
Matrix Completion from Noisy Entries

Raghu­nan­dan H. Keshavan*, Andrea Montanari*[†] and Sewoong Oh*

Abstract

Given a matrix M of low-rank, we consider the problem of reconstructing it from noisy observations of a small, random subset of its entries. The problem arises in a variety of applications, from collaborative filtering (the ‘Netflix problem’) to structure-from-motion and positioning. We study a low complexity algorithm introduced in [1], based on a combination of spectral techniques and manifold optimization, that we call here OPTSPACE. We prove performance guarantees that are order-optimal in a number of circumstances.

1 Introduction

Spectral techniques are an authentic workhorse in machine learning, statistics, numerical analysis, and signal processing. Given a matrix M , its largest singular values –and the associated singular vectors– ‘explain’ the most significant correlations in the underlying data source. A low-rank approximation of M can further be used for low-complexity implementations of a number of linear algebra algorithms [2].

In many practical circumstances we have access only to a sparse subset of the entries of an $m \times n$ matrix M . It has recently been discovered that, if the matrix M has rank r , and unless it is too ‘structured’, a small random subset of its entries allow to reconstruct it *exactly*. This result was first proved by Candés and Recht [3] by analyzing a convex relaxation introduced by Fazel [4]. A tighter analysis of the same convex relaxation was carried out in [5]. A number of iterative schemes to solve the convex optimization problem appeared soon thereafter [6, 7, 8] (also see [9] for a generalization).

In an alternative line of work, the authors of [1] attacked the same problem using a combination of spectral techniques and manifold optimization: we will refer to their algorithm as OPTSPACE. OPTSPACE is intrinsically of low complexity, the most complex operation being computing r singular values and the corresponding singular vectors of a sparse $m \times n$ matrix. The performance guarantees proved in [1] are comparable with the information theoretic lower bound: roughly $nr \max\{r, \log n\}$ random entries are needed to reconstruct M exactly (here we assume m of order n). A related approach was also developed in [10], although without performance guarantees for matrix completion.

The above results crucially rely on the assumption that M is *exactly* a rank r matrix. For many applications of interest, this assumption is unrealistic and it is therefore important to investigate their robustness. Can the above approaches be generalized when the underlying data is ‘well approximated’ by a rank r matrix? This question was addressed in [11] within the convex relaxation approach of [3]. The present paper proves a similar robustness result for OPTSPACE. Remarkably the guarantees we obtain are order-optimal in a variety of circumstances, and improve over the analogous results of [11].

*Department of Electrical Engineering, Stanford University

[†]Departments of Statistics, Stanford University

1.1 Model definition

Let M be an $m \times n$ matrix of rank r , that is

$$M = U\Sigma V^T. \quad (1)$$

where U has dimensions $m \times r$, V has dimensions $n \times r$, and Σ is a diagonal $r \times r$ matrix. We assume that each entry of M is perturbed, thus producing an ‘approximately’ low-rank matrix N , with

$$N_{ij} = M_{ij} + Z_{ij}, \quad (2)$$

where the matrix Z will be assumed to be ‘small’ in an appropriate sense.

Out of the $m \times n$ entries of N , a subset $E \subseteq [m] \times [n]$ is revealed. We let N^E be the $m \times n$ matrix that contains the revealed entries of N , and is filled with 0’s in the other positions

$$N_{ij}^E = \begin{cases} N_{ij} & \text{if } (i, j) \in E, \\ 0 & \text{otherwise.} \end{cases} \quad (3)$$

The set E will be uniformly random given its size $|E|$.

1.2 Algorithm

For the reader’s convenience, we recall the algorithm introduced in [1], which we will analyze here. The basic idea is to minimize the cost function $F(X, Y)$, defined by

$$F(X, Y) \equiv \min_{S \in \mathbb{R}^{r \times r}} \mathcal{F}(X, Y, S), \quad (4)$$

$$\mathcal{F}(X, Y, S) \equiv \frac{1}{2} \sum_{(i,j) \in E} (N_{ij} - (XSY^T)_{ij})^2. \quad (5)$$

Here $X \in \mathbb{R}^{n \times r}$, $Y \in \mathbb{R}^{m \times r}$ are orthogonal matrices, normalized by $X^T X = m\mathbf{1}$, $Y^T Y = n\mathbf{1}$.

Minimizing $F(X, Y)$ is an *a priori* difficult task, since F is a non-convex function. The key insight is that the singular value decomposition (SVD) of N^E provides an excellent initial guess, and that the minimum can be found with high probability by standard gradient descent after this initialization. Two caveats must be added to this description: (1) In general the matrix N^E must be ‘trimmed’ to eliminate over-represented rows and columns; (2) For technical reasons, we consider a slightly modified cost function to be denoted by $\tilde{F}(X, Y)$.

