

Low-rank Matrix Completion with Noisy Observations: a Quantitative Comparison

Raghunandan H. Keshavan, Andrea Montanari[†] and Sewoong Oh
Electrical Engineering and [†]Statistics Department
Stanford University, Stanford, CA 94304
{raghuram,montanari,swoh}@stanford.edu

Abstract—We consider a problem of significant practical importance, namely, the reconstruction of a low-rank data matrix from a small subset of its entries. This problem appears in many areas such as collaborative filtering, computer vision and wireless sensor networks. In this paper, we focus on the matrix completion problem in the case when the observed samples are corrupted by noise. We compare the performance of three state-of-the-art matrix completion algorithms (OptSpace, ADMiRA and FPCA) on a single simulation platform and present numerical results. We show that in practice these efficient algorithms can be used to reconstruct real data matrices, as well as randomly generated matrices, accurately.

I. INTRODUCTION

We consider the problem of reconstructing an $m \times n$ low rank matrix M from a small set of observed entries possibly corrupted by noise. This problem is of considerable practical interest and has many applications. One example is collaborative filtering, where users submit rankings for small subsets of, say, movies, and the goal is to infer the preference of unrated movies for a recommendation system [1]. It is believed that the movie-rating matrix is approximately low-rank, since only a few factors contribute to a user’s preferences. Other examples of matrix completion include the problem of inferring 3-dimensional structure from motion [2] and triangulation from incomplete data of distances between wireless sensors [3].

A. Prior and related work

On the theoretical side, most recent work focuses on algorithms for exactly recovering the unknown low-rank matrix and providing an upper bound on the number of observed entries that guarantee successful recovery with high probability, when the observed set is drawn uniformly at random over all subsets of the same size. The main assumptions of this *matrix completion problem with exact observations* is that the matrix M to be recovered has rank $r \ll m, n$ and that the observed entries are known exactly. Adopting techniques from compressed sensing, Candès and Recht introduced a convex relaxation to the NP-hard problem which is to find a minimum rank matrix matching the observed entries [4]. They introduced the concept of incoherence property and proved that for a matrix M of rank r which has the incoherence property, solving the convex relaxation correctly recovers the unknown matrix, with high probability, if the number of observed entries $|E|$ satisfies, $|E| \geq Crn^{1.2} \log n$.

Recently [5] improved the bound to $|E| \geq Crn \max\{\log n, r\}$ with an extra condition that the matrix has bounded condition number, where the condition number of a matrix is defined as the ratio between the largest singular value and the smallest singular value of M . We introduced an efficient algorithm called OPTSPACE, based on spectral methods followed by a local manifold optimization. For a bounded rank $r = O(1)$, the performance bound of OPTSPACE is order optimal [5]. Candès and Tao proved a similar bound $|E| \geq Cnr(\log n)^6$ with a stronger assumption on the original matrix M , known as the *strong incoherence condition* but without any assumption on the condition number of the matrix M [6]. For any value of r , it is only suboptimal by a poly-logarithmic factor.

While most theoretical work focus on proving bounds for the exact matrix completion problem, a more interesting and practical problem is when the matrix M is only approximately low rank or when the observation is corrupted by noise. The main focus of this *matrix completion with noisy observations* is to design an algorithm to find an $m \times n$ low-rank matrix \widehat{M} that best approximates the original matrix M and provide a bound on the root mean squared error (RMSE) given by,

$$\text{RMSE} = \frac{1}{\sqrt{mn}} \|M - \widehat{M}\|_F. \quad (1)$$

Candès and Plan introduced a generalization of the convex relaxation from [4] to the noisy case, and provided a bound on the RMSE [7]. More recently, a bound on the RMSE achieved by the OPTSPACE algorithm with noisy observations was obtained in [8]. This bound is order optimal in a number of situations and improves over the analogous result in [7]. Detailed comparison of these two results are provided in Section II-D.

On the practical side, directly solving the convex relaxation introduced in [4] requires solving a Semidefinite Program (SDP), the complexity of which grows proportional to n^3 . Recently, many authors have proposed efficient algorithms for solving the low-rank matrix completion problem. These include Accelerated Proximal Gradient (APG) algorithm [9], Fixed Point Continuation with Approximate SVD (FPCA) [10], Atomic Decomposition for Minimum Rank Approximation (ADMiRA) [11], SOFT-IMPUTE [12], Subspace Evolution and Transfer (SET) [13], Singular Value Projection (SVP) [14], and OPTSPACE [5]. In this paper, we provide numeri-

cal comparisons of the performance of three state-of-the-art algorithms, namely, OPTSPACE, ADMIRA and FPCA, and show that these efficient algorithms can be used to reconstruct real data matrices, as well as randomly generated matrices, accurately.

