that implies our necessary and sufficient condition.

Theorem 3.1 *Let \mathcal{A} be a linear map of $n \times n$ matrices into \mathbb{R}^m . Suppose that for every Y in the null-space of \mathcal{A} and any projection operators onto r -dimensional subspaces P and Q that*

$$\|(I - P)Y(I - Q)\|_* \geq \|PYQ\|_*. \quad (3.1)$$

Then for every matrix Z with row and column spaces equal to the range of Q and P respectively,

$$\|Z + Y\|_* \geq \|Z\|_*$$

for all Y in the null-space of \mathcal{A} . In particular, if 3.1 holds for every pair of projection operators P and Q , then for every Y in the null space of \mathcal{A} and for every decomposition $Y = Y_1 + Y_2$ where Y_1 has rank r and Y_2 has rank greater than r , it holds that

$$\|Y_1\|_* \leq \|Y_2\|_*.$$

We will need the following lemma

Lemma 3.2 *For any block partitioned matrix*

$$X = \begin{bmatrix} A & B \\ C & D \end{bmatrix}$$

we have $\|X\|_ \geq \|A\|_* + \|D\|_*$.*

Proof This lemma follows from the dual description of the nuclear norm:

$$\|X\|_* = \sup \left\{ \left\langle \begin{bmatrix} Z_{11} & Z_{12} \\ Z_{21} & Z_{22} \end{bmatrix}, \begin{bmatrix} A & B \\ C & D \end{bmatrix} \right\rangle \mid \left\| \begin{bmatrix} Z_{11} & Z_{12} \\ Z_{21} & Z_{22} \end{bmatrix} \right\| = 1 \right\}. \quad (3.2)$$

and similarly

$$\|A\|_* + \|D\|_* = \sup \left\{ \left\langle \begin{bmatrix} Z_{11} & 0 \\ 0 & Z_{22} \end{bmatrix}, \begin{bmatrix} A & B \\ C & D \end{bmatrix} \right\rangle \mid \left\| \begin{bmatrix} Z_{11} & 0 \\ 0 & Z_{22} \end{bmatrix} \right\| = 1 \right\}. \quad (3.3)$$

Since (3.2) is a supremum over a larger set than (3.3), the claim follows. ■

Theorem 3.1 now trivially follows

Proof [of Theorem 3.1] Without loss of generality, we may choose coordinates such that P and Q both project onto the space spanned by first r standard basis vectors. Then we may partition Y as

$$Y = \begin{bmatrix} Y_{11} & Y_{12} \\ Y_{21} & Y_{22} \end{bmatrix}$$

and write, using Lemma 3.2,

$$\|Y - Z\|_* - \|Z\|_* = \left\| \begin{bmatrix} Y_{11} - Z & Y_{12} \\ Y_{21} & Y_{22} \end{bmatrix} \right\|_* - \|Z\|_* \geq \|Y_{11} - Z\|_* + \|Y_{22}\|_* - \|Z\|_* \geq \|Y_{22}\|_* - \|Y_{11}\|_*$$

which is non-negative by assumption. Note that if the theorem holds for all projection operators P and Q whose range has dimension r , then $\|Z + Y\|_* \geq \|Z\|_*$ for all matrices Z of rank r and hence the second part of the theorem follows. ■

3.2 Proof of the Weak Bound

Now we can turn to the proof of Theorem 1.2. The key observation in proving this lemma is the following characterization of the null-space of \mathcal{A} provided by Stojnic et al [25]

Lemma 3.3 *The null space of \mathcal{A} is identically distributed to the span of $n^2(1 - \mu)$ matrices G_i where each G_i is sampled i.i.d. from $\mathfrak{G}(n)$.*

This is nothing more than a statement that the null-space of \mathcal{A} is a random subspace. However, when we parameterize elements in this subspace as linear combinations of Gaussian vectors, we can leverage Comparison Theorems for Gaussian processes to yield our bounds.

Let $M = n^2(1 - \mu)$ and let G_1, \dots, G_M be i.i.d. samples from $\mathfrak{G}(n)$. Let X_0 be a matrix of rank βn . Let P_{X_0} and Q_{X_0} denote the projections onto the column and row spaces of X_0 respectively. By theorem 3.1 and Lemma 3.3, we need to show that for all $v \in \mathbb{R}^M$,

$$\|(I - P_{X_0}) \left(\sum_{i=1}^M v_i G_i \right) (I - Q_{X_0})\|_* \geq \|P_{X_0} \left(\sum_{i=1}^M v_i G_i \right) Q_{X_0}\|_*. \quad (3.4)$$

That is, $\sum_{i=1}^M v_i G_i$ is an arbitrary element of the null space of \mathcal{A} , and this equation restates the sufficient condition provided by Theorem 3.1. Now it is clear by homogeneity that we can restrict our attention to those $v \in \mathbb{R}^M$ with norm 1. The following crucial lemma characterizes when the expected value of this difference is nonnegative

Lemma 3.4 *Let $r = \beta n$ and suppose P and Q are projection operators onto r -dimensional subspaces of \mathbb{R}^n . For $i = 1, \dots, M$ let G_i be sampled from $\mathfrak{G}(n)$. Then*

$$\begin{aligned} \mathbb{E} \left[\inf_{\|v\|_{\ell_2}=1} \|(I - P) \left(\sum_{i=1}^M v_i G_i \right) (I - Q)\|_* - \|P \left(\sum_{i=1}^M v_i G_i \right) Q\|_* \right] \\ \geq \left(\frac{8}{3\pi} + o(1) \right) \left((1 - \beta)^{3/2} - \beta^{3/2} \right) n^{3/2} - \sqrt{Mn}. \end{aligned} \quad (3.5)$$

We will prove this Lemma and a similar inequality required for the proof the Strong Bound in Section 3.4 below. But we now show how using this Lemma and a concentration of measure argument, we prove Theorem 1.2.

