

Subspace Expanders and Matrix Rank Minimization

Amin Khajehnejad, Samet Oymak and Babak Hassibi
California Institute of Technology, Pasadena CA 91125 *

February 22, 2011

Abstract

Matrix rank minimization (RM) problems recently gained extensive attention due to numerous applications in machine learning, system identification and graphical models. In RM problem, one aims to find the matrix with the lowest rank that satisfies a set of linear constraints. The existing algorithms include nuclear norm minimization (NNM) and singular value thresholding. Thus far, most of the attention has been on i.i.d. Gaussian measurement operators. In this work, we introduce a new class of measurement operators, and a novel recovery algorithm, which is notably faster than NNM. The proposed operators are based on what we refer to as subspace expanders, which are inspired by the well known expander graphs based measurement matrices in compressed sensing. We show that given an $n \times n$ PSD matrix of rank r , it can be uniquely recovered from a minimal sampling of $O(nr)$ measurements using the proposed structures, and the recovery algorithm can be cast as matrix inversion after a few initial processing steps.

1 Introduction

In the rank minimization problem, one would like to find a matrix X with the lowest rank that satisfies a set of linear constraints $\mathcal{A}(X)$, often of smaller size than the number of matrix entries. This problem in its full generality is NP-hard. However, a number of recent papers have demonstrated that under certain circumstances, and when the unknown matrix is sufficiently low rank, RM can be solved via convex optimization programs, mainly nuclear norm minimization [1, 2, 3, 15]. Random linear Gaussian measurements are considered as a standard choice, for which certain recovery guarantees can be asymptotically proven for NNM. The two main conditions that certify the success of NNM are Restricted Isometry Property (RIP) [2] and null space conditions [3]. Based on the analysis initially developed by Stojnic for sparse recovery, tight analytical thresholds for NNM have been recently found by Oymak et al. and Chandrasekaran et al. using a “escape through the mesh” analysis of the null space conditions when measurement operators are i.i.d. Gaussian [15, 16, 17]. In parallel, there have also been promising

*This work was supported in part by the National Science Foundation under grants CCF-0729203, CNS-0932428 and CCF-1018927, by the Office of Naval Research under the MURI grant N00014-08-1-0747, and by Caltech’s Lee Center for Advanced Networking.

results on matrix completion problem as well, whereby one observes a subset of the entries of a low rank matrix, rather than linear combinations [4, 5].

The RM problem is often regarded as a generalization of compressed sensing(CS), wherein the aim is to recover a sparse vector from a set of ill-posed linear measurements, represented by a wide measurement matrix $A^{m \times n}$ with $m < n$ [12]. It is now well understood that although certain random ensembles of measurement matrices (e.g. Gaussian, partial Fourier, etc.) are legitimate choices for CS, carefully designed matrices can lead to additional benefits for the sparse compression/recovery. A few examples are faster encoding time and recovery algorithms for sparse matrices and in particular expander graphs [6, 7, 8, 9], higher recovery thresholds for measurement matrices based on Reed-Solomon codes [11, 14] etc. To the best of our knowledge, this point has not been studied in RM problem, where only random measurement ensembles (mostly Gaussian) have been studied. Do alternative recovery algorithms for RM problem exist, with success guarantees on certain classes of carefully designed linear measurement operators?

In this paper, we introduce a new class of measurement operators, along with a novel recovery algorithm that is provably successful for the proposed operators, and is faster than NNM. The premise of the recovery algorithm relies on the measurement operator $\mathcal{A}(\cdot) : \mathbb{R}^{n \times n} \rightarrow \mathbb{R}^{m \times m}$ having three key properties: Hermitian, low density and rank expansion. Low density property means that $\mathcal{A}(X)$ can be described as the linear combination of d linear operators $\mathcal{A}(X) = \sum_{i=1}^d \mathcal{A}_i(X)$, where d is a constant, and for each $1 \leq i \leq d$, the rank of $\mathcal{A}_i(X)$ is not larger than the rank of X . In fact, the low density operators that we introduce are characterized by only $O(mn)$ variables, as opposed to full i.i.d Gaussian linear measurements that require $(mn)^2$ variables. The interpretation of expansion is that $\mathcal{A}(\cdot)$ maps a sufficiently low dimensional subspace to a higher dimensional subspace, which is equivalent to mapping every positive semidefinite matrix X to a positive semidefinite matrix $\mathcal{A}(X)$ with rank greater than $c \cdot \text{rank}(X)$, where $c > 1$ is a constant. The contributions of this paper are thus threefold. We first prove that low density rank expander operators exist. We then provide a uniqueness result, in the sense that a sufficiently low rank PSD matrix X is the unique PSD solution of the equations given by a high quality rank expander. We further propose a new recovery algorithm and provide theoretical recovery guarantees when the suggested rank expander measurement operators are exploited. Although most of the results in this paper are stressed for the case of PSD matrices, we briefly mention a generalization of our results to the case of Hermitian matrices. This includes both the existence of high quality expanders operating on Hermitian matrices, and the solidity of the recovery algorithm. The validity of the proposed algorithm is verified by numerical simulations.