B. Outline

The organization of this paper is as follows. In Section 2, we describe the matrix completion problem and efficient algorithms to solve the matrix completion problem when the observations are corrupted by noise. Section 3 discusses the results of numerical simulations and compares the performance of three matrix completion algorithms with respect to speed and accuracy.

II. THE MODEL DEFINITION AND ALGORITHMS

A. Model definition

The matrix M has dimensions $m \times n$, and we define $\alpha \equiv m/n$ to denote the ratio. In the following we assume, without loss of generality, $\alpha \geq 1$. We assume that the matrix M has exact low rank $r \ll n$, that is, there exist matrices U of dimensions $m \times r$, V of dimensions $n \times r$, and a diagonal matrix Σ of dimensions $r \times r$, such that

$$M = U\Sigma V^T.$$

Notice that for a given matrix M , the factors (U, V, Σ) are not unique. Further, each entry of M is perturbed, thus producing an ‘approximately’ low-rank matrix N , with

$$N_{ij} = M_{ij} + Z_{ij},$$

where the matrix Z accounts for the noise.

Out of the $m \times n$ entries of N , a subset $E \subseteq [m] \times [n]$ is observed. Let N^E be the $m \times n$ observed matrix with all the observed values, such that

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

Our goal is to find a low rank estimation $\widehat{M}(N^E, E)$ of the original matrix M from the observed noisy matrix N^E and the set of observed indices E .

B. Algorithms

In the case when there is no noise, that is $N_{ij} = M_{ij}$, solving the following optimization problem will recover the original matrix correctly, if the number of observed entries $|E|$ is large enough.

$$\begin{aligned} & \text{minimize} \quad \text{rank}(X) \\ & \text{subject to} \quad \mathcal{P}_E(X) = \mathcal{P}_E(M), \end{aligned} \quad (2)$$

where $X \in \mathbb{R}^{m \times n}$ is the variable matrix, $\text{rank}(X)$ is the rank of matrix X , and $\mathcal{P}_E(\cdot)$ is the projector operator defined as

$$\mathcal{P}_E(M)_{ij} = \begin{cases} M_{ij} & \text{if } (i, j) \in E, \\ 0 & \text{otherwise.} \end{cases} \quad (3)$$

This problem finds the matrix with the minimum rank that matches all the observations. Notice that the solution of

problem (2) is optimal. If this problem does not recover the correct matrix M then there exists at least one other rank- r matrix that matches all the observations and no other algorithm can distinguish which one is the correct solution. However, this optimization problem is NP-hard and all known algorithms require doubly exponential time in n [4].

In compressed sensing, minimizing the l_1 norm of a vector is the tightest convex relaxation of minimizing the l_0 norm, or equivalently minimizing the number of non-zero entries, for sparse signal recovery. We can adopt this idea to matrix completion, where $\text{rank}(\cdot)$ of a matrix corresponds to l_0 norm of a vector, and nuclear norm to l_1 norm [4], where the nuclear norm of a matrix is defined as the sum of its singular values.

$$\begin{aligned} & \text{minimize} \quad \|X\|_* \\ & \text{subject to} \quad \mathcal{P}_E(X) = \mathcal{P}_E(M), \end{aligned} \quad (4)$$

where $\|X\|_*$ denotes the nuclear norm of X .

In this paper, we are interested in the more practical case when the observations are contaminated by noise or the original matrix to be reconstructed is only approximately low rank. In this case, the constraint $\mathcal{P}_E(X) = \mathcal{P}_E(M)$ must be relaxed. This results in either the problem [7], [10], [9], [12]

$$\begin{aligned} & \text{minimize} \quad \|X\|_* \\ & \text{subject to} \quad \|\mathcal{P}_E(X) - \mathcal{P}_E(N)\|_F \leq \Theta, \end{aligned} \quad (5)$$

or its Lagrangian version

$$\text{minimize} \quad \mu \|X\|_* + \frac{1}{2} \|\mathcal{P}_E(X) - \mathcal{P}_E(N)\|_F^2. \quad (6)$$

In the following, we briefly explain the objective of the three state-of-the-art matrix completion algorithms based on the relaxation, namely, FPCA, ADMIRA, and OPTSPACE.

