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

Yingjie Bi

Industrial Engineering and Operations Research
University of California, Berkeley
yingjiebi@berkeley.edu

Haixiang Zhang

Department of Mathematics
University of California, Berkeley
haixiang.zhang@berkeley.edu

Javad Lavaei

Industrial Engineering and Operations Research
University of California, Berkeley
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.

1 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 et al., 2009) and motion detection (Zhou et al., 2013; Fattahi and Sojoudi, 2020), and in engineering problems, such as power system state estimation (Zhang et al., 2018c).

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:

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

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 nonsquare 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 methods, 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 used in the literature such as Bhojanapalli et al. (2016b); Ge et al. (2017); Zhu et al. (2018); Zhang et al. (2018b,a, 2019); Ha et al. (2020); Zhang and Zhang (2020); Bi and Lavaei (2020); Zhang et al. (2021); Zhang (2021).

Definition 1 (Recht et al. (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 nonexistence of spurious local minima are presented in Zhang et al. (2021); Zhang (2021). Zhang et al. (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 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. As surveyed in Section 2 below, the RIP constant to ensure linear convergence in the previous results is much smaller than the RIP constant needed to ensure the absence of spurious local minima. The gap between these two types of bounds will be addressed in this paper.

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 above two aspects for the low-rank matrix recovery problem separately.

1.1 Summary of main contributions

For the local convergence, we prove in Section 3 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 in Section 2, 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 any δ , we 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 is also derived for the asymmetric problem (5). Based on these results, local linear convergence is then established. The enlarged local regularity regions identified in this work are useful for two reasons: i) to reduce the number of iterations needed by various saddle-escaping algorithms to reach the local convergence region, ii) to mitigate the difficulty of finding an initial point that will converge to the ground truth in case the RIP constant is close to 1 and global convergence cannot be established due to the existence of spurious solutions.

For the global convergence, in Section 4, we first study the symmetric problem (4) with linear measurements by proving that the objective g_s satisfies the strict saddle property if the function f_s has the δ -RIP $_{2r}$ property for $\delta < 1/2$. Note that this result is sharp, because in Zhang et al. (2018a)

a counterexample has been found with the existence of spurious local minima under $\delta = 1/2$. Using the above strict saddle property and the local convergence result proven in Section 3, 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 arbitrary tolerance ϵ . Moreover, the global linear convergence is also proven for both the symmetric and asymmetric problems for arbitrary objectives and an arbitrary rank r under the δ -RIP $_{2r}$ condition with $\delta < 1/3$.

1.2 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

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

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 $f(M)$ in the following 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 matrix 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 is restricted Lipschitz continuous only over a bounded region. Here, for simplicity we assume that these properties hold for all low-rank matrices.

For the asymmetric problem, it can be verified that the function F in (7) satisfies the $2\delta/(1 + \delta)$ -RIP $_{2r}$ property after scaling (see Zhang et al. (2021)). Furthermore, if $M^* = U^* V^{*T}$ is a balanced factorization of the ground truth M^* , then

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

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, the asymmetric case, linear measurement and general nonlinear function)

Paper	Objective	Assumption	Radius of Local Regularity Region
Bhojanapalli et al. (2016a)	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	$\delta_{2r} \leq \delta$	$1.29 \frac{\sqrt{1+2\delta-3\delta^2}}{1+\delta} \sigma_r(M^*)^{1/2}$

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 convert 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 $2\sigma_r(M^*)$.

2 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 et al., 2020), or probabilistic results for randomly generated measurements (Zheng and Lafferty, 2015). Moreover, Li and Lin (2020); Zhou et al. (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 refer 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, 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 et al. (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 et al. (2017) leads to the global linear 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. Our results requiring the δ -RIP $_{2r}$ property with $\delta < 1/2$ for symmetric problems with linear measurements and the δ -RIP $_{2r}$ property with $\delta < 1/3$ for both symmetric and asymmetric problems with general objectives significantly improve all previous results on the global linear convergence.