OPTSPACE(matrix N^E)

- 1: Trim N^E , and let \tilde{N}^E be the output;
 - 2: Compute the rank- r projection of \tilde{N}^E , $\mathsf{T}_r(\tilde{N}^E) = X_0 S_0 Y_0^T$;
 - 3: Minimize $\tilde{F}(X, Y)$ through gradient descent, with initial condition (X_0, Y_0) .
-

We may note here that the rank of the matrix M , if not known, can be reliably estimated from \tilde{N}^E . We refer to the journal version of this paper for further details.

The various steps of the above algorithm are defined as follows.

Trimming. We say that a row is ‘over-represented’ if it contains more than $2|E|/m$ revealed entries (i.e. more than twice the average number of revealed entries). Analogously, a column is over-represented if it contains more than $2|E|/n$ revealed entries. The trimmed matrix \tilde{N}^E is obtained from N^E by setting to 0 over-represented rows and columns. \tilde{M}^E and \tilde{Z}^E are defined similarly. Hence, $\tilde{N}^E = \tilde{M}^E + \tilde{Z}^E$.

Rank- r projection. Let

$$\tilde{N}^E = \sum_{i=1}^{\min(m,n)} \sigma_i x_i y_i^T, \quad (6)$$

be the singular value decomposition of \tilde{N}^E , with singular vectors $\sigma_1 \geq \sigma_2 \geq \dots$. We then define

$$\mathsf{T}_r(\tilde{N}^E) = \frac{mn}{|E|} \sum_{i=1}^r \sigma_i x_i y_i^T. \quad (7)$$

Apart from an overall normalization, $\mathsf{T}_r(\tilde{N}^E)$ is the best rank- r approximation to \tilde{N}^E in Frobenius norm.

Minimization. The modified cost function \tilde{F} is defined as

$$\tilde{F}(X, Y) = F(X, Y) + \rho G(X, Y) \quad (8)$$

$$\equiv F(X, Y) + \rho \sum_{i=1}^m G_1 \left(\frac{\|X^{(i)}\|^2}{3\mu_0 r} \right) + \rho \sum_{j=1}^n G_1 \left(\frac{\|Y^{(j)}\|^2}{3\mu_0 r} \right), \quad (9)$$

where $X^{(i)}$ denotes the i -th row of X , and $Y^{(j)}$ the j -th row of Y . See Section 1.3 below for the definition of μ_0 . The function $G_1 : \mathbb{R}^+ \rightarrow \mathbb{R}$ is such that $G_1(z) = 0$ if $z \leq 1$ and $G_1(z) = e^{(z-1)^2} - 1$ otherwise. Further, we can choose $\rho = \Theta(n\epsilon)$.

Let us stress that the regularization term is mainly introduced for our proof technique to work (and a broad family of functions G_1 would work as well). In numerical experiments we did not find any performance loss in setting $\rho = 0$.

One important feature of OPTSPACE is that $F(X, Y)$ and $\tilde{F}(X, Y)$ are regarded as functions of the r -dimensional subspaces of \mathbb{R}^m and \mathbb{R}^n generated (respectively) by the columns of X and Y . This interpretation is justified by the fact that $F(X, Y) = F(XA, YB)$ for any two orthogonal matrices $A, B \in \mathbb{R}^{r \times r}$ (the same property holds for \tilde{F}). The set of r dimensional subspaces of \mathbb{R}^m is a differentiable Riemannian manifold $\mathsf{G}(m, r)$ (the Grassman manifold). The gradient descent algorithm is applied to the function $\tilde{F} : \mathsf{M}(m, n) \equiv \mathsf{G}(m, r) \times \mathsf{G}(n, r) \rightarrow \mathbb{R}$. For further details on optimization by gradient descent on matrix manifolds we refer to [12, 13].

1.3 Main results

Our first result shows that, in great generality, the rank- r projection of \tilde{N}^E provides a reasonable approximation of M . Throughout this paper, without loss of generality, we assume $\alpha \equiv m/n \geq 1$.

Theorem 1.1. *Let $N = M + Z$, where M has rank r and $|M_{ij}| \leq M_{\max}$ for all $(i, j) \in [m] \times [n]$, and assume that the subset of revealed entries $E \subseteq [m] \times [n]$ is uniformly random with size $|E|$. Then there exists numerical constants C and C' such that*

$$\frac{1}{\sqrt{mn}} \|M - \mathsf{T}_r(\tilde{N}^E)\|_F \leq C M_{\max} \left(\frac{nr\alpha^{3/2}}{|E|} \right)^{1/2} + C' \frac{n\sqrt{r\alpha}}{|E|} \|\tilde{Z}^E\|_2, \quad (10)$$

with probability larger than $1 - 1/n^3$.