FPCA, introduced in [10], is an efficient algorithm for solving the convex relaxation, which is a nuclear norm regularized least squares problem in (6). Following the same line of argument given in [7], we choose $\mu = \sqrt{2np}\sigma$, where $p = |E|/mn$ and σ^2 is the variance of each entry in Z .

ADMIRA, introduced in [11], is an efficient algorithm which is based on the atomic decomposition and extends the idea of the Compressive Sampling Matching Pursuit (CoSaMP) [15]. ADMIRA is an iterative method for solving the following rank- r matrix approximation problem.

$$\begin{aligned} & \text{minimize} \quad \|\mathcal{P}_E(X) - \mathcal{P}_E(N)\|_F \\ & \text{subject to} \quad \text{rank}(X) \leq r. \end{aligned} \quad (7)$$

One drawback of ADMIRA is that it requires the prior knowledge of the rank of the original matrix M . In the following numerical simulations, for fair comparison, we first run a rank estimation algorithm to guess the rank of the original matrix and use the estimated rank in ADMIRA. The rank estimation algorithm is explained in the next section.

OPTSPACE, introduced in [5], is a novel and efficient algorithm based on the spectral method followed by a local optimization, which consists of the following three steps.

1. Trim the matrix N^E .

2. Compute the rank- r projection of the trimmed observation matrix.

3. Minimize $\|\mathcal{P}_E(XSY^T) - \mathcal{P}_E(N)\|_F^2$ through gradient descent, using the rank- r projection as the initial guess.

In the trimming step, we set to zero all columns in N^E with the number of samples larger than $2|E|/n$ and set to zero all rows with the number of samples larger than $2|E|/m$. In the second step, the rank- r projection of a matrix A is defined as

$$\mathcal{P}_r(A) = \frac{mn}{|E|} \sum_{i=1}^r \sigma_i x_i y_i^T, \quad (8)$$

where the SVD of A is given by $A = \sum_{i=1}^n \sigma_i x_i y_i^T$. The basic idea is that the rank- r projection of the trimmed observation matrix provides an excellent initial guess, so that the standard gradient descent provides a good estimate after this initialization. Note that we need to estimate the target rank r . To estimate the target rank r for ADMIRA and OPTSPACE, we used the following simple rank estimation procedure.

C. Rank estimation algorithm

Let \tilde{N}^E be the trimmed version of N^E . By singular value decomposition of the trimmed matrix, we have

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

where x_i and y_i are the left and right singular vectors corresponding to i th singular value σ_i . Then, the following cost function is defined in terms of the singular values.

$$R(i) = \frac{\sigma_{i+1} + \sigma_1 \sqrt{\frac{i\sqrt{mn}}{|E|}}}{\sigma_i}.$$

Finally, the estimated rank is the index i that minimizes the cost function $R(i)$.

The idea behind this algorithm is that if enough entries of N are revealed and there is little noise then there is a clear separation between the first r singular values, which reveal the structure of the matrix M to be reconstructed, and the spurious ones [5]. Hence, σ_{i+1}/σ_i is minimum when i is the correct rank r . The second term is added to ensure the robustness of the algorithm.

D. Comparison of the performance guarantees

Performance guarantees for matrix completion problem with noisy observations are proved in [7] and [8]. Theorem 7 of [7] shows the following bound on the performance of solving convex relaxation (5) under some constraints on the matrix M known as the strong incoherence property.

$$\text{RMSE} \leq 7\sqrt{\frac{n}{|E|}} \|\mathcal{P}_E(Z)\|_F + \frac{2}{n\sqrt{\alpha}} \|\mathcal{P}_E(Z)\|_F, \quad (9)$$

where RMSE is defined in Eq. (1). The constant in front of the first term is in fact slightly smaller than 7 in [7], but in any case larger than $4\sqrt{2}$.

Theorem 1.2 of [8] shows the following bound on the performance of OPTSPACE under the assumptions that M

is incoherent and has a bounded condition number $\kappa = \sigma_1(M)/\sigma_r(M)$, where the condition number of a matrix is defined as the ratio between the largest singular value $\sigma_1(M)$ and the smallest singular value $\sigma_r(M)$ of M .

$$\text{RMSE} \leq C\kappa^2 \sqrt{\alpha r} \frac{n}{|E|} \|\mathcal{P}_E(Z)\|_2, \quad (10)$$

for some numerical constant C .

