

Local and Global Linear Convergence of General Low-Rank Matrix Recovery Problems

Yingjie Bi,¹ Haixiang Zhang,² Javad Lavaei¹

¹ Industrial Engineering and Operations Research, University of California, Berkeley

² Department of Mathematics, University of California, Berkeley
yingjiebi@berkeley.edu, haixiang_zhang@berkeley.edu, lavaei@berkeley.edu

Abstract

We study the convergence rate of gradient-based local search methods for solving low-rank matrix recovery problems with general objectives in both symmetric and asymmetric cases, under the assumption of the restricted isometry property. First, we develop a new technique to verify the Polyak–Łojasiewicz inequality in a neighborhood of the global minimizers, which leads to a local linear convergence region for the gradient descent method. Second, based on the local convergence result and a sharp strict saddle property proven in this paper, we present two new conditions that guarantee the global linear convergence of the perturbed gradient descent method. The developed local and global convergence results provide much stronger theoretical guarantees than the existing results. As a by-product, this work significantly improves the existing bounds on the RIP constant required to guarantee the non-existence of spurious solutions.

Introduction

The low-rank matrix recovery problem is to recover an unknown low-rank ground truth matrix from certain measurements. This problem has a variety of applications in machine learning, such as recommendation systems (Koren, Bell, and Volinsky 2009) and motion detection (Zhou, Yang, and Yu 2013; Fattahi and Sojoudi 2020), and in engineering problems, such as power system state estimation (Zhang, Madani, and Lavaei 2018).

In this paper, we consider two variants of the low-rank matrix recovery problem with a general measurement model represented by an arbitrary smooth function. The first variant is the symmetric problem, in which the ground truth $M^* \in \mathbb{R}^{n \times n}$ is a symmetric and positive semidefinite matrix with $\text{rank}(M^*) = r$, and M^* is a global minimizer of some loss function f_s . Then, M^* can be recovered by solving the optimization problem:

$$\begin{aligned} \min \quad & f_s(M) \\ \text{s. t.} \quad & \text{rank}(M) \leq r, \\ & M \succeq 0, M \in \mathbb{R}^{n \times n}. \end{aligned} \quad (1)$$

Note that minimizing $f_s(M)$ over positive semidefinite matrices without the rank constraint would often lead to finding a solution with the highest-rank possible rather than the

rank-constrained solution M^* . The second variant of the low-rank matrix recovery problem to be studied is the asymmetric problem, in which $M^* \in \mathbb{R}^{n \times m}$ is a possibly non-square matrix with $\text{rank}(M^*) = r$, and it is a global minimizer of some loss function f_a . Similarly, M^* can be recovered by solving

$$\begin{aligned} \min \quad & f_a(M) \\ \text{s. t.} \quad & \text{rank}(M) \leq r, \\ & M \in \mathbb{R}^{n \times m}. \end{aligned} \quad (2)$$

As a special case, the loss function f_s or f_a can be induced by linear measurements. In this situation, we are given a linear operator $\mathcal{A} : \mathbb{R}^{n \times n} \rightarrow \mathbb{R}^p$ or $\mathcal{A} : \mathbb{R}^{n \times m} \rightarrow \mathbb{R}^p$, where p denotes the number of measurements. To recover M^* from the vector $d = \mathcal{A}(M^*)$, the function $f_s(M)$ or $f_a(M)$ is often chosen to be

$$\frac{1}{2} \|\mathcal{A}(M) - d\|^2. \quad (3)$$

Besides, there are many natural choices for the loss function, such as a nonlinear model associated with the 1-bit matrix recovery (Davenport et al. 2014).

The symmetric problem (1) can be transformed into an unconstrained optimization problem by factoring M as XX^T with $X \in \mathbb{R}^{n \times r}$, which leads to the following equivalent formulation:

$$\min_{X \in \mathbb{R}^{n \times r}} f_s(XX^T). \quad (4)$$

In the asymmetric case, one can similarly factor M as UV^T with $U \in \mathbb{R}^{n \times r}$ and $V \in \mathbb{R}^{m \times r}$. Note that $(UP, V(P^{-1})^T)$ gives another possible factorization of M for any invertible matrix $P \in \mathbb{R}^{r \times r}$. To reduce the redundancy, a regularization term is usually added to the objective function to enforce that the factorization is balanced, i.e., $U^T U = V^T V$ is satisfied (Tu et al. 2016). Since every factorization can be converted into a balanced one by selecting an appropriate P , the original asymmetric problem (2) is equivalent to

$$\min_{U \in \mathbb{R}^{n \times r}, V \in \mathbb{R}^{m \times r}} f_a(UV^T) + \frac{\phi}{4} \|U^T U - V^T V\|_F^2, \quad (5)$$

where $\phi > 0$ is an arbitrary constant.

To handle the symmetric and asymmetric problems in a unified way, we will use the same notation X to denote the

matrix of decision variables in both cases. In the symmetric case, X is obtained from the equation $M = XX^T$. In the asymmetric case, X is defined as

$$X = \begin{bmatrix} U \\ V \end{bmatrix} \in \mathbb{R}^{(n+m) \times r}.$$

To rewrite the asymmetric problem (5) in terms of X , we apply the technique in Tu et al. (2016) by defining an auxiliary function $F : \mathbb{R}^{(n+m) \times (n+m)} \rightarrow \mathbb{R}$ as

$$F \left(\begin{bmatrix} N_{11} & N_{12} \\ N_{21} & N_{22} \end{bmatrix} \right) = \frac{1}{2} (f_a(N_{12}) + f_a(N_{21}^T)) + \frac{\phi}{4} (\|N_{11}\|_F^2 + \|N_{22}\|_F^2 - \|N_{12}\|_F^2 - \|N_{21}\|_F^2), \quad (6)$$

in which the argument of the function F is partitioned into four blocks, denoted as $N_{11} \in \mathbb{R}^{n \times n}$, $N_{12} \in \mathbb{R}^{n \times m}$, $N_{21} \in \mathbb{R}^{m \times n}$, $N_{22} \in \mathbb{R}^{m \times m}$. The problem (5) then reduces to

$$\min_{X \in \mathbb{R}^{(n+m) \times r}} F(XX^T), \quad (7)$$

which is a special case of the symmetric problem (4). Henceforth, the objective functions of the two problems will be referred as to $g_s(X) = f_s(XX^T)$ and $g_a(X) = F(XX^T)$, respectively.