First note, that if we plug in $M = (1 - \mu)n^2$ and divide the right hand side by $n^{3/2}$, the right hand side of (3.5) is non-negative if (1.5) holds. To bound the probability that (3.4) is non-negative, we employ a powerful concentration inequality for the Gaussian distribution bounding deviations of smoothly varying functions from their expected value.

To quantify what we mean by smoothly varying, recall that a function f is *Lipshitz* with respect to the Euclidean norm if there exists a constant L such that $|f(x) - f(y)| \leq L\|x - y\|_{\ell_2}$ for all x and y . The smallest such constant L is called the *Lipshitz constant* of the map f . If f is Lipshitz, it cannot vary too rapidly. In particular, note that if f is differentiable and Lipshitz, then L is a bound on the norm of the gradient of f . The following theorem states that the deviations of a Lipshitz function applied to a Gaussian random variable have Gaussian tails.

Theorem 3.5 *Let x be a normally distributed random vector and let f be a function with Lipshitz constant L . Then*

$$\mathbb{P}[|f(x) - \mathbb{E}[f(x)]| \geq t] \leq 2 \exp \left(-\frac{t^2}{2L^2} \right).$$

See [15] for a proof of this theorem with slightly weaker constants and several references for more complicated proofs that give rise to this concentration inequality. The following Lemma bounds the Lipschitz constant of interest

Lemma 3.6 For $i = 1, \dots, M$, let $X_i \in \mathbb{R}^{n_1 \times n_1}$ and $Y_i \in \mathbb{R}^{n_2 \times n_2}$. Define the function

$$F_I(X_1, \dots, X_M, Y_1, \dots, Y_M) = \inf_{\|v\|_{\ell_2}=1} \left\| \sum_{i=1}^M v_i X_i \right\|_* - \left\| \sum_{i=1}^M v_i Y_i \right\|_*.$$

Then the Lipschitz constant of F_I is at most $\sqrt{n_1 + n_2}$.

The proof of this lemma is straightforward and can be found in the appendix. Using Theorem 3.5 and Lemmas 3.4 and 3.6, we can now bound

$$\begin{aligned} \mathbb{P} \left[\inf_{\|v\|_{\ell_2}=1} \left\| (I - P_{X_0}) \left(\sum_{i=1}^M v_i G_i \right) (I - Q_{X_0}) \right\|_* - \left\| P_{X_0} \left(\sum_{i=1}^M v_i G_i \right) Q_{X_0} \right\|_* \leq t n^{3/2} \right] \\ \leq \exp \left(-\frac{1}{2} \left(\frac{8}{3\pi} \left((1 - \beta)^{3/2} - \beta^{3/2} \right) - \sqrt{1 - \mu} - t \right)^2 n^2 + o(n^2) \right). \end{aligned} \quad (3.6)$$

Setting $t = 0$ completes the proof of Theorem 1.2. We will use this concentration inequality with a non-zero t to prove the Strong Bound.

3.3 Proof of the Strong Bound

The proof of the Strong Bound is similar to that of the Weak Bound except we prove that (3.4) holds for *all* operators P and Q that project onto r -dimensional subspaces. Our proof will require an ϵ -net for the projection operators—a set of points such that any projection operator is within ϵ of some element in the set. We will show that if a slightly stronger bound than (3.4) holds on the ϵ -net, then (3.4) holds for all choices of row and column spaces.

Let us first examine how (3.4) changes when we perturb P and Q . Let P, Q, P' and Q' all be projection operators onto r -dimensional subspaces. Let W be some $n \times n$ matrix and observe that

$$\begin{aligned} & \left\| (I - P')W(I - Q') \right\|_* - \left\| P'WQ' \right\|_* - \left(\left\| (I - P)W(I - Q) \right\|_* - \left\| PWQ \right\|_* \right) \\ & \leq \left\| (I - P)W(I - Q) - (I - P')W(I - Q') \right\|_* + \left\| PWQ - P'WQ' \right\|_* \\ & \leq \left\| (I - P)W(I - Q) - (I - P')W(I - Q) \right\|_* + \left\| (I - P')W(I - Q) - (I - P')W(I - Q') \right\|_* \\ & \quad + \left\| PWQ - P'WQ \right\|_* + \left\| P'WQ - P'WQ' \right\|_* \\ & \leq \|P - P'\| \|W\|_* \|I - Q\| + \|I - P'\| \|W\|_* \|Q - Q'\| + \|P - P'\| \|W\|_* \|Q\| + \|P'\| \|W\|_* \|Q - Q'\| \\ & \leq 2(\|P - P'\| + \|Q - Q'\|) \|W\|_*. \end{aligned}$$

Here, the first and second lines follow from the triangle inequality, the third line follows because $\|AB\|_* \leq \|A\| \|B\|_*$, and the fourth line follows because $P, P', Q,$ and Q' are all projection operators. Rearranging this inequality gives

$$\left\| (I - P')W(I - Q') \right\|_* - \left\| P'WQ' \right\|_* \geq \left\| (I - P)W(I - Q) \right\|_* - \left\| PWQ \right\|_* - 2(\|P - P'\| + \|Q - Q'\|) \|W\|_*.$$

As we have just discussed, if we can prove that with overwhelming probability

$$\|(I - P)W(I - Q)\|_* - \|PWQ\|_* - 4\epsilon\|W\|_* \geq 0 \quad (3.7)$$

for all P and Q in an ϵ -net for the projection operators onto r -dimensional subspaces, we will have proved the Strong Bound.

To proceed, we need to know the size of an ϵ -net. The following bound on such a net is due to Szarek.

Theorem 3.7 (Szarek [27]) *Consider the space of all projection operators on \mathbb{R}^n projecting onto r dimensional subspaces endowed with the metric*

$$d(P, P') = \|P - P'\|$$

Then there exists an ϵ -net in this metric space with cardinality at most $(\frac{3\pi}{2\epsilon})^{r(n-r/2-1/2)}$.