2 Basic Definitions and Lemmas

Let \mathbb{S}^n denote the space of Hermitian matrices of size $n \times n$, and \mathbb{S}_+^n denote the set of positive semidefinite (PSD) matrices. An orthogonal projection is a matrix $P \in \mathbb{S}_+^n$ with $P^2 = P$. We say that $U \in \mathbb{R}^{n_1 \times n_2}$ is a partial unitary matrix if $U^T U = I$, i.e. the columns of U form an orthonormal set. Notice that $U U^T$

is an orthogonal projection. Let $\eta_+(X), \eta_-(X), \eta_0(X)$ denote the number of positive, negative and zero eigenvalues of X , respectively. Also for a Hermitian matrix X , let X_- and X_+ denote the PSD matrices induced by the negative and positive eigenvalues of X respectively, i.e. $X = X_+ - X_-$.

For a given matrix $X \in \mathbb{R}^{n_1 \times n_2}$, $\lambda_i(X)$ and $\sigma_i(X)$ denotes i 'th largest eigenvalue and the i 'th largest singular value, respectively. The nuclear norm, spectral norm and Frobenius norm operators are denoted by $\|\cdot\|_*$, $\|\cdot\|$ and $\|\cdot\|_F$, respectively, and are defined by $\|X\|_* = \sum_{i=1}^{\min\{n_1, n_2\}} \sigma_i(X)$, $\|X\| = \sigma_1(X)$ and $\|X\|_* = \left(\sum_{i=1}^{\min\{n_1, n_2\}} \sigma_i^2(X)\right)^{1/2}$. In addition, we define $\text{Span}_C(X), \text{Span}_R(X)$ to be the linear spaces spanned by the columns and rows of X , respectively.

A function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is called L -Lipschitz if $|f(x) - f(y)| \leq L\|x - y\|_{\ell_2}$, for every x, y . For hermitian matrices A, B , $A \succeq B$ means that $A - B$ is positive semidefinite. For a linear operator $\mathcal{A}(\cdot)$ acting on a linear space, we denote the null space of \mathcal{A} by $\mathcal{N}(\mathcal{A})$, i.e. $W \in \mathcal{N}(\mathcal{A})$ iff $\mathcal{A}(W) = 0$. We denote by $\mathcal{G}(d_1, d_2)$ the ensemble of real $d_1 \times d_2$ matrices in which the entries are i.i.d. $\mathcal{N}(0, 1)$ (zero-mean, unit variance Gaussian).

The following lemmas are crucial to the technical discussions of this paper. Their proofs are skipped due to space considerations.

Lemma 2.1. *Let $f(X)$ be a function on matrices in the following form: $f(X) = \sum_{i=1}^m a_i \sigma_i(X)$ for some real constants $\{a_i\}_{i=1}^m$. Then $f(X)$ is a $\sqrt{\sum_{i=1}^m a_i^2}$ Lipschitz function of X .*

Lemma 2.2. (A Gaussian concentration inequality, [20])

Let \mathbf{x} be drawn from $\mathcal{G}(n, 1)$ and $f : \mathbb{R}^n \rightarrow \mathbb{R}$ be a function with Lipschitz constant L . Then, we have the following concentration inequality

$$\mathbb{P}(|f(x) - \mathbb{E}f(x)| \geq t) \leq 2 \exp\left(-\frac{t^2}{2L^2}\right) \quad (1)$$

Lemma 2.3. (Weyl's Inequalities, [10])

Let $A, B \in \mathbb{S}^n$. Then:

$$\lambda_j(A + B) \leq \lambda_i(A) + \lambda_{j-i+1}(B) \quad \forall i \leq j \quad (2)$$

$$\lambda_j(A + B) \geq \lambda_i(A) + \lambda_{j-i+n}(B) \quad \forall j \leq i \quad (3)$$