Although the assumptions on the above two theorems are not directly comparable, as far as the error bounds are concerned, the bound (10) improves over the bound (9) in several respects: The bound (10) does not have the second term in the bound (9) which actually grows with the number of observed entries; The bound (10) decreases as $n/|E|$ rather than $(n/|E|)^{1/2}$; The bound (10) is proportional to the operator norm of the noise matrix $\|\mathcal{P}_E(Z)\|_2$ instead of the Frobenius norm $\|\mathcal{P}_E(Z)\|_F \geq \|\mathcal{P}_E(Z)\|_2$. For E uniformly random, one expects $\|\mathcal{P}_E(Z)\|_F$ to be roughly of order $\|\mathcal{P}_E(Z)\|_2 \sqrt{n}$. For instance, if the entries of Z are i.i.d. Gaussian with bounded variance σ , $\|\mathcal{P}_E(Z)\|_F = \Theta(\sqrt{|E|})$ while $\|\mathcal{P}_E(Z)\|_2$ is of order $\sqrt{|E|/n}$.

In the following, we numerically compare the performances of three efficient algorithms, OPTSPACE, ADMIRA and FPCA, for solving the matrix completion problem, with real data matrices as well as randomly generated matrices.

III. NUMERICAL RESULTS

In this section, we present numerical comparisons between three approximate low-rank matrix completion algorithms: OPTSPACE, ADMIRA and FPCA. The performance of each algorithm is compared in terms of the relative root mean squared error defined as in Eq. (1). for randomly generated matrices in Section III-A and real data matrices in Section III-B. We used MATLAB implementations of the algorithms and tested them on a 3.0 GHz Desktop computer with 2 GB RAM. FPCA is available from www.columbia.edu/~sm2756/FPCA.htm and OPTSPACE is available from www.stanford.edu/~raghuram/optspace/.

A. Numerical results with randomly generated matrices

For numerical simulations with randomly generated matrices, we use $n \times n$ test matrices M of rank r generated as $M = UV^T$, where U and V are $n \times r$ matrices with each entry being sampled independently from a standard Gaussian distribution $\mathcal{N}(0, 1)$, unless specified otherwise. Each entry is revealed independently with probability ϵ/n , so that on an average $n\epsilon$ entries are revealed. The observation is corrupted by added noise matrix Z , so that the observation for the index (i, j) is $M_{ij} + Z_{ij}$.

In the standard scenario, we typically make the following three assumptions on the noise matrix Z . (1) The noise Z_{ij} does not depend on the value of the matrix M_{ij} . (2) The entries of Z , $\{Z_{ij}\}$, are independent. (3) The distribution of each entries of Z is Gaussian. The above matrix completion algorithms are expected to be especially effective under this standard scenario for the following two reasons. First, the




Fig. 1. The RMSE (above) and the computation time in seconds (below) as a function of the average number of observed entries per row ϵ for SNR=4 under the standard scenario.

squared error objective function that the algorithms minimize is well suited for the Gaussian noise. Second, the independence of Z_{ij} 's ensure that the noise matrix is almost full rank and the singular values are evenly distributed. This implies that for a given noise power $\|Z\|_F$, the spectral norm $\|Z\|_2$ is much smaller than $\|Z\|_F$. In the following, we fix $m = n = 500$ and $r = 4$, and study how the performance changes with different noise. Each of the simulation results is averaged over 10 instances and is shown with respect to two basic parameters, the average number of revealed entries per row ϵ and the signal-to-noise ratio, $\text{SNR} = \mathbb{E}[\|M\|_F^2] / \mathbb{E}[\|Z\|_F^2]$.

1) *Standard scenario:* In this standard scenario, the noise Z_{ij} 's are distributed as i.i.d. Gaussian $\mathcal{N}(0, \sigma^2)$. Note that the SNR is equal to $4/\sigma^2$. There is a basic trade-off between two metrics of interest: the accuracy of the estimation is measured using RMSE and the computation complexity is measured by the running time in seconds.

In order to interpret the simulation results, they are compared to the RMSE achieved by the oracle and a simple rank- r projection algorithm defined as Eq. (8). The rank- r projection algorithm simply computes $\mathcal{P}_r(N^E)$. The oracle has prior knowledge of the linear subspace spanned by $\{UX^T + YV^T : X \in \mathbb{R}^{m \times r}, Y \in \mathbb{R}^{n \times r}\}$, and the RMSE of the oracle estimate

is $\sigma\sqrt{(2nr - r^2)/n\epsilon}$ [7].

Figure 1 shows the performance and the computation time for each of the algorithms with respect to ϵ under the standard scenario for fixed SNR= 4. For most values of ϵ , the simple rank- r projection has the worst performance. However, when all the entries are revealed and the noise is i.i.d. Gaussian, the rank- r projection coincides with the oracle bound, which in this simulation corresponds to the value $\epsilon = 500$. Note that the behavior of the performance curves of FPCA, ADMIRA, and OPTSPACE with respect to ϵ is similar to the oracle bound, which is proportional to $1/\sqrt{\epsilon}$.