With this in hand, we now calculate the probability that for a given P and Q in the ϵ -net,

$$\inf_{\|v\|_{\ell_2}=1} \|(I - P) \left(\sum_{i=1}^M v_i G_i \right) (I - Q)\|_* - \|P \left(\sum_{i=1}^M v_i G_i \right) Q\|_* \geq 4\epsilon \sup_{\|v\|_{\ell_2}=1} \left\| \sum_{i=1}^M v_i G_i \right\|_*. \quad (3.8)$$

As we will show in Section 3.4, we can upper bound the right hand side of this inequality using a similar bound as in Lemma 3.4.

Lemma 3.8 *For $i = 1, \dots, M$ let G_i be sampled from $\mathfrak{G}(n)$. Then*

$$\mathbb{E} \left[\sup_{\|v\|_{\ell_2}=1} \left\| \sum_{i=1}^M v_i G_i \right\|_* \right] \leq \left(\frac{8}{3\pi} + o(1) \right) n^{3/2} + \sqrt{Mn}. \quad (3.9)$$

Moreover, we prove the following in the appendix.

Lemma 3.9 *For $i = 1, \dots, M$, let $X_i \in \mathbb{R}^{n \times n}$ and define the function*

$$F_S(X_1, \dots, X_M) = \sup_{\|v\|_{\ell_2}=1} \left\| \sum_{i=1}^M v_i X_i \right\|_*.$$

Then the Lipschitz constant of F_S is at most \sqrt{n} .

Using Lemmas 3.8 and 3.9 combined with Theorem 3.5, we have that

$$\mathbb{P} \left[4\epsilon \sup_{\|v\|_{\ell_2}=1} \left\| \sum_{i=1}^M v_i G_i \right\|_* \geq tn^{3/2} \right] \leq \exp \left(-\frac{1}{2} \left(\frac{8}{3\pi} - \sqrt{1-\mu} + o(1) - \frac{t}{4\epsilon} \right)^2 n^2 \right), \quad (3.10)$$

and if we set the exponents of (3.6) and (3.10) equal to each other and solve for t , we find after some algebra and the union bound

$$\begin{aligned}
& \mathbb{P} \left[\inf_{\|v\|_{\ell_2}=1} \left\| (I - P) \left(\sum_{i=1}^M v_i G_i \right) (I - Q) \right\|_* - \left\| P \left(\sum_{i=1}^M v_i G_i \right) Q \right\|_* \geq 4\epsilon \sup_{\|v\|_{\ell_2}=1} \left\| \sum_{i=1}^M v_i G_i \right\|_* \right] \\
& \geq \mathbb{P} \left[\inf_{\|v\|_{\ell_2}=1} \left\| (I - P) \left(\sum_{i=1}^M v_i G_i \right) (I - Q) \right\|_* - \left\| P \left(\sum_{i=1}^M v_i G_i \right) Q \right\|_* > tn^{3/2} > 4\epsilon \sup_{\|v\|_{\ell_2}=1} \left\| \sum_{i=1}^M v_i G_i \right\|_* \right] \\
& \geq 1 - \mathbb{P} \left[\inf_{\|v\|_{\ell_2}=1} \left\| (I - P) \left(\sum_{i=1}^M v_i G_i \right) (I - Q) \right\|_* - \left\| P \left(\sum_{i=1}^M v_i G_i \right) Q \right\|_* < tn^{3/2} \right] \\
& \quad - \mathbb{P} \left[4\epsilon \sup_{\|v\|_{\ell_2}=1} \left\| \sum_{i=1}^M v_i G_i \right\|_* > tn^{3/2} \right] \\
& \geq 1 - 2 \exp \left(-\frac{1}{2} \left(\frac{8}{3\pi} \frac{(1 - \beta)^{3/2} - \beta^{3/2} - 4\epsilon}{1 + 4\epsilon} - \sqrt{1 - \mu} \right)^2 n^2 + o(n^2) \right).
\end{aligned}$$

Now, let Ω be an ϵ -net for the set of projection operators discussed above. Again by the union bound, we have that

$$\begin{aligned}
& \mathbb{P} \left[\forall P, Q \inf_{\|v\|_{\ell_2}=1} \left\| (I - P) \left(\sum_{i=1}^M v_i G_i \right) (I - Q) \right\|_* - \left\| P \left(\sum_{i=1}^M v_i G_i \right) Q \right\|_* \geq 4\epsilon \sup_{\|v\|_{\ell_2}=1} \left\| \sum_{i=1}^M v_i G_i Q \right\|_* \right] \\
& \leq 1 - 2 \exp \left(- \left\{ \frac{1}{2} \left(\frac{8}{3\pi} \frac{(1 - \beta)^{3/2} - \beta^{3/2} - 4\epsilon}{1 + 4\epsilon} - \sqrt{1 - \mu} \right)^2 + \beta(2 - \beta) \log \left(\frac{3\pi}{2\epsilon} \right) \right\} n^2 + o(n^2) \right).
\end{aligned} \tag{3.11}$$

Finding the parameters μ , β , and ϵ that make the terms multiplying n^2 negative completes the proof of the Strong Bound.

3.4 Comparison Theorems for Gaussian Processes and the Proofs of Lemmas 3.4 and 3.8

Both of the two following Comparison Theorems provide sufficient conditions for when the expected supremum or infimum of one Gaussian process is greater to that of another. Elementary proofs of both of these Theorems and several other Comparison Theorems can be found in §3.3 of [15].