3 Rank Expanders and Proposed Operators

Definition 1. *Let $\mathcal{A} : \mathbb{R}^{n \times n} \rightarrow \mathbb{R}^{m \times m}$ be a linear operator with $m < n$. For $d > 0, 0 \leq \epsilon < 1, 1 \leq r_0 \leq n$, we say that \mathcal{A} is an unbalanced (ϵ, d, r_0, n) -rank expander, if it satisfies the following conditions:*

1. *For every $X \in \mathbb{S}_+^n$, $\mathcal{A}(X) \in \mathbb{S}_+^m$*
2. *For every $X \in \mathbb{S}^n$, $\text{rank}(\mathcal{A}(X)) \leq d \cdot \text{rank}(X)$*
3. *For all orthogonal projections P with $\text{rank}(P) = r \leq r_0$, $rd \geq \text{rank}(\mathcal{A}(P)) > (1 - \epsilon)rd$*

This definition is inspired by the definition of unbalanced expander graphs, that maintain similar properties with respect to positive vectors (instead of PSD matrices) and with ℓ_0 -norm (instead of rank). An unbalanced d -regular (k, ϵ) -expander graph is a bipartite graph with n nodes on the left and m nodes on the right, and regular degree d for left hand side nodes, such that every subset S of left nodes with $|S| \leq r_0$ has a neighborhood $N(S)$ of size at least $|N(S)| \geq (1 - \epsilon)|S|d$. Unbalanced expander graphs have been proven to have elegant properties that make them suitable for sparse vector recovery (a.k.a. compressed sensing) in addition to being useful as parity check matrices for error correcting codes. With that in mind, one might be inspired to generalize the notion of expander graphs to subspace (rank) expanders, in order to obtain operators that can be used in low rank matrix recovery. The following lemma is immediate.

Lemma 3.1. *If $\mathcal{A}(\cdot)$ is an (ϵ, d, r_0, n) -rank expander then, for every $X \in \mathbb{S}_+^n$ with rank $r \leq r_0$, we have $rd \geq \text{rank}(\mathcal{A}(X)) \geq (1 - \epsilon)rd$*

We now move on to describe the proposed measurement structures. Afterwards, we prove that these constructions indeed result in rank expanders and that the expansion property allows us to find alternative fast reconstruction algorithms for the RM problem.

3.1 Proposed Measurement Operator

Let $G_1, \dots, G_d \in \mathbb{R}^{m \times n}$ be matrices to be specified later. The proposed measurement operator $\mathcal{A}(\cdot)$ has the following low density form:

$$\mathcal{A}(X) = \sum_{i=1}^d G_i X G_i^T \quad (4)$$

Where $X \in \mathbb{R}^{n \times n}$. We will prove that upon appropriate choices of G_i 's, $\mathcal{A}(\cdot)$ is an unbalanced rank expander. It is easy to check that with this choice of $\mathcal{A}(\cdot)$, conditions 1 and 2 of Definition 1 are immediately satisfied. Furthermore, as long as the $G_i X G_i^T$'s are almost incoherent, one would expect their ranks to add up. In particular, it is easy to show that when X is fixed and $\{G_i\}_{i=1}^d$ are drawn i.i.d. from $\mathcal{G}(m, n)$, we have

$$\mathbb{P}(\text{rank}(\mathcal{A}(X)) = \min\{d \times \text{rank}(X), m\}) = 1. \quad (5)$$

However, the challenge of condition 3 is in the fact that the rank expansion property must hold for every sufficiently low rank X . Let $X \in \mathbb{S}_+^n$ and $X^{1/2}$ denote an arbitrary square root of X (i.e. $X = X^{1/2} X^{1/2}$). Note that $\mathcal{A}(X)$ can be written in the following form:

$$\mathcal{A}(X) = \begin{pmatrix} G_1 X^{1/2} & \dots & G_d X^{1/2} \end{pmatrix} \begin{pmatrix} G_1 X^{1/2} & \dots & G_d X^{1/2} \end{pmatrix}^T \quad (6)$$

It then follows that

$$\text{rank}(\mathcal{A}(X)) = \text{rank} \begin{pmatrix} G_1 X^{1/2} & G_2 X^{1/2} & \dots & G_d X^{1/2} \end{pmatrix} \quad (7)$$

For analyzing the rank expansion property, we can thus limit ourselves to the form (7).