Projection onto rank- r matrices through SVD is pretty standard (although trimming is crucial for achieving the above guarantee). The key point here is that a much better approximation is obtained by minimizing the cost $\tilde{F}(X, Y)$ (step 3 in the pseudocode above), provided M satisfies an appropriate incoherence condition. Let $M = U\Sigma V^T$ be a low rank matrix, and assume, without loss of generality, $U^T U = m\mathbf{1}$ and $V^T V = n\mathbf{1}$. We say that M is (μ_0, μ_1) -incoherent if the following conditions hold.

A1. For all $i \in [m], j \in [n]$ we have, $\sum_{k=1}^r U_{ik}^2 \leq \mu_0 r, \sum_{k=1}^r V_{jk}^2 \leq \mu_0 r$.

A2. There exists μ_1 such that $|\sum_{k=1}^r U_{ik}(\Sigma_k/\Sigma_1)V_{jk}| \leq \mu_1 r^{1/2}$.

Theorem 1.2. *Let $N = M + Z$, where M is a (μ_0, μ_1) -incoherent matrix of rank r , and assume that the subset of revealed entries $E \subseteq [m] \times [n]$ is uniformly random with size $|E|$. Further, let $\Sigma_{\min} = \Sigma_1 \leq \dots \leq \Sigma_r = \Sigma_{\max}$ with $\Sigma_{\max}/\Sigma_{\min} \equiv \kappa$. Let \tilde{M} be the output of OPTSPACE on input N^E . Then there exists numerical constants C and C' such that if*

$$|E| \geq Cn\sqrt{\alpha}\kappa^2 \max \{ \mu_0 r \sqrt{\alpha} \log n; \mu_0^2 r^2 \alpha \kappa^4; \mu_1^2 r^2 \alpha \kappa^4 \}, \quad (11)$$

then, with probability at least $1 - 1/n^3$,

$$\frac{1}{\sqrt{mn}} \|\widehat{M} - M\|_F \leq C' \kappa^2 \frac{n\sqrt{\alpha r}}{|E|} \|Z^E\|_2. \quad (12)$$

provided that the right-hand side is smaller than Σ_{\min} .

Apart from capturing the effect of additive noise, these two theorems improve over the work of [1] even in the noiseless case. Indeed they provide quantitative bounds in finite dimensions, while the results of [1] were only asymptotic.

1.4 Noise models

In order to make sense of the above results, it is convenient to consider a couple of simple models for the noise matrix Z :

Independent entries model. We assume that Z 's entries are independent random variables, with zero mean $\mathbb{E}\{Z_{ij}\} = 0$ and sub-gaussian tails. The latter means that

$$\mathbb{P}\{|Z_{ij}| \geq x\} \leq 2e^{-\frac{x^2}{2\sigma^2}}, \quad (13)$$

for some bounded constant σ^2 .

Worst case model. In this model Z is arbitrary, but we have an uniform bound on the size of its entries: $|Z_{ij}| \leq Z_{\max}$.

The basic parameter entering our main results is the operator norm of \widetilde{Z}^E , which is bounded as follows.

Theorem 1.3. *If Z is a random matrix drawn according to the independent entries model, then there is a constant C such that,*

$$\|\widetilde{Z}^E\|_2 \leq C\sigma \left(\frac{\sqrt{\alpha}|E| \log |E|}{n} \right)^{1/2}, \quad (14)$$

with probability at least $1 - 1/n^3$.

If Z is a matrix from the worst case model, then

$$\|\widetilde{Z}^E\|_2 \leq \frac{2|E|}{n\sqrt{\alpha}} Z_{\max}, \quad (15)$$

for any realization of E .

Note that for $|E| = \Omega(n \log n)$, no row or column is over-represented with high probability. It follows that in the regime of $|E|$ for which the conditions of Theorem 1.2 are satisfied, we have $Z^E = \widetilde{Z}^E$. Then, among the other things, this result implies that for the independent entries model the right-hand side of our error estimate, Eq. (12), is with high probability smaller than Σ_{\min} , if $|E| \geq Cr\alpha^{3/2}n \log n \kappa^4 (\sigma/\Sigma_{\min})^2$. For the worst case model, the same statement is true if $Z_{\max} \leq \Sigma_{\min}/C\sqrt{r}\kappa^2$.