Theorem 3.10 (Slepian's Lemma [24]) *Let X and Y be Gaussian random variables in \mathbb{R}^N such that*

$$\begin{cases} \mathbb{E}[X_i X_j] \leq \mathbb{E}[Y_i Y_j] & \text{for all } i \neq j \\ \mathbb{E}[X_i^2] = \mathbb{E}[Y_i^2] & \text{for all } i \end{cases}$$

Then

$$\mathbb{E}[\max_i Y_i] \leq \mathbb{E}[\max_i X_i].$$

Theorem 3.11 (Gordan [12, 13]) Let $X = (X_{ij})$ and $Y = (Y_{ij})$ be Gaussian random vectors in $\mathbb{R}^{N_1 \times N_2}$ such that

$$\begin{cases} \mathbb{E}[X_{ij}X_{ik}] \leq \mathbb{E}[Y_{ij}Y_{ik}] & \text{for all } i, j, k \\ \mathbb{E}[X_{ij}X_{lk}] \geq \mathbb{E}[Y_{ij}Y_{lk}] & \text{for all } i \neq l \text{ and } j, k \\ \mathbb{E}[X_{ij}^2] = \mathbb{E}[X_{ij}^2] & \text{for all } j, k \end{cases}$$

Then

$$\mathbb{E}[\min_i \max_j Y_{ij}] \leq \mathbb{E}[\min_i \max_j X_{ij}].$$

The following two lemmas follow from applications of these Comparison Theorems. We prove them in more generality than necessary for the current work because both Lemmas are interesting in their own right. Let $\|\cdot\|_p$ be any norm on $D \times D$ matrices and let $\|\cdot\|_d$ be its associated dual norm (See Section 1.2). Let us define the quantity $\sigma(\|G\|_p)$ as

$$\sigma(\|G\|_p) = \sup_{\|Z\|_d=1} \|Z\|_F, \quad (3.12)$$

and note that by this definition, we have

$$\sigma(\|G\|_p) = \sup_{\|Z\|_d=1} \mathbb{E}[\langle G, Z \rangle^2]^{1/2}$$

motivating the notation.

This first Lemma is now a straightforward consequence of Slepian's Lemma

Lemma 3.12 Let $\Delta > 0$ and let g be a Gaussian random vector in \mathbb{R}^M . Let G, G_1, \dots, G_M be sampled i.i.d. from $\mathfrak{G}(D)$. Then

$$\mathbb{E} \left[\sup_{\|v\|_{\ell_2}=1} \sup_{\|Y\|_d=1} \Delta \langle g, v \rangle + \left\langle \sum_{i=1}^M v_i G_i, Y \right\rangle \right] \leq \mathbb{E}[\|G\|_p] + \sqrt{M(\Delta^2 + \sigma(\|G\|_p)^2)}.$$

Proof

We follow the strategy used to prove Theorem 3.20 in [15]. Let G, G_1, \dots, G_M be sampled i.i.d. from $\mathfrak{G}(D)$ and $g \in \mathbb{R}^M$ be a Gaussian random vector and let γ be a zero-mean, unit-variance Gaussian random variable. For $v \in \mathbb{R}^M$ and $Y \in \mathbb{R}^{D \times D}$ define

$$\begin{aligned} Q_L(v, Y) &= \Delta \langle g, v \rangle + \left\langle \sum_{i=1}^M v_i G_i, Y \right\rangle + \sigma(\|G\|_p) \gamma \\ Q_R(v, Y) &= \langle G, Y \rangle + \sqrt{\Delta^2 + \sigma(\|G\|_p)^2} \langle g, v \rangle. \end{aligned}$$

Now observe that for any unit vectors in \mathbb{R}^M v, \hat{v} and any $D \times D$ matrices Y, \hat{Y} with dual norm 1

$$\begin{aligned} & \mathbb{E}[Q_L(v, Y)Q_L(\hat{v}, \hat{Y})] - \mathbb{E}[Q_R(v, Y)Q_R(\hat{v}, \hat{Y})] \\ &= \Delta^2 \langle v, \hat{v} \rangle + \langle v, \hat{v} \rangle \langle Y, \hat{Y} \rangle + \sigma(\|G\|_p)^2 - \langle Y, \hat{Y} \rangle - (\Delta^2 + \sigma(\|G\|_p)^2) \langle v, \hat{v} \rangle \\ &= (\sigma(\|G\|_p)^2 - \langle Y, \hat{Y} \rangle)(1 - \langle v, \hat{v} \rangle). \end{aligned}$$

The difference in expectation is thus equal to zero if $v = \hat{v}$ and is greater than or equal to zero if $v \neq \hat{v}$. Hence, by Slepian's Lemma and a compactness argument (see Proposition A.1 in the Appendix),

$$\mathbb{E} \left[\sup_{\|v\|_{\ell_2}=1} \sup_{\|Y\|=1} Q_L(v, Y) \right] \leq \mathbb{E} \left[\sup_{\|v\|_{\ell_2}=1} \sup_{\|Y\|=1} Q_R(v, Y) \right]$$

which proves the Lemma. \blacksquare

The following lemma can be proved in a similar fashion

Lemma 3.13 *Let $\|\cdot\|_p$ be a norm on $\mathbb{R}^{D_1 \times D_1}$ with dual norm $\|\cdot\|_d$ and let $\|\cdot\|_b$ be a norm on $\mathbb{R}^{D_2 \times D_2}$. Let g be a Gaussian random vector in \mathbb{R}^M . Let G_0, G_1, \dots, G_M be sampled i.i.d. from $\mathfrak{G}(D_1)$ and G'_1, \dots, G'_M be sampled i.i.d. from $\mathfrak{G}(D_2)$. Then*

$$\begin{aligned} \mathbb{E} \left[\inf_{\|v\|_{\ell_2}=1} \inf_{\|Y\|_b=1} \sup_{\|Z\|_d=1} \left\langle \sum_{i=1}^M v_i G_i, Z \right\rangle + \left\langle \sum_{i=1}^M v_i G'_i, Y \right\rangle \right] \\ \geq \mathbb{E} [\|G_0\|_p] - \mathbb{E} \left[\sup_{\|v\|_{\ell_2}=1} \sup_{\|Y\|_b=1} \sigma(\|G\|_p) \langle g, v \rangle + \left\langle \sum_{i=1}^M v_i G'_i, Y \right\rangle \right]. \end{aligned}$$