3.2 Existence of Rank Expanders

Our goal is to prove the existence of high quality (small ϵ) rank expanders for certain regimes of d, ϵ, r_0 and n . Based on Lemma 3.1, we can restrict our attention to X being an orthogonal projection of rank at most r_0 . Our analysis is for the case when G_i 's are chosen i.i.d. from $\mathcal{G}(m, n)$ ¹. The main existence theorem is the following:

Theorem 3.1. Existence of Rank Expander *For any $0 < \epsilon < 1$ there are constants C_1 and C_2 so that for any n and $r_0 \leq n$, whenever $m = \sqrt{C_1 C_2 n r_0}$ and $d = \sqrt{\frac{C_2 n}{C_1 r_0}}$ and $\{G_i\}_{i=1}^d$'s are independent instances of $\mathcal{G}(m, n)$, the operator $\mathcal{A}(X) = \sum_{i=1}^d G_i X G_i^T$ is an (ϵ, d, r_0, n) expander with probability at least $1 - \exp(-\Omega(n))$.*

Before explaining the technicalities involved in the proof of the above theorem, consider the following argument. Given $\mathcal{A} : \mathbb{R}^{n \times n} \rightarrow \mathbb{R}^{m \times m}$, suppose for all $X \in \mathbb{S}_+^n$ with $\text{rank}(X) \leq r^*$, X can be uniquely decoded from $\mathcal{A}(X)$, say by exhaustive search. It then follows that \mathcal{A} has to be injective on the restricted domain $\{X \in \mathbb{S}_+^n : \text{rank}(X) \leq r^*\}$. It will soon become apparent in the sequel that given an (ϵ, d, r_0, n) expander, for $r^* = r_0/2$ this condition holds. A simple argument counting the degrees of freedom of the low rank domain and the corresponding range of A reveals that the problem parameters should satisfy the following relationship.

$$m = \Omega(\sqrt{n r_0}), \quad m d = \Omega(n), \quad d r_0 = O(m) \quad (8)$$

In fact, it turns out that Theorem 3.1 is true as long as (8) holds asymptotically, which implies the optimality of the number of measurements in the suggested expander operators. For the proof, we set $m = C_1 d r_0$ and $d m = C_2 n$ where $C_1 > 1$, $C_2 > 1$ will be the constants in Theorem 3.1.

Proof sketch of Theorem 3.1. The proof is based on three major technical steps.

Step 1: We consider an ϵ_0 cover with operator norm $\|\cdot\|$ over the set of orthogonal projections of rank $r \leq r_0$. From [13], we know that there is such a cover of size at most $M = (C_0/\epsilon_0)^{nr}$, which we denote by $\{U_i U_i^T\}_{i=1}^M \in \mathbb{S}_+^n$, with $U_i \in \mathbb{R}^{n \times r}$. Also, we first focus on a particular rank $r \leq r_0$, and later union bound the undesirable probability over all values of $r \leq r_0$.

Step 2: Now consider a U_i from the ϵ_0 cover. Denote $\mathcal{B}(U_i) = [G_1 U_i \dots G_d U_i] \in \mathbb{R}^{m \times dr}$. Since $\text{rank}(\mathcal{A}(U_i U_i^T)) = \text{rank}(\mathcal{B}(U_i))$, we can focus on $\mathcal{B}(U_i)$. Note that since U_i is a fixed partial unitary, due to the unitary invariance of i.i.d. Gaussian matrices, $\mathcal{B}(U_i)$ has i.i.d. Gaussian distribution. Now define the function $f(X) = \sum_{i=(1-\epsilon)dr+1}^{dr} \sigma_i(X)$, for $X \in \mathbb{R}^{m \times rd}$. Notice that $f(X) > 0$ implies $\text{rank}(f(X)) > (1 - \epsilon)dr$, because it means some of the smallest ϵdr singular values of X are nonzero. On the other hand $f(X)$ is a linear function of the singular values of X , and thus it satisfies the Lipschitz condition of Lemma 2.1. Since $\mathcal{B}(U_i)$ is Gaussian, we can apply Lemma 2.2 to get the following concentration bound

$$\mathbb{P}(f(\mathcal{B}(U_i)) < \delta \frac{e}{\sqrt{m}}) < \exp(-\frac{(C_3 - \delta)^2 e}{2}) \quad (9)$$