The unconstrained problems (4) and (5) are often solved by local search algorithms, such as the gradient descent method, due to their efficiency in handling large-scale problems. Since the objective functions $g_s(X)$ and $g_a(X)$ are nonconvex, local search methods may converge to a spurious (non-global) local minimum. To guarantee the absence of such spurious solutions, the restricted isometry property (RIP) defined below is the most common condition imposed on the functions f_s and f_a (Bhojanapalli, Neyshabur, and Srebro 2016; Ge, Jin, and Zheng 2017; Zhu et al. 2018; Zhang et al. 2018b,a; Zhang, Sojoudi, and Lavaei 2019; Ha, Liu, and Barber 2020; Zhang and Zhang 2020; Bi and Lavaei 2020; Zhang, Bi, and Lavaei 2021; Zhang 2021).

Definition 1 (Recht, Fazel, and Parrilo (2010); Zhu et al. (2018)). A twice continuously differentiable function $f_s : \mathbb{R}^{n \times n} \rightarrow \mathbb{R}$ satisfies the *restricted isometry property* of rank $(2r_1, 2r_2)$ for a constant $\delta \in [0, 1)$, denoted as δ -RIP $_{2r_1, 2r_2}$, if

$$(1 - \delta)\|K\|_F^2 \leq [\nabla^2 f_s(M)](K, K) \leq (1 + \delta)\|K\|_F^2$$

holds for all matrices $M, K \in \mathbb{R}^{n \times n}$ with $\text{rank}(M) \leq 2r_1$ and $\text{rank}(K) \leq 2r_2$. In the case when $r_1 = r_2 = r$, the notation RIP $_{2r, 2r}$ will be simplified as RIP $_{2r}$. A similar definition can be also made for the asymmetric loss function f_a .

The state-of-the-art results on the non-existence of spurious local minima are presented in Zhang, Bi, and Lavaei (2021); Zhang (2021). Zhang, Bi, and Lavaei (2021) shows that the problem (4) or (5) is devoid of spurious local minima if i) the associated function f_s or f_a satisfies the δ -RIP $_2$ property with $\delta < 1/2$ in case $r = 1$, ii) the function f_s or f_a satisfies the δ -RIP $_{2r}$ property with $\delta \leq 1/3$ in case $r > 1$. Zhang (2021) further shows that a special case of the symmetric problem (4) does not have spurious local minima if

iii) f_s is in the form (3) given by linear measurements and satisfies the δ -RIP $_{2r}$ property with $\delta < 1/2$. The absence of spurious local minima under the above conditions does not automatically imply the existence of numerical algorithms with a fast convergence to the ground truth. In this paper, we will significantly strengthen the above two results by establishing the global linear convergence for problems with an arbitrary rank r and a general loss function f_s or f_a under the almost same RIP assumption as above, i.e., $\delta < 1/2$ for the symmetric problem and $\delta < 1/3$ for the asymmetric problem.

One common approach to establish fast convergence is to first show that the objective function has favorable regularity properties, such as strong convexity, in a neighborhood of the global minimizers, which guarantees that common iterative algorithms will converge to a global minimizer at least linearly if they are initialized in this neighborhood. Second, given the local convergence result, certain algorithms can be utilized to reach the above neighborhood from an arbitrary initial point. Note that randomization and stochasticity are often needed in those algorithms to avoid saddle points that are far from the ground truth, such as random initialization (Lee et al. 2016) or random perturbation during the iterations (Ge et al. 2015; Jin et al. 2017). In this paper, we deal with the two above-mentioned aspects for the low-rank matrix recovery problem separately.

Notations and Conventions

In this paper, I_n denotes the identity matrix of size $n \times n$, $A \otimes B$ denotes the Kronecker product of matrices A and B , and $A \succeq 0$ means that A is a symmetric and positive semidefinite matrix. $\sigma_i(A)$ denotes the i -th largest singular value of the matrix A . $\mathbf{A} = \text{vec}(A)$ is the vector obtained from stacking the columns of a matrix A . For a vector \mathbf{A} of dimension n^2 , its symmetric matricization $\text{mat}_S(\mathbf{A})$ is defined as $(A + A^T)/2$ with A being the unique matrix satisfying $\mathbf{A} = \text{vec}(A)$. For two matrices A and B of the same size, their inner product is denoted as $\langle A, B \rangle = \text{tr}(A^T B)$ and $\|A\|_F = \sqrt{\langle A, A \rangle}$ denotes the Frobenius norm of A . Given a matrix M and a set \mathcal{Z} of matrices, define

$$\text{dist}(X, \mathcal{Z}) = \min_{Z \in \mathcal{Z}} \|X - Z\|_F.$$

Moreover, $\|v\|$ denotes the Euclidean norm of a vector v . The action of the Hessian $\nabla^2 f(M)$ of a matrix function f on any two matrices K and L is given by

$$[\nabla^2 f(M)](K, L) = \sum_{i,j,k,l} \frac{\partial^2 f}{\partial M_{ij} \partial M_{kl}}(M) K_{ij} L_{kl}.$$

Summary of Main Contributions

For the local convergence, we prove that a regularity property named the Polyak-Łojasiewicz (PL) inequality always holds in a neighborhood of the global minimizers. The PL inequality is significantly weaker than the regularity condition used in previous works to study the local convergence of the low-rank matrix recovery problem, while it still guarantees a linear convergence to the ground truth. Hence, as will be compared with the prior literature, not only are the

obtained local regularity regions remarkably larger than the existing ones, but also they require significantly weaker RIP assumptions. Specifically, if f_s satisfies the δ -RIP $_{2r}$ property for an arbitrary δ , we will show that there exists some constant $\mu > 0$ such that the objective function g_s of the symmetric problem (4) satisfies the PL inequality

$$\frac{1}{2} \|\nabla g_s(X)\|_F^2 \geq \mu(g_s(X) - f_s(M^*))$$

for all X in the region

$$\{X \in \mathbb{R}^{n \times r} \mid \text{dist}(X, \mathcal{Z}) \leq \tilde{C}\}$$

with

$$\tilde{C} < \sqrt{2(\sqrt{2}-1)\sqrt{1-\delta^2}\sigma_r(M^*)}^{1/2}.$$

Here, $\text{dist}(X, \mathcal{Z})$ is the Frobenius distance between the matrix X and the set \mathcal{Z} of global minimizers of the problem (4). A similar result will also be derived for the asymmetric problem (5). Based on these results, local linear convergence can then be established. Compared with the previous results, our new results are advantageous for two reasons. First, the weaker RIP assumptions imposed by our results allow them to be applicable to a much broader class of problems, especially those problems with nonlinear measurements where the RIP constant of the loss function f_s or f_a varies at different points. In this case, the region in which the RIP constant is below the previous bounds may be significantly small or even empty, while the region satisfying our bounds is much larger since the radius of the region is increased by more than a constant factor. Second, when the RIP constant is large and global convergence cannot be established due to the existence of spurious solutions, the enlarged local regularity regions identified in this work can reduce the sample complexity to find the correct initial point converging to the ground truth. This has a major practical value in problems like data analytics in power systems (Jin et al. 2021) in which there is a fundamental limit to the number of measurements due to the physics of the network.