Due to space constraints, the proof of Theorem 1.3 will be given in the journal version of this paper.

1.5 Comparison with related work

Let us begin by mentioning that a statement analogous to our preliminary Theorem 1.1 was proved in [14]. Our result however applies to any number of revealed entries, while the one of [14] requires $|E| \geq (8 \log n)^4 n$ (which for $n \leq 5 \cdot 10^8$ is larger than n^2).

As for Theorem 1.2, we will mainly compare our algorithm with the convex relaxation approach recently analyzed in [11]. Our basic setting is indeed the same, while the algorithms are rather different.

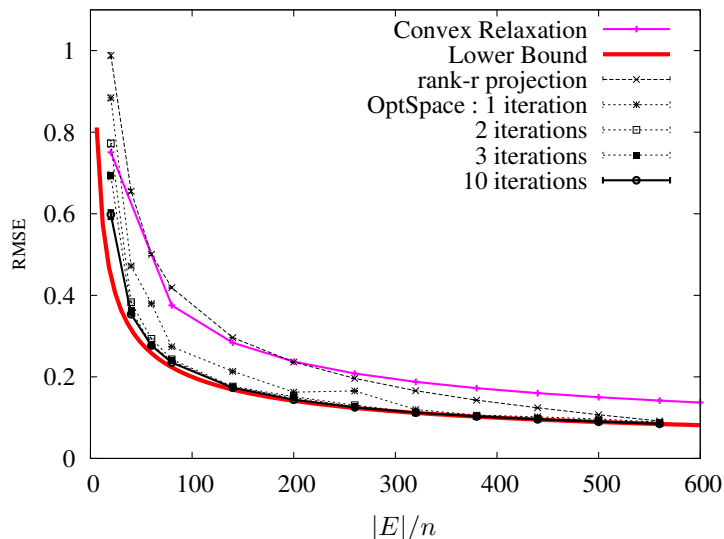


Figure 1: Root mean square error achieved by OPTSPACE for reconstructing a random rank-2 matrix, as a function of the number of observed entries $|E|$, and of the number of line minimizations. The performance of nuclear norm minimization and an information theory lower bound are also shown.

Figure 1 compares the average root mean square error for the two algorithms as a function of $|E|$. Here M is a random rank $r = 2$ matrix of dimension $m = n = 600$, generated by letting $M = \tilde{U}\tilde{V}^T$ with $\tilde{U}_{ij}, \tilde{V}_{ij}$ i.i.d. $N(0, 20/\sqrt{n})$. The noise is distributed according to the independent entries model with $Z_{ij} \sim N(0, 1)$. This example is taken from [11] Figure 2, from which we took the data for the convex relaxation approach, as well as the information theory lower bound. After one iteration, OPTSPACE has a smaller root mean square error than [11], and in about 10 iterations it becomes indistinguishable from the information theory lower bound.

Next let us compare our main result with the performance guarantee in [11], Theorem 7. Let us stress that we require some bound on the condition number κ , while the analysis of [11, 5] requires a stronger incoherence assumption. As far as the error bound is concerned, [11] proved

$$\frac{1}{\sqrt{mn}} \|\widehat{M} - M\|_F \leq 7 \sqrt{\frac{n}{|E|}} \|Z^E\|_F + \frac{2}{n\sqrt{\alpha}} \|Z^E\|_F. \quad (16)$$

(The constant in front of the first term is in fact slightly smaller than 7 in [11], but in any case larger than $4\sqrt{2}$).

Theorem 1.2 improves over this result in several respects: (1) We do not have the second term on the right hand side of (16), that actually increases with the number of observed entries; (2) Our error decreases as $n/|E|$ rather than $(n/|E|)^{1/2}$; (3) The noise enters Theorem 1.2 through the operator norm $\|Z^E\|_2$ instead of its Frobenius norm $\|Z^E\|_F \geq \|Z^E\|_2$. For E uniformly random, one expects $\|Z^E\|_F$ to be roughly of order $\|Z^E\|_2\sqrt{n}$. For instance, within the independent entries model with bounded variance σ , $\|Z^E\|_F = \Theta(\sqrt{|E|})$ while $\|Z^E\|_2$ is of order $\sqrt{|E|/n}$ (up to logarithmic terms).

2 Some notations

The matrix M to be reconstructed takes the form (1) where $U \in \mathbb{R}^{m \times r}$, $V \in \mathbb{R}^{n \times r}$. We write $U = [u_1, u_2, \dots, u_r]$ and $V = [v_1, v_2, \dots, v_r]$ for the columns of the two factors, with $\|u_i\| = \sqrt{m}$, $\|v_i\| = \sqrt{n}$, and $u_i^T u_j = 0$, $v_i^T v_j = 0$ for $i \neq j$ (there is no loss of generality in this, since normalizations can be absorbed by redefining Σ).