¹Existence of expanders using other ensembles of matrices G , and in particular sparse matrices, shall remain as an interesting open problem

where $C_3 = 1 - \sqrt{1/C_1}$ and $e = \epsilon dr m$. Here δ serves as a safety margin, in order to account for the perturbation $P - U_i U_i^T$, when we consider a certain orthogonal projection P , which we know is the proximity of some $U_i U_i^T$ of the cover, i.e. $\|P - U_i U_i^T\| < \epsilon_0$. In other words, as will be shown in the next step, lower bounding $f(\mathcal{B}(U_i))$ certifies that $\text{rank}(\mathcal{A}(P))$ will also be large, and thus (9) is basically determines a lower bound on the probability of failure. The exponent of the righthand side of (9) is $e = \epsilon dr m = \epsilon C_2 r n$ which is proportional to the exponent of the size of the cover $\log(M) = O(rn)$. Consequently, with careful choices of parameters, using a union bound on failures, we can make sure $f(\mathcal{B}(U)) \geq \delta \frac{e}{\sqrt{m}}$ for all $U \in \{U_i\}$ w.h.p. In particular we need:

$$(C_3 - \delta)^2 \epsilon C_2 > (\log C_0 - \log \epsilon_0) \quad (10)$$

Step 3: Now it remains to show that if $E = P - P_i$ is a perturbation on P_i which makes $\text{rank}(P) \leq (1 - \epsilon)dr$, then $\|E\|$ has to be large because $f(\mathcal{B}(U_i))$ is large. In particular, showing $\|E\| > \epsilon_0$ will finish the proof since we know that $\|P - P_i\| \leq \epsilon_0$. In order to show this step, we make use of Lemma 2.3 to find $\lambda_{i+(1-\epsilon)dr}(\mathcal{A}(P_i)) \leq \lambda_i(\mathcal{A}(E))$ and hence to deduce that $\|\mathcal{B}(E_+^{1/2})\|_* \geq f(\mathcal{B}(U_i))$. We carry out some more arguments to upper bound $\|\mathcal{B}(E_+^{1/2})\|_*$ in terms of $\|E\|$ to get a contradiction as long as $\delta\sqrt{\epsilon} > 2\sqrt{\epsilon_0}$ (*) holds. Finally, we conclude that whenever (*) and the condition (10) are satisfied, with high probability $f(\mathcal{B}(\sqrt{P})) > 0$ for all projections P with rank r , which implies $\text{rank}(\mathcal{A}(P)) > (1 - \epsilon)dr$ as desired. In particular, sufficiently large values of C_1, C_2 will do the job. Also, inequalities (8) will similarly work since increasing C_1, C_2 only improve the conditions. ■

4 Fast Recovery Algorithm

Before presenting the main algorithm, we will provide some results about low rank positive semidefinite matrix recovery. Specifically, we emphasize on a uniqueness result concerning PDS matrices and the rank expander operators. Suppose a matrix $X_0 \in \mathbb{S}_+^n$ and a linear operator $\mathcal{A}(\cdot)$ are given, and we ask under what conditions X_0 is the unique (PSD) inverse image of $\mathcal{A}(X_0)$, i.e. it is possible to recover X_0 by simply characterizing the set $\{X \mid X \succeq 0, \mathcal{A}(X) = \mathcal{A}(X_0)\}$. The following lemma which is adopted from [18, 15] provides an answer to this question.

Lemma 4.1. *Any PSD matrix X of rank at most r is the unique PSD inverse image of $\mathcal{A}(X)$, if and only if every nonzero Hermitian $W \in \mathcal{N}(\mathcal{A})$ has at least $r + 1$ negative eigenvalues.*

Now, we explain how rank expanders can facilitate the existence of the condition in Lemma 4.1.

Lemma 4.2. *Let $\mathcal{A}(\cdot)$ be an (ϵ, d, r_0, n) -rank expander with $\epsilon < 1/2$. Then for every nonzero Hermitian $W \in \mathcal{N}(\mathcal{A})$ we have that $\eta_-(W) > r_0/2$.*

Proof. Any $W \in \mathcal{N}(\mathcal{A})$, write $W = W_+ - W_-$. Since $\mathcal{A}(W) = 0$ we have $B = \mathcal{A}(W_+) = \mathcal{A}(W_-)$. Assume $\eta_-(W) \leq r_0/2$. Let $r = \min\{\eta_-(W), \eta_+(W)\}$. Then $dr \geq \text{rank}(B)$. On the other hand

Algorithm 1 Reconstruct a low rank PSD matrix X from under-determined linear measurements $Y = \sum_{i=1}^d A_i X A_i^*$.