Among the three algorithms, FPCA has the largest RMSE, and OPTSPACE is very close to the oracle bound for all values of ϵ . Note that when all the values are revealed, ADMIRA is an efficient way of implementing rank- r projection, and the performances are expected to be similar. This is confirmed by the observation that for $\epsilon \geq 400$ the two curves are almost identical. One of the reasons why the RMSE of FPCA does not decrease with ϵ for large values of ϵ is that FPCA overestimates the rank and returns estimated matrices with rank much higher than r , whereas the rank estimation algorithm used for ADMIRA and OPTSPACE always returned the correct rank r for $\epsilon \geq 80$.

The second figure in Figure 1 shows the average running time of the algorithms with respect to ϵ . Note that due to the large difference between the running time of three algorithms, the time is displayed in log scale. For most of the simulations, ADMIRA had shortest running time and FPCA the longest, and the gap was noticeably large as clearly shown in the figure. For FPCA and OPTSPACE, the computation time increased with ϵ , whereas ADMIRA had relatively stable computation time independent of ϵ .

Figure 2 show the performance and computation time for each of the algorithms against the SNR within the standard scenario for fixed $\epsilon = 40$. The behavior of the performance curves of ADMIRA and OPTSPACE are similar to the oracle bound which is linear in σ which, in the standard scenario, is equal to $2/\sqrt{\text{SNR}}$. The performance of the rank- r projection algorithm is determined by two factors. One is the added noise which is linear in σ and the other is the error caused by the erased entries which is constant independent of SNR. These two factors add up, whence the performance curve of the rank- r projection follows. The reason the RMSE of FPCA does not decrease with SNR for values of SNR less than 1 is not that the estimates are good but rather the estimated entries gets very small and the resulting RMSE is close to $\sqrt{\mathbb{E}[\|M\|_F^2/n^2]}$, which is 2 in this simulation, regardless of the noise power. When there is no noise, which corresponds to the value $1/\text{SNR} = 0$, FPCA and OPTSPACE both recover the original matrix correctly for this chosen value of $\epsilon = 40$. For all three algorithms, the computation time is larger for smaller noise, and the reason is that it takes more iterations until the stopping criterion is met. Also, for most of the simulations with different SNR, ADMIRA had shortest running time and FPCA the longest.

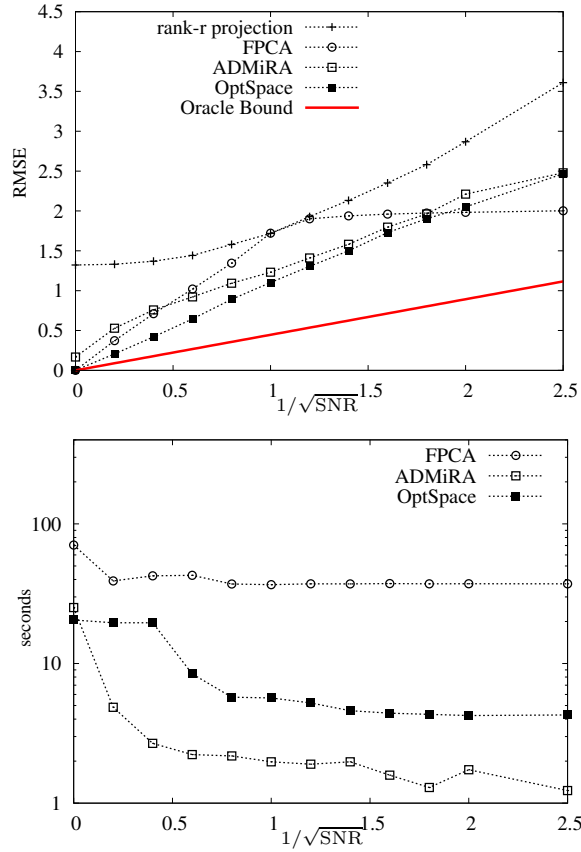


Fig. 2. The RMSE (above) and the computation time in seconds (below) as a function of $1/\text{SNR}$ for fixed $\epsilon = 40$ under the standard scenario.




Fig. 3. The RMSE as a function of the average number of observed entries per row ϵ for fixed $\text{SNR}=4$ within the multiplicative noise model.