We shall write $\Sigma = \text{diag}(\Sigma_1, \dots, \Sigma_r)$ with $\Sigma_1 \geq \Sigma_2 \geq \dots \geq \Sigma_r > 0$. The maximum and minimum singular values will also be denoted by $\Sigma_{\max} = \Sigma_1$ and $\Sigma_{\min} = \Sigma_r$. Further, the maximum size of an entry of M is $M_{\max} \equiv \max_{ij} |M_{ij}|$.

Probability is taken with respect to the uniformly random subset $E \subseteq [m] \times [n]$ given $|E|$ and (eventually) the noise matrix Z . Define $\epsilon \equiv |E|/\sqrt{mn}$. In the case when $m = n$, ϵ corresponds to the average number of revealed entries per row or column. Then it is convenient to work with a model in which each entry is revealed independently with probability ϵ/\sqrt{mn} . Since, with high probability $|E| \in [\epsilon\sqrt{\alpha}n - A\sqrt{n\log n}, \epsilon\sqrt{\alpha}n + A\sqrt{n\log n}]$, any guarantee on the algorithm performances that holds within one model, holds within the other model as well if we allow for a vanishing shift in ϵ . We will use C, C' etc. to denote universal numerical constants.

Given a vector $x \in \mathbb{R}^n$, $\|x\|$ will denote its Euclidean norm. For a matrix $X \in \mathbb{R}^{n \times n'}$, $\|X\|_F$ is its Frobenius norm, and $\|X\|_2$ its operator norm (i.e. $\|X\|_2 = \sup_{u \neq 0} \|Xu\|/\|u\|$). The standard scalar product between vectors or matrices will sometimes be indicated by $\langle x, y \rangle$ or $\langle X, Y \rangle$, respectively. Finally, we use the standard combinatorics notation $[N] = \{1, 2, \dots, N\}$ to denote the set of first N integers.

3 Proof of Theorem 1.1

As explained in the introduction, the crucial idea is to consider the singular value decomposition of the trimmed matrix \tilde{N}^E instead of the original matrix N^E . Apart from a trivial rescaling, these singular values are close to the ones of the original matrix M .

Lemma 3.1. *There exists a numerical constant C such that, with probability greater than $1 - 1/n^3$,*

$$\left| \frac{\sigma_q}{\epsilon} - \Sigma_q \right| \leq CM_{\max} \sqrt{\frac{\alpha}{\epsilon}} + \frac{1}{\epsilon} \|\tilde{Z}^E\|_2, \quad (17)$$

where it is understood that $\Sigma_q = 0$ for $q > r$.

Proof. For any matrix A , let $\sigma_q(A)$ denote the q th singular value of A . Then, $\sigma_q(A+B) \leq \sigma_q(A) + \sigma_1(B)$, whence

$$\begin{aligned} \left| \frac{\sigma_q}{\epsilon} - \Sigma_q \right| &\leq \left| \sigma_q(\tilde{M}^E)/\epsilon - \Sigma_q \right| + \sigma_1(\tilde{Z}^E)/\epsilon \\ &\leq CM_{\max} \sqrt{\frac{\alpha}{\epsilon}} + \frac{1}{\epsilon} \|\tilde{Z}^E\|_2, \end{aligned}$$

where the second inequality follows from the following Lemma as shown in [1].

Lemma 3.2 (Keshavan, Montanari, Oh, 2009 [1]). *There exists a numerical constant C such that, with probability larger than $1 - 1/n^3$,*

$$\frac{1}{\sqrt{mn}} \left\| M - \frac{\sqrt{mn}}{\epsilon} \tilde{M}^E \right\|_2 \leq CM_{\max} \sqrt{\frac{\alpha}{\epsilon}}. \quad (18)$$

□

We will now prove Theorem 1.1.