- 1: **Input:**
 - 2: Constant integer $d \geq 1$.
 - 3: Matrices $A_i \in \mathbb{R}^{m \times n}$, $1 \leq i \leq d$, and $Y \in \mathbb{R}^{m \times m}$.
 - 4: **Output:**
 - 5: Low rank PSD matrix X .
 - 6: **Initialize**
 - 7: Compute $Y = S \Sigma S^*$, with S full column rank (SVD).
 - 8: Set $P = I - S S^*$.
 - 9: Set $Q := \text{Null}((P A_1)^T, \dots, (P A_d)^T)^T$.
 - 10: Compute $B_i = A_i Q$, and set $M = \sum_{i=1}^d B_i \otimes B_i$.
 - 11: Find $X \in \mathbb{R}^{n \times n}$ with $\text{vec}(X) = (Q \otimes Q) M^\dagger \text{vec}(Y)$.
-

$\text{rank}(W_+ + W_-) = \text{rank}(W_+) + \text{rank}(W_-) \geq 2r$ hence we have $\text{rank}(\mathcal{A}(W_+ + W_-)) \geq 2(1 - \epsilon)dr$ since $2r \leq r_0$. Note that $\mathcal{A}(W_+ + W_-) = \mathcal{A}(W_+) + \mathcal{A}(W_-) = 2B$. It follows that $dr \geq \text{rank}(B) = \text{rank}(2B) \geq 2(1 - \epsilon)dr \implies 1 \geq 2(1 - \epsilon) \iff \epsilon \geq 1/2$, which is a contradiction. ■

The combination of Lemmas 4.1 and 4.2 suggests that by using an (ϵ, d, r_0, n) -rank expander with $\epsilon < 1/2$ as a measurement operator, one can guarantee that every PSD matrix X_0 of rank at most $r_0/2$ is the unique PSD solution to the measurements $\mathcal{A}(X_0)$. Therefore, every program (e.g. SDP) that can identify a point in the feasible set $\{X \mid X \succeq 0, \mathcal{A}(X) = \mathcal{A}(X_0)\}$ successfully returns X_0 . Quite Interestingly, for the case of high quality expanders (small ϵ), we propose an alternative method for identifying one feasible point, which by the token of the aforementioned uniqueness argument, can successfully recover low rank PSD matrices. The algorithm is based on using the expansion property of the suggested linear operator, and the fact that the original matrix is low rank. The key point is that for rank expanders, the original under-determined system can be equivalently transformed to an over-determined linear system, after a few simple processing steps, that mostly involves taking SVD and finding null spaces. The new system of linear equation can then be solved by matrix inversion. This routine is described in Algorithm 1, and is in fact the generalization of a positive sparse vector recovery algorithm elaborated in [9]. The next theorem provides a guarantee for the success of the proposed algorithm.

Theorem 4.1. PSD Recovery. *If the operator $\mathcal{A}(X) = \sum_{i=1}^d A_i X A_i^T$ is a (ϵ, d, r_0, n) -rank expander with $\epsilon < 1/2$, then for every $k \leq r_0(1 - \epsilon)$, every PSD matrix X of rank k can be perfectly recovered from $\mathcal{A}(X)$ using Algorithm 1.*

In order to prove Theorem 4.1, we first prove the following lemma.

Lemma 4.3. *Suppose that the operator $\mathcal{A}(X) = \sum_{i=1}^d A_i X A_i^T$ is an unbalanced (ϵ, d, r_0, n) -rank expander with $k/(1 - \epsilon) < m < n$, where A_i 's are $m \times n$. Also suppose that $X \in \mathbb{S}^n$ has rank $k \leq r_0(1 - \epsilon)/2$. Further, let \mathcal{S} be the linear space of all $n \times 1$ vectors u such that $\text{Span}_C\{\mathcal{A}(uu^T)\} \in \text{Span}_C\{\mathcal{A}(X)\}$. Then $\dim(\mathcal{S}) < k/(1 - \epsilon)$.*

Proof. If $\dim(\mathcal{S}) \geq k/(1 - \epsilon)$, then we can find an orthonormal matrix U of size $n \times k/(1 - \epsilon)$ such that all of its columns are in \mathcal{S} . Therefore, by definition $\text{Span}_C\{\mathcal{A}(UU^T)\} \in \text{Span}_C\{\mathcal{A}(X)\}$, and thus $\text{rank}(\mathcal{A}(UU^T)) \leq \text{rank}(\mathcal{A}(X)) \leq k \cdot d$. However, since $k/(1 - \epsilon) < r_0$, from the definition of rank-expander, we must have $\text{rank}(\mathcal{A}(UU^T)) > d(1 - \epsilon)\text{rank}(UU^T) = k \cdot d$, which is a contradiction. ■