3 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 et al. (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 at all matrices in this neighborhood.

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

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

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).

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

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

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 Theorem 1, we obtain the following 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 side of (13) is nonnegative, we have $0 \leq 1 - \mu\eta \leq 1$. Using the transformation from asymmetric problems to symmetric problems, one can obtain parallel results for the asymmetric problem (5) as stated 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 initialization 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

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

4 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 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). The details together with the proofs are 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 with linear measurements, 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 $\delta < 1/2$ and the loss function f_s in the form (3) given by linear measurements. For every $C > 0$, there exists some $\kappa > 0$ such that for each $X \in \mathbb{R}^{n \times r}$ the two conditions (14) will imply $\|XX^T - ZZ^T\|_F < C$.

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 $\delta < 1/2$ and the loss function f_s in the form (3) given by linear measurements. 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.

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 in the problem.

For general symmetric and asymmetric problems with arbitrary objectives and rank r , we will utilize the following result that serves a similar purpose as Lemma 4, which is a consequence of Theorem 6 and Theorem 7 in Zhang et al. (2021).

Lemma 6. Suppose that $\delta < 1/3$. For every $C > 0$, there exists some $\kappa > 0$ such that for each $X \in \mathbb{R}^{n \times r}$ the two conditions (14) will imply $\|XX^T - ZZ^T\|_F < C$ for the symmetric problem (4). A similar result holds for the asymmetric problem (5) if M^* is replaced with \tilde{M}^* .

Similar to the relationship between Lemma 4 and Theorem 5, Lemma 6 can be leveraged to obtain the following global linear convergence result.

Theorem 7. Given $\delta < 1/3$, 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 $\|\hat{X}\hat{X}^T - M^*\|_F \leq \epsilon$ for the symmetric problem (4) or $\|\hat{X}\hat{X}^T - \tilde{M}^*\|_F \leq \epsilon$ for the asymmetric problem (5) with high probability in $O(\log(1/\epsilon))$ number of iterations.

5 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 Section 3 and Section 4.

In the experiments, 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 p 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

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

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 shows 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



(a) A symmetric problem with $r = 1$, $n = 40$, $p = 120$ and δ estimated to be 0.49.

(b) An asymmetric problem with $r = 5$, $n = 10$, $m = 8$, $p = 220$ and δ estimated to be 0.32.

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.

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, 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 become very close to a saddle point. However, such perturbation is necessary theoretically to deal with pathological cases.

6 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 with linear measurements under the δ -RIP $_{2r}$ property with $\delta < 1/2$. Then, we present two conditions that guarantee the global convergence of the perturbed gradient method in linear time. Compared with the existing results, these conditions are remarkably weaker and can be applied to a larger class of low-rank matrix recovery problems.

References

- Bhojanapalli, S., Kyrillidis, A., and Sanghavi, S. (2016a). Dropping convexity for faster semi-definite optimization. In *29th Annual Conference on Learning Theory*, volume 49 of *Proceedings of Machine Learning Research*, pages 530–582.
- Bhojanapalli, S., Neyshabur, B., and Srebro, N. (2016b). 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.
- Davenport, M. A., Plan, Y., van den Berg, E., and Wootters, M. (2014). 1-bit matrix completion. *Inf. Inference*, 3(3):189–223.

- 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*, pages 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*, pages 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*, pages 1724–1732.
- Karimi, H., Nutini, J., and Schmidt, M. (2016). Linear convergence of gradient and proximal-gradient methods under the Polyak–Lojasiewicz condition. In *Machine Learning and Knowledge Discovery in Databases*, volume 9851 of *Lecture Notes in Computer Science*, pages 795–811.
- 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*, pages 1246–1257.
- Li, H. and Lin, Z. (2020). Provable accelerated gradient method for nonconvex low rank optimization. *Mach. Learn.*, 109(1):103–134.
- Park, D., Kyriillidis, 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*, pages 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, pages 12583–12592.
- Zhang, H., Bi, Y., and Lavaei, J. (2021). General low-rank matrix optimization: Geometric analysis and sharper bounds. arXiv:2104.10356.
- Zhang, R. Y. (2021). Sharp global guarantees for nonconvex low-rank matrix recovery in the over-parameterized 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*, pages 5862–5871.