Proof Define the functionals

$$\begin{aligned} P_L(v, Y, Z) &= \left\langle \sum_{i=1}^M v_i G_i, Z \right\rangle + \left\langle \sum_{i=1}^M v_i G'_i, Y \right\rangle + \gamma \sigma(\|G_0\|_p) \\ P_R(v, Y, Z) &= \langle G_0, Z \rangle + \sigma(\|G_0\|_p) \langle g, v \rangle + \left\langle \sum_{i=1}^M v_i G'_i, Y \right\rangle. \end{aligned}$$

Let v and \hat{v} be unit vectors in \mathbb{R}^M , Y and \hat{Y} be $D_2 \times D_2$ matrices with $\|Y\|_b = \|\hat{Y}\|_b = 1$, and Z and \hat{Z} be $D_1 \times D_1$ matrices with $\|Z\|_d = \|\hat{Z}\|_d = 1$. Then we have

$$\begin{aligned} &\mathbb{E}[P_L(v, Y, Z)P_L(\hat{v}, \hat{Y}, \hat{Z})] - \mathbb{E}[P_R(v, Y, Z)P_L(\hat{v}, \hat{Y}, \hat{Z})] \\ &= \langle v, \hat{v} \rangle \langle Z, \hat{Z} \rangle + \langle v, \hat{v} \rangle \langle Y, \hat{Y} \rangle + \sigma(\|G_0\|_p)^2 - \langle Z, \hat{Z} \rangle - \sigma(\|G_0\|_p)^2 \langle v, \hat{v} \rangle - \langle v, \hat{v} \rangle \langle Y, \hat{Y} \rangle \\ &= (\sigma(\|G_0\|_p)^2 - \langle Z, \hat{Z} \rangle)(1 - \langle v, \hat{v} \rangle). \end{aligned}$$

The difference in expectations is greater than or equal to zero and equal to zero when $v = \hat{v}$ and $Y = \hat{Y}$. Hence, by Gordan's Lemma and a compactness argument,

$$\mathbb{E} \left[\inf_{\|v\|_{\ell_2}=1} \inf_{\|Y\|_b=1} \sup_{\|Z\|_d=1} Q_L(v, Y, Z) \right] \geq \mathbb{E} \left[\inf_{\|v\|_{\ell_2}=1} \inf_{\|Y\|_b=1} \sup_{\|Z\|_d=1} Q_R(v, Y, Z) \right]$$

completing the proof. \blacksquare

Together with Lemmas 3.12 and 3.13, we can prove the Lemma 3.4.

Proof [of Lemma 3.4] For $i = 1, \dots, M$, let $G_i \in \mathfrak{G}((1 - \beta)n)$ and $G'_i \in \mathfrak{G}(\beta n)$. Then

$$\begin{aligned}
& \mathbb{E} \left[\inf_{\|v\|_{\ell_2}=1} \left\| \sum_{i=1}^M v_i G_i \right\|_* - \left\| \sum_{i=1}^M v_i G'_i \right\|_* \right] \\
&= \mathbb{E} \left[\inf_{\|v\|_{\ell_2}=1} \inf_{\|Y\|=1} \sup_{\|Z\|=1} \left\langle \sum_{i=1}^M v_i G_i, Z \right\rangle + \left\langle \sum_{i=1}^M v_i G'_i, Y \right\rangle \right] \\
&\geq \mathbb{E} [\|G_0\|_*] - \mathbb{E} \left[\sup_{\|v\|_{\ell_2}=1} \sup_{\|Y\|=1} \sigma(\|G\|_*) \langle g, v \rangle + \left\langle \sum_{i=1}^M v_i G'_i, Y \right\rangle \right] \\
&\geq \mathbb{E} [\|G_0\|_*] - \mathbb{E} [\|G'_0\|_*] - \sqrt{M} \sqrt{\sigma(\|G\|_*)^2 + \sigma(\|G'\|_*)^2}
\end{aligned}$$

where the first inequality follows from Lemma 3.13, and the second inequality follows from Lemma 3.12.

Now we only need to plug in the expected values of the nuclear norm and the quantity $\sigma(\|G\|_*)$. Let G be sampled from $\mathfrak{G}(D)$. Then

$$\mathbb{E}\|G\|_* = D\mathbb{E}\sigma_i = \frac{8}{3\pi} D^{3/2} + q(D) \quad (3.13)$$

where $q(D)/D^{3/2} = o(1)$. The constant in front of the $D^{3/2}$ comes from integrating $\sqrt{\lambda}$ against the Marčenko-Pastur distribution (see, e.g., [17, 2]):

$$\frac{1}{2\pi} \int_0^4 \sqrt{4-t} dt = \frac{8}{3\pi} \approx 0.85.$$

Secondly, a straightforward calculation reveals

$$\sigma(\|G\|_*) = \sup_{\|H\| \leq 1} \|G\|_F = \sqrt{D}.$$

Plugging these values in with the appropriate dimensions completes the proof. ■

Proof [of Lemma 3.8] This lemma immediately follows from applying Lemma 3.12 with $\Delta = 0$ and from the calculations at the end of the proof above. It is also an immediate consequence of Lemma 3.21 from [15]. ■

4 Numerical Experiments

We now show that these asymptotic estimates hold even for small values of n . We conducted a series of experiments for a variety of the matrix sizes n , ranks r , and numbers of measurements m . As in the previous section, we let $\beta = \frac{r}{n}$ and $\mu = \frac{m}{n^2}$. For a fixed n , we constructed random recovery scenarios for low-rank $n \times n$ matrices. For each n , we varied μ between 0 and 1 where the matrix is completely determined. For a fixed n and μ , we generated all possible ranks such that $\beta(2 - \beta) \leq \mu$. This cutoff was chosen because beyond that point there would be an infinite set of matrices of rank r satisfying the m equations.