Proof. (Theorem 1.1) For any matrix A of rank at most $2r$, $\|A\|_F \leq \sqrt{2r}\|A\|_2$, whence

$$\begin{aligned} \frac{1}{\sqrt{mn}} \|M - \mathsf{T}_r(\tilde{N}^E)\|_F &\leq \frac{\sqrt{2r}}{\sqrt{mn}} \left\| M - \frac{\sqrt{mn}}{\epsilon} \left(\tilde{N}^E - \sum_{i \geq r+1} \sigma_i x_i y_i^T \right) \right\|_2 \\ &\leq \frac{\sqrt{2r}}{\sqrt{mn}} \left(\left\| M - \frac{\sqrt{mn}}{\epsilon} \tilde{M}^E \right\|_2 + \frac{\sqrt{mn}}{\epsilon} \|\tilde{Z}^E\|_2 + \frac{\sqrt{mn}}{\epsilon} \sigma_{r+1} \right) \\ &\leq 2CM_{\max} \sqrt{2\alpha r/\epsilon} + (2\sqrt{2r}/\epsilon) \|\tilde{Z}^E\|_2 \\ &\leq C' M_{\max} \left(\frac{nr\alpha^{3/2}}{|E|} \right)^{1/2} + 2\sqrt{2} \left(\frac{n\sqrt{r\alpha}}{|E|} \right) \|\tilde{Z}^E\|_2. \end{aligned}$$

This proves our claim. □

4 Proof of Theorem 1.2

Recall that the cost function is defined over the Riemannian manifold $M(m, n) \equiv G(m, r) \times G(n, r)$. The proof of Theorem 1.2 consists in controlling the behavior of F in a neighborhood of $\mathbf{u} = (U, V)$ (the point corresponding to the matrix M to be reconstructed). Throughout the proof we let $\mathcal{K}(\mu)$ be the set of matrix couples $(X, Y) \in \mathbb{R}^{m \times r} \times \mathbb{R}^{n \times r}$ such that $\|X^{(i)}\|^2 \leq \mu r$, $\|Y^{(j)}\|^2 \leq \mu r$ for all i, j

4.1 Preliminary remarks and definitions

Given $\mathbf{x}_1 = (X_1, Y_1)$ and $\mathbf{x}_2 = (X_2, Y_2) \in M(m, n)$, two points on this manifold, their distance is defined as $d(\mathbf{x}_1, \mathbf{x}_2) = \sqrt{d(X_1, X_2)^2 + d(Y_1, Y_2)^2}$, where, letting $(\cos \theta_1, \dots, \cos \theta_r)$ be the singular values of $X_1^T X_2 / m$,

$$d(X_1, X_2) = \|\theta\|_2. \quad (19)$$

Given S achieving the minimum in Eq. (4), it is also convenient to introduce the notations

$$d_-(\mathbf{x}, \mathbf{u}) \equiv \sqrt{\Sigma_{\min}^2 d(\mathbf{x}, \mathbf{u})^2 + \|S - \Sigma\|_F^2}, \quad (20)$$

$$d_+(\mathbf{x}, \mathbf{u}) \equiv \sqrt{\Sigma_{\max}^2 d(\mathbf{x}, \mathbf{u})^2 + \|S - \Sigma\|_F^2}. \quad (21)$$

4.2 Auxiliary lemmas and proof of Theorem 1.2

The proof is based on the following two lemmas that generalize and sharpen analogous bounds in [1] (for proofs we refer to the journal version of this paper).

Lemma 4.1. *There exists numerical constants C_0, C_1, C_2 such that the following happens. Assume $\epsilon \geq C_0 \mu_0 r \sqrt{\alpha} \max\{\log n; \mu_0 r \sqrt{\alpha} (\Sigma_{\min} / \Sigma_{\max})^4\}$ and $\delta \leq \Sigma_{\min} / (C_0 \Sigma_{\max})$. Then,*

$$F(\mathbf{x}) - F(\mathbf{u}) \geq C_1 n \epsilon \sqrt{\alpha} d_-(\mathbf{x}, \mathbf{u})^2 - C_1 n \sqrt{r \alpha} \|Z^E\|_2 d_+(\mathbf{x}, \mathbf{u}), \quad (22)$$

$$F(\mathbf{x}) - F(\mathbf{u}) \leq C_2 n \epsilon \sqrt{\alpha} \Sigma_{\max}^2 d(\mathbf{x}, \mathbf{u})^2 + C_2 n \sqrt{r \alpha} \|Z^E\|_2 d_+(\mathbf{x}, \mathbf{u}), \quad (23)$$

for all $\mathbf{x} \in M(m, n) \cap \mathcal{K}(4\mu_0)$ such that $d(\mathbf{x}, \mathbf{u}) \leq \delta$, with probability at least $1 - 1/n^4$. Here $S \in \mathbb{R}^{r \times r}$ is the matrix realizing the minimum in Eq. (4).