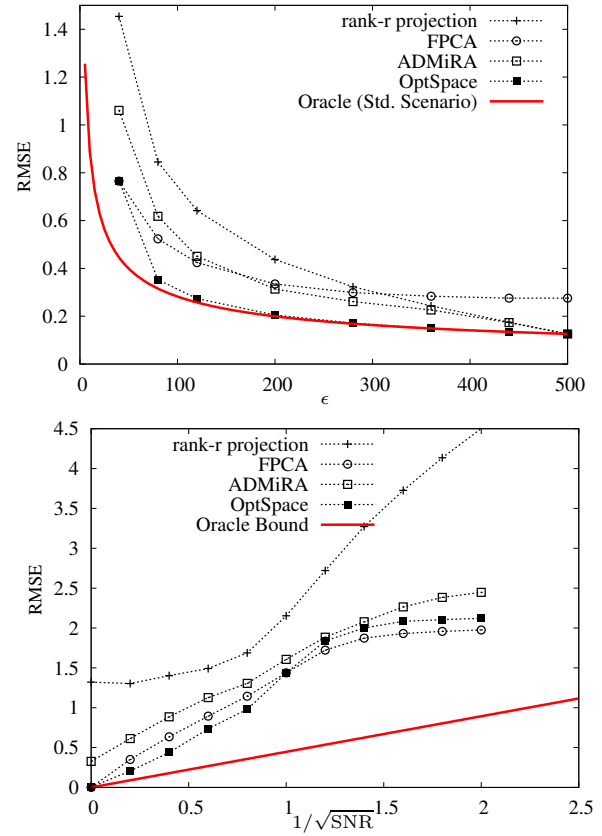


Fig. 4. The RMSE as a function of the average number of observed entries per row ϵ for fixed $\text{SNR}=4$ with outliers (above) and the RMSE as a function of $1/\text{SNR}$ for fixed $\epsilon = 40$ with outliers (below).

2) *Multiplicative Gaussian noise*: In sensor network localization [16], where the entries of the matrix corresponds to the pair-wise distances between the sensors, the observation noise is oftentimes assumed to be multiplicative. In formulae, $Z_{ij} = \xi_{ij}M_{ij}$, where ξ_{ij} 's are distributed as i.i.d. Gaussian with zero mean. The variance of ξ_{ij} 's are chosen to be $1/r$ so that the resulting noise power is one. Note that in this case, Z_{ij} 's are mutually dependent through M_{ij} 's and the values of the noise also depend on the value of the matrix entry M_{ij} .

Figure 3 shows the RMSE with respect to ϵ under multiplicative Gaussian noise. The RMSE of the rank- r projection for $\epsilon = 40$ is larger than 1.5 and is omitted in the figure. The bottommost line corresponds to the oracle performance under standard scenario, and is displayed here, and all of the following figures, to serve as a reference for comparison. The main difference with respect to Figure 1 is that all the performance curves are larger under multiplicative noise. For the same value of SNR , it is more difficult to distinguish the noise from the original matrix, since the noise is now correlated with the matrix M .

3) *Outliers*: In structure from motion [2], the entries of the matrix corresponds to the position of points of interest in 2-dimensional images captured by cameras in different angles and locations. However, due to failures in the feature extraction




Fig. 5. The RMSE as a function of the average number of observed entries per row ϵ for fixed SNR= 4 with quantization.




Fig. 6. The RMSE as a function of the average number of observed entries per row ϵ for fixed SNR= 4 with ill-conditioned matrices.

algorithm, some of the observed positions are corrupted by large noise where as most of the observations are noise free. To account for such outliers, we use the following model.

$$Z_{ij} = \begin{cases} a & \text{with probability } 1/200, \\ -a & \text{w.p. } 1/200, \\ 0 & \text{w.p. } 99/100. \end{cases}$$

The value of a is chosen according to the target SNR= $400/a^2$. This is clearly independent of the matrix entries and Z_{ij} 's are mutually independent, but the distribution is now non-Gaussian.

Figure 4 shows the performance of the algorithms with respect to ϵ and the SNR with outliers. Comparing the first figure to Figure 1, we can see that the performance for large value of ϵ is less affected by outliers compared to the small values of ϵ . The second figure clearly shows how the performance degrades for non-Gaussian noise when the number of samples is small. The algorithms minimize the squared error $\|\mathcal{P}_E(X) - \mathcal{P}_E(N)\|_F^2$ as in (6) and (7). For outliers, a suitable algorithm would be to minimize the l_1 -norm of the errors instead of the l_2 -norm. Hence, for this simulation outliers, we can see that the performance of the rank- r projection, ADMiRA and OPTSPACE is worse than the Gaussian noise case. However, the performance of FPCA is almost the same as in the standard scenario.