For each (n, μ, β) triple, we repeated the following procedure 10 times. A matrix of rank r was generated by choosing two random $n \times r$ factors Y_L and Y_R with i.i.d. random entries and setting




Figure 2: Random rank recovery experiments for (a) $n = 30$ and (b) $n = 40$. The color of each cell reflects the empirical recovery rate. White denotes perfect recovery in all experiments, and black denotes failure for all experiments. In both frames, we plot the Weak Bound (1.5), showing that the predicted recovery regions are contained within the empirical regions, and the boundary between success and failure is well approximated for large values of β .

$Y_0 = Y_L Y_R^*$. A matrix \mathbf{A} was sampled from the Gaussian ensemble with m rows and n^2 columns. Then the nuclear norm minimization

$$\begin{aligned} & \text{minimize} && \|X\|_* \\ & \text{subject to} && \mathbf{A} \text{vec } X = \mathbf{A} \text{vec } Y_0 \end{aligned}$$

was solved using the freely available software SeDuMi [26] using the semidefinite programming formulation described in [21]. On a 2.0 GHz Laptop, each semidefinite program could be solved in less than two minutes for 40×40 dimensional X . We declared Y_0 to be recovered if

$$\|X - Y_0\|_F / \|Y_0\|_F < 10^{-3}.$$

Figure 2 displays the results of these experiments for $n = 30$ and 40 . The color of the cell in the figures reflects the empirical recovery rate of the 10 runs (scaled between 0 and 1). White denotes perfect recovery in all experiments, and black denotes failure for all experiments. It is remarkable to note that not only are the plots very similar for $n = 30$ and $n = 40$, but that the Weak Bound falls completely within the white region and is an excellent approximation of the boundary between success and failure for large β .

References

- [1] A. Argyriou, C. A. Micchelli, and M. Pontil, “Convex multi-task feature learning,” *Machine Learning*, 2008, published online first at <http://www.springerlink.com/>.
- [2] Z. D. Bai, “Methodologies in spectral analysis of large dimensional random matrices,” *Statistica Sinica*, vol. 9, no. 3, pp. 611–661, 1999.

- [3] R. Baraniuk, M. Davenport, R. DeVore, and M. Wakin, “A simple proof of the restricted isometry property for random matrices,” *Constructive Approximation*, 2008, to Appear. Preprint available at <http://dsp.rice.edu/cs/jlcs-v03.pdf>.
- [4] C. Beck and R. D’Andrea, “Computational study and comparisons of LFT reducibility methods,” in *Proceedings of the American Control Conference*, 1998.
- [5] E. J. Candès, J. Romberg, and T. Tao, “Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information,” *IEEE Trans. Inform. Theory*, vol. 52, no. 2, pp. 489–509, 2006.
- [6] E. J. Candès and T. Tao, “Decoding by linear programming,” *IEEE Transactions on Information Theory*, vol. 51, no. 12, pp. 4203–4215, 2005.
- [7] D. L. Donoho and J. Tanner, “Neighborliness of randomly projected simplices in high dimensions,” *Proc. Natl. Acad. Sci. USA*, vol. 102, no. 27, pp. 9452–9457, 2005.
- [8] —, “Sparse nonnegative solution of underdetermined linear equations by linear programming,” *Proc. Natl. Acad. Sci. USA*, vol. 102, no. 27, pp. 9446–9451, 2005.
- [9] L. El Ghaoui and P. Gahinet, “Rank minimization under LMI constraints: A framework for output feedback problems,” in *Proceedings of the European Control Conference*, 1993.
- [10] M. Fazel, “Matrix rank minimization with applications,” Ph.D. dissertation, Stanford University, 2002.
- [11] M. Fazel, H. Hindi, and S. Boyd, “A rank minimization heuristic with application to minimum order system approximation,” in *Proceedings of the American Control Conference*, 2001.
- [12] Y. Gordan, “Some inequalities for gaussian processes and applications,” *Israel Journal of Math*, vol. 50, pp. 265–289, 1985.
- [13] —, “Gaussian processes and almost spherical sections of convex bodies,” *Annals of Probability*, vol. 16, pp. 180–188, 1988.
- [14] R. A. Horn and C. R. Johnson, *Topics in Matrix Analysis*. New York: Cambridge University Press, 1991.
- [15] M. Ledoux and M. Talagrand, *Probability in Banach Spaces*. Berlin: Springer-Verlag, 1991.
- [16] N. Linial, E. London, and Y. Rabinovich, “The geometry of graphs and some of its algorithmic applications,” *Combinatorica*, vol. 15, pp. 215–245, 1995.
- [17] V. A. Marčenko and L. A. Pastur, “Distributions of eigenvalues for some sets of random matrices,” *Math. USSR-Sbornik*, vol. 1, pp. 457–483, 1967.
- [18] R. Meka, P. Jain, C. Caramanis, and I. S. Dhillon, “Rank minimization via online learning,” in *Proceedings of the International Conference on Machine Learning*, 2008.
- [19] M. Mesbahi and G. P. Papavassilopoulos, “On the rank minimization problem over a positive semidefinite linear matrix inequality,” *IEEE Transactions on Automatic Control*, vol. 42, no. 2, pp. 239–243, 1997.
- [20] P. A. Parrilo and S. Khatri, “On cone-invariant linear matrix inequalities,” *IEEE Trans. Automat. Control*, vol. 45, no. 8, pp. 1558–1563, 2000.
- [21] B. Recht, M. Fazel, and P. Parrilo, “Guaranteed minimum rank solutions of matrix equations via nuclear norm minimization,” Submitted. Preprint Available at <http://www.ist.caltech.edu/~brecht/publications.html>.