Corollary 4.2. *There exist a constant C such that, under the hypotheses of Lemma 4.1*

$$\|S - \Sigma\|_F \leq C \Sigma_{\max} d(\mathbf{x}, \mathbf{u}) + C \frac{\sqrt{r}}{\epsilon} \|Z^E\|_2. \quad (24)$$

Further, for an appropriate choice of the constants in Lemma 4.1, we have

$$\sigma_{\max}(S) \leq 2 \Sigma_{\max} + C \frac{\sqrt{r}}{\epsilon} \|Z^E\|_2, \quad (25)$$

$$\sigma_{\min}(S) \geq \frac{1}{2} \Sigma_{\min} - C \frac{\sqrt{r}}{\epsilon} \|Z^E\|_2. \quad (26)$$

Lemma 4.3. *There exists numerical constants C_0, C_1, C_2 such that the following happens. Assume $\epsilon \geq C_0 \mu_0 r \sqrt{\alpha} (\Sigma_{\max} / \Sigma_{\min})^2 \max\{\log n; \mu_0 r \sqrt{\alpha} (\Sigma_{\max} / \Sigma_{\min})^4\}$ and $\delta \leq \Sigma_{\min} / (C_0 \Sigma_{\max})$. Then,*

$$\|\text{grad } \tilde{F}(\mathbf{x})\|^2 \geq C_1 n \epsilon^2 \Sigma_{\min}^4 \left[d(\mathbf{x}, \mathbf{u}) - C_2 \frac{\sqrt{r} \Sigma_{\max}}{\epsilon \Sigma_{\min}} \frac{\|Z^E\|_2}{\Sigma_{\min}} \right]_+^2, \quad (27)$$

for all $\mathbf{x} \in M(m, n) \cap \mathcal{K}(4\mu_0)$ such that $d(\mathbf{x}, \mathbf{u}) \leq \delta$, with probability at least $1 - 1/n^4$. (Here $[a]_+ \equiv \max(a, 0)$.)

We can now turn to the proof of our main theorem.

Proof. (Theorem 1.2). Let $\delta = \Sigma_{\min}/C_0\Sigma_{\max}$ with C_0 large enough so that the hypotheses of Lemmas 4.1 and 4.3 are verified.

Call $\{\mathbf{x}_k\}_{k \geq 0}$ the sequence of pairs $(X_k, Y_k) \in \mathcal{M}(m, n)$ generated by gradient descent. By assumption, the following is true with a large enough constant C :

$$\|Z^E\|_2 \leq \frac{\epsilon}{C\sqrt{r}} \left(\frac{\Sigma_{\min}}{\Sigma_{\max}} \right)^2 \Sigma_{\min}. \quad (28)$$

Further, by using Corollary 4.2 in Eqs. (22) and (23) we get

$$F(\mathbf{x}) - F(\mathbf{u}) \geq C_1 n \epsilon \sqrt{\alpha} \Sigma_{\min}^2 \{d(\mathbf{x}, \mathbf{u})^2 - \delta_{0,-}^2\}, \quad (29)$$

$$F(\mathbf{x}) - F(\mathbf{u}) \leq C_2 n \epsilon \sqrt{\alpha} \Sigma_{\max}^2 \{d(\mathbf{x}, \mathbf{u})^2 + \delta_{0,+}^2\}, \quad (30)$$

where

$$\delta_{0,-} \equiv C \frac{\sqrt{r}\Sigma_{\max}}{\epsilon\Sigma_{\min}} \frac{\|Z^E\|_2}{\Sigma_{\min}}, \quad \delta_{0,+} \equiv C \frac{\sqrt{r}\Sigma_{\max}}{\epsilon\Sigma_{\min}} \frac{\|Z^E\|_2}{\Sigma_{\max}}. \quad (31)$$

By Eq. (28), we can assume $\delta_{0,+} \leq \delta_{0,-} \leq \delta/10$.

For $\epsilon \geq C\alpha\mu_1^2 r^2 (\Sigma_{\max}/\Sigma_{\min})^4$ as per our assumptions, using Eq. (28) in Theorem 1.1, together with the bound $d(\mathbf{u}, \mathbf{x}_0) \leq \|M - X_0 S Y_0^T\|_{F/n\sqrt{\alpha}\Sigma_{\min}}$, we get

$$d(\mathbf{u}, \mathbf{x}_0) \leq \frac{\delta}{10}. \quad (32)$$

We make the following claims :

1. $\mathbf{x}_k \in \mathcal{K}(4\mu_0)$ for all k .