For the global convergence analysis, we first study the symmetric problem (4) with an arbitrary objective and an arbitrary rank r and prove that the objective g_s satisfies the strict saddle property if the function f_s has the δ -RIP $_{2r}$ property with $\delta < 1/2$. Note that this result is sharp, because in Zhang et al. (2018a) a counterexample has been found that contains spurious local minima under $\delta = 1/2$. Using the above strict saddle property and the local convergence result, we show that the perturbed gradient descent method with local improvement will find an approximate solution X satisfying $\|XX^T - M^*\|_F \leq \epsilon$ in $O(\log 1/\epsilon)$ number of iterations for an arbitrary tolerance ϵ . Moreover, the convergence result for symmetric problems also implies the global linear convergence for asymmetric problems under the δ -RIP $_{2r}$ condition with $\delta < 1/3$.

Assumptions

The assumptions required in this work will be introduced below. To avoid using different notations for the symmetric and asymmetric problems, we use the universal notation $f(M)$ henceforth to denote either $f_s(M)$ or $f_a(M)$. Similarly, M^* denotes the ground truth in either of the cases.

Assumption 1. The function f is twice continuously differentiable. In addition, its gradient ∇f is ρ_1 -restricted Lipschitz continuous for some constant ρ_1 , i.e., the inequality

$$\|\nabla f(M) - \nabla f(M')\|_F \leq \rho_1 \|M - M'\|_F$$

holds for all matrices M and M' with $\text{rank}(M) \leq r$ and $\text{rank}(M') \leq r$. The Hessian of the function f is also ρ_2 -restricted Lipschitz continuous for some constant ρ_2 , i.e., the inequality

$$|[\nabla^2 f(M) - \nabla^2 f(M')](K, K)| \leq \rho_2 \|M - M'\|_F \|K\|_F^2$$

holds for all matrices M, M', K with $\text{rank}(M) \leq r$, $\text{rank}(M') \leq r$ and $\text{rank}(K) \leq 2r$.

Assumption 2. The function f satisfies the δ -RIP $_{2r}$ property. Furthermore, ρ_1 in Assumption 1 is chosen to be large enough such that $\rho_1 \geq 1 + 2\delta$.

Assumption 3. The ground truth M^* satisfies $\|M^*\|_F \leq D$, and the initial point X_0 of the local search algorithm also satisfies $\|X_0 X_0^T\|_F \leq D$, where D is a constant given by the prior knowledge (every large enough D satisfies this assumption).

Assumption 4. In the asymmetric problem (5), the coefficient ϕ of the regularization term is chosen to be $\phi = (1 - \delta)/2$.

Note that the results of this paper still hold if the gradient and Hessian of the function f are restricted Lipschitz continuous only over a bounded region. Here, for simplicity we assume that these properties hold for all low-rank matrices.

As mentioned before Definition 1, the RIP-related Assumption 2 is a widely used assumption in studying the landscape of low-rank matrix recovery problems, which is satisfied in a variety of problems, such as those for which f is given by a sufficiently large number of random Gaussian linear measurements (Candès and Plan 2011). Moreover, in the case when the function f does not satisfy the RIP assumption globally, it often satisfies RIP in a neighborhood of the global minimizers, and the theorems in this paper can still be applied to obtain local convergence results.

For the asymmetric problem, it can be verified that, by choosing the coefficient ϕ of the regularization term as in Assumption 4, the function F in (7) satisfies the $2\delta/(1 + \delta)$ -RIP $_{2r}$ property after scaling (see Zhang, Bi, and Lavaei (2021)). Other values of ϕ can also lead to the RIP property on F , but the specific value in Assumption 4 is the one minimizing the RIP constant. Furthermore, if $M^* = U^* V^{*T}$ is a balanced factorization of the ground truth M^* , then

$$\tilde{M}^* = \begin{bmatrix} U^* \\ V^* \end{bmatrix} [U^{*T} \quad V^{*T}] \in \mathbb{R}^{(n+m) \times (n+m)} \quad (8)$$

is called the augmented ground truth, which is obviously a global minimizer of the transformed asymmetric problem (7). \tilde{M}^* is independent of the factorization (U^*, V^*) , and

$$\|\tilde{M}^*\|_F = 2\|M^*\|_F \leq 2D, \quad \sigma_r(\tilde{M}^*) = 2\sigma_r(M^*).$$

We include the proofs of the above statements in Appendix A for completeness. In addition, we prove in Appendix A that the gradient and Hessian of the function g_s in

the symmetric problem (4) and those of the function g_a in the asymmetric problem (5) share the same Lipschitz property over a bounded region. Using the above observations, one can translate any results developed for symmetric problems to similar results for asymmetric problems by simply replacing δ with $2\delta/(1 + \delta)$, D with $2D$, and $\sigma_r(M^*)$ with $2\sigma_r(M^*)$.

Related Works

The low-rank matrix recovery problem has been investigated in numerous papers. In this section, we focus on the existing results related to the linear convergence for the factored problems (4) and (5) solved by local search methods.