Proof of Theorem 4.1. Let \mathcal{S} be as defined in Lemma 4.3 with $\dim(\mathcal{S}) = r$, and $Q^{n \times r}$ be a basis for \mathcal{S} , and let $X = X^{1/2}X^{T/2}$. It is easy to check that the columns of $X^{1/2}$ are all in \mathcal{S} , and thus $X^{1/2} = QV^{1/2}$, for some unknown $V^{1/2}$ of size $r \times k$. Therefore, we can write:

$$Y = \mathcal{A}(\mathbf{X}) = \sum_{i=1}^d A_i Q V^{1/2} V^{T/2} Q^T A_i^T \quad (11)$$

or equivalently

$$\text{vec}(Y) = \left(\sum_{i=1}^d B_i \otimes B_i^T \right) \text{vec}(V) \quad (12)$$

where $B_i = A_i Q$. (12) is a linear system of equations with m^2 equations and r^2 unknowns. Moreover, from Lemma 4.3, we know that $r < k/(1 - \epsilon) < m$. In addition, from Lemma 4.2, the solution to the system of linear equations in (12) is unique and must be $V = X$, since otherwise $X - V$ is in the null space of $\mathcal{A}(\cdot)$, and has at most $r < r_0/2$ negative eigenvalues. Therefore, (12) is an over-determined linear system and can be solved by matrix inversion as given in line 11 of Algorithm 1. ■

5 Extension to Hermitians

In this section, we briefly extend the results of the previous sections to the case of Hermitian matrices. Specifically, we prove the existence of expanders for Hermitian matrices (rather than only PSD matrices, which was the case discussed previously), and then state a theorem certifying the success of Algorithm 1 for these classes of low rank matrices.

5.1 Expansion

Lemma 5.1. Expansion for Hermitians. *Assume \mathcal{A} is a (ϵ, d, r_0, n) expander and let X be Hermitian with $\text{rank}(X) \leq r_0$. Then, $\text{rank}(\mathcal{A}(X)) \geq (1 - 4\epsilon)d \cdot \text{rank}(X)$.*

The formal proof of the above lemma is skipped, but it is mostly based on some dimension counting arguments and the fact that $\text{rank}(A + B) = \text{rank}(A) + \text{rank}(B)$ is equivalent to $\text{Span}_C(A) \cap \text{Span}_C(B) = \emptyset$ and $\text{Span}_R(A) \cap \text{Span}_R(B) = \emptyset$ (see e.g. [19]).

5.2 Recovery

Assume $X \in \mathbb{S}$ with $\text{rank}(X) = r \leq r_0$. With Gaussian measurements, $\text{rank}(\mathcal{A}(X)) = rd$ almost surely, due to (5) and the result of [19]. Notice that if the eigenvalue decomposition of X is $X = \sum_{i=1}^r \lambda_i u_i u_i^T$, then $\mathcal{A}(X) = \sum_{i=1}^r \lambda_i \mathcal{A}(u_i u_i^T)$, and using $\text{rank}(\mathcal{A}(X)) = r \times d$ and $\text{rank}(\mathcal{A}(u_i u_i^T)) \leq d$, it follows that

$\text{Span}_C(\mathcal{A}(u_i u_i^T)) \subset \text{Span}_C(\mathcal{A}(X))$ for all $i \leq r$. Consequently, similar to the case of PSD matrices, we need to find the space of u such that $\text{Span}_C(\mathcal{A}(u u^T)) \subset \text{Span}_C(\mathcal{A}(X))$. All technical steps follow identically and similar to the PDS case, and we can assert that sufficiently low rank Hermitian solutions to (12) are unique, as follows.

Theorem 5.1. Hermitian Recovery. *Let $X_0 \in \mathbb{S}$ with $\text{rank}(X_0) \leq r_0(1 - \epsilon)$. Suppose $\mathcal{A}(\cdot)$ is as described in Theorem 3.1. With probability at least $1 - \exp(-\Omega(n))$, X_0 can be perfectly recovered from $\mathcal{A}(X_0)$ by Algorithm 1.*

Note that although the same algorithm works for both cases of PSD and Hermitian matrices, there is a significant difference. In Theorem 4.1, “all” $X \in \mathbb{S}_+^n$ with sufficiently rank low are recoverable, whereas a similar fact is true for “almost all” Hermitian matrices of low rank. These notions are often distinguished in the literature by the terms “strong” recovery and “weak” recovery, respectively. Furthermore, note that for the case of PSD recovery, one can alternatively use convex optimization to find the unique PSD solution. However for the recovery of Hermitians, we do not know of any other method but our proposed Algorithm 1.