Indeed without loss of generality we can assume $\mathbf{x}_0 \in \mathcal{K}(3\mu_0)$ (because otherwise we can rescale those lines of X_0, Y_0 that violate the constraint). Therefore $\tilde{F}(\mathbf{x}_0) = F(\mathbf{x}_0) \leq 4C_2 n \epsilon \sqrt{\alpha} \Sigma_{\max}^2 \delta^2 / 100$. On the other hand $\tilde{F}(\mathbf{x}) \geq \rho(e^{1/9} - 1)$ for $\mathbf{x} \notin \mathcal{K}(4\mu_0)$. Since $\tilde{F}(\mathbf{x}_k)$ is a non-increasing sequence, the thesis follows provided we take $\rho \geq C_2 n \epsilon \sqrt{\alpha} \Sigma_{\min}^2$.

2. $d(\mathbf{x}_k, \mathbf{u}) \leq \delta/10$ for all k .

Assuming $\epsilon \geq C\alpha\mu_1^2 r^2 (\Sigma_{\max}/\Sigma_{\min})^6$, we have $d(\mathbf{x}_0, \mathbf{u})^2 \leq (\Sigma_{\min}^2/C'\Sigma_{\max}^2)(\delta/10)^2$. Also assuming Eq. (28) with large enough C we can show the following. For all \mathbf{x}_k such that $d(\mathbf{x}_k, \mathbf{u}) \in [\delta/10, \delta]$, we have $\tilde{F}(\mathbf{x}) \geq F(\mathbf{x}) \geq F(\mathbf{x}_0)$. This contradicts the monotonicity of $\tilde{F}(\mathbf{x})$, and thus proves the claim.

Since the cost function is twice differentiable, and because of the above, the sequence $\{\mathbf{x}_k\}$ converges to

$$\Omega = \{\mathbf{x} \in \mathcal{K}(4\mu_0) \cap \mathcal{M}(m, n) : d(\mathbf{x}, \mathbf{u}) \leq \delta, \text{grad } \tilde{F}(\mathbf{x}) = 0\}. \quad (33)$$

By Lemma 4.3 for any $\mathbf{x} \in \Omega$,

$$d(\mathbf{x}, \mathbf{u}) \leq C \frac{\sqrt{r}\Sigma_{\max}}{\epsilon\Sigma_{\min}} \frac{\|Z^E\|_2}{\Sigma_{\min}} \quad (34)$$

which implies the thesis using Corollary 4.2. \square

Acknowledgements

This work was partially supported by a Terman fellowship, the NSF CAREER award CCF-0743978 and the NSF grant DMS-0806211.

References

- [1] R. H. Keshavan, A. Montanari, and S. Oh. Matrix completion from a few entries. [arXiv:0901.3150](#), January 2009.
- [2] A. Frieze, R. Kannan, and S. Vempala. Fast monte-carlo algorithms for finding low-rank approximations. *J. ACM*, 51(6):1025–1041, 2004.
- [3] E. J. Candès and B. Recht. Exact matrix completion via convex optimization. [arXiv:0805.4471](#), 2008.
- [4] M. Fazel. *Matrix Rank Minimization with Applications*. PhD thesis, Stanford University, 2002.
- [5] E. J. Candès and T. Tao. The power of convex relaxation: Near-optimal matrix completion. [arXiv:0903.1476](#), 2009.
- [6] J-F Cai, E. J. Candès, and Z. Shen. A singular value thresholding algorithm for matrix completion. [arXiv:0810.3286](#), 2008.
- [7] S. Ma, D. Goldfarb, and L. Chen. Fixed point and Bregman iterative methods for matrix rank minimization. [arXiv:0905.1643](#), 2009.
- [8] K. Toh and S. Yun. An accelerated proximal gradient algorithm for nuclear norm regularized least squares problems. <http://www.math.nus.edu.sg/~matys>, 2009.
- [9] J. Wright, A. Ganesh, S. Rao, and Y. Ma. Robust principal component analysis: Exact recovery of corrupted low-rank matrices. [arXiv:0905.0233](#), 2009.
- [10] K. Lee and Y. Bresler. Admira: Atomic decomposition for minimum rank approximation. [arXiv:0905.0044](#), 2009.
- [11] E. J. Candès and Y. Plan. Matrix completion with noise. [arXiv:0903.3131](#), 2009.
- [12] A. Edelman, T. A. Arias, and S. T. Smith. The geometry of algorithms with orthogonality constraints. *SIAM J. Matr. Anal. Appl.*, 20:303–353, 1999.
- [13] P.-A. Absil, R. Mahony, and R. Sepulchrer. *Optimization Algorithms on Matrix Manifolds*. Princeton University Press, 2008.
- [14] D. Achlioptas and F. McSherry. Fast computation of low-rank matrix approximations. *J. ACM*, 54(2):9, 2007.