The major previous results on the local regularity property are summarized in Table 1. In this table, each number in the last column reported for the existing works denotes the radius R such that their respective objective functions g satisfy the (α, β) -regularity condition

$$\langle \nabla g(X), X - \mathcal{P}_{\mathcal{Z}}(X) \rangle \geq \frac{\alpha}{2} \text{dist}(X, \mathcal{Z})^2 + \frac{1}{2\beta} \|\nabla g(X)\|_F^2$$

for all matrices X with $\text{dist}(X, \mathcal{Z}) \leq R$. Here, \mathcal{Z} is the set of global minimizers, and $\mathcal{P}_{\mathcal{Z}}(X)$ is a global minimizer $Z \in \mathcal{Z}$ that is the closest to X . The (α, β) -regularity condition is slightly weaker than the strong convexity condition, and it can lead to linear convergence on the same region. In Table 1, we do not include specialized results that are only applicable to a specific objective (Jin et al. 2017; Hou, Li, and Zhang 2020), or probabilistic results for randomly generated measurements (Zheng and Lafferty 2015). Moreover, Li and Lin (2020); Zhou, Cao, and Gu (2020) used the accelerated gradient descent to obtain a faster convergence rate, but their convergence regions are even smaller than the ones based on the (α, β) -regularity condition as listed in Table 1. Each number in the last column reported for our results refers to the radius of the region satisfying the PL inequality, which is a weaker condition than the (α, β) -regularity condition while offering the same convergence rate guarantee. It can be observed that we have identified far larger regions than the existing ones under weaker RIP assumptions by replacing the (α, β) -regularity condition with the PL inequality.

Regarding the existing global convergence results for the low-rank matrix recovery problem, Jin et al. (2017) established the global linear convergence for the symmetric problem in a very specialized case with f_s being a quadratic loss function, and Tu et al. (2016) proposed the Procrustes flow method with the global linear convergence for the linear measurement case under the assumption that the function f_s satisfies the $1/10$ -RIP $_{6r}$ property for symmetric problems or the function f_a satisfies the $1/25$ -RIP $_{6r}$ property for asymmetric problems under a careful initialization. Zhao, Wang, and Liu (2015) established the global linear convergence for asymmetric problems with linear measurements under the assumption that f_a satisfies δ -RIP $_{2r}$ with $\delta \leq O(1/r)$ using alternating exact minimization over variables U and V in (5). In addition, the strict saddle property proven in Ge, Jin, and Zheng (2017) leads to the global lin-

ear convergence of perturbed gradient methods for the linear measurement case under the $1/10$ -RIP $_{2r}$ assumption for symmetric problems and the $1/20$ -RIP $_{2r}$ assumption for asymmetric problems. Later, Zhu et al. (2018) proved a weaker strict saddle property under the $1/5$ -RIP $_{2r,4r}$ assumption for symmetric problems with general objectives, while Li, Zhu, and Tang (2017) proved the same weaker property under the $1/5$ -RIP $_{2r,4r}$ assumption for asymmetric problems with general objectives and nuclear norm regularization. Our results requiring the δ -RIP $_{2r}$ property with $\delta < 1/2$ for symmetric problems with general objectives and the δ -RIP $_{2r}$ property with $\delta < 1/3$ for asymmetric problems with general objectives depend on significantly weaker RIP assumptions and thus can be applied to a broader class of problems, which is a major improvement over all previous results on the global linear convergence.

Besides local search methods for the factored problems, there are other approaches for tackling the low-rank matrix recovery. Earlier works such as Candès and Recht (2009); Recht, Fazel, and Parrilo (2010) solved the original nonconvex problems based on convex relaxations. Although they can achieve good performance guarantees under the RIP assumptions, they are not suitable for large-scale problems. Other approaches for solving the low-rank matrix recovery include applying the inertial proximal gradient descent method directly to the original objective functions without factoring the decision variable M (Dutta et al. 2020). However, it may converge to an arbitrary critical point, while in this paper we show that RIP-based local search methods can guarantee the global convergence to a global minimum.

Local Convergence

In this section, we present the local regularity results for problems (4) and (5), which state that the functions g_s and g_a satisfy the PL inequality locally, leading to local linear convergence results for the gradient descent method. The proofs are delegated to Appendix B.

First, we consider the symmetric problem (4). The development of the local PL inequality for this problem is enlightened by the high-level idea behind the proof of the absence of spurious local minima in Zhang, Sojoudi, and Lavaei (2019); Zhang and Zhang (2020); Bi and Lavaei (2020). The objective is to find a function f_s^* corresponding to the worst-case scenario, meaning that it satisfies the δ -RIP $_{2r}$ property with the smallest possible δ while the PL inequality is violated at a particular matrix X . This is achieved by designing a semidefinite program parameterized by X with constraints implied by the δ -RIP $_{2r}$ property and the negation of the PL inequality. Denote the optimal value of the semidefinite program by $\delta_f^*(X)$. If a given function f_s satisfies δ -RIP $_{2r}$ with $\delta < \delta_f^*(X)$ for all $X \in \mathbb{R}^{n \times r}$ in a neighborhood of the global minimizers, it can be concluded that the PL inequality holds for all matrices in this neighborhood.

Lemma 1. *Consider the symmetric problem (4) and an arbitrary positive number \tilde{C} satisfying*

$$\tilde{C} < \sqrt{2(\sqrt{2} - 1)\sqrt{1 - \delta^2}\sigma_r(M^*)^{1/2}}. \quad (9)$$

Paper	Objective	Assumption	Radius of Local Regularity Region
Bhojanapalli, Kyrillidis, and Sanghavi (2016)	S/G	f_s Convex, $\delta_{2r} \leq \delta$	$\frac{1}{100} \frac{1-\delta}{1+\delta} \frac{\sigma_r(M^*)}{\sigma_1(M^*)} \sigma_r(M^*)^{1/2}$
Tu et al. (2016)	S/L	$\delta_{6r} \leq 1/10$	$\frac{1}{4} \sigma_r(M^*)^{1/2}$
Tu et al. (2016)	A/L	$\delta_{6r} \leq 1/25$	$\frac{1}{4} \sigma_r(M^*)^{1/2}$
Park et al. (2018)	A/G	f_a Convex, $\delta_{2r} \leq \delta$	$\frac{\sqrt{2}}{10} \sqrt{\frac{1-\delta}{1+\delta}} \sigma_r(M^*)^{1/2}$
Zhu et al. (2021)	A/G	$\delta_{2r,4r} \leq 1/50$	$\sigma_r(M^*)^{1/2}$
Ours	S/G	$\delta_{2r} \leq \delta$	$0.91\sqrt{1-\delta^2} \sigma_r(M^*)^{1/2}$
Ours	A/G	$\delta_{2r} \leq \delta$	$1.29 \frac{\sqrt{1+2\delta-3\delta^2}}{1+\delta} \sigma_r(M^*)^{1/2}$

Table 1: Previous local regularity results for the low-rank matrix recovery problems and the comparison with our results (“S”, “A”, “L” and “G” stand for the symmetric case, asymmetric case, linear measurement and general nonlinear function).

