Title

Max-norm optimization for robust matrix recovery

Journal

Mathematical Programming, 167(1)

ISSN

0025-5610

Authors

Fang, Ethan X
Liu, Han
Toh, Kim-Chuan
et al.

Publication Date

2018

DOI

10.1007/s10107-017-1159-y

Peer reviewed

Math. Program., Ser. B (2018) 167:5–35
https://doi.org/10.1007/s10107-017-1159-y

FULL LENGTH PAPER

Max-norm optimization for robust matrix recovery

Ethan X. Fang[1] · **Han Liu**[2] · **Kim-Chuan Toh**[3] · **Wen-Xin Zhou**[4]

Received: 2 September 2015 / Accepted: 2 May 2017 / Published online: 30 May 2017

Abstract This paper studies the matrix completion problem under arbitrary sampling schemes. We propose a new estimator incorporating both max-norm and nuclear-norm regularization, based on which we can conduct efficient low-rank matrix recovery using a random subset of entries observed with additive noise under general non-uniform and unknown sampling distributions. This method significantly relaxes the uniform sampling assumption imposed for the widely used nuclear-norm penalized approach, and makes low-rank matrix recovery feasible in more practical settings. Theoretically, we prove that the proposed estimator achieves fast rates of convergence under different settings. Computationally, we propose an alternating direction method of multipliers

K.-C. Toh: Research supported in part by Ministry of Education Academic Research Fund R-146-000-194-112.

✉ Ethan X. Fang
xxf13@psu.edu

Han Liu
hanliu@princeton.edu

Kim-Chuan Toh
mattohkc@nus.edu.sg

Wen-Xin Zhou
stevewenxin@hotmail.com

[1] Department of Statistics, Department of Industrial and Manufacturing Engineering, Pennsylvania State University, University Park, PA 16802, USA

[2] Department of Operations Research and Financial Engineering, Princeton University, Princeton, NJ 08544, USA

[3] Department of Mathematics, National University of Singapore, 10 Lower Kent Ridge Road, Singapore 119076, Singapore

[4] Department of Mathematics, University of California, San Diego, La Jolla, CA 92093, USA

algorithm to efficiently compute the estimator, which bridges a gap between theory and practice of machine learning methods with max-norm regularization. Further, we provide thorough numerical studies to evaluate the proposed method using both simulated and real datasets.

Mathematics Subject Classification 90C25 · 90C29 · 15A60

1 Introduction

We consider the matrix completion problem, which aims to reconstruct an unknown matrix based on a small number of entries contaminated by additive noise. This problem has drawn significant attention over the past decade due to its wide applications, including collaborative filtering (the well-known Netflix problem) [4,27], multi-task learning [1–3], sensor-network localization [5] and system identification [24]. Specifically, our goal is to recover an unknown matrix $M^0 \in \mathbb{R}^{d_1 \times d_2}$ based on a subset of its entries observed with noise, say $\{Y_{i_t, j_t}\}_{t=1}^n$. In general, the problem of recovering a partially observed matrix is ill-posed, as the unobserved entries can take any values without further assumption. However, in many applications mentioned above, it is natural to impose the condition that the target matrix is of either exact or approximately low-rank, which avoids the ill-posedness and makes the recovery possible.

To obtain a low-rank estimate of the matrix, a straightforward approach is to consider the rank minimization problem

$$\min_{M \in \mathbb{R}^{d_1 \times d_2}} \operatorname{rank}(M), \text{ subject to } \|Y_\Omega - M_\Omega\|_F \leq \delta, \tag{1.1}$$

where $\Omega = \{(i_t, j_t) : t = 1, \ldots, n\}$ is the index set of observed entries, and $\delta > 0$ is a tuning parameter. This method directly searches for a matrix of the lowest rank with reconstruction error controlled by δ. However, the optimization problem (1.1) is computationally intractable due to its nonconvexity. A commonly used alternative is the following convex relaxation of (1.1):

$$\min_{M \in \mathbb{R}^{d_1 \times d_2}} \|M\|_*, \text{ subject to } \|Y_\Omega - M_\Omega\|_F \leq \delta, \tag{1.2}$$

where $\|\cdot\|_*$ denotes the nuclear-norm (also known as the trace-norm, Ky Fan-norm or Schatten 1-norm), and it is defined as the sum of singular values of a matrix. Low-rank matrix recovery based on nuclear-norm regularization has been extensively studied in both noiseless and noisy cases [9,10,19,21,26,30–32]. Furthermore, various computational algorithms have been proposed to solve this problem. For example, Cai et al. [6] propose a singular value thresholding algorithm which is equivalent to the gradient method for solving the dual of a regularized version of (1.2); Toh and Yun [37] propose an accelerated proximal gradient method to solve a least squares version of (1.2); Liu and Vandenberghe [24] exploit an interior-point method; Chen et al. [11] adopt an alternating direction method of multipliers approach to solve (1.2).

Though significant progress has been made, it remains unclear whether the nuclear-norm is the best convex relaxation for the rank minimization problem (1.1). Recently,

	Movie 1	Movie 2	Movie 3	Movie 4	Movie 5
User 1	5			4	
User 2		3		2	
User 3	1		4		5
User 4		2			1
User 5			1	3	

(a)

	Movie 1	Movie 2	Movie 3	Movie 4	Movie 5
User 1					1
User 2		2		3	4
User 3	1		4		5
User 4		2			1
User 5	4	2	1	3	1

(b)

Fig. 1 **a** The theoretical guarantee of the nuclear-norm estimator assumes each entry is equally likely to be observed. **b** In practice, some entries related to some popular movies or some active users, such as Movie 5 or User 5, are more likely to be sampled than others. Thus, the uniform sampling assumption is violated

some disadvantages of the nuclear-norm regularization have been noted. For instance, the theoretical guarantee of the nuclear-norm regularization relies on an assumption that the indices of the observed entries are uniformly sampled. That is, each entry is equally likely to be observed as illustrated in Fig. 1a. This assumption is restrictive in applications. Taking the well-known Netflix problem as an example, our goal is to reconstruct a movie-user rating matrix, in which each row represents a user and each column represents a movie. The (k, ℓ)-th entry of the rating matrix represents the k-th user's rating for the ℓ-th movie. In practice, we only observe a small proportion of the entries. In this example, the uniform sampling assumption is arguably violated due to the following reasons: (1) Some users are more active than others, and they rate more movies than others. (2) Some movies are more popular than others and are rated by more users. As a consequence, the entries from certain columns or rows are more likely to be observed. See Fig. 1b for a simple illustration. To sum up, the sampling distribution can be highly non-uniform in real world applications.

To relax or even avoid the unrealistic uniform sampling assumption, several recent papers propose to use the matrix max-norm as a convex surrogate for the rank. Srebro and Salakhutdinov [35] observe from empirical comparisons that the max-norm regularized approach outperforms the nuclear-norm based one for matrix completion and collaborative filtering under non-uniform sampling schemes. Lee et al. [22] and Jalali and Srebro [17] demonstrate the advantage of using max-norm regularizer over nuclear-norm in some other applications. More recently, Cai and Zhou [7] prove that the max-norm regularized estimator is minimax rate-optimal (over a class of approximately low-rank matrices) under non-uniform sampling schemes.

Though the max-norm approach possesses attractive theoretical properties, efficiently solving large-scale max-norm optimization problem remains challenging and prevents the wide adoption of max-norm regularizer. As we shall see later, despite the fact that the max-norm is a convex regularizer and can be formulated as a semidefinite programming problem, classical methods such as interior-point methods are only scalable to moderate dimensions, while the problem of practical interest is of large dimensions. In recent work, Lee et al. [22] and Shen et al. [33] propose first-order algorithms for a nonconvex relaxation of the problem. However, these methods are sensitive to the choice of initial points and stepsizes, and are only capable of producing stationary solutions, whose statistical properties remain open due to the nonconvexity.

Meanwhile, although the max-norm estimator is adaptive to general sampling schemes, it was shown in Cai and Zhou [7] that if the target matrix is of exact low-rank, and the sampling scheme is uniform, the max-norm estimator only achieves a sub-optimal rate compared to the nuclear-norm estimator. Specifically, letting $\hat{M}_{\max}$ and $\hat{M}_*$ be the estimators using max-norm and nuclear-norm regularizers, we have

$$(d_1 d_2)^{-1}\|\hat{M}_{\max} - M^0\|_F^2 = \mathcal{O}_{\mathbb{P}}(n^{-1/2}\sqrt{rd}) \quad \text{and} \quad (d_1 d_2)^{-1}\|\hat{M}_* - M^0\|_F^2 = \mathcal{O}_{\mathbb{P}}(n^{-1} rd \log d),$$

where r is the rank of M^0 and $d = d_1 + d_2$. To compare, under the uniform sampling scheme, the nuclear-norm regularized method achieves the optimal rate of convergence (up to a logarithmic factor) and is computationally more scalable.

To achieve the advantages of both regularizers, we propose a new estimator using a hybrid regularizer. Meanwhile, we propose an efficient alternating direction method of multipliers (ADMM) algorithm to solve the optimization problem. Our method includes the max-norm regularizer as a special case, and the proposed algorithm is scalable to modestly large dimensions. The contribution of this paper is two-fold: First, we propose an estimator for matrix completion under general sampling scheme, which achieves optimal rate of convergence in the exact low-rank case and is adaptive to different sampling schemes. Second, we provide an efficient algorithm to solve the corresponding max-norm plus nuclear-norm penalized optimization problem. We illustrate the efficiencies of the proposed methods and algorithms by numerical experiments on both simulated and real datasets.

Notation Throughout this paper, we adopt the following notations. For any positive integer d, $[d]$ denotes the set of integers $\{1, 2, \ldots, d\}$. For a vector $v = (v_1, \ldots, v_d)^T \in \mathbb{R}^d$ and a positive number $p \in (0, \infty)$, we denote $\|u\|_p$ as the ℓ_p-norm, i.e., $\|u\|_p = \left(\sum_{i=1}^d |u_i|^p\right)^{1/p}$. Also, we let $\|u\|_\infty = \max_{i=1,\ldots,d} |u_i|$. For a matrix $M = (M_{k\ell}) \in \mathbb{R}^{d_1 \times d_2}$, let $\|M\|_F = \left(\sum_{k=1}^{d_1}\sum_{\ell=1}^{d_2} M_{k\ell}^2\right)^{1/2}$ be the Frobenius-norm, and we denote the matrix elementwise ℓ_∞-norm by $\|M\|_\infty = \max_{k,l} |M_{k\ell}|$. Given the ℓ_p and ℓ_q norms on $\mathbb{R}^{d_1}$ and $\mathbb{R}^{d_2}$, we define the corresponding $\|\cdot\|_{p,q}$ operator-norm, where $\|M\|_{p,q} = \sup_{\|x\|_p=1} \|Mx\|_q$. For examples, $\|M\| = \|M\|_{2,2}$ is the spectral-norm, and $\|M\|_{2,\infty} = \max_{k=1,\ldots,d_1} \left(\sum_{\ell=1}^{d_2} M_{k\ell}^2\right)^{1/2}$ is the maximum row norm of M. We denote by $a \asymp b$ if $c_1 b \le a \le c_2 b$ for two constants c_1 and c_2.

Paper Organization The rest of this paper is organized as follows. In Sect. 2, we review the max-norm approach and formulate the problem. In Sect. 3, we propose the algorithm. In Sect. 4, we provide theoretical analysis of the estimator. We provide extensive numerical studies in Sect. 5, and we conclude the paper in Sect. 6.

2 Preliminaries and problem formulation

In this section, we first introduce the concept of the matrix max-norm [23]. Next, we propose a new estimator which involves both max-norm and nuclear-norm regularizers.

Definition 2.1 The max-norm of a matrix $M \in \mathbb{R}^{d_1 \times d_2}$ is defined as

$$\|M\|_{\max} = \min_{M=UV^T} \|U\|_{2,\infty} \|V\|_{2,\infty},$$

where the minimum is over all factorizations $M = UV^T$ for $U \in \mathbb{R}^{d_1 \times k}$, $V \in \mathbb{R}^{d_2 \times k}$ for $k = 1, \ldots, \min(d_1, d_2)$, and $\|U\|_{2,\infty}$, $\|V\|_{2,\infty}$ denote the operator-norms of $U : \ell_2^k \to \ell_\infty^{d_1}$ and $V : \ell_2^k \to \ell_\infty^{d_2}$.

We briefly compare the max-norm and nuclear-norm regularizers. We refer to Srebro and Shraibman [36] and Cai and Zhou [7] for more detailed discussions. Recall that the nuclear-norm of the matrix M is defined as

$$\|M\|_* = \min \Big\{ \sum_j |\sigma_j| : M = \sum_j \sigma_j u_j v_j^T, u_j \in \mathbb{R}^{d_1}, v_j \in \mathbb{R}^{d_2}, \|u_j\|_2 = \|v_j\|_2 = 1 \Big\}.$$

From the definition, the nuclear-norm encourages low-rank approximation with factors in the ℓ_2-space. On the other hand, it is known [18] that the max-norm has a similar interpretation by replacing the constraints in the ℓ_2-space by those in the ℓ_∞-space:

$$\|M\|_{\max} \asymp \min \Big\{ \sum_j |\sigma_j| : M = \sum_j \sigma_j u_j v_j^T, u_j \in \mathbb{R}^{d_1}, v_j \in \mathbb{R}^{d_2}, \|u_j\|_\infty = \|v_j\|_\infty = 1 \Big\},$$

where the factor of equivalence is the Grothendieck's constant $K \in (1.67, 1.79)$. More specifically, a consequence of Grothendieck's inequality is that $K_G^{-1} \|M\|_{1\to\infty} \le \|M\|_{\max} \le \|M\|_{1\to\infty}$ [36], where $\|M\|_{1\to\infty} := \max_{u \in \mathbb{R}^{d_2} : \|u\|_1 \le 1} \|Mu\|_\infty$ for any $M \in \mathbb{R}^{d_1 \times d_2}$. This gives some intuition on why the max-norm regularizer could outperform the nuclear-norm regularizer when the matrix entries are uniformly bounded. This scenario indeed stands in many applications. For example, in the Netflix problem or the low-rank correlation matrix estimation problem, the entries of the unknown matrix are either ratings or correlation coefficients, and are uniformly bounded.