4) *Quantization noise*: One common model for noise is the quantization noise. For a regular quantization, we choose a parameter a and quantize the matrix entries to the nearest value in $\{\dots, -a/2, a/2, 3a/2, 5a/2, \dots\}$. The parameter a is chosen carefully such that the resulting SNR is 4. The performance for this quantization is expected to be worse than the multiplicative noise case, since now the noise is deterministic and completely depends on the matrix entries M_{ij} , whereas in the multiplicative noise model it was random.

Figure 5 shows the performance against ϵ within quantization noise. The overall behavior of the performance curves is similar to Figure 1, but all the curves are shifted up. Note that the bottommost line is the oracle performance in the standard

scenario which is the same in all the figures. Compared to Figure 3, for the same values of SNR= 4, quantization is much more damaging than the multiplicative noise as expected.

5) *Ill conditioned matrices*: In this simulation, we look at how the performance degrades under the standard scenario if the matrix M is ill-conditioned. M is generated as $M = \sqrt{4/166} U \text{diag}([1, 4, 7, 10]) V^T$, where U and V are generated as in the standard scenario. The resulting matrix has condition number 10 and the normalization constant $\sqrt{4/166}$ is chosen such that $\mathbb{E}[\|M\|_F]$ is the same as in the standard case.

Figure 6 shows the performance with respect to ϵ with ill-conditioned matrix M . The performance of OPTSPACE is similar to that of ADMiRA for many values of ϵ . However, a modification of OPTSPACE called INCREMENTAL OPTSPACE achieves a better performance in this case of ill-conditioned matrix. The INCREMENTAL OPTSPACE algorithm starts from finding a rank-1 approximation from N^E and incrementally finds higher rank approximations and has more robust performance when M is ill-conditioned, but is computationally more expensive.

B. Numerical results with real data matrices

In this section, we consider the low-rank matrix completion problems in the context of recommender systems, based on two real data sets : the Jester joke data set [17] and the Movielens data set [18]. The Jester joke data set contains 4.1×10^6 ratings for 100 jokes from 73,421 users. ¹ Since the number of users is large compared to the number of jokes, we randomly select $n_u \in \{100, 1000, 2000, 4000\}$ users for comparison purposes. As in [10], we randomly choose two ratings for each user as a test set, and this test set, which we denote by T , is used in computing the prediction error in Normalized Mean Absolute Error (NMAE). The Mean

¹The dataset is available at <http://www.ieor.berkeley.edu/~goldberg/jester-data/>

Absolute Error (MAE) is defined as in [10], [19].

$$MAE = \frac{1}{|T|} \sum_{(i,j) \in T} |M_{ij} - \widehat{M}_{ij}|,$$

where M_{ij} is the original rating in the data set and \widehat{M}_{ij} is the predicted rating for user i and item j . The Normalized Mean Absolute Error (NMAE) is defined as

$$NMAE = \frac{MAE}{M_{\max} - M_{\min}},$$

where M_{\max} and M_{\min} are upper and lower bounds for the ratings. In the case of Jester joke, all the ratings are in $[-10, 10]$ which implies that $M_{\max} = 10$ and $M_{\min} = -10$.

n_u	n_s	OPTSPACE	FPCA	ADMIRA
100	7484	0.17674	0.20386	0.18194
1000	73626	0.15835	0.16114	0.16194
2000	146700	0.15747	0.16101	0.16286
4000	290473	0.15918	0.16291	0.16317
943	80000	0.18638	0.19018	0.24276

The numerical results for Jester joke data set using INCREMENTAL OPTSPACE, FPCA and ADMIRA are presented in the first four columns of the table above. The number of jokes m is fixed at 100 and the number of users n_u and the number of samples n_s is given in the first two columns. The resulting NMAE of each algorithm is shown in the table. To get an idea of how good the predictions are, consider the case where each missing entry is predicted with a random number drawn uniformly at random in $[-10, 10]$ and the actual rating is also a random number with same distribution. After a simple computation, we can see that the resulting NMAE of the random prediction is 0.333. As another comparison, for the same data set with $n_u = 18000$, simple nearest neighbor algorithm and EIGENTASTE both yield NMAE of 0.187 [19]. The NMAE of INCREMENTAL OPTSPACE is lower than these simple algorithms even for $n_u = 100$ and tends to decrease with n_u .