There exists a constant $\mu > 0$ such that the PL inequality

$$\frac{1}{2} \|\nabla g_s(X)\|_F^2 \geq \mu(g_s(X) - f_s(M^*))$$

holds for all matrices in the region

$$\{X \in \mathbb{R}^{n \times r} \mid \text{dist}(X, \mathcal{Z}) \leq \tilde{C}\}, \quad (10)$$

where \mathcal{Z} is the set of global minimizers of the problem (4).

In the above, note that $\sigma_r(M^*)$ and thus \tilde{C} are always positive because M^* is assumed to be rank r . Both the (α, β) -regularity condition used in the prior literature and the PL inequality deployed here guarantee a linear convergence if it is already known that the trajectory at all iterations remains within the region in which the associated condition holds. However, there is a key difference between these two conditions. The (α, β) -regularity condition ensures that $\text{dist}(X, \mathcal{Z})$ is nonincreasing during the iterations under a sufficiently small step size, and thus the trajectory never leaves the local neighborhood. In contrast, the weaker PL inequality may not be able to guarantee this property. To resolve this issue, in our convergence proof we will adopt a different distance function given by $\|XX^T - M^*\|_F$. By Taylor’s formula and the definition of the δ -RIP $_{2r}$ property, we have

$$\begin{aligned} \frac{1-\delta}{2} \|M - M^*\|_F^2 &\leq f_s(M) - f_s(M^*) \\ &\leq \frac{1+\delta}{2} \|M - M^*\|_F^2, \end{aligned} \quad (11)$$

for all matrices $M \in \mathbb{R}^{n \times n}$ with $\text{rank}(M) \leq r$. Therefore, if $M, M' \in \mathbb{R}^{n \times n}$ are two matrices such that $f_s(M) \leq f_s(M')$, then the inequality (11) implies that

$$\|M - M^*\|_F \leq \sqrt{\frac{1+\delta}{1-\delta}} \|M' - M^*\|_F. \quad (12)$$

Therefore, the distance function $\|XX^T - M^*\|_F$ is almost nonincreasing if the function value $g_s(X)$ does not increase. Combining this idea with the local PL inequality proved in Lemma 1, we obtain the next local convergence result.

Theorem 2. For the symmetric problem (4), the gradient descent method converges to the optimal solution linearly if the initial point X_0 satisfies

$$\|X_0 X_0^T - M^*\|_F < 2(\sqrt{2} - 1)(1 - \delta)\sigma_r(M^*)$$

and the step size η satisfies

$$1/\eta \geq 12\rho_1 r^{1/2} \left(\sqrt{\frac{1+\delta}{1-\delta}} \|X_0 X_0^T - M^*\|_F + D \right).$$

Specifically, there exists some constant $\mu > 0$ (which depends on X_0 but not on η) such that

$$\|X_t X_t^T - M^*\|_F \leq (1 - \mu\eta)^{t/2} \sqrt{\frac{1+\delta}{1-\delta}} \|X_0 X_0^T - M^*\|_F, \quad \forall t \in \{0, 1, \dots\}, \quad (13)$$

where X_t denotes the output of the algorithm at iteration t .

Note that since the left-hand side of (13) is nonnegative, we have $0 \leq 1 - \mu\eta \leq 1$. As a remark, although our bound on the step size η in Theorem 2 seems complex, it essentially says that η needs to be small, and the upper bound on the acceptable values of the step size can be explicitly calculated out routinely after all the parameters of the problem are given. Furthermore, using the transformation from asymmetric problems to symmetric problems, one can obtain parallel results for the asymmetric problem (5) as below.

Theorem 3. Consider the asymmetric problem (5). The PL inequality is satisfied in the region

$$\{X \in \mathbb{R}^{(n+m) \times r} \mid \text{dist}(X, \mathcal{Z}) \leq \tilde{C}\},$$

where \mathcal{Z} denotes the set of global minimizers and

$$\tilde{C} < 2\sqrt{\sqrt{2}-1} \frac{\sqrt{1+2\delta-3\delta^2}}{1+\delta} \sigma_r(M^*)^{1/2}.$$

Moreover, local linear convergence is guaranteed for the gradient descent method if the initial point X_0 satisfies

$$\|X_0 X_0^T - \tilde{M}^*\|_F < 4(\sqrt{2}-1) \frac{1-\delta}{1+\delta} \sigma_r(M^*)$$

and the step size η satisfies

$$1/\eta \geq 12\rho_1 r^{1/2} \left(\sqrt{\frac{1+3\delta}{1-\delta}} \|X_0 X_0^T - \tilde{M}^*\|_F + 2D \right).$$

Global Convergence

Having developed local convergence results, the next step is to design an algorithm whose trajectory will eventually enter the local convergence region from any initial point. The major challenge is to deal with the saddle points outside the local regularity region. One common approach is the perturbed gradient descent method, which adds random noise to jump out of a neighborhood of a strict saddle point. Using the symmetric problem as an example, the basic idea is to first use the analysis in Jin et al. (2017) to show that the perturbed gradient descent method will successfully find a matrix X that approximately satisfies the first-order and second-order necessary optimality conditions, i.e.,

$$\|\nabla g_s(X)\|_F \leq \kappa, \quad \lambda_{\min}(\nabla^2 g_s(X)) \geq -\kappa, \quad (14)$$

after a certain number of iterations where the number depends on κ . Here, $\lambda_{\min}(\nabla^2 g_s(X))$ denotes the minimum eigenvalue of the matrix \mathbf{G} that satisfies the equation

$$(\text{vec}(U))^T \mathbf{G} \text{vec}(V) = [\nabla^2 g_s(X)](U, V),$$

for all $U, V \in \mathbb{R}^{n \times r}$. The second step is to prove the strict saddle property for the problem, which means that for appropriate values of κ the two conditions in (14) imply that $\|X X^T - M^*\|_F$ is so small that X is in the local convergence region given by Theorem 2. After this iteration, the algorithm switches to the simple gradient descent method. This two-phase algorithm is commonly called the perturbed gradient descent method with local improvement (Jin et al. 2017), whose details are given by Algorithm 1 in Appendix C. The proofs in this section are also given in Appendix C.