6 Simulation Results

Numerical simulations were performed to verify the validity of Alg. 1. We used $n = 50$, and linear measurement operators in the form of (4) with $d = 2, 3, 4$, and two types of distributions for G_i ’s: 1) i.i.d Gaussian matrices, and 2) sparse matrices where every row of each G_i has exactly one 1 in a random location. We did not explicitly prove that the sparse constructions expanders. However, such low density structures are of high practical interests. The resulting curves of successful recovery thresholds are given in Figure 1. In all of our simulations for Gaussian matrices, the transition between successful recovery and failure was very sharp, i.e. either failed all times or succeeded, depending on the number of measurements and the rank of X . Figure 1 illustrates the empirical transition phase or the recovery threshold of Alg. 1. The corresponding curves for sparse matrices are also shown in Figure 1. The transition for sparse matrices was not as sharp as Gaussians. On the same curves, the performance of the standard trace minimization with nonnegativity (PSD) constraint is also displayed. Observe that the performance of the proposed algorithm is very comparable to the convex relaxation method. In addition, the curves for sparse measurements collapse into the same curves as the dense measurements for sufficiently large m . In practice, Alg. 1 is extremely faster than NNM. To give an example, for simulations in MATLAB on a 2.4 GHz Intel Core i5 with 4 GB RAM, and for the case of $d = 2, m = 39, k = 10$ with Gaussian G_i ’s, each reconstruction took on average 34s for the NNM, while it was only about 0.05s for Alg. 1. For sparse structures, it took about 1.8s for NNM and 0.05s for Alg.1. NNM was solved via the SEDUMI toolbox.

References

- [1] M. Fazel, “Matrix Rank Minimization with Applications,” Ph.D. thesis 2002.




Figure 1: Empirical recovery thresholds of Alg. 1 and NNM for 50×50 matrices with linear operators. Measurements are $m \times m$ and k is the rank.

- [2] B. Recht, M. Fazel, P.A. Parrilo, “Guaranteed Minimum-Rank Solutions of Linear Matrix Equations via Nuclear Norm Minimization,” SIAM Review, 52(3),471-501 (2010).
- [3] B. Recht, W. Xu and B. Hassibi, “Necessary and Sufficient Conditions for Success of the Nuclear Norm Heuristic for Rank Minimization,” CDC 2008.
- [4] E. J. Candes and B. Recht, “Exact Matrix Completion Via Convex Optimization,” Found. of Comput. Math., 9, 717-772.
- [5] E. J. Candes and T. Tao, “The power of Convex Relaxation: Near-optimal Matrix Completion,” IEEE Trans. Inform. Theory, 56(5), 2053-2080.
- [6] W. Xu, B Hassibi, “Efficient Compressive Sensing with Deterministic Guarantees Using Expander Graphs,” ITW 2007.
- [7] R. Berinde, P. Indyk, “Sparse Recovery Using Sparse Matrices”, Tech. Rep., 2008.
- [8] R. Berinde, P. Indyk, “Sequential Sparse Matching Pursuit,” Allerton 2009.
- [9] M. A. Khajehnejad, W. Xu, A. G. Dimakis, B. Hassibi, “Sparse Recovery of Positive Signals with Minimal Expansion”, IEEE Tran. on Signal Proc., 59(1),196-208.
- [10] R. Bhatia, “Matrix Analysis (Graduate Texts in Mathematics)”. Springer, 1996.
- [11] F. Parvaresh and B. Hassibi, “Explicit Measurements with Almost Optimal Thresholds for Compressed Sensing,” ICASPP 2008.
- [12] <http://dsp.rice.edu/cs>.

- [13] S.J. Szarek, “Metric Entropy of Homogeneous Spaces. In: Quantum Probability”, arXiv:math/9701213v1.
- [14] M. Akcakaya, V. Tarokh, “A Frame Construction and A Universal Distortion Bound for Sparse Representations,” IEEE Tran. Sig. Proc. 56(6), 2008.
- [15] S. Oymak and B. Hassibi, “New Null Space Results and Recovery Thresholds for Matrix Rank Minimization,” arXiv:1011.6326v1.
- [16] V. Chandrasekaran, B. Recht, P. Parrilo and A. Willsky, “The Convex Geometry of Linear Inverse Problems,” arXiv:1012.0621v1.
- [17] M. Stojnic, “Various Thresholds for ℓ_1 -Optimization in Compressed Sensing,” , arXiv:0907.3666v1.
- [18] W. Xu, “On the Uniqueness of Positive Semidefinite Matrix Solution under Compressed Observations,” ISIT 2010.
- [19] D. Callan, “When Is ‘Rank’ Additive?” The College Math. J. , 29(2),145-147, 1998.
- [20] M. Ledoux, M. Talagrand, “Probability in Banach Spaces,” Springer-Verlag, 1991.