Looking at a complete matrix where all the entries are known can bring some insight into the structure of real data matrices. With Jester joke data set, we deleted all users containing missing entries, and generated a complete matrix M with 14,116 users and 100 jokes. The distribution of the singular values of M is shown in Figure 7. We must point out that this rating matrix is not low-rank or even approximately low-rank, although it is common to make such assumptions. This is one of the difficulties in dealing with real data. The other aspect is that the samples are not drawn uniformly at random as commonly assumed in [6], [5].

Numerical simulation results on the Movielens data set is also shown in the last row of the above table. The data set contains 100,000 ratings for 1,682 movies from 942 users.² We use 80,000 randomly chosen ratings to estimate the 20,000 ratings in the test set, which is called *u1.base*

²The dataset is available at <http://www.grouplens.org/node/73>




Fig. 7. Distribution of the singular values of the complete sub matrix in the Jester joke data set.

and *u1.test*, respectively, in the movielens data set. In the last row of the above table, we compare the resulting NMAE using INCREMENTAL OPTSPACE, FPCA and ADMIRA.

REFERENCES

- [1] “Netflix prize,” <http://www.netflixprize.com/>.
- [2] P. Chen and D. Suter, “Recovering the missing components in a large noisy low-rank matrix: application to sfm,” *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 26, no. 8, pp. 1051–1063, Aug. 2004.
- [3] S. Oh, A. Karbasi, and A. Montanari, “Sensor network localization from local connectivity: performance analysis for the MDS-MAP algorithm,” 2009, <http://infoscience.epfl.ch/record/140635>.
- [4] E. J. Candès and B. Recht, “Exact matrix completion via convex optimization,” 2008, [arXiv:0805.4471](https://arxiv.org/abs/0805.4471).
- [5] R. H. Keshavan, A. Montanari, and S. Oh, “Matrix completion from a few entries,” January 2009, [arXiv:0901.3150](https://arxiv.org/abs/0901.3150).
- [6] E. J. Candès and T. Tao, “The power of convex relaxation: Near-optimal matrix completion,” 2009, [arXiv:0903.1476](https://arxiv.org/abs/0903.1476).
- [7] E. J. Candès and Y. Plan, “Matrix completion with noise,” 2009, [arXiv:0903.3131](https://arxiv.org/abs/0903.3131).
- [8] R. H. Keshavan, A. Montanari, and S. Oh, “Matrix completion from noisy entries,” June 2009, [arXiv:0906.2027](https://arxiv.org/abs/0906.2027).
- [9] K. Toh and S. Yun, “An accelerated proximal gradient algorithm for nuclear norm regularized least squares problems,” 2009, <http://www.math.nus.edu.sg/~matys>.
- [10] S. Ma, D. Goldfarb, and L. Chen, “Fixed point and Bregman iterative methods for matrix rank minimization,” 2009, [arXiv:0905.1643](https://arxiv.org/abs/0905.1643).
- [11] K. Lee and Y. Bresler, “Admira: Atomic decomposition for minimum rank approximation,” 2009, [arXiv:0905.0044](https://arxiv.org/abs/0905.0044).
- [12] R. Mazumder, T. Hastie, and R. Tibshirani, “Spectral regularization algorithms for learning large incomplete matrices,” 2009, http://www-stat.stanford.edu/~hastie/Papers/SVD_JMLR.pdf.
- [13] W. Dai and O. Milenkovic, “Set: an algorithm for consistent matrix completion,” 2009, [arXiv:0909.2705](https://arxiv.org/abs/0909.2705).
- [14] R. Meka, P. Jain, and I. S. Dhillon, “Guaranteed rank minimization via singular value projection,” 2009, [arXiv:0909.5457](https://arxiv.org/abs/0909.5457).
- [15] D. Needell and J. A. Tropp, “Cosamp: Iterative signal recovery from incomplete and inaccurate samples,” *Applied and Computational Harmonic Analysis*, vol. 26, no. 3, pp. 301–321, Apr 2008. [Online]. Available: <http://arxiv.org/abs/0803.2392>
- [16] Z. Wang, S. Zheng, S. Boyd, and Y. Ye, “Further relaxations of the sdp approach to sensor network localization,” Tech. Rep., 2006.
- [17] “Jester jokes,” <http://eigentaste.berkeley.edu/user/index.php>.
- [18] “Movielens,” <http://www.movielens.org>.
- [19] K. Goldberg, T. Roeder, D. Gupta, and C. Perkins, “Eigentaste: A constant time collaborative filtering algorithm,” pp. 133–151, July 2001.