- [22] B. Recht, W. Xu, and B. Hassibi, “Necessary and sufficient conditions for success of the nuclear norm heuristic for rank minimization,” in *Proceedings of the 47th IEEE Conference on Decision and Control*, 2008.
- [23] J. D. M. Rennie and N. Srebro, “Fast maximum margin matrix factorization for collaborative prediction,” in *Proceedings of the International Conference of Machine Learning*, 2005.
- [24] D. Slepian, “The one-sided barrier problem for gaussian noise,” *Bell System Technical Journal*, vol. 41, pp. 463–501, 1962.
- [25] M. Stojnic, W. Xu, and B. Hassibi, “Compressed sensing - probabilistic analysis of a null-space characterization,” in *IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, 2008.
- [26] J. F. Sturm, “Using SeDuMi 1.02, a MATLAB toolbox for optimization over symmetric cones,” *Optimization Methods and Software*, vol. 11-12, pp. 625–653, 1999.
- [27] S. J. Szarek, “Metric entropy of homogeneous spaces,” in *Quantum probability (Gdańsk, 1997)*, ser. Banach Center Publ. Warsaw: Polish Acad. Sci., 1998, vol. 43, pp. 395–410, preprint available at arXiv:math/9701213v1.
- [28] K. Q. Weinberger and L. K. Saul, “Unsupervised learning of image manifolds by semidefinite programming,” *International Journal of Computer Vision*, vol. 70, no. 1, pp. 77–90, 2006.

A Appendix

A.1 Rank-deficient case of Theorem 1.1

As promised above, here is the completion of the proof of Theorem 1.1

Proof In an appropriate basis, we may write

$$X_0 = \begin{bmatrix} X_{11} & 0 \\ 0 & 0 \end{bmatrix} \quad \text{and} \quad X_* - X_0 = Y = \begin{bmatrix} Y_{11} & Y_{12} \\ Y_{21} & Y_{22} \end{bmatrix}.$$

If Y_{11} and $Y_{22} - Y_{21}Y_{11}^{-1}Y_{12}$ have full rank, then all our previous arguments apply. Thus, assume that at least one of them is not full rank. Nonetheless, it is always possible to find an *arbitrarily small* $\epsilon > 0$ such that

$$Y_{11} + \epsilon I \quad \text{and} \quad \begin{bmatrix} Y_{11} + \epsilon I & Y_{12} \\ Y_{21} & Y_{22} + \epsilon I \end{bmatrix}$$

are full rank. This, of course, is equivalent to having $Y_{22} + \epsilon I - Y_{21}(Y_{11} + \epsilon I)^{-1}Y_{12}$ full rank. We

can write

$$\begin{aligned}
\|X_*\|_* &= \|X_0 + X_* - X_0\|_* \\
&= \left\| \begin{bmatrix} X_{11} & 0 \\ 0 & 0 \end{bmatrix} + \begin{bmatrix} Y_{11} & Y_{12} \\ Y_{21} & Y_{22} \end{bmatrix} \right\|_* \\
&\geq \left\| \begin{bmatrix} X_{11} - \epsilon I & 0 \\ 0 & Y_{22} - Y_{21}(Y_{11} + \epsilon I)^{-1}Y_{12} \end{bmatrix} \right\|_* - \left\| \begin{bmatrix} Y_{11} + \epsilon I & Y_{12} \\ Y_{21} & Y_{21}(Y_{11} + \epsilon I)^{-1}Y_{12} \end{bmatrix} \right\|_* \\
&= \|X_{11} - \epsilon I\|_* + \left\| \begin{bmatrix} 0 & 0 \\ 0 & Y_{22} - Y_{21}(Y_{11} + \epsilon I)^{-1}Y_{12} \end{bmatrix} \right\|_* - \left\| \begin{bmatrix} Y_{11} + \epsilon I & Y_{12} \\ Y_{21} & Y_{21}(Y_{11} + \epsilon I)^{-1}Y_{12} \end{bmatrix} \right\|_* \\
&\geq \|X_0\|_* - r\epsilon + \left\| \begin{bmatrix} \epsilon I - \epsilon I & 0 \\ 0 & Y_{22} - Y_{21}(Y_{11} + \epsilon I)^{-1}Y_{12} \end{bmatrix} \right\|_* - \left\| \begin{bmatrix} Y_{11} + \epsilon I & Y_{12} \\ Y_{21} & Y_{21}(Y_{11} + \epsilon I)^{-1}Y_{12} \end{bmatrix} \right\|_* \\
&\geq \|X_0\|_* - 2r\epsilon + \left\| \begin{bmatrix} -\epsilon I & 0 \\ 0 & Y_{22} - Y_{21}(Y_{11} + \epsilon I)^{-1}Y_{12} \end{bmatrix} \right\|_* - \left\| \begin{bmatrix} Y_{11} + \epsilon I & Y_{12} \\ Y_{21} & Y_{21}(Y_{11} + \epsilon I)^{-1}Y_{12} \end{bmatrix} \right\|_* \\
&\geq \|X_0\|_* - 2r\epsilon,
\end{aligned}$$

where the last inequality follows from the condition of part 1 and noting that

$$X_0 - X_* = \begin{bmatrix} -\epsilon I & 0 \\ 0 & Y_{22} - Y_{21}(Y_{11} + \epsilon I)^{-1}Y_{12} \end{bmatrix} + \begin{bmatrix} Y_{11} + \epsilon I & Y_{12} \\ Y_{21} & Y_{21}(Y_{11} + \epsilon I)^{-1}Y_{12} \end{bmatrix},$$

lies in the null space of $\mathcal{A}(\cdot)$ and the first matrix above has rank more than r . But, since ϵ can be arbitrarily small, this implies that $X_0 = X_*$. \blacksquare

A.2 Lipschitz Constants of F_I and F_S

We begin with the proof of Lemma 3.9 and then use this to estimate the Lipschitz constant in Lemma 3.6.