In this section, we will present two conditions that guarantee the global linear convergence of the above algorithm. For symmetric problems, the next lemma provides the strict saddle property and fulfills the purpose for the second step mentioned above. Its proof is a generalization of the one for the absence of spurious local minima under the same assumption in Zhang (2021).

Lemma 4. *Consider the symmetric problem (4) with $\delta < 1/2$. For every $C > 0$, there exists some $\kappa > 0$ such that for every $X \in \mathbb{R}^{n \times r}$ the two conditions given in (14) will imply $\|X X^T - Z Z^T\|_F < C$.*

The remaining step is to show that the trajectory of the perturbed gradient descent method will always belong to a bounded region in which the gradient and Hessian of the objective g_s are Lipschitz continuous (see Appendix C). Combining the above results with Theorem 3 in Jin et al. (2017), we can obtain the following global linear convergence result.

Theorem 5. *Consider the symmetric problem (4) with $\delta < 1/2$. For every $\epsilon > 0$, the perturbed gradient descent method with local improvement under a suitable step size η and perturbation size w finds a solution \hat{X} satisfying $\|\hat{X} \hat{X}^T - M^*\|_F \leq \epsilon$ with high probability in $O(\log(1/\epsilon))$ number of iterations. Here, η and w are defined in Algorithm 1 in Appendix C.*

In the above theorem, the order $O(\log(1/\epsilon))$ of the convergence rate is determined by the number of iterations spent in the second phase of the algorithm, because the number of iterations in the first phase is independent of ϵ . Note that we only show the relationship between the number of iterations and ϵ , but the convergence rate also depends on the initial point X_0 and the loss function f_s . Moreover, although not being related to the final convergence rate, Theorem 3 in Jin et al. (2017) also shows that the number of iterations in the first phase is polynomial with respect to the problem size.

For asymmetric problems with arbitrary objectives and rank r , if we apply the transformation from asymmetric problems to symmetric problems and replace δ in Theorem 5 with $2\delta/(1+\delta)$, Theorem 5 immediately implies the following global linear convergence result.

Theorem 6. *Consider the asymmetric problem (5) with $\delta < 1/3$. For every $\epsilon > 0$, the perturbed gradient descent method with local improvement under a suitable step size η and perturbation size w finds a solution \hat{X} satisfying $\|\hat{X} \hat{X}^T - \tilde{M}^*\|_F \leq \epsilon$ with high probability in $O(\log(1/\epsilon))$ number of iterations.*

Numerical Illustration

In this section, we conduct numerical experiments to demonstrate the behavior of the perturbed gradient descent algorithm for solving low-rank matrix recovery problems. The linear convergence rate observed for the examples below supports our theoretical analyses in the previous two sections.

In the first experiment, we consider the loss function (3) induced by a linear operator \mathcal{A} with

$$\mathcal{A}(M) = (\langle A_1, M \rangle, \dots, \langle A_p, M \rangle).$$

Here, each entry of A_i is independently generated from the standard Gaussian distribution. As shown in Candès and Plan (2011), such linear operator \mathcal{A} satisfies RIP with high probability if the number of measurements is large enough. Since it is NP-hard to check whether the resulting loss function f_s or f_a satisfies the δ -RIP $_{2r}$ for certain δ , the δ parameter is estimated as follows: For the symmetric problem (4), we first generate 10^4 random matrices $X \in \mathbb{R}^{n \times 2r}$ with each entry independently selected from the standard Gaussian distribution, and then find the proper scaling factor $a \in \mathbb{R}$ and the smallest δ such that

$$(1-\delta)\|X X^T\|_F^2 \leq \|a\mathcal{A}(X X^T)\|^2 \leq (1+\delta)\|X X^T\|_F^2$$

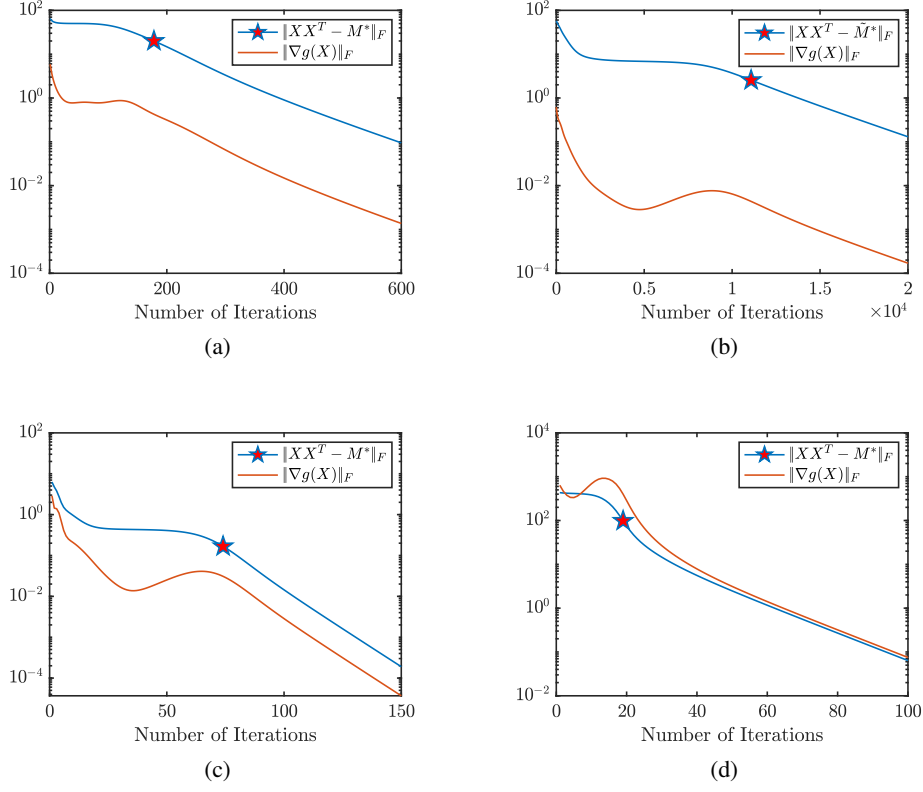


Figure 1: The trajectory of the perturbed gradient descent method for solving the low-rank matrix recovery problem. The marker in each figure shows the boundary of the local convergence region provided by Theorem 2. (a) A symmetric linear problem with $r = 1$, $n = 40$, $p = 120$ and δ estimated to be 0.49. (b) An asymmetric linear problem with $r = 5$, $n = 10$, $m = 8$, $p = 220$ and δ estimated to be 0.32. (c) The 1-bit matrix recovery problem with $r = 5$, $n = 10$. (d) The 1-bit matrix recovery problem with $r = 2$, $n = 600$.

holds for all generated matrices X . The δ parameter for the asymmetric problem (5) can be estimated similarly. After that, the ground truth $M^* = XX^T$ or $M^* = UV^T$ is generated randomly with each entry of X or (U, V) independently selected from the standard Gaussian distribution. The initial point is generated in the same way.