As mentioned in Sect. 1, the advantages of using the max-norm over the nuclear-norm are well illustrated in the literature from both theoretical and practical perspectives. Specifically, we consider the matrix completion problem in a general sampling scheme. Let $M^0 \in \mathbb{R}^{d_1 \times d_2}$ denote the unknown matrix to be recovered. Assume that we are given a random index set Ω of size n:

$$\Omega = \big\{(i_t, j_t) : t = 1, \ldots, n\big\} \subset \big([d_1] \times [d_2]\big)^n,$$

where $[d_i] = \{1, 2, \ldots, d_i\}$ for $i = 1, 2$. We further assume that the samples of the indices are drawn independently from a general sampling distribution $\Pi = \{\pi_{k\ell}\}_{k \in [d_1], \ell \in [d_2]}$ on $[d_1] \times [d_2]$. Note that we consider the sampling scheme with replacement, i.e., we assume $\mathbb{P}\{(i_t, j_t) = (k, \ell)\} = \pi_{k\ell}$ for all $t \in [n]$ and all $(k, \ell) \in [d_1] \times [d_2]$. For example, the sampling scheme is uniform if $\pi_{k\ell} = (d_1 d_2)^{-1}$ for all $(k, \ell) \in [d_1] \times [d_2]$. Given the sampled index set Ω, we further observe noisy entries $\{Y_{i_t, j_t}\}_{t \in [n]}$:

$$Y_{i_t,j_t} = M^0_{i_t,j_t} + \sigma\xi_t, \text{ for } t = 1, \ldots, n,$$

where $\sigma > 0$ denotes the noise level, and ξ_t's are independent and identically distributed random variables with $\mathbb{E}(\xi_t) = 0$ and $\mathbb{E}(\xi_t^2) = 1$.

Using the max-norm regularization, Cai and Zhou [7] propose to construct an estimator

$$\hat{M}_{\max} = \underset{M\in\mathbb{R}^{d_1\times d_2}}{\operatorname{argmin}} \frac{1}{n}\sum_{t=1}^{n}\left(Y_{i_t,j_t} - M_{i_t,j_t}\right)^2, \text{ subject to } M \in \mathcal{K}(\alpha, R), \tag{2.1}$$

where $\mathcal{K}(\alpha, R) = \left\{M \in \mathbb{R}^{d_1\times d_2} : \|M\|_\infty \le \alpha, \|M\|_{\max} \le R\right\}$ with α being a prespecified upper bound for the elementwise ℓ_∞-norm of M^0 and $R > 0$ a tuning parameter. Note that, in many real world applications, we have a tight upper bound on the magnitudes of all the entries of M^0 in advance. This condition enforces that M^0 should not be too "spiky", and a loose upper bound may jeopardize the estimation accuracy [26]. Also, the recent work by Lee et al. [22] argues that the max-norm regularizer produces better empirical results on low-rank matrix recovery for uniformly bounded data.

Cai and Zhou [7] provide theoretical guarantees for the max-norm regularizer (2.1). Specifically, under the approximately low-rank assumption that $\|M^0\|_{\max} \le R$, we have,

$$\frac{1}{d_1 d_2}\|\hat{M}_{\max} - M^0\|_F^2 = \mathcal{O}_{\mathbb{P}}\left(\sqrt{\frac{R^2 d}{n}}\right),$$

where $d = d_1 + d_2$. This rate matches the minimax lower bound over all approximately low-rank matrices even under non-uniform sampling schemes. See Cai and Zhou [7] for more details.

The optimization problem (2.1) is computationally challenging. Cai and Zhou [7] employ a first-order method proposed in Lee et al. [22]. In particular, Lee et al. [22] and Shen et al. [33] consider first-order methods based on rewriting problem (2.1) into the following form:

$$\begin{aligned}&\min_{U,V} \frac{1}{n}\sum_{t=1}^{n}\left(U_{i_t}^T V_{j_t} - Y_{i_t,j_t}\right)^2,\\ &\text{subject to } \|U\|_{2,\infty}^2 \le R,\ \|V\|_{2,\infty}^2 \le R,\ \max_{(k,\ell)\in[d_1]\times[d_2]} |U_k^T V_\ell| \le \alpha,\end{aligned}$$

where U_i and V_j denote the i-th row of U and the j-th row of V, respectively. Then, Lee et al. [22] and Shen et al. [33] consider different efficient first-order methods to solve this problem. However, the problem is nonconvex, and the convergence behaviors of those methods on such a nonconvex problem are generally sensitive to the choice of the initial point and stepsize selection. More seriously, the algorithms mentioned can only guarantee local stationary solutions, which may not necessarily possess the nice

theoretical properties for the solution to problem (2.1). More recently, Orabona et al. [29] solve the optimization problem (2.1) without the uniform-boundedness constraint. However, it is unclear how to extend their algorithms to solve the problem (2.1) with the ℓ_∞-norm constraint.

In the next section, we aim to solve the max-norm penalized optimization problem

$$\min_{M\in\mathbb{R}^{d_1\times d_2}} \frac{1}{n}\sum_{t=1}^{n}\left(Y_{i_t,j_t}-M_{i_t,j_t}\right)^2+\lambda\|M\|_{\max},\ \text{subject to}\ \|M\|_\infty\le\alpha, \tag{2.2}$$

where $\lambda > 0$ is a tuning parameter. By convexity and strong duality, the problem (2.2) is equivalent to (2.1) for a properly chosen λ. Specifically, for any R specified in (2.1), there exists a λ such that the solutions to the two problems coincide.

As discussed in Sect. 1, a major drawback of the max-norm penalized estimator (2.1) is that if the underlying true matrix M^0 is of exact low-rank, and when the sampling scheme is indeed uniform, the max-norm regularizer does not perform as well as the nuclear-norm regularizer. Since the underlying structure of M^0 and the sampling scheme are unknown, it is difficult to choose the better approach in practice. To overcome this issue, we propose the following hybrid estimator which is expected to be more flexible and adaptive:

$$\widehat{M}:=\operatorname*{argmin}_{M\in\mathbb{R}^{d_1\times d_2}} \frac{1}{n}\sum_{t=1}^{n}\left(Y_{i_t,j_t}-M_{i_t,j_t}\right)^2+\lambda\|M\|_{\max}+\mu\|M\|_*,\ \text{subject to}\ \|M\|_\infty\le\alpha \tag{2.3}$$

where μ is a nonnegative tuning parameter. The addition of the nuclear-norm penalization is motivated by the fact that the nuclear-norm also serves as a convex surrogate for the rank of the estimator. Thus, the addition of the nuclear-norm encourages the estimator to be low rank or approximately low rank as compared to the max-norm estimator in (2.2). However, note that our primary goal here is not to find a low-rank estimator but one which approximates the underlying matrix M^0 at near optimal recovery and is robust against the unknown sampling scheme. It is worth mentioning that the use of the sum of two norms in matrix recovery has been considered in other contexts. For example, in robust principal component analysis Candès et al. [8], the sum of the nuclear and ℓ_1 norms is used in the recovery of the low-rank and sparse components of a given superposition. In Doan and Vavasis [12], a similar combination of the two norms (denoted as $\|\cdot\|_{1,*} := \|X\|_1 + \theta\|X\|_*$ for a given matrix X and a parameter θ) is used to find hidden sparse rank-one matrices in a given matrix. The geometry of the unit $\|\cdot\|_{1,*}$-norm ball is further analyzed in Drusvyatskiy et al. [13]. It is interesting to note that (2.3) is the first time that the sum of the max-norm and nuclear norm is considered in matrix recovery.

In Sect. 3, we propose an efficient algorithm to solve (2.3), which includes (2.2) as a special case by taking $\mu = 0$. Section 4 provides theoretical justification for the hybrid estimator $\widehat{M}$ in (2.3). In particular, it achieves fast rate of convergence under the "ideal" situation, and is robust against non-uniform sampling schemes. To sum up, this estimator possesses the advantages of both the max-norm and nuclear-norm regularizers. Section 5 provides empirical results of the algorithm.

3 Algorithm

In this section, we propose a new algorithm to solve the problem (2.3). The key step is to reformulate the problem to expose the structure.

3.1 Algorithmic framework

We first review that the max-norm regularized problem (2.2) can be equivalently formulated as a semidefinite programming (SDP) problem. By Definition 2.1, it is unclear how to efficiently compute the max-norm of a given matrix. By Srebro et al. [34], the max-norm of a matrix A can be computed via solving the following SDP problem:

$$\|A\|_{\max} = \min R, \text{ subject to } \begin{pmatrix} W_1 & A \\ A^T & W_2 \end{pmatrix} \succeq 0, \ \|\mathrm{diag}(W_1)\|_\infty \le R, \ \|\mathrm{diag}(W_2)\|_\infty \le R.$$

Thus, the max-norm penalized problem (2.2) can be formulated as an SDP problem that

$$\begin{aligned} \min_{Z \in \mathbb{R}^{d\times d}} \ & \frac{1}{2}\sum_{t=1}^{n}\left(Y_{i_t,j_t} - Z^{12}_{i_t,j_t}\right)^2 + \lambda \, \|\mathrm{diag}(Z)\|_\infty, \\ \text{subject to } \ & \|Z^{12}\|_\infty \le \alpha, \ Z \succeq 0, \end{aligned} \tag{3.1}$$

where $d = d_1 + d_2$, and

$$Z = \begin{pmatrix} Z^{11} & Z^{12} \\ (Z^{12})^T & Z^{22} \end{pmatrix}, \ Z^{11} \in \mathbb{R}^{d_1\times d_1}, \ Z^{12} \in \mathbb{R}^{d_1\times d_2} \text{ and } Z^{22} \in \mathbb{R}^{d_2\times d_2}.$$

One may observe that the problem (3.1) does not explicitly encourage the optimal solutions to be low-rank matrices, although such a property is desirable in many practical applications such as collaborative filtering. Thus, we propose to add the regularization term involving $\langle I, \ Z\rangle$, which is the convex surrogate for the rank of the positive semidefinite matrix Z, to the objective function in (3.1) to obtain the following hybrid optimization problem:

$$\begin{aligned} \min_{Z \in \mathbb{R}^{d\times d}} \ & \frac{1}{2}\sum_{t=1}^{n}\left(Y_{i_t,j_t} - Z^{12}_{i_t,j_t}\right)^2 + \lambda \, \|\mathrm{diag}(Z)\|_\infty + \mu\langle I, \ Z\rangle, \\ \text{subject to } \ & \|Z^{12}\|_\infty \le \alpha, \ Z \succeq 0, \end{aligned} \tag{3.2}$$

where $\mu \ge 0$ is a tuning parameter. Note that the estimator in Cai and Zhou [7] is constructed by solving a special case of this problem by setting $\mu = 0$.

Remark 3.1 The problem (3.2) is equivalent to the problem (2.3). To see this, by Lemma 1 of Fazel et al. [15], there exists an SDP formulation of the trace-norm such that $\|M\|_* \le t$ if and only if there exist matrices $Z^{11} \in \mathbb{R}^{d_1\times d_1}$, $Z^{12} \in \mathbb{R}^{d_1\times d_2}$ and $Z^{22} \in \mathbb{R}^{d_2\times d_2}$ satisfying

$$\begin{pmatrix} Z^{11} & Z^{12} \\ (Z^{12})^T & Z^{22} \end{pmatrix} \succeq 0, \quad \text{and } \mathrm{Trace}\big(Z^{11}\big) + \mathrm{Trace}\big(Z^{22}\big) \leq 2t.$$

The optimization problem (3.2) is computationally challenging. Directly solving the problem by generic interior-point method based SDP solvers is not computationally scalable. This is because the problem of interest is often of high dimensions, and the ℓ_∞-norm constraint in (3.2) induces a large number of constraints in the SDP. In addition, the feasible set is in a very complex form as it involves both the positive semidefinite and ℓ_∞-norm constraints. Although gradient projection methods are the most straightforward methods to use, the complicated feasible set also makes them difficult to be applied. This is because applying such a method requires projecting the intermediate solution to the feasible set, but it is unclear how to efficiently compute the projection.

To solve the problem efficiently, we consider an equivalent form of (3.2) below. As we shall see immediately, this formulation is crucial for efficiently solving the problem:

$$\min_{X,Z} \ \mathcal{L}(Z) + \mu\langle I,\ X\rangle,\ \text{subject to } X \succeq 0,\ Z \in \mathcal{P},\ X - Z = 0, \tag{3.3}$$

where the function $\mathcal{L}(Z)$ and the set $\mathcal{P}$ are defined as follows:

$$\mathcal{L}(Z) = \frac{1}{2}\sum_{t=1}^{n}\left(Y_{i_t,j_t} - Z^{12}_{i_t,j_t}\right)^2 + \lambda\,\|\mathrm{diag}(Z)\|_\infty, \quad \mathcal{P} = \{Z \in \mathcal{S}^d : \|Z^{12}\|_\infty \leq \alpha\}, \tag{3.4}$$

and $\mathbb{S}^d$ denotes the set of symmetric matrices in $\mathbb{R}^{d\times d}$.

Intuitively, the advantage of formulating the problem (3.2) into the form of (3.3) is that we "split" the complicated feasible set of (3.3) into two parts. In particular, X and Z in (3.3) enforce the positive semidefinite constraint and the ℓ_∞-norm constraints, respectively. The motivation of this splitting is that though projection onto the feasible set of (3.2), which contains both the semidefinite and ℓ_∞-norm constrains, is difficult, we can efficiently compute the projection onto the positive semidefinite set or the ℓ_∞-constraint set individually. As a result, adopting an alternating direction approach, in each step, we only need to project X onto the positive semidefinite cone, and control the ℓ_∞-norm of Z. Meanwhile, we impose an additional constraint $X - Z = 0$ to ensure the feasibility of both X and Z to the problem (3.2).

To solve (3.3), we consider the augmented Lagrangian function of (3.3) defined by

$$L(X, Z; W) = \mathcal{L}(Z) + \mu\langle I,\ X\rangle + \langle W, X - Z\rangle + \frac{\rho}{2}\|X - Z\|_F^2,\ X \in \mathbb{S}^d_+,\ Z \in \mathcal{P},$$

where W is the dual variable, and $\mathbb{S}^d_+ = \{A \in \mathcal{S}^d : A \succeq 0\}$ is the positive semidefinite cone.

Then, we apply the ADMM algorithm to solve the problem (3.3). The algorithm runs iteratively, at the t-th iteration, we update $(X, Z; W)$ by

$$\begin{aligned} X^{t+1} &= \operatorname*{argmin}_{X\in\mathbb{S}^d_+} L(X, Z^t; W^t) = \Pi_{\mathbb{S}^d_+}\left\{Z^t - \rho^{-1}(W^t + \mu I)\right\}, \\ Z^{t+1} &= \operatorname*{argmin}_{Z\in\mathcal{P}} L(X^{t+1}, Z; W^t) = \operatorname*{argmin}_{Z\in\mathcal{P}} \mathcal{L}(Z) + \frac{\rho}{2}\|Z - X^{t+1} - \rho^{-1}W^t\|_F^2, \\ W^{t+1} &= W^t + \tau\rho(X^{t+1} - Z^{t+1}), \end{aligned} \tag{3.5}$$

where $\tau \in (0, (1+\sqrt{5})/2)$ is a step-length parameter which is typically chosen to be 1.618. Here, $\Pi_{\mathbb{S}^d_+}(A)$ denotes the projection of the matrix $A \in \mathbb{S}^d$ onto the semidefinite cone $\mathbb{S}^d_+$. The worst-case $\mathcal{O}(t^{-1})$ rate of convergence of ADMM method is shown, for example, in Fang et al. [14].

3.2 Solving subproblems

For fast implementations of the algorithm (3.5), it is important to solve the X- and Z-subproblems of (3.5) efficiently. For the X-subproblem, we have that the minimizer is obtained by truncating all the negative eigenvalues of the matrix $Z^t - \rho^{-1}W^t$ to 0's by Eckart–Young Theorem [38]. Moreover, the following proposition provides a solution to the Z-subproblem in (3.5), which can be computed efficiently.

Proposition 3.2 *Let $\Omega = \{(i_t, j_t)\}_{t=1}^n$ be the index set of observed entries in M^0. For a given matrix $C \in \mathbb{R}^{d\times d}$, we have*

$$\mathcal{Z}(C) = \operatorname*{argmin}_{Z\in\mathcal{P}} \mathcal{L}(Z) + \frac{\rho}{2}\|Z - C\|_F^2,$$

where

$$\begin{aligned} \mathcal{Z}(C) &= \begin{pmatrix} \mathcal{Z}^{11}(C) & \mathcal{Z}^{12}(C) \\ \mathcal{Z}^{12}(C)^T & \mathcal{Z}^{22}(C) \end{pmatrix}, \\ \mathcal{Z}^{12}_{k\ell}(C) &= \begin{cases} \Pi_{[-\alpha,\alpha]}\left(\frac{Y_{k\ell}+\rho C^{12}_{k\ell}}{1+\rho}\right), & \text{if } (k,\ell)\in\Omega, \\ \Pi_{[-\alpha,\alpha]}(C^{12}_{k\ell}), & \text{otherwise,} \end{cases} \\ \mathcal{Z}^{11}_{k\ell}(C) &= C^{11}_{k\ell} \quad \text{if } k\neq\ell, \quad \mathcal{Z}^{22}_{k\ell}(C) = C^{22}_{k\ell} \quad \text{if } k\neq\ell, \\ \operatorname{diag}\left\{\mathcal{Z}(C)\right\} &= \operatorname*{argmin}_{z\in\mathbb{R}^d} \lambda\|z\|_\infty + \frac{\rho}{2}\|\operatorname{diag}(C) - z\|_2^2, \end{aligned} \tag{3.6}$$

and $\Pi_{[a,b]}(x) = \min\{b, \max(a, x)\}$ projects $x \in \mathbb{R}$ to the interval $[a, b]$.

Proof By the definition of $\mathcal{L}(Z)$ in (3.2), we have

$$\mathcal{Z}(C) = \operatorname*{argmin}_{Z\in\mathcal{P}} \frac{1}{2}\sum_{t=1}^n \left(Z^{12}_{i_t,j_t} - Y_{i_t,j_t}\right)^2 + \lambda\|\operatorname{diag}(Z)\|_\infty^2 + \frac{\rho}{2}\|Z - C\|_F^2.$$

This optimization problem is equivalent to

$$\begin{aligned}\min_{Z^{12}} \Big\{ \frac{1}{2} \sum_{t=1}^{n} \Big(Z^{12}_{i_t,j_t} - Y_{i_t,j_t} \Big)^2 + \rho \|Z^{12} - C^{12}\|_F^2 : \|Z^{12}\|_\infty \le \alpha \Big\} \\ + \min_{Z^{11}} \Big\{ \frac{\rho}{2} \|Z^{11}_{k\ell} - C^{11}_{k\ell}\|_F^2 : k \ne \ell \Big\} + \min_{Z^{22}} \Big\{ \frac{\rho}{2} \|Z^{22}_{k\ell} - C^{11}_{k\ell}\|_F^2 : k \ne \ell \Big\} \\ + \min_{\mathrm{diag}(Z)} \Big\{ \lambda \|\mathrm{diag}(Z)\|_\infty + \frac{\rho}{2} \|\mathrm{diag}(C) - \mathrm{diag}(Z)\|_2^2 \Big\}.\end{aligned} \tag{3.7}$$

For the first term of the above optimization problem, utilizing its separable structure, it is equivalent to

$$\sum_{(j,k)\in S} \min_{|Z^{12}_{k\ell}| \le \alpha} \Big\{ \frac{1}{2} \big(Z^{12}_{k\ell} - Y_{k\ell} \big)^2 + \rho \big(Z^{12}_{k\ell} - C^{12}_{k\ell} \big)^2 \Big\} + \rho \sum_{(j,k)\notin S} \min_{|Z^{12}_{k\ell}| \le \alpha} \big(Z^{12}_{k\ell} - C^{12}_{k\ell} \big)^2,$$

from which we see that its minimizer is given by $\mathcal{Z}^{12}(C)$.

In addition, the optimality of $\mathcal{Z}^{11}(C)$ and $\mathcal{Z}^{22}(C)$ are obtained by considering the remaining terms of (3.7), which concludes the proof. □

Note that in (3.7), we need to solve the following optimization problem

$$\min_{z\in\mathbb{R}^d} \beta \|z\|_\infty + \frac{1}{2} \|c - z\|_2^2, \tag{3.8}$$

where $c = (c_1, \ldots, c_d)^T = \mathrm{diag}(C)$ and $\beta = \lambda/\rho$. A direct approach to solve this problem is to reformulate it into a linearly constrained quadratic programming problem. In the next lemma, we show that it actually admits a closed-form solution. For ease of presentation, we assume without loss of generality that $c_1 \ge c_2 \ge \ldots \ge c_d \ge 0$.

Lemma 3.3 *Suppose that $c_1 \ge c_2 \ge \cdots \ge c_d \ge 0$. The solution to the optimization problem (3.8) is of the form*

$$z^* = (t^*, \ldots, t^*, c_{k^*+1}, \ldots, c_d)^T,$$

where $t^* = \frac{1}{k^*} \sum_{i=1}^{k^*} (c_i - \beta)$ and k^* is the index such that $c_{k^*+1} < \frac{1}{k^*} \Big(\sum_{i=1}^{k^*} c_i - \beta \Big) \le c_{k^*}$. If no such k^* exists, then $z^* = (t^*, \ldots, t^*)^T$, where $t^* = \frac{1}{d} \sum_{i=1}^{d} (c_i - \beta)$.

Proof Let $z = (z_1, \ldots, z_d)^T$. By the assumption that $c_1 \ge c_2 \ge \ldots \ge c_d \ge 0$, one can prove by contradiction that the optimal solution to (3.8) must satisfy the property that $z_1 \ge z_2 \ge \cdots \ge z_d \ge 0$. It is clear that (3.8) is equivalent to the following convex minimization problem:

$$\min_{z,t} \Big\{ \beta t + \frac{1}{2} \|c - z\|^2 : z_i \ge 0,\ z_i - t \le 0,\ i = 1, \ldots, d \Big\}, \tag{3.9}$$

whose KKT optimality conditions are given by

$$
\begin{aligned}
&z - c - \mu + \hat{\mu} = 0,\\
&\beta - \textstyle\sum_{j=1}^{d} \hat{\mu}_j = 0,\\
&\mu_j \geq 0,\ z_j \geq 0,\ \mu_j z_j = 0,\ j = 1, \ldots, d,\\
&\hat{\mu}_j \geq 0,\ z_j - t \leq 0,\ \hat{\mu}_j (z_j - t) = 0,\ j = 1, \ldots, d.
\end{aligned}
$$

Define

$$
t_k = \frac{1}{k} \left(\sum_{i=1}^{k} c_i - \beta \right), \quad k = 1, \ldots, d.
$$

Let k^* be the index such that $c_{k^*+1} < t_{k^*} \leq c_{k^*}$. If no such k^* exists, i.e., $c_i < t_i$ for all $i = 1, \ldots, d$, then set $k^* = d$. Now one can verify that the point $(z^*, t^*, \mu^*, \hat{\mu}^*)$ defined below satisfies the KKT conditions:

$$
\mu^* = 0,\ \ t^* = t_{k^*},\ \ z_i^* = \begin{cases} t^* & \text{for } i = 1, \ldots, k^*, \\ c_i & \text{for } i = k^* + 1, \ldots, d, \end{cases} \qquad \hat{\mu}_i^* = \begin{cases} c_i - t^* & \text{for } i = 1, \ldots, k^*, \\ 0 & \text{for } i = k^* + 1, \ldots, d. \end{cases}
$$

Hence z^* is the optimal solution to (3.8). This completes the proof. □

Remark 3.4 We avoid presenting the general case of $c = (c_1, \ldots, c_d)^T$ for simplicity. The solution in the general case can be derived similarly, and we implement the algorithm for the general case in later numerical studies.

The algorithm for solving problem (3.2) is summarized in Algorithm 1.

Algorithm 1 Solving max-norm optimization problem (3.2) by the ADMM

Input: $X^0, Z^0, W^0, Y_\Omega, \lambda, \mu, \alpha, \rho, \tau, t = 0$.
while Stopping criterion is not satisfied. **do**
Update $X^{t+1} \leftarrow \Pi_{\mathcal{S}_+^d} \{ Z^t - \rho^{-1}(W^t + \mu I) \}$.
Update $Z^{t+1} \leftarrow \mathcal{Z}(X^{t+1} + \rho^{-1} W^t)$ by (3.6).
Update $W^{t+1} \leftarrow W^t + \tau\rho(X^{t+1} - Z^{t+1})$.
$t \leftarrow t + 1$.
end while
Output: $\hat{Z} = Z^t$, $\hat{M} = \hat{Z}^{12} \in \mathbb{R}^{d_1 \times d_2}$.

Remark 3.5 Taking a closer look at Algorithm 1, we see that the equivalent reformulation (3.3) of the original problem (3.2) brings us computational efficiency. In particular, all sub-problems can be solved efficiently. Among them, the most computationally expensive step is the X-update step as we need to compute an eigenvalue decomposition of the matrix $Z^t - \rho^{-1} W^t$, which has the complexity of $\mathcal{O}(d^3)$. Meanwhile, we point out that if a rank-r solution to the X-subproblem is desired, the computational complexity can be reduced to $\mathcal{O}(rd^2)$.

Remark 3.6 Note that if the user requires an exact low rank solution, solving the X-subproblem can be further accelerated. In particular, we can apply the Eckart–Young Theorem and project the solution onto the nearest face for the target rank, see e.g., Oliveira et al. [28] where this idea is applied to the SDP relaxation of the quadratic assignment problem with nonnegativity constraints added.

In addition to the algorithm for the regularized max-norm minimization problem (2.2), we also provide the algorithm for solving the constrained version (2.1) in Appendix 7. We focus our discussions on the regularized version since it is computationally more challenging.

3.3 Stopping conditions

In this section, we discuss the stopping conditions for Algorithm 1. Denote by $\delta_{\mathcal{C}}(\cdot)$ the indicator function over a given set $\mathcal{C}$ such that $\delta_{\mathcal{C}}(x) = 0$ if $x \in \mathcal{C}$ and $\delta_{\mathcal{C}}(x) = \infty$ if $x \notin \mathcal{C}$. The optimality conditions for (3.3) are given as follows:

$$0 \in \partial \delta_{\mathcal{S}^d_+}(X) + \mu I + W, \quad 0 \in \partial \delta_{\mathcal{P}}(Z) + \nabla \mathcal{L}(Z) - W, \quad X - Z = 0, \tag{3.10}$$

where W is the Lagrangian multiplier associated with the equality constraint $X - Z = 0$. Here $\partial \delta_{\mathcal{S}^d_+}(X)$ denotes the subdifferential of $\delta_{\mathcal{S}^d_+}(\cdot)$ at X; similarly for $\partial \delta_{\mathcal{P}}(Z)$.

By the optimality conditions of X^{t+1} and Z^{t+1} in (3.5), we have that

$$0 \in \partial \delta_{\mathcal{S}^d_+}\left(X^{t+1}\right) + \rho\left\{X^{t+1} - Z^t + \rho^{-1}\left(W^t + \mu I\right)\right\}$$

if and only if

$$\rho\left(Z^t - Z^{t+1}\right) + W^{t+1} - \widetilde{W}^{t+1} \in \partial \delta_{\mathcal{S}^d_+}\left(X^{t+1}\right) + \mu I + W^{t+1},$$

and

$$0 \in \partial \delta_{\mathcal{P}}\left(Z^{t+1}\right) + \nabla \mathcal{L}\left(Z^{t+1}\right) + \rho\left(Z^{t+1} - X^{t+1} - \rho^{-1} W^t\right)$$

if and only if

$$\widetilde{W}^{t+1} - W^{t+1} \in \partial \delta_{\mathcal{P}}\left(Z^{t+1}\right) + \nabla \mathcal{L}\left(Z^{t+1}\right) - W^{t+1},$$

where $\widetilde{W}^{t+1} = W^t + \rho(X^{t+1} - Z^{t+1})$. Observe that the iterate $(X^{t+1}, Z^{t+1}, W^{t+1})$ generated from Algorithm 1 is an accurate approximate optimal solution to (3.5) if the residual

$$\eta^{t+1} := \max\left\{R_P^{t+1}, R_D^{t+1}\right\}$$

is small, where

$$R_P^{t+1} = \|X^{t+1} - Z^{t+1}\|_F, \quad R_D^{t+1} = \max\big\{\|\rho(Z^t - Z^{t+1}) + W^{t+1} - \widetilde{W}^{t+1}\|_F,\ \|W^{t+1} - \widetilde{W}^{t+1}\|_F\big\},$$

denote the primal and dual residuals. In the practical implementation, we let the algorithm stop when $\eta^{t+1} \le 10^{-4}$ or when the number of iterations exceeds 200.

3.4 Practical implementations

We should mention that tuning the parameter ρ properly in the ADMM method is critical for the method to converge at a reasonable rate. In our implementation, starting with the initial value of 0.1 for ρ, we adaptively tune the parameter at every tenth iterations based on the following criterion:

$$\begin{cases} \text{set } \rho \leftarrow 0.7\rho & \text{if } \|R_P^{t+1}\| < 0.5\|R_D^{t+1}\|, \\ \text{set } \rho \leftarrow 1.3\rho & \text{if } \|R_D^{t+1}\| < 0.5\|R_P^{t+1}\|. \end{cases}$$

The basic idea is to balance the progress of $\|R_P^{t+1}\|$ and $\|R_D^{t+1}\|$ so that the stopping criterion $\eta^{t+1} \le 10^{-4}$ can be attained within a small number of iterations.

Another important computational issue which we need to address is to cut down the cost of computing the full eigenvalue decomposition in the X-update step in Algorithm 1. Given a matrix $G \in \mathcal{S}^d$, we observe that to compute the projection $\Pi_{\mathcal{S}_+^d}(G)$, we need only the eigen-pairs corresponding to the positive eigenvalues of G. Thus in our implementation, we use the LAPACK subroutine `dsyevx.f` to compute only a partial eigenvalue decomposition of G if we know that the number of positive eigenvalues of G is substantially smaller than d, say less than 10% of d. Such a partial eigenvalue decomposition is typically cheaper than a full eigenvalue decomposition when the number of eigenvalues of interest is much smaller than the dimension d. For Algorithm 1, at the $(t+1)$-th iteration, we estimate the potential number of positive eigenvalues of $G^t := Z^t - \rho^{-1}(W^t + \mu I)$ (and hence the rank of X^{t+1}) based on the rank of the previously computed iterate X^t. Such an estimation is usually accurate when the sequence of iterates $\{(X^t, Y^t, Z^t)\}$ starts to converge. During the initial phase of Algorithm 1, we do not have a good estimate on the rank of X^{t+1}, and we compute the projection based on the full eigenvalue decomposition of G^t.

To further reduce the cost of computing X^{t+1} in Algorithm 1, we employ a heuristic strategy to truncate the small positive eigenvalues of G^t to 0's. That is, if there is a group of positive eigenvalues of G^t with magnitudes which are significantly larger than the remaining positive eigenvalues, we compute X^{t+1} using only the eigen-pairs corresponding to the large positive eigenvalues of G^t. Such a strategy can significantly reduce the cost of computing X^{t+1} since the number of large eigenvalues of G^t is typically small in a low-rank matrix completion problem. A surprising bonus of adopting such a cost cutting heuristic is that the recovery error can actually become 30–50% smaller, despite the fact that the computed X^{t+1} now is only an approximate solution of the X-update subproblem in Algorithm 1. One possible explanation for such a phenomenon is that the truncation of small positive eigenvalues of X^{t+1} to 0's

actually has a debiasing effect to eliminate the attenuation of the singular values of the recovered matrix due to the presence of the convex regularization term. In the case of compressed sensing, such a debiasing effect has been explained in Figueiredo et al. [16].

4 Theoretical properties

In this section, we provide theoretical guarantees for the hybrid estimator (2.3). To facilitate our discussions, we introduce the following notations. Let $X_1, \ldots, X_n$ be i.i.d. copies of a random matrix X with distribution $\Pi = (\pi_{k\ell})_{k\in[d_1],\ell\in[d_2]}$ on the set $\mathcal{X} = \left\{e_k(d_1)e_\ell(d_2)^T, k = 1, \ldots, d_1, \ell = 1, \ldots, d_2\right\}$, i.e., $\mathbb{P}\{X = e_k(d_1)e_\ell(d_2)^T\} = \pi_{k\ell}$, where $e_k(d)$ are the canonical basis vectors in $\mathbb{R}^d$. By definition,

$$\|A\|_{L_2(\Pi)}^2 = \mathbb{E}\langle A, X\rangle^2 = \sum_{k=1}^{d_1}\sum_{\ell=1}^{d_2}\pi_{k\ell}A_{k\ell}^2 \tag{4.1}$$

for all matrices $A = (A_{k\ell})_{1\le k\le d_1, 1\le\ell\le d_2} \in \mathbb{R}^{d_1\times d_2}$. Moreover, let

$$\pi_{k\cdot} = \sum_{\ell=1}^{d_2}\pi_{k\ell} \quad \text{and} \quad \pi_{\cdot\ell} = \sum_{k=1}^{d_1}\pi_{k\ell}$$

be, respectively, the probabilities of observing an element from the k-th row and the ℓ-th column.

Considering the exact low-rank matrix recovery, i.e., $\text{rank}(M^0) \le r_0$, the first part of the next theorem shows that the estimator (2.3) achieves a fast rate of convergence under some "ideal" situations, and the second part indicates that it is also robust against non-uniform sampling schemes. For ease of presentation, we conduct the analysis by considering a constrained form of (2.3), namely,

$$\hat{M} := \underset{M\in\mathcal{K}(\alpha,R)}{\text{argmin}} \frac{1}{n}\sum_{t=1}^{n}\left(Y_{i_t,j_t} - M_{i_t,j_t}\right)^2 + \mu\|M\|_*, \tag{4.2}$$

where $\mathcal{K}(\alpha, R) = \{M \in \mathbb{R}^{d_1\times d_2} : \|M\|_\infty \le \alpha \text{ and } \|M\|_{\max} \le R\}$. Our proof partly follows the arguments in Cai and Zhou [7]. The major technical challenge here is to carefully balance the tuning parameters R and μ in (4.2) to achieve the desired recovery results for both the uniform and non-uniform sampling schemes.

Theorem 4.1 *Assume that* $\|M^0\|_\infty \le \alpha$, $\text{rank}(M^0) \le r_0$ *and that* $\xi_1, \ldots, \xi_n$ *are i.i.d.* $N(0, 1)$ *random variables. The sampling distribution* Π *is such that* $\min_{(k,\ell)\in[d_1]\times[d_2]}\pi_{k\ell} \ge (\nu d_1 d_2)^{-1}$ *for some* $\nu \ge 1$. *Choose* $R \ge \alpha r_0^{1/2}$ *in* (4.2) *and write* $d = d_1 + d_2$.

(i) Let $\mu = c_1(dn)^{-1/2}$ for some constant $c_1 > 0$. Then, for a sample size $2 < n \le d_1 d_2$, the estimator $\hat{M}$ given at (4.2) satisfies

$$\frac{1}{d_1 d_2}\|\hat{M} - M^0\|_F^2 \lesssim \max\left\{\nu^2 \frac{r_0 d}{n} + \nu(\alpha \vee \sigma)\sqrt{\frac{R^2 d}{n}},\ \nu\alpha^2 \frac{\log d}{n}\right\} \tag{4.3}$$

with probability at least $1 - 3d^{-1}$.

(ii) Let $\mu = c_2\,\sigma \max_{(k,\ell)\in[d_1]\times[d_2]}(\pi_{k\cdot} \vee \pi_{\cdot\ell})^{1/2}$ for some constant $c_2 > 0$. Then, for a sample size $2 < n \le d_1 d_2$, the estimator $\hat{M}$ given at (4.2) satisfies

$$\begin{aligned}&\frac{1}{d_1 d_2}\|\hat{M} - M^0\|_F^2 \\ &\lesssim \max\left\{\nu^2(\alpha \vee \sigma)^2 \max_{k,\ell}(\pi_{k\cdot} \vee \pi_{\cdot\ell})\frac{r_0 d_1 d_2 \log d}{n},\ \nu\alpha^2\sqrt{\frac{\log d}{n}}\right\}\end{aligned} \tag{4.4}$$

with probability at least $1 - 3d^{-1}$.

Proof of Theorem 4.1 Recall that $Y_{i_t,j_t} = M^0_{i_t,j_t} + \sigma\xi_t = \langle X_t,\ M^0\rangle + \sigma\xi_t$ for $t = 1, \ldots, n$. By the optimality of $\hat{M}$ in (4.2), we have that

$$\frac{1}{n}\sum_{i=1}^{n}\langle X_i, \hat{M} - M^0\rangle^2 \le \frac{2\sigma}{n}\sum_{i=1}^{n}\xi_i\langle X_i, \hat{M} - M^0\rangle + \mu(\|M^0\|_* - \|\hat{M}\|_*). \tag{4.5}$$

For each matrix $A \in \mathbb{R}^{d_1\times d_2}$, denote by $u_j(A)$ and $v_j(A)$ the left and right orthonormal singular vectors of A, i.e., $A = \sum_{j=1}^{r}\sigma_j(A)u_j(A)v_j^T(A)$, where $r = \mathrm{rank}(A)$ and $\sigma_1(A) \ge \cdots \ge \sigma_r(A) > 0$ are the singular values of A. Let $S_1(A)$ and $S_2(A)$ be, respectively, the linear span of $\{u_j(A)\}$ and $\{v_j(A)\}$. Consequently, following the proof of Theorem 3 in Klopp [20] we have

$$\|M^0\|_* - \|\hat{M}\|_* \le \|\mathcal{P}_{M^0}(\hat{M} - M^0)\|_* - \|\mathcal{P}^{\perp}_{M^0}(\hat{M} - M^0)\|_*, \tag{4.6}$$

where $\mathcal{P}_A(B) = P_{S_1(A)}B + P^{\perp}_{S_1(A)}BP_{S_2(A)}$, $\mathcal{P}^{\perp}_A(B) = P^{\perp}_{S_1(A)}BP^{\perp}_{S_2(A)}$, and P_S denotes the projector onto the linear subspace S.

(i) Looking at the inequality (4.5), it follows from (6.6) in Cai and Zhou [7] that with probability greater than $1 - d^{-1}$,

$$\frac{1}{n}\sum_{i=1}^{n}\langle X_i, \hat{M} - M^0\rangle^2 \le 24\sigma\sqrt{\frac{R^2 d}{n}} + \mu\sqrt{2r_0}\|\hat{M} - M^0\|,$$

which, together with (6.13) of Cai and Zhou [7], implies that with probability at least $1 - 3d^{-1}$,

$$\begin{aligned}
&\frac{1}{2\nu d_1 d_2}\|\hat{M} - M^0\|_F^2 \le \frac{1}{2}\|\hat{M} - M^0\|_{L_2(\Pi)}^2 \\
&\quad \le \max\left\{\mu\sqrt{2r_0}\|\hat{M} - M^0\|_F + C_1(\alpha+\sigma)\sqrt{\frac{R^2 d}{n}},\ C_2\alpha^2\frac{\log d}{n}\right\} \\
&\quad \le \max\left\{\frac{1}{4\nu d_1 d_2}\|\hat{M} - M^0\|_F^2 + 2\nu r_0 d_1 d_2\mu^2 + C_1(\alpha+\sigma)\sqrt{\frac{R^2 d}{n}},\ C_2\alpha^2\frac{\log d}{n}\right\},
\end{aligned}$$

where $C_5, C_6 > 0$ are absolute constants. This proves (4.3) by rearranging the constants.

(ii) First we assume that the regularization parameter μ satisfies $\mu \ge 3\|\Sigma_\xi\|$, where $\Sigma_\xi := n^{-1}\sum_{i=1}^n \xi_i X_i \in \mathbb{R}^{d_1\times d_2}$. By (4.6) and the inequality $|\langle A, B\rangle| \le \|A\| \cdot \|B\|_*$ which holds for all matrices A and B where $\|A\|$ is the spectral norm, the right-hand side of (4.5) is bounded by

$$\begin{aligned}
&\left(2\sigma\|\Sigma_\xi\| + \mu\right)\|\mathcal{P}_{M^0}\left(\hat{M} - M^0\right)\|_* + \left(2\sigma\|\Sigma_\xi\| - \mu\right)\|\mathcal{P}_{M^0}^\perp\left(\hat{M} - M^0\right)\|_* \\
&\quad \le \frac{5}{3}\mu\|\mathcal{P}_{M^0}\left(\hat{M} - M^0\right)\|_* \le \frac{5}{3}\mu\sqrt{2r_0}\|\hat{M} - M^0\|_F
\end{aligned} \tag{4.7}$$

whenever $\mu \ge 3\sigma\|\Sigma_\xi\|$, where $r_0 = \operatorname{rank}(M^0)$.

Let $\varepsilon_1, \ldots, \varepsilon_n$ be i.i.d. Rademacher random variables. Then, it follows from Lemmas 12 and 13 in Klopp [20] that with probability greater than $1 - 2d^{-1}$,

$$\begin{aligned}
&\frac{1}{2\nu d_1 d_2}\|\hat{M} - M^0\|_F^2 \\
&\quad \le \frac{1}{2}\|\hat{M} - M^0\|_{L_2(\Pi)}^2 \\
&\quad \le \max\left\{\frac{5}{3}\mu\sqrt{2r_0}\|\hat{M} - M^0\|_F + C_3\nu\alpha^2 r_0 d_1 d_2(\mathbb{E}\|\Sigma_\varepsilon\|)^2,\ C_4\alpha^2\sqrt{\frac{\log d}{n}}\right\} \\
&\quad \le \max\left[\frac{1}{4\nu d_1 d_2}\|\hat{M} - M^0\|_F^2 + \nu r_0 d_1 d_2\left\{6\mu^2 + C_3\alpha^2(\mathbb{E}\|\Sigma_\varepsilon\|)^2\right\},\ C_4\alpha^2\sqrt{\frac{\log d}{n}}\right],
\end{aligned}$$

where $C_3, C_4 > 0$ are absolute constants and $\Sigma_\varepsilon := n^{-1}\sum_{i=1}^n \varepsilon_i X_i$.

It remains to consider the quantities $\|\Sigma_\xi\|$ and $\mathbb{E}\|\Sigma_\varepsilon\|$. For $\|\Sigma_\xi\|$ with Gaussian multipliers $\xi_1, \ldots, \xi_n$, applying Lemma 5 in Klopp [20] yields that, for every $n > 0$,

$$\|\Sigma_\xi\| \le C_5 \max_{(k,\ell)\in[d_1]\times[d_2]}(\pi_{k\cdot} \vee \pi_{\cdot\ell})^{1/2}\sqrt{\frac{\log d}{n}} + C_6\frac{(\log d)^{3/2}}{n}$$

holds with probability at least $1 - d^{-1}$, where $C_5, C_6 > 0$ are absolute constants. Furthermore, by Corollary 8.2 in Mackey et al. [25],

$$\mathbb{E}\|\Sigma_\varepsilon\| \le \max_{(k,\ell)\in[d_1]\times[d_2]} (\pi_{k\cdot} \vee \pi_{\cdot\ell})^{1/2} \sqrt{\frac{3\log d}{n}} + \frac{\log d}{n}.$$

Together, the previous three displays prove (4.4). □

Remark 4.1 It is known that both the trace-norm and the max-norm serve as semidefinite relaxations of the rank. In the context of approximately low-rank matrix reconstruction, we consider two types of convex relaxations for low-rankness. For any $\alpha, R > 0$, define the matrix classes

$$\mathcal{M}_{\max}(\alpha, R) = \left\{ M \in \mathbb{R}^{d_1\times d_2} : \|M\|_\infty \le \alpha, \|M\|_{\max} \le R \right\} \tag{4.8}$$

and

$$\mathcal{M}_{\mathrm{tr}}(\alpha, R) = \left\{ M \in \mathbb{R}^{d_1\times d_2} : \|M\|_\infty \le \alpha, (d_1 d_2)^{-1/2}\|M\|_* \le R \right\}. \tag{4.9}$$

For any integer $1 \le r \le \min(d_1, d_2)$, set $\mathcal{M}_{\mathrm{r}}(\alpha, r) = \{M \in \mathbb{R}^{d_1\times d_2} : \|M\|_\infty \le \alpha, \mathrm{rank}\,(M) \le r\}$ and note that $\mathcal{M}_{\mathrm{r}}(\alpha, r) \subset \mathcal{M}_{\max}(\alpha, \alpha r^{1/2}) \subset \mathcal{M}_{\mathrm{tr}}(\alpha, \alpha r^{1/2})$. The following results [7] provide recovery guarantees for approximately low-rank matrix completion in the sense that the target matrix M^0 either belongs to $\mathcal{M}_{\max}(\alpha, R)$ or $\mathcal{M}_{\mathrm{tr}}(\alpha, R)$ for some $\alpha, R > 0$. As before, set $d = d_1 + d_2$.

(i) Assume that $M^0 \in \mathcal{M}_{\max}(\alpha, R)$ and $\xi_1, \ldots, \xi_n$ are i.i.d. $N(0, 1)$ random variables. Then, for a sample size $2 < n \le d_1 d_2$, the max-norm constrained least squares estimator $\hat{M}_{\max} := \mathrm{argmin}_{M\in\mathcal{K}(\alpha,R)} \frac{1}{n}\sum_{t=1}^n \left(Y_{i_t,j_t} - M_{i_t,j_t}\right)^2$ satisfies

$$\|\hat{M}_{\max} - M^0\|_{L_2(\Pi)}^2 \lesssim (\alpha \vee \sigma)\sqrt{\frac{R^2 d}{n}} + \frac{\alpha^2 \log d}{n}$$

with probability at least $1 - 3d^{-1}$.

(ii) Assume that $M^0 \in \mathcal{M}_{\mathrm{tr}}(\alpha, R)$, $\xi_1, \ldots, \xi_n$ are i.i.d. $N(0, 1)$ random variables and that the sampling distribution Π is uniform on $\mathcal{X}$. Then, for a sample size $2 < n \le d_1 d_2$, the trace-norm penalized estimator $\hat{M}_{\mathrm{tr}} := \mathrm{argmin}_{M:\|M\|_\infty\le\alpha} \frac{1}{n}\sum_{t=1}^n \left(Y_{i_t,j_t} - M_{i_t,j_t}\right)^2 + \mu\|M\|_*$ with $\mu \asymp \sigma\sqrt{\frac{\log d}{dn}}$ satisfies

$$\frac{1}{d_1 d_2}\|\hat{M}_{\mathrm{tr}} - M^0\|_F^2 \lesssim (\alpha \vee \sigma)\sqrt{\frac{R^2 d \log d}{n}} + \frac{\alpha^2 \log d}{n}$$

with probability at least $1 - 3d^{-1}$.

Remark 4.2 When the underlying matrix has exactly rank r, i.e., $M^0 \in \mathcal{M}_r(\alpha, r)$, it is known that using the trace-norm regularized approach leads to a mean square error of order $\mathcal{O}(n^{-1}rd\log d)$. Under the uniform sampling scheme, the trace-norm regularized method is the most preferable one as it achieves optimal rate of convergence (up to a logarithmic factor) and is computationally attractive, although from a practical point of view, the uniform sampling assumption is controversial. In comparison with the result in Cai and Zhou [7], which is suboptimal under the uniform sampling scheme and exact low-rank setting, here we established near optimal recovery results (up to a logarithmic factor) under such a setting, and we can still guarantee recoveries under non-uniform sampling schemes.

An important message we wish to convey is that, when learning in a non-uniform world, the underlying sampling distribution also contributes to the recovery guarantee. More specifically, Part (ii) of Theorem 4.1 sheds light on how the sampling distribution affects the recovery error bound. The optimal rate of convergence in the class of low-rank matrices is also achieved by $\hat{M}$ when the sampling scheme is uniform. From (4.3) and (4.4), we see that the actual performance of the hybrid estimator $\hat{M}$ depends heavily on the sampling distribution and so is the optimal choice of the regularization parameter μ.

5 Numerical experiments

We compare the nuclear-norm, max-norm and hybrid regularizers for matrix completion on an iMac with Intel i5 Processor at 2.7 GHz with 16 GB memory. We test different methods on simulated and real datasets. All the tuning parameters are chosen by data splitting.

5.1 Simulated datasets

We first test the methods on simulated data, where we consider three sampling schemes. In the first scheme, the indices of observed entries are uniformly sampled without replacement, while in the other two schemes, the indices are sampled non-uniformly. Specifically, in all three schemes, we let the target matrix M^0 be generated by $M^0 = M_L M_R^T$, where M_L and M_R are two $d_t \times r$ matrices, and each entry is sampled independently from a standard normal distribution $N(0, 1)$. Thus, $M^0 \in \mathbb{R}^{d_t \times d_t}$ is a rank r matrix. In all three settings, as listed in Tables 1, 2 and 3. we consider different combinations of dimensions, ranks and sampling ratios (SR), where $\mathrm{SR} = n/d_t^2$. We compare the matrix recovery results using the nuclear-norm, max-norm penalized estimators and the hybrid estimator. For the nuclear-norm approach, we compute the estimator by adopting the accelerated proximal-gradient method discussed in Toh and Yun [37]. For the max-norm and hybrid approaches, we compute the estimator by solving problem (3.2) using Algorithm 3.2, where μ in (3.2) is set to 0 when we compute the max-norm penalized estimator.

In Scheme 1, we uniformly sample the entries. In Schemes 2 and 3, we conduct non-uniform sampling schemes in the following way. Denote by $\pi_{k\ell}$ the probability

Table 1 Averaged relative error and running time in seconds for different methods under uniform sampling scheme

σ	d_t	(r, SR)	Nuclear		Max		Hybrid	
			RE	Time	RE	Time	RE	Time
0	500	(5, 0.10)	1.1×10^{-3}	6.9	4.1×10^{-2}	12.0	4.0×10^{-2}	12.5
		(5, 0.15)	7.7×10^{-4}	7.0	3.8×10^{-2}	11.1	2.9×10^{-2}	13.4
		(10, 0.10)	5.5×10^{-2}	8.0	1.2×10^{-1}	11.4	2.9×10^{-2}	13.4
		(10, 0.15)	1.3×10^{-3}	8.6	3.8×10^{-2}	11.7	2.3×10^{-2}	12.0
	1000	(5, 0.10)	8.8×10^{-4}	44.8	2.9×10^{-2}	110.4	1.9×10^{-2}	115.1
		(5, 0.15)	6.6×10^{-4}	43.3	1.9×10^{-2}	111.3	1.8×10^{-2}	114.3
		(5, 0.20)	5.5×10^{-4}	44.6	1.8×10^{-2}	112.4	6.7×10^{-3}	120.0
		(10, 0.10)	1.5×10^{-3}	44.4	2.9×10^{-2}	108.7	2.0×10^{-2}	121.7
		(10, 0.15)	1.0×10^{-3}	45.8	2.0×10^{-2}	112.8	1.3×10^{-2}	117.8
		(10, 0.20)	8.3×10^{-4}	45.5	1.5×10^{-2}	110.8	8.9×10^{-3}	117.3
	1500	(5, 0.10)	8.1×10^{-4}	162.8	2.3×10^{-2}	385.4	1.2×10^{-2}	408.2
		(5, 0.15)	6.3×10^{-4}	158.3	1.7×10^{-2}	396.9	1.1×10^{-2}	406.6
		(5, 0.20)	5.3×10^{-4}	158.1	1.3×10^{-2}	410.9	5.6×10^{-3}	405.3
		(10, 0.10)	1.3×10^{-3}	165.9	2.0×10^{-2}	413.8	1.5×10^{-2}	413.3
		(10, 0.15)	9.5×10^{-4}	160.8	1.4×10^{-2}	410.1	1.3×10^{-2}	423.2
		(10, 0.20)	7.8×10^{-4}	161.0	1.2×10^{-2}	395.1	7.0×10^{-3}	398.2
0.01	500	(5, 0.10)	7.4×10^{-2}	6.4	6.4×10^{-2}	10.5	6.3×10^{-2}	12.3
		(5, 0.15)	5.4×10^{-3}	6.4	4.8×10^{-2}	11.4	4.3×10^{-2}	13.1
		(10, 0.10)	1.7×10^{-1}	6.3	5.2×10^{-2}	10.9	6.6×10^{-2}	11.9
		(10, 0.15)	7.8×10^{-2}	6.5	4.0×10^{-2}	11.2	4.8×10^{-2}	14.2
	1000	(5, 0.10)	4.8×10^{-2}	47.1	3.9×10^{-2}	101.7	3.6×10^{-2}	119.8
		(5, 0.15)	4.5×10^{-2}	47.5	2.8×10^{-2}	106.8	3.3×10^{-2}	116.6
		(5, 0.20)	4.7×10^{-2}	47.6	2.6×10^{-2}	117.3	2.6×10^{-2}	119.8
		(10, 0.10)	6.2×10^{-2}	47.1	4.3×10^{-2}	106.1	4.2×10^{-2}	116.7
		(10, 0.15)	4.9×10^{-2}	47.2	3.3×10^{-2}	105.9	3.0×10^{-2}	120.2
		(10, 0.20)	4.5×10^{-2}	47.7	2.7×10^{-2}	112.2	3.2×10^{-3}	120.3
	1500	(5, 0.10)	4.2×10^{-2}	161.2	2.9×10^{-2}	377.9	2.9×10^{-2}	406.1
		(5, 0.15)	4.1×10^{-2}	167.5	2.4×10^{-2}	408.7	2.8×10^{-2}	409.3
		(5, 0.20)	4.4×10^{-2}	153.4	2.1×10^{-2}	412.9	2.1×10^{-2}	415.6
		(10, 0.10)	5.0×10^{-3}	166.9	3.3×10^{-2}	397.2	3.3×10^{-2}	404.6
		(10, 0.15)	4.7×10^{-3}	160.8	2.6×10^{-2}	395.4	2.5×10^{-2}	424.2
		(10, 0.20)	4.3×10^{-3}	150.6	2.1×10^{-2}	401.9	2.0×10^{-2}	380.7

Under noiseless and noisy settings, for the nuclear norm approach, we set $\mu = 1 \times 10^{-4}\|Y_\Omega\|_F$ and $2 \times 10^{-4}\|Y_\Omega\|_F$. For the max-norm approach, we set $\lambda = 2\|Y_\Omega\|_F$ and $0.05\|Y_\Omega\|_F$. For the hybrid approach, and we set $\lambda = 0.01\|Y_\Omega\|_F$ and $0.8\|Y_\Omega\|_F$, $\mu = 0.02\lambda$ and $1 \times 10^{-4}\lambda$

Table 2 Averaged relative error and running time in seconds for different methods using non-uniformly sampled data as in Scheme 2

σ	d_t	(r, SR)	Nuclear		Max		Hybrid	
			RE	Time	RE	Time	RE	Time
0	500	(5, 0.10)	7.4×10^{-1}	7.6	2.2×10^{-1}	12.5	1.2×10^{-1}	15.8
		(5, 0.15)	6.1×10^{-1}	7.8	9.6×10^{-2}	13.1	6.1×10^{-2}	15.7
		(10, 0.10)	7.7×10^{-1}	7.5	2.1×10^{-1}	12.9	1.6×10^{-1}	16.1
		(10, 0.15)	6.1×10^{-1}	8.5	6.1×10^{-2}	13.0	7.6×10^{-2}	15.7
	1000	(5, 0.10)	7.4×10^{-1}	45.2	2.2×10^{-1}	97.0	1.1×10^{-1}	113.9
		(5, 0.15)	6.1×10^{-1}	48.2	1.2×10^{-1}	104.0	4.3×10^{-2}	113.1
		(5, 0.20)	6.2×10^{-1}	45.4	1.1×10^{-1}	105.6	3.5×10^{-2}	105.0
		(10, 0.10)	7.5×10^{-1}	45.8	1.9×10^{-1}	97.3	8.8×10^{-2}	113.8
		(10, 0.15)	6.0×10^{-1}	47.6	5.9×10^{-2}	105.2	4.1×10^{-2}	109.7
		(10, 0.20)	6.0×10^{-1}	44.6	6.1×10^{-2}	108.8	4.3×10^{-2}	108.2
	1500	(5, 0.10)	7.5×10^{-1}	143.2	2.3×10^{-1}	388.7	1.0×10^{-1}	372.3
		(5, 0.15)	6.0×10^{-1}	147.2	1.3×10^{-1}	398.0	6.2×10^{-2}	389.0
		(5, 0.20)	6.0×10^{-1}	138.5	1.1×10^{-1}	397.6	2.2×10^{-2}	358.8
		(10, 0.10)	7.5×10^{-1}	143.2	1.4×10^{-1}	360.0	7.4×10^{-2}	386.1
		(10, 0.15)	6.0×10^{-1}	142.3	5.9×10^{-2}	392.3	2.8×10^{-2}	380.2
		(10, 0.20)	6.0×10^{-1}	137.1	9.9×10^{-2}	395.2	2.4×10^{-2}	359.4
0.01	500	(5, 0.10)	7.4×10^{-1}	7.5	2.2×10^{-1}	15.1	1.3×10^{-1}	16.2
		(5, 0.15)	6.1×10^{-1}	8.3	1.0×10^{-1}	14.9	7.1×10^{-2}	16.2
		(10, 0.10)	7.7×10^{-1}	8.7	2.4×10^{-1}	15.5	1.7×10^{-1}	16.2
		(10, 0.15)	6.2×10^{-1}	8.3	8.0×10^{-2}	15.2	8.6×10^{-2}	16.5
	1000	(5, 0.10)	7.4×10^{-1}	44.5	2.2×10^{-1}	117.9	1.0×10^{-1}	118.2
		(5, 0.15)	6.1×10^{-1}	47.0	1.2×10^{-1}	116.9	5.2×10^{-2}	120.8
		(5, 0.20)	6.2×10^{-1}	46.7	1.1×10^{-1}	120.7	4.3×10^{-2}	123.0
		(10, 0.10)	7.5×10^{-1}	45.6	2.0×10^{-1}	117.3	9.3×10^{-2}	122.9
		(10, 0.15)	6.1×10^{-1}	47.3	6.5×10^{-2}	119.3	5.3×10^{-2}	123.3
		(10, 0.20)	6.1×10^{-1}	46.3	6.3×10^{-2}	123.2	5.0×10^{-2}	120.5
	1500	(5, 0.10)	7.5×10^{-1}	152.6	2.3×10^{-1}	395.6	7.2×10^{-2}	396.9
		(5, 0.15)	6.0×10^{-1}	156.3	1.2×10^{-1}	382.0	5.3×10^{-2}	394.2
		(5, 0.20)	6.0×10^{-1}	162.4	1.1×10^{-1}	396.3	3.0×10^{-2}	398.2
		(10, 0.10)	7.5×10^{-1}	154.5	1.4×10^{-1}	403.2	7.3×10^{-2}	406.1
		(10, 0.15)	6.0×10^{-1}	158.7	5.9×10^{-2}	396.5	4.3×10^{-2}	399.1
		(10, 0.20)	6.0×10^{-1}	157.7	9.5×10^{-2}	405.4	3.6×10^{-2}	400.3

For the nuclear norm approach, we set $\mu = 2 \times 10^{-4}\|Y_\Omega\|_F$. For the max-norm approach, we set $\lambda = 0.1\|Y_\Omega\|_F$. For the hybrid approach, and we set $\lambda = 0.2\|Y_\Omega\|_F$, $\mu = 2 \times 10^{-4}\lambda$

Table 3 Averaged relative error and running time in seconds for the nuclear-norm and max-norm penalized matrix completion using non-uniformly sampled data as in Scheme 3

σ	d_t	(r, SR)	Nuclear		Max		Hybrid	
			RE	Time	RE	Time	RE	Time
0	500	(5, 0.10)	7.4×10^{-1}	7.2	2.6×10^{-1}	14.7	1.9×10^{-1}	17.8
		(5, 0.15)	7.2×10^{-1}	7.3	1.9×10^{-1}	14.8	8.6×10^{-2}	16.7
		(10, 0.10)	8.0×10^{-1}	7.3	3.9×10^{-1}	13.9	3.2×10^{-1}	17.6
		(10, 0.15)	7.1×10^{-1}	7.4	1.5×10^{-1}	14.6	1.1×10^{-1}	17.9
	1000	(5, 0.10)	7.4×10^{-1}	42.4	2.4×10^{-1}	120.6	1.5×10^{-1}	121.6
		(5, 0.15)	7.1×10^{-1}	42.1	1.9×10^{-1}	115.7	7.9×10^{-2}	119.9
		(5, 0.20)	6.2×10^{-1}	44.2	1.2×10^{-1}	118.2	3.9×10^{-2}	119.8
		(10, 0.10)	7.5×10^{-1}	42.9	1.9×10^{-1}	110.5	1.3×10^{-1}	119.9
		(10, 0.15)	7.1×10^{-1}	42.8	1.4×10^{-1}	115.7	6.6×10^{-2}	119.0
		(10, 0.20)	6.0×10^{-1}	44.1	7.0×10^{-2}	118.7	3.7×10^{-2}	119.6
	1500	(5, 0.10)	7.5×10^{-1}	142.1	2.4×10^{-1}	391.7	1.6×10^{-1}	380.7
		(5, 0.15)	7.1×10^{-1}	143.8	2.1×10^{-1}	385.4	7.5×10^{-2}	386.4
		(5, 0.20)	6.0×10^{-1}	146.6	1.1×10^{-1}	385.0	2.9×10^{-2}	387.9
		(10, 0.10)	7.5×10^{-1}	143.1	1.7×10^{-1}	372.9	1.1×10^{-1}	377.9
		(10, 0.15)	7.1×10^{-1}	144.2	1.6×10^{-2}	390.4	3.9×10^{-2}	388.5
0.01	500	(5, 0.10)	7.5×10^{-1}	7.5	4.1×10^{-2}	13.7	4.0×10^{-2}	15.4
		(5, 0.15)	7.2×10^{-1}	7.8	3.8×10^{-2}	13.7	2.9×10^{-2}	15.1
		(10, 0.10)	8.0×10^{-1}	7.5	1.2×10^{-1}	12.9	2.9×10^{-2}	16.1
		(10, 0.15)	7.1×10^{-1}	7.8	3.8×10^{-2}	13.8	2.3×10^{-2}	16.3
		(10, 0.20)	6.2×10^{-1}	8.5	2.8×10^{-2}	13.8	2.1×10^{-2}	16.2
	1000	(5, 0.10)	7.4×10^{-1}	44.4	2.4×10^{-1}	115.9	1.5×10^{-1}	118.3
		(5, 0.15)	7.1×10^{-1}	45.6	1.9×10^{-1}	117.6	7.7×10^{-2}	119.1
		(5, 0.20)	6.2×10^{-1}	47.8	1.1×10^{-1}	117.1	4.4×10^{-2}	120.0
		(10, 0.10)	7.5×10^{-1}	44.6	2.0×10^{-1}	112.3	1.4×10^{-1}	118.0
		(10, 0.15)	7.1×10^{-1}	45.6	1.4×10^{-1}	117.3	6.6×10^{-2}	117.6
		(10, 0.20)	6.1×10^{-1}	48.3	7.0×10^{-2}	113.4	4.7×10^{-2}	119.4
	1500	(5, 0.10)	7.5×10^{-1}	148.2	2.4×10^{-1}	381.7	1.6×10^{-1}	386.9
		(5, 0.15)	7.1×10^{-1}	150.4	2.1×10^{-1}	396.8	6.5×10^{-2}	396.1
		(5, 0.20)	6.0×10^{-1}	156.2	1.1×10^{-1}	396.9	3.2×10^{-2}	390.0
		(10, 0.10)	7.5×10^{-1}	148.6	1.7×10^{-1}	401.5	1.1×10^{-1}	396.9
		(10, 0.15)	7.1×10^{-1}	151.4	1.6×10^{-1}	405.3	4.8×10^{-2}	389.2
		(10, 0.20)	6.0×10^{-1}	160.1	8.0×10^{-2}	398.4	3.7×10^{-2}	393.1

The parameters are chosen to be the same as those in Table 2

that the (k, ℓ)-th entry is sampled. For each $(k, \ell) \in [d_t] \times [d_t]$, let $\pi_{k\ell} = p_k p_\ell$, where we let p_k (and p_ℓ) be

$$p_k = \begin{cases} 2p_0 & \text{if } k \le \frac{d_t}{10} \\ 4p_0 & \text{if } \frac{d_t}{10} < k \le \frac{d_t}{5} \\ p_0 & \text{otherwise,} \end{cases} \text{ for Scheme 2, and } p_k = \begin{cases} 3p_0 & \text{if } k \le \frac{d_t}{10} \\ 9p_0 & \text{if } \frac{d_t}{10} < k \le \frac{d_t}{5} \\ p_0 & \text{otherwise,} \end{cases} \text{ for Scheme 3,}$$

and p_0 is a normalizing constant such that $\sum_{k=1}^{d_t} p_k = 1$.

In the implementation of Algorithm 1, we set the tuning parameter λ to be proportional to $\|Y_\Omega\|_F$, as suggested by Toh and Yun [37], where Y_Ω denotes the partially observed matrix. From the theoretical analysis in Sect. 4, we have that the parameter μ should be smaller than λ by a factor of about $(d_1 d_2)^{-1/2} = d_t^{-1}$ in the hybrid approach.

To evaluate the matrix recovery results, we adopt the metric of relative error (RE) defined by

$$\text{RE} = \frac{\|\hat{M} - M^0\|_F}{\|M^0\|_F},$$

where $\hat{M}$ is the output solution by the algorithm. We consider different settings of d_t, r and SR. We run simulations under each setting for five different instances. We first consider the noiseless cases. The averaged relative errors and running times are summarized in the upper halves of Tables 1, 2 and 3, corresponding to Schemes 1, 2 and 3, respectively. In Table 1, where uniformly sampled data is considered, we find that the nuclear-norm approach obtains the best recovery results. Meanwhile, we find that the hybrid approach performs significantly better than the pure max-norm approach. This observation is consistent with the existing theoretical result that max-norm regularization does not perform as well as nuclear-norm regularization if the observed entries are indeed uniformly sampled, and the proposed hybrid approach significantly boosts the performance of the max-norm regularized method without specifying data generating schemes. In Tables 2 and 3, where non-uniform sampling distributions are considered, we observe that both max-norm regularized and hybrid approaches significantly outperform the nuclear-norm approach, especially when the sampling ratio is low. This observation matches the theoretical analysis in Sect. 4 and Cai and Zhou [7]. We also find that the hybrid approach always outperforms the max-norm approach. This is because, while the max-norm approach is robust, the additional nuclear norm penalization helps to fully utilize the underlying low-rank structure in our generating schemes.

Next, we consider settings with noises, where we use the same sampling schemes as in Schemes 1, 2 and 3, and for each sampled entry, we observe a noisy sample:

$$Y_{i_t, j_t} = M^0_{i_t, j_t} + \sigma \xi_t \cdot \|M^0\|_\infty, \text{ where } \sigma = 0.01 \text{ and } \xi_t \sim N(0, 1).$$

We report the averaged relative errors and running times in the lower halves of Tables 1, 2 and 3. As expected, under non-uniform sampling schemes, the max-norm and hybrid approaches produce better recovery results than the nuclear-norm approach, and the hybrid approach outperforms the max-norm approach. Surprisingly, we find

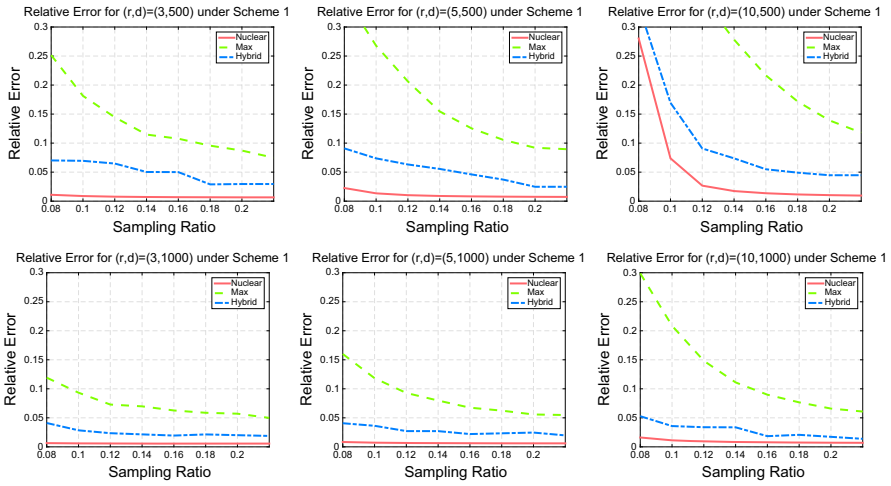

Fig. 2 Relative Errors under different settings for the noisy case under Scheme 1 (uniform sampling)

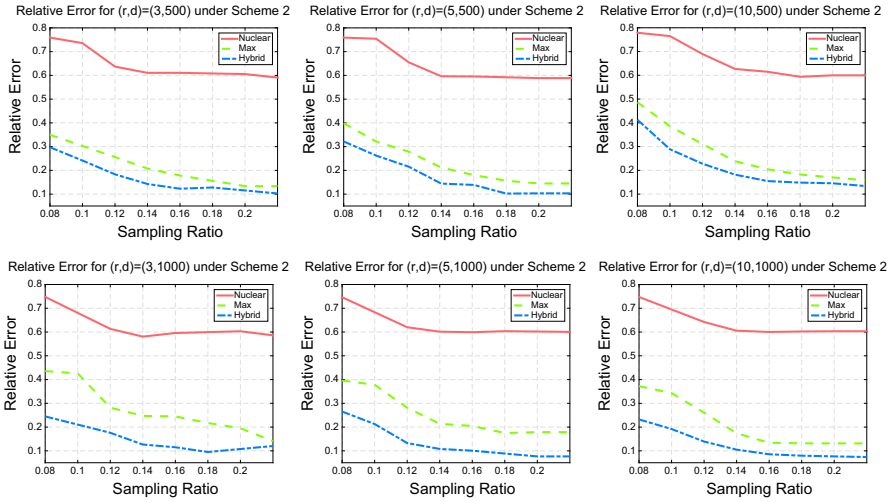

Fig. 3 Relative Errors under different settings for the noisy case under Scheme 2 (non-uniform sampling)

that under the uniform sampling scheme, the max-norm and hybrid approaches also outperform the nuclear-norm approach in the noisy setting. These observations provide further evidences that the max-norm and hybrid approaches are more robust to noises and sampling schemes than the nuclear-norm approach in practice.

In addition, for the noisy setting, we plot how the relative errors decrease as sampling ratios increase under the three schemes. Specifically, for SR $= 0.08, 0.10, \ldots, 0.22$, $r = 3, 5, 10$ and $d_t = 500$ and 1000, we plot the averaged relative errors over five repetitions in Figs. 2, 3 and 4. Under the uniform sampling scheme, Fig. 2 shows that the nuclear-norm approach provides the best recovery results while the hybrid approach performs much better than the max-norm approach. Under non-uniform

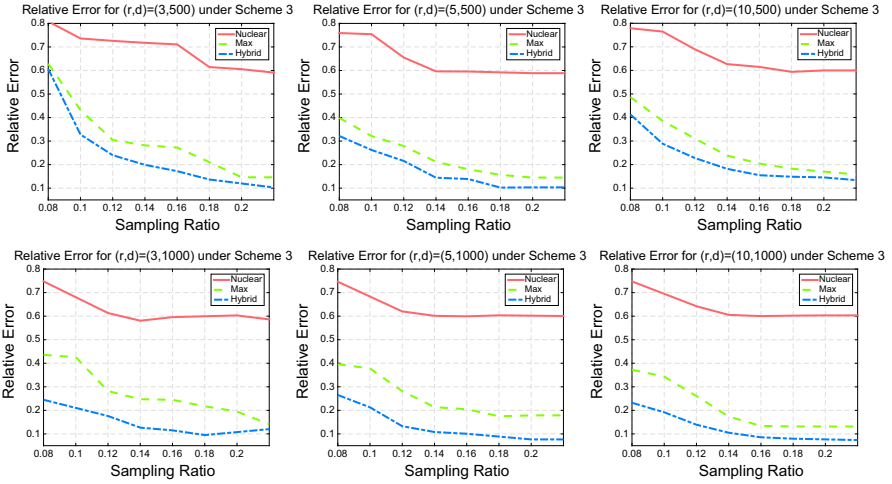

Fig. 4 Relative Errors under different settings for the noisy case under Scheme 3 (non-uniform sampling)

sampling schemes, Figs. 3 and 4 demonstrate that the hybrid approach has the best performance, while the nuclear-norm approach gives the poorest results.

5.2 Real datasets

In this subsection, we test our methods using some real datasets. We first consider the well-known Jester joke dataset. This dataset contains more than 4.1 million ratings for 100 jokes from 73,421 users, and it is publicly available through http://www.ieor.berkeley.edu/~goldberg/jester-data/. The whole Jester joke dataset contains three sub-datasets, which are: (1) jester-1: 24,983 users who rate 36 or more jokes; (2) jester-2: 23,500 users who rate 36 or more jokes; (3) jester-3: 24,938 users who rate between 15 and 35 jokes. More detailed descriptions can be found in Toh and Yun [37] and Chen et al. [11], where the nuclear-norm based approach is used to study this dataset.

Due to the large number of users, as in Chen et al. [11], we randomly select n_u users' ratings from the datasets. Since many entries are unknown, we cannot compute the relative error as we did for the simulated data. Instead, we take the metric of the normalized mean absolute error (NMAE) to measure the accuracy of the estimator $\hat{M}$:

$$\text{NMAE} = \frac{\sum_{(j,k)\notin\Omega} |\hat{M}_{jk} - M^0_{jk}|}{|\Omega|(r_{\max} - r_{\min})},$$

where $r_{\min}$ and $r_{\max}$ denote the lower and upper bounds for the ratings, respectively. In the Jester joke dataset, the range is $[-10, 10]$. Thus, we have $r_{\max} - r_{\min} = 20$.

In each iteration, we first randomly select n_u users, and then randomly permute the ratings from the users to generate $M^0 \in \mathbb{R}^{n_u \times 100}$. Next, we adopt the generating scheme used in Scheme 2 in the previous subsection to generate a set of observed

Table 4 Averaged normalized mean absolute error and running time in seconds for different methods using Jester joke dataset

Example	(n_u, SR)	Nuclear		Max		Hybrid	
		NMAE	Time	NMAE	Time	NMAE	Time
jester-1	(1000, 0.15)	0.210	4.82	0.197	110.55	0.200	87.13
	(1000, 0.20)	0.209	4.83	0.194	111.79	0.203	89.98
	(1000, 0.25)	0.204	5.12	0.188	111.36	0.197	89.02
	(1500, 0.15)	0.210	5.93	0.194	302.47	0.201	250.07
	(1500, 0.20)	0.206	6.08	0.192	307.70	0.195	255.29
	(1500, 0.25)	0.204	6.39	0.185	305.91	0.194	254.66
	(2000, 0.15)	0.212	7.06	0.192	647.25	0.196	566.84
	(2000, 0.20)	0.208	7.30	0.188	671.73	0.192	547.89
	(2000, 0.25)	0.205	7.45	0.183	640.75	0.192	558.02
jester-2	(1000, 0.15)	0.211	4.86	0.199	109.15	0.196	86.34
	(1000, 0.20)	0.207	5.01	0.192	110.40	0.193	87.81
	(1000, 0.25)	0.204	4.89	0.188	110.41	0.187	90.07
	(1500, 0.15)	0.212	5.86	0.197	313.01	0.198	247.26
	(1500, 0.20)	0.210	6.10	0.192	313.39	0.193	260.84
	(1500, 0.25)	0.205	6.34	0.189	322.05	0.187	255.88
	(2000, 0.15)	0.213	6.99	0.197	633.97	0.198	577.32
	(2000, 0.20)	0.208	7.50	0.194	644.04	0.193	562.32
	(2000, 0.25)	0.204	7.42	0.187	687.24	0.188	576.56
jester-3	(1000, 0.15)	0.227	4.27	0.221	97.82	0.218	83.18
	(1000, 0.20)	0.220	4.41	0.212	103.28	0.212	84.02
	(1000, 0.25)	0.221	4.54	0.213	105.48	0.212	84.90
	(1500, 0.15)	0.225	5.47	0.218	272.30	0.215	237.38
	(1500, 0.20)	0.220	5.54	0.212	280.34	0.212	240.19
	(1500, 0.25)	0.218	5.69	0.208	284.05	0.211	241.21
	(2000, 0.15)	0.226	6.46	0.216	585.71	0.218	521.87
	(2000, 0.20)	0.222	6.59	0.217	606.53	0.212	525.93
	(2000, 0.25)	0.218	6.70	0.211	614.04	0.210	526.78

For the nuclear norm approach, we set $\mu = 2 \times 10^{-4}\|Y_S\|_F$. For the max-norm approach, we set $\lambda = 0.5\|Y_S\|_F$. For the hybrid approach, we set $\lambda = 0.8\|Y_S\|_F$, $\mu = 1 \times 10^{-4}\lambda$

indices Ω. Note that we can only observe the entry (j, k) if $(j, k) \in \Omega$, and $M^0_{j,k}$ is available. Thus, the actual sampling ratio is less than the input SR. We consider different settings of n_u and SR, and we report the averaged NMAE and running times in Table 4 after running each setting five times. It can be seen that the max-norm and hybrid approaches outperform the nuclear-norm approach in all cases. This provides strong evidences that the proposed estimator and algorithm could be useful in practice.

Meanwhile, we observe that the running times for solving the max-norm penalized optimization problems are significantly longer than that for solving the nuclear-norm penalized problem. This is because solving max-norm penalized optimization

Table 5 Averaged normalized mean absolute error and running time in seconds for different methods using Movie-100K dataset

SR	Nuclear		Max		Hybrid	
	NMAE	Time	NMAE	Time	NMAE	Time
0.10	0.243	108.4	0.231	266.8	0.232	292.2
0.15	0.235	112.5	0.222	274.9	0.223	288.9
0.20	0.233	112.1	0.213	263.4	0.220	286.2
0.25	0.223	123.8	0.208	285.5	0.215	294.7

The parameters are chosen to be the same as those in Table 4

Table 6 Averaged normalized mean absolute error and running time in seconds for different methods using Movie-1M dataset

n	SR	Nuclear		Max		Hybrid	
		NMAE	Time	NMAE	Time	NMAE	Time
1500	0.10	0.248	154.7	0.235	377.6	0.236	409.2
	0.15	0.238	154.1	0.222	318.3	0.229	410.9
	0.20	0.233	153.9	0.216	329.8	0.223	401.9
	0.25	0.225	210.7	0.208	473.3	0.218	506.2
2000	0.10	0.244	357.8	0.227	733.2	0.230	956.9
	0.15	0.234	363.5	0.214	725.7	0.213	946.0
	0.20	0.230	365.6	0.206	782.6	0.206	946.3
	0.25	0.220	391.9	0.199	744.4	0.210	950.7

The parameters are chosen to be the same as those in Table 4

problems is intrinsically more difficult than solving nuclear-norm penalized ones. Specifically, in Algorithm 1, the most computationally expensive step is to compute a full eigenvalue decomposition of a matrix of size $d_1 + d_2$ by $d_1 + d_2$ during the X-update step. As a comparison, in nuclear-norm regularized optimizations, we only need to compute a singular value decomposition of a matrix of size d_1 by d_2. In the Jester joke dataset, since $d_2 \ll d_1$, singular value decomposition takes the advantage of a small d_2, but the computational cost of the max-norm approach is dominated by the large d_1. In practical matrix completion problems, the computational efficiency is sometimes not the top priority, but more attention is placed on reducing the reconstruction error. Thus, depending on the specific application, the max-norm and hybrid approaches provide useful complements to the nuclear-norm approach.

We also consider the MovieLens data. The dataset is available through http://www.grouplens.org. We first implement the proposed methods on the Movie-100K dataset, which contains 100,000 ratings for 1682 movies by 943 users. The ratings range from $r_{\min} = 1$ to $r_{\max} = 5$. In this experiment, we first randomly permute the rows and columns of the matrix, and then sample the observed entries as in Scheme 2 in the previous subsection. Table 5 reports the averaged NMAE and running times of different methods. Next, we implement the methods on the Movie-1M dataset. This dataset contains 1,000,209 ratings of 3900 movies made by 6040 users. We randomly select n users and n movies to conduct the tests, where $n = 1500$ or 2000. We report the results in Table 6. From Tables 5 and 6, we observe that the max-norm and hybrid

approaches lead to better matrix recovery results than the nuclear-norm approach in all cases. In addition, we observe that the differences between running times of the max-norm and nuclear-norm approaches are less significant than those in the Jester joke problem. This is because d_1 and d_2 are of the same order in the MovieLens example. Therefore, in practice, if the computational efficiency is the top priority, and if $d_1 \ll d_2$ or $d_1 \gg d_2$, the nuclear-norm approach is preferable. While if controlling the reconstruction accuracy attracts more concern, we recommend the proposed hybrid approach.