Proof [of Lemma 3.9] Note that the function F_S is convex as we can write as a supremum of a collection of convex functions

$$F_S(X_1, \dots, X_M) = \sup_{\|v\|_{\ell_2}=1} \sup_{\|Z\|<1} \left\langle \sum_{i=1}^M v_i X_i, Z \right\rangle. \quad (\text{A.1})$$

The Lipschitz constant L is bounded above by the maximal norm of a subgradient of this convex function. That is, if we denote $\bar{X} := (X_1, \dots, X_M)$, then we have

$$L \leq \sup_{\bar{X}} \sup_{\bar{Z} \in \partial F_S(\bar{X})} \left(\sum_{i=1}^M \|Z_i\|_F^2 \right)^{1/2}.$$

Now, by (A.1), a subgradient of F_S at \bar{X} is given of the form $(v_1 Z, v_2 Z, \dots, v_M Z)$ where v has norm 1 and Z has operator norm 1. For any such subgradient

$$\sum_{i=1}^M \|v_i Z\|_F^2 = \|Z\|_F^2 \leq n$$

bounding the Lipschitz constant as desired. \blacksquare

Proof [of Lemma 3.6] For $i = 1, \dots, M$, let $X_i, \hat{X}_i \in \mathbb{R}^{n_1 \times n_1}$, and $Y_i, \hat{Y}_i \in \mathbb{R}^{n_2 \times n_2}$. Let

$$w^* = \arg \min_{\|w\|_{\ell_2}=1} \left\| \sum_{i=1}^M w_i \hat{X}_i \right\|_* - \left\| \sum_{i=1}^M w_i \hat{Y}_i \right\|_*.$$

Then we have that

$$\begin{aligned} & F_I(X_1, \dots, X_M, Y_1, \dots, Y_M) - F_I(\hat{X}_1, \dots, \hat{X}_M, \hat{Y}_1, \dots, \hat{Y}_M) \\ &= \left(\inf_{\|v\|_{\ell_2}=1} \left\| \sum_{i=1}^M v_i X_i \right\|_* - \left\| \sum_{i=1}^M v_i Y_i \right\|_* \right) - \left(\inf_{\|w\|_{\ell_2}=1} \left\| \sum_{i=1}^M w_i \hat{X}_i \right\|_* - \left\| \sum_{i=1}^M w_i \hat{Y}_i \right\|_* \right) \\ &\leq \left\| \sum_{i=1}^M w_i^* X_i \right\|_* - \left\| \sum_{i=1}^M w_i^* Y_i \right\|_* - \left\| \sum_{i=1}^M w_i^* \hat{X}_i \right\|_* + \left\| \sum_{i=1}^M w_i^* \hat{Y}_i \right\|_* \\ &\leq \left\| \sum_{i=1}^M w_i^* (X_i - \hat{X}_i) \right\|_* + \left\| \sum_{i=1}^M w_i^* (Y_i - \hat{Y}_i) \right\|_* \\ &\leq \sup_{\|w\|_{\ell_2}=1} \left\| \sum_{i=1}^M w_i (X_i - \hat{X}_i) \right\|_* + \left\| \sum_{i=1}^M w_i (Y_i - \hat{Y}_i) \right\|_* = \sup_{\|w\|_{\ell_2}=1} \left\| \sum_{i=1}^M w_i \tilde{X}_i \right\|_* + \left\| \sum_{i=1}^M w_i \tilde{Y}_i \right\|_* \end{aligned}$$

where $\tilde{X}_i = X_i - \hat{X}_i$ and $\tilde{Y}_i = Y_i - \hat{Y}_i$. This last expression is a convex function of \tilde{X}_i and \tilde{Y}_i ,

$$\sup_{\|w\|_{\ell_2}=1} \left\| \sum_{i=1}^M w_i \tilde{X}_i \right\|_* + \left\| \sum_{i=1}^M w_i \tilde{Y}_i \right\|_* = \sup_{\|w\|_{\ell_2}=1} \sup_{\|Z_X\| < 1} \sup_{\|Z_Y\| < 1} \left\langle \sum_{i=1}^M w_i \tilde{X}_i, Z_X \right\rangle + \left\langle \sum_{i=1}^M w_i \tilde{Y}_i, Z_Y \right\rangle$$

with Z_X $n_1 \times n_2$ and Z_Y $n_2 \times n_2$. Using an identical argument as the one presented in the proof of Lemma 3.9, we have that a subgradient of this expression is of the form

$$(w_1 Z_X, w_2 Z_X, \dots, w_M Z_X, w_1 Z_Y, w_2 Z_Y, \dots, w_M Z_Y)$$

where w has norm 1 and Z_X and Z_Y have operator norms 1, and thus

$$\sum_{i=1}^M \|w_i Z_X\|_F^2 + \|w_i Z_Y\|_F^2 = \|Z_X\|_F^2 + \|Z_Y\|_F^2 \leq n_1 + n_2$$

completing the proof. ■

A.3 Compactness Argument for Comparison Theorems

Proposition A.1 *Let Ω be a compact metric space with distance function ρ . Suppose that f and g are real-valued function on Ω such that f is continuous and for any finite subset $X \subset \Omega$*

$$\max_{x \in X} f(x) \leq \max_{x \in X} g(x).$$

Then

$$\sup_{x \in \Omega} f(x) \leq \sup_{x \in \Omega} g(x).$$

Proof Let $\epsilon > 0$. Since f is continuous and Ω is compact, f is uniformly continuous on Ω . That is, there exists a $\delta > 0$ such that for all $x, y \in \Omega$, $\rho(x, y) < \delta$ implies $|f(x) - f(y)| < \epsilon$. Let X_δ be a δ -net for Ω . Then, for any $x \in \Omega$, there is a y in the δ -net with $\rho(x, y) < \delta$ and hence

$$f(x) \leq f(y) + \epsilon \leq \sup_{z \in X_\delta} f(z) + \epsilon \leq \sup_{z \in X_\delta} g(z) + \epsilon \leq \sup_{z \in \Omega} g(z) + \epsilon.$$

Since this holds for all $x \in \Omega$ and $\epsilon > 0$, this completes the proof. ■