- Zhang, Y., Madani, R., and Lavaei, J. (2018c). 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*, pages 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.

A Properties of the factored objectives

We first study the smoothness properties for the gradient and Hessian of the objective function g_s in the symmetric problem (4). The following lemma is borrowed from the proof of Theorem 7 in Bi and Lavaei (2020).

Lemma 8. *If Q is a quadratic form satisfying δ -RIP $_{2r}$, then*

$$|[Q](K, L) - \langle K, L \rangle| \leq 2\delta \|K\|_F \|L\|_F,$$

for all matrices $K, L \in \mathbb{R}^{n \times n}$ of rank at most $2r$.

Lemma 9. *For a given constant R greater than D , the gradient ∇g_s of the function g_s in the symmetric problem (4) is $8\rho_1 r^{1/2} R$ -restricted Lipschitz continuous and the Hessian $\nabla^2 g_s$ is $4\rho_1 r^{1/4} R^{1/2} (2r^{1/2} R \rho_2 / \rho_1 + 3)$ -restricted Lipschitz continuous over the region*

$$\mathcal{D} = \{X \in \mathbb{R}^{n \times r} \mid \|XX^T\|_F \leq R\}.$$

Proof. For every $U \in \mathcal{D}$, we have

$$\|U\|_F = \sqrt{\sum_{i=1}^r \sigma_i(U)^2} \leq \sqrt[4]{r \sum_{i=1}^r \sigma_i(U)^4} = \sqrt[4]{r \sum_{i=1}^r \lambda_i(UU^T)^2} = r^{1/4} \|UU^T\|_F^{1/2} \leq r^{1/4} R^{1/2}. \quad (15)$$

Furthermore, for every $U, V \in \mathcal{D}$, it holds that

$$\|UU^T - VV^T\|_F = \|U(U - V)^T + (U - V)V^T\|_F \leq 2r^{1/4} R^{1/2} \|U - V\|_F.$$

To prove that the gradient ∇g_s is Lipschitz continuous, one can write

$$\begin{aligned} \|\nabla g_s(U) - \nabla g_s(V)\|_F &= 2\|\nabla f_s(UU^T)U - \nabla f_s(VV^T)V\|_F \\ &\leq 2\|\nabla f_s(UU^T)U - \nabla f_s(VV^T)U\|_F + 2\|\nabla f_s(VV^T)(U - V)\|_F \\ &\leq 2\rho_1 \|UU^T - VV^T\|_F \|U\|_F + 2\rho_1 \|VV^T - M^*\|_F \|U - V\|_F \\ &\leq 4\rho_1 r^{1/2} R \|U - V\|_F + 4\rho_1 R \|U - V\|_F \\ &\leq 8\rho_1 r^{1/2} R \|U - V\|_F. \end{aligned}$$

Similarly, for every $W \in \mathbb{R}^{n \times r}$, we have

$$\begin{aligned} &[\nabla^2 g_s(U)](W, W) - [\nabla^2 g_s(V)](W, W) \\ &= [\nabla^2 f_s(UU^T)](UW^T + WU^T, UW^T + WU^T) \\ &\quad - [\nabla^2 f_s(VV^T)](VW^T + WV^T, VW^T + WV^T) \\ &\quad + 2\langle \nabla f_s(UU^T) - \nabla f_s(VV^T), WW^T \rangle \\ &= [\nabla^2 f_s(UU^T) - \nabla^2 f_s(VV^T)](UW^T + WU^T, UW^T + WU^T) \\ &\quad + [\nabla^2 f_s(VV^T)](UW^T + WU^T, UW^T + WU^T) \\ &\quad - [\nabla^2 f_s(VV^T)](VW^T + WV^T, VW^T + WV^T) \\ &\quad + 