Figure 1(a) and (b) show the difference between the obtained solution and the ground truth together with the norm of the gradient of the objective function at different iterations. The convergence behavior clearly divides into two stages. The convergence rate is sublinear initially and then switches to linear when the current point moves into the local region associated with the PL inequality. In Figure 1(a) and (b), the marker shows the first time when the current point falls into the local convergence region provide in Theorem 2 or Theorem 3. It can be seen that these theorems predict the boundary of the transition from a sublinear convergence rate to the linear convergence rate fairly tightly. After this point, $O(\log(1/\epsilon))$ additional iterations are needed to find an approximate solution with accuracy ϵ . On the other hand, the occasion when perturbation needs to be added is rare in practice since it is unlikely for the trajectory to be very close to a saddle point. However, such perturbation is necessary theoretically to deal with pathological cases.

Second, we consider the 1-bit matrix recovery (Davenport et al. 2014) with full measurements, which is a nonlinear low-rank matrix recovery problem. In this problem, there is an unknown symmetric ground truth matrix $\hat{M} \in \mathbb{R}^{n \times n}$ with $\hat{M} \succeq 0$ and $\text{rank}(\hat{M}) = r$. One is allowed to take independent measurements on every entry \hat{M}_{ij} , where each measurement value is a binary random variable whose distribution is given by $Y_{ij} = 1$ with probability $\sigma(\hat{M}_{ij})$ and $Y_{ij} = 0$ otherwise. Here, $\sigma(x)$ is commonly chosen to be the sigmoid function $e^x/(e^x + 1)$. After a number of measurements are taken, let y_{ij} be the percentage of the measurements on the (i, j) -th entry that are equal to 1. The goal is to find the maximum likelihood estimator for the ground truth \hat{M} , which can be formulated as finding the global minimizer M^* of the problem (4) with

$$f_s(M) = - \sum_{i=1}^n \sum_{j=1}^n (y_{ij} M_{ij} - \log(1 + e^{M_{ij}})).$$

To establish the RIP condition for the function f_s above, consider its Hessian $\nabla^2 f_s(M)$ that is given by

$$[\nabla^2 f_s(M)](K, L) = \sum_{i=1}^n \sum_{j=1}^n \sigma'(M_{ij}) K_{ij} L_{ij},$$

for every $M, K, L \in \mathbb{R}^{n \times n}$. On the region

$$\{M \in \mathbb{R}^{n \times n} \mid |M_{ij}| \leq 2.29, \forall i, j = 1, \dots, n\}, \quad (15)$$

we have $1/12 < \sigma'(M_{ij}) \leq 1/4$, and thus the function f_s satisfies the δ -RIP $_{2r}$ property with $\delta < 1/2$.

Note that due to the noisy measurements the global minimizer M^* is not equal to \hat{M} in general. However, for demonstration purposes we should know M^* a priori, and hence we consider the case when the number of measurements is large enough such that $y_{ij} = \sigma(\hat{M}_{ij})$ and $M^* = \hat{M}$. In Figure 1(c), the ground truth and the initial point are generated randomly in the region (15). Here, we can observe a similar two-stage convergence behavior as in the example with linear measurements. We experiment on the same problem with a larger matrix size n as shown in Figure 1(d), which also gives similar results.

Conclusion

In this paper, we study the local and global convergence behaviors of gradient-based local search methods for solving low-rank matrix recovery problems in both symmetric and asymmetric cases. First, we present a novel method to identify a local region in which the PL inequality is satisfied, which is significantly larger than the region associated with the regularity conditions proven in the prior literature. This leads to a linear convergence result for the gradient descent method over a large local region. Second, we develop the strict saddle property for symmetric problems under the δ -RIP $_{2r}$ property with $\delta < 1/2$. Then, we prove the global linear convergence of the perturbed gradient descent method for symmetric problems under the δ -RIP $_{2r}$ property with $\delta < 1/2$, and the same convergence property can also be guaranteed for asymmetric problems with $\delta < 1/3$. Compared with the existing results, these conditions are remarkably weaker and can be applied to a larger class of problems.

Acknowledgments

This work was supported by grants from AFOSR, ARO, ONR, and NSF.

Appendix

Appendix is available in the full version of this paper at <https://arxiv.org/abs/2104.13348>.

References

Bhojanapalli, S.; Kyrillidis, A.; and Sanghavi, S. 2016. Dropping Convexity for Faster Semi-definite Optimization. In *29th Annual Conference on Learning Theory*, volume 49 of *Proceedings of Machine Learning Research*, 530–582.

Bhojanapalli, S.; Neyshabur, B.; and Srebro, N. 2016. Global Optimality of Local Search for Low Rank Matrix Recovery. In *Advances in Neural Information Processing Systems*, volume 29.

Bi, Y.; and Lavaei, J. 2020. Global and Local Analyses of Nonlinear Low-Rank Matrix Recovery Problems. ArXiv:2010.04349.

Candès, E. J.; and Plan, Y. 2011. Tight Oracle Inequalities for Low-Rank Matrix Recovery From a Minimal Number of Noisy Random Measurements. *IEEE Trans. Inf. Theory*, 57(4): 2342–2359.

Candès, E. J.; and Recht, B. 2009. Exact Matrix Completion via Convex Optimization. *Found. Comput. Math.*, 9(6): 717–772.

Davenport, M. A.; Plan, Y.; van den Berg, E.; and Wootters, M. 2014. 1-Bit Matrix Completion. *Inf. Inference*, 3(3): 189–223.