Remark 5.1 Note that the improvement from the hybrid and max-norm approaches over the nuclear-norm approach is about 5%, which looks marginal. However, a 5% improvement can be significant in practice as the nuclear-norm approach is widely recognized as a highly efficient approach. In the earlier Netflix competition, it is seen that the results from top teams (where nuclear-norm approach is used as part of the algorithms) are all very close, and a 5% improvement can be significant for practitioners. See http://www.research.att.com/articles/featured_stories/2010_05/201005_netflix2_article.html?fbid=pgKJkRJ5mbi. In addition, though the nuclear-norm approach is computationally more efficient, we note that in this particular application, computation efficiency is not of the highest priority, and the modest sacrifice of computational cost is tolerable here.

6 Conclusions

We propose a new matrix completion method using a hybrid nuclear- and max-norm regularizer. Compared with the standard nuclear-norm based approach, our method is adaptive under different sampling schemes and achieves fast rates of convergence. To handle the computational challenge, we propose the first scalable algorithm with provable convergence guarantee. This bridges the gap between theory and practice of the max-norm approach. In addition, we provide thorough numerical results to backup the developed theory. This work paves the way for more potential machine learning applications of max-norm regularization.

A possible future direction is to further improve the computational efficiency. The most computationally expensive component in Algorithm 1 is the X-update step, in which an eigenvalue decomposition is needed. By solving some approximate version of this subproblem, it is possible to further boost the empirical performance and solve problems of larger sizes.

7 Extensions

In this section, we consider solving the max-norm constrained version of the optimization problem (2.3). In particular, we consider

$$\min_{M \in \mathbb{R}^{d_1 \times d_2}} \frac{1}{2} \sum_{t=1}^{n} \left(Y_{i_t, j_t} - M_{i_t, j_t} \right)^2 + \langle M, I \rangle, \text{ subject to } \|M\|_\infty \leq \alpha, \|M\|_{\max} \leq R. \tag{7.1}$$

This problem can be formulated as an SDP problem as follows:

$$\begin{aligned} &\min_{Z\in\mathbb{R}^{d\times d}} \frac{1}{2}\sum_{t=1}^{n}(Y_{i_t,j_t}-Z^{12}_{i_t,j_t})^2+\mu\langle I,\ Z\rangle,\\ &\text{subject to } \ \|Z^{12}\|_\infty\le\alpha,\ \|\mathrm{diag}(Z)\|_\infty\le R,\ \ Z\succeq 0. \end{aligned} \tag{7.2}$$

Let the loss function be

$$\mathcal{L}(Z)=\frac{1}{2}\sum_{t=1}^{n}\left(Y_{i_t,j_t}-Z^{12}_{i_t,j_t}\right)^2+\mu\langle I,Z\rangle.$$

We define the set

$$\mathcal{P}=\{Z\in\mathcal{S}^d:\mathrm{diag}(Z)\ge 0,\|Z^{11}\|_\infty\le R,\|Z^{22}\|_\infty\le R,\|Z^{12}\|_\infty<\alpha\}.$$

Thus, we have an equivalent formulation of (7.2) below, which is more conducive for computation:

$$\min_{X,Z}\mathcal{L}(Z)+\mu\langle X,I\rangle,\ \text{subject to } X\succeq 0,\ Z\in\mathcal{P},X-Z=0. \tag{7.3}$$

We consider the augmented Lagrangian function of (7.3) defined by

$$L(X,Z;W)=\mathcal{L}(Z)+\langle W,X-Z\rangle+\frac{\rho}{2}\|X-Z\|_F^2,\ X\in\mathcal{S}^d_+,\ Z\in\mathcal{P},$$

where W is the dual variable. Then, it is natural to apply the ADMM to solve the problem (7.3). At the t-th iteration, we update $(X,Z;W)$ by

$$\begin{aligned} X^{t+1}&=\operatorname*{argmin}_{X\in\mathcal{S}^d_+}L(X,Z^t;W^t)=\Pi_{\mathcal{S}^d_+}\left\{Z^t-\rho^{-1}(W^t+\mu I)\right\},\\ Z^{t+1}&=\operatorname*{argmin}_{Z\in\mathcal{P}}L(X^{t+1},Z;W^t)=\operatorname*{argmin}_{Z\in\mathcal{P}}\mathcal{L}(Z)+\frac{\rho}{2}\|Z-X^{t+1}-\rho^{-1}W^t\|_F^2,\\ W^{t+1}&=W^t+\tau\rho(X^{t+1}-Z^{t+1}), \end{aligned} \tag{7.4}$$

The next proposition provides a closed-form solution for the Z-subproblem in (7.4).

Proposition 7.1 *Denote the observed set of indices of M^0 by $\Omega=\{(i_t,j_t)\}_{t=1}^n$. For a given matrix $C\in\mathbb{R}^{d\times d}$, we have*

$$\mathcal{Z}(C)=\operatorname*{argmin}_{Z\in\mathcal{P}}\mathcal{L}(Z)+\frac{\rho}{2}\|Z-C\|_F^2, \tag{7.5}$$

where

$$\mathcal{Z}(C) = \begin{pmatrix} \mathcal{Z}^{11}(C) & \mathcal{Z}^{12}(C) \\ \mathcal{Z}^{12}(C)^T & \mathcal{Z}^{22}(C) \end{pmatrix}$$

$$\mathcal{Z}^{12}_{k\ell}(C) = \begin{cases} \Pi_{[-\alpha,\alpha]}\left(\frac{Y_{k\ell}+\rho C^{12}_{k\ell}}{\rho}\right), & \text{if } (k,\ell) \in S, \\ \Pi_{[-\alpha,\alpha]}(C^{12}_{k\ell}), & \text{otherwise,} \end{cases}$$

$$\mathcal{Z}^{11}_{k\ell}(C) = \begin{cases} \Pi_{[-R,R]}\left(C^{11}_{k\ell}\right) & \text{if } k \neq \ell, \\ \Pi_{[0,R]}\left(C^{11}_{k\ell}\right) & \text{if } k = \ell, \end{cases} \quad \mathcal{Z}^{22}_{k\ell}(C) = \begin{cases} \Pi_{[-R,R]}\left(C^{22}_{k\ell}\right) & \text{if } k \neq \ell, \\ \Pi_{[0,R]}\left(C^{22}_{k\ell}\right) & \text{if } k = \ell, \end{cases}$$

and $\Pi_{[a,b]}(x) = \min\{b, \max(a, x)\}$ *projects* $x \in \mathbb{R}$ *to the interval* $[a, b]$.

We summarize the algorithm for solving the problem (7.2) below.

Algorithm 2 Solving max-norm optimization problem (7.2) by the ADMM

Initialize X^0, Z^0, W^0, ρ, λ.
Input: X^0, Z^0, W^0, Y_Ω, λ, R, α, ρ, τ, $t = 0$.
  while Stopping criterion is not satisfied. **do**
    Update $X^{t+1} \leftarrow \Pi_{\mathcal{S}^d_+}(Z^t - \rho^{-1}(W^t + \mu I))$.
    Update $Z^{t+1} \leftarrow \mathcal{Z}(X^{t+1} + \rho^{-1}W^t)$ by (7.5).
    Update $W^{t+1} \leftarrow W^t + \tau\rho(X^{t+1} - Z^{t+1})$.
    $t \leftarrow t + 1$.
  end while
Output: $\hat{Z} = Z^t$, $\hat{M} = \hat{Z}^{12} \in \mathbb{R}^{d_1 \times d_2}$.

References

1. Abernethy, J., Bach, F., Evgeniou, T., Vert, J.-P.: A new approach to collaborative filtering: operator estimation with spectral regularization. J. Mach. Learn. Res. **10**, 803–826 (2009)
2. Amit, Y., Fink, M., Srebro, N., Ullman, S.: Uncovering shared structures in multiclass classification. In: Proceedings of the 24th International Conference on Machine Learning. ACM (2007)
3. Argyriou, A., Evgeniou, T., Pontil, M.: Convex multi-task feature learning. Mach. Learn. **73**, 243–272 (2008)
4. Bennett, J., Lanning, S.: The Netflix prize. In: Proceedings of KDD Cup and Workshop. http://www.cs.uic.edu/~liub/KDD-cup-2007/proceedings.html (2007)
5. Biswas, P., Liang, T., Toh, K., Wang, T., Ye, Y.: Semidefinite programming approaches for sensor network localization with noisy distance measurements. IEEE Trans. Autom. Sci. Eng. **3**, 360–371 (2006)
6. Cai, J.-F., Candès, E.J., Shen, Z.: A singular value thresholding algorithm for matrix completion. SIAM J. Optim. **20**, 1956–1982 (2010)
7. Cai, T.T., Zhou, W.-X.: Matrix completion via max-norm constrained optimization. Electron. J. Stat. **10**, 1493–1525 (2016)
8. Candès, E.J., Li, X., Ma, Y., Wright, J.: Robust principal component analysis? J. ACM **58**, 1–37 (2009)
9. Candès, E.J., Recht, B.: Exact matrix completion via convex optimization. Found. Comput. Math. **9**, 717–772 (2009)
10. Candès, E.J., Tao, T.: The power of convex relaxation: near-optimal matrix completion. IEEE Trans. Inf. Theory **56**, 2053–2080 (2010)

11. Chen, C., He, B., Yuan, X.: Matrix completion via an alternating direction method. IMA J. Numer. Anal. **32**, 227–245 (2012)
12. Doan, X., Vavasis, S.: Finding approximately rank-one submatrices with the nuclear norm and ℓ_1-norm. SIAM J. Optim. **23**, 2502–2540 (2013)
13. Drusvyatskiy, D., Vavasis, S., Wolkowicz, H.: Extreme point inequalities and geometry of the rank sparsity ball. Math. Program. **152**, 521–544 (2015)
14. Fang, E.X., He, B., Liu, H., Yuan, X.: Generalized alternating direction method of multipliers: new theoretical insights and applications. Math. Prog. Comp. **7**, 149–187 (2015)
15. Fazel, M., Hindi, H., Boyd, S.P.: A rank minimization heuristic with application to minimum order system approximation. In: Proceedings of the American Control Conference, vol. **6**. IEEE (2001)
16. Figueiredo, M., Nowak, R., Wright, S.: Gradient projection for sparse reconstruction: application to compressed sensing and other inverse problems. IEEE J. Sel. Top. Signal Process. **1**, 586598 (2007)
17. Jalali, A., Srebro, N.: Clustering using max-norm constrained optimization. In: Proceedings of the 29th International Conference on Machine Learning (ICML-12) (2012)
18. Jameson, G.J.O.: Summing and Nuclear Norms in Banach Space Theory, vol. 8. Cambridge University Press, Cambridge (1987)
19. Keshavan, R.H., Montanari, A., Oh, S.: Matrix completion from noisy entries. J. Mach. Learn. Res. **11**, 2057–2078 (2010)
20. Klopp, O.: Noisy low-rank matrix completion with general sampling distribution. Bernoulli **20**, 282–303 (2014)
21. Koltchinskii, V., Lounici, K., Tsybakov, A.B.: Nuclear-norm penalization and optimal rates for noisy low-rank matrix completion. Ann. Statist. **39**, 2302–2329 (2011)
22. Lee, J., Recht, B., Srebro, N., Tropp, J., Salakhutdinov, R.: Practical large-scale optimization for max-norm regularization. In: Proceedings of the Advances in Neural Information Processing Systems (2010)
23. Linial, N., Mendelson, S., Schechtman, G., Shraibman, A.: Complexity measures of sign matrices. Combinatorica **27**, 439–463 (2007)
24. Liu, Z., Vandenberghe, L.: Interior-point method for nuclear norm approximation with application to system identification. SIAM J. Matrix Anal. Appl. **31**, 1235–1256 (2009)
25. Mackey, L., Jordan, M.I., Chen, R.Y., Farrell, B., Tropp, J.A.: Matrix concentration inequalities via the method of exchangeable pairs. Ann. Probab. **42**, 906–945 (2014)
26. Negahban, S., Wainwright, M.J.: Restricted strong convexity and weighted matrix completion: optimal bounds with noise. J. Mach. Learn. Res. **13**, 1665–1697 (2012)
27. Netflix (2006). Netflix problem. http://www.netflixprize.com
28. Oliveira, D.E., Wolkowicz, H., Xu, Y.: ADMM for the SDP relaxation of the QAP. arXiv preprint arXiv:1512.05448 (2015)
29. Orabona, F., Argyriou, A., Srebro, N.: PRISMA: Proximal iterative smoothing algorithm. arXiv preprint arXiv:1206.2372 (2012)
30. Recht, B.: A simpler approach to matrix completion. J. Mach. Learn. Res. **12**, 3413–3430 (2011)
31. Recht, B., Fazel, M., Parrilo, P.A.: Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization. SIAM Rev. **52**, 471–501 (2010)
32. Rohde, A., Tsybakov, A.B.: Estimation of high-dimensional low-rank matrices. Ann. Stat. **39**, 887–930 (2011)
33. Shen, J., Xu, H., Li, P.: Online optimization for max-norm regularization. Mach. Learn. **106**, 419–457 (2017)
34. Srebro, N., Rennie, J., Jaakkola, T.S.: Maximum-margin matrix factorization. In: Proceedings of the Advances in Neural Information Processing Systems (2005)
35. Srebro, N., Salakhutdinov, R.R.: Collaborative filtering in a non-uniform world: learning with the weighted trace norm. In: Proceedings of the Advances in Neural Information Processing Systems (2010)
36. Srebro, N., Shraibman, A.: Rank, trace-norm and max-norm. In: Proceedings of the 18th Annual Conference on Learning Theory (2005)
37. Toh, K.-C., Yun, S.: An accelerated proximal gradient algorithm for nuclear norm regularized linear least squares problems. Pac. J. Optim. **6**, 615–640 (2010)
38. Trefethen, L.N., Bau III, D.: Numerical Linear Algebra, vol. 50. SIAM, Philadelphia (1997)