Dutta, A.; Hanzely, F.; Liang, J.; and Richtárik, P. 2020. Best Pair Formulation & Accelerated Scheme for Non-Convex Principal Component Pursuit. *IEEE Trans. Signal Process.*, 68: 6128–6141.

Fattahi, S.; and Sojoudi, S. 2020. Exact Guarantees on the Absence of Spurious Local Minima for Non-negative Rank-1 Robust Principal Component Analysis. *J. Mach. Learn. Res.*, 21(59): 1–51.

Ge, R.; Huang, F.; Jin, C.; and Yuan, Y. 2015. Escaping From Saddle Points — Online Stochastic Gradient for Tensor Decomposition. In *Proceedings of the 28th Conference on Learning Theory*, volume 40 of *Proceedings of Machine Learning Research*, 797–842.

Ge, R.; Jin, C.; and Zheng, Y. 2017. No Spurious Local Minima in Nonconvex Low Rank Problems: A Unified Geometric Analysis. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, 1233–1242.

Ha, W.; Liu, H.; and Barber, R. F. 2020. An Equivalence Between Critical Points for Rank Constraints Versus Low-Rank Factorizations. *SIAM J. Optim.*, 30(4): 2927–2955.

Hou, T. Y.; Li, Z.; and Zhang, Z. 2020. Fast Global Convergence for Low-Rank Matrix Recovery via Riemannian Gradient Descent With Random Initialization. ArXiv:2012.15467.

Jin, C.; Ge, R.; Netrapalli, P.; Kakade, S. M.; and Jordan, M. I. 2017. How to Escape Saddle Points Efficiently. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, 1724–1732.

Jin, M.; Lavaei, J.; Sojoudi, S.; and Baldick, R. 2021. Boundary Defense Against Cyber Threat for Power System State Estimation. *IEEE Trans. Inf. Forensics Secur.*, 16: 1752–1767.

Koren, Y.; Bell, R.; and Volinsky, C. 2009. Matrix Factorization Techniques for Recommender Systems. *Computer*, 42(8): 30–37.

Lee, J. D.; Simchowitz, M.; Jordan, M. I.; and Recht, B. 2016. Gradient Descent Only Converges to Minimizers. In *29th Annual Conference on Learning Theory*, volume 49 of *Proceedings of Machine Learning Research*, 1246–1257.

Li, H.; and Lin, Z. 2020. Provable Accelerated Gradient Method for Nonconvex Low Rank Optimization. *Mach. Learn.*, 109(1): 103–134.

Li, Q.; Zhu, Z.; and Tang, G. 2017. Geometry of Factored Nuclear Norm Regularization. ArXiv:1704.01265.

- Park, D.; Kyriallidis, A.; Caramanis, C.; and Sanghavi, S. 2018. Finding Low-Rank Solutions via Nonconvex Matrix Factorization, Efficiently and Provably. *SIAM J. Imaging Sci.*, 11(4): 2165–2204.
- Recht, B.; Fazel, M.; and Parrilo, P. A. 2010. Guaranteed Minimum-Rank Solutions of Linear Matrix Equations via Nuclear Norm Minimization. *SIAM Rev.*, 52(3): 471–501.
- Tu, S.; Boczar, R.; Simchowitz, M.; Soltanolkotabi, M.; and Recht, B. 2016. Low-Rank Solutions of Linear Matrix Equations via Procrustes Flow. In *Proceedings of the 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, 964–973.
- Zhang, G.; and Zhang, R. Y. 2020. How Many Samples Is a Good Initial Point Worth in Low-Rank Matrix Recovery? In *Advances in Neural Information Processing Systems*, volume 33, 12583–12592.
- Zhang, H.; Bi, Y.; and Lavaei, J. 2021. General Low-Rank Matrix Optimization: Geometric Analysis and Sharper Bounds. In *Advances in Neural Information Processing Systems*, volume 34.
- Zhang, R. Y. 2021. Sharp Global Guarantees for Nonconvex Low-Rank Matrix Recovery in the Overparameterized Regime. ArXiv:2104.10790.
- Zhang, R. Y.; Jozs, C.; Sojoudi, S.; and Lavaei, J. 2018a. How Much Restricted Isometry Is Needed in Nonconvex Matrix Recovery? In *Advances in Neural Information Processing Systems*, volume 31.
- Zhang, R. Y.; Sojoudi, S.; and Lavaei, J. 2019. Sharp Restricted Isometry Bounds for the Inexistence of Spurious Local Minima in Nonconvex Matrix Recovery. *J. Mach. Learn. Res.*, 20(114): 1–34.
- Zhang, X.; Wang, L.; Yu, Y.; and Gu, Q. 2018b. A Primal-Dual Analysis of Global Optimality in Nonconvex Low-Rank Matrix Recovery. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, 5862–5871.
- Zhang, Y.; Madani, R.; and Lavaei, J. 2018. Conic Relaxations for Power System State Estimation With Line Measurements. *IEEE Control Netw. Syst.*, 5(3): 1193–1205.
- Zhao, T.; Wang, Z.; and Liu, H. 2015. A Nonconvex Optimization Framework for Low Rank Matrix Estimation. In *Advances in Neural Information Processing Systems*, volume 28.
- Zheng, Q.; and Lafferty, J. 2015. A Convergent Gradient Descent Algorithm for Rank Minimization and Semidefinite Programming From Random Linear Measurements. In *Advances in Neural Information Processing Systems*, volume 28.
- Zhou, D.; Cao, Y.; and Gu, Q. 2020. Accelerated Factored Gradient Descent for Low-Rank Matrix Factorization. In *Proceedings of the 23th International Conference on Artificial Intelligence and Statistics*, volume 108 of *Proceedings of Machine Learning Research*, 4430–4440.
- Zhou, X.; Yang, C.; and Yu, W. 2013. Moving Object Detection by Detecting Contiguous Outliers in the Low-Rank Representation. *IEEE Trans. Pattern Anal. Mach. Intell.*, 35(3): 597–610.
- Zhu, Z.; Li, Q.; Tang, G.; and Wakin, M. B. 2018. Global Optimality in Low-Rank Matrix Optimization. *IEEE Trans. Signal Process.*, 66(13): 3614–3628.
- Zhu, Z.; Li, Q.; Tang, G.; and Wakin, M. B. 2021. The Global Optimization Geometry of Low-Rank Matrix Optimization. *IEEE Trans. Inf. Theory*, 67(2): 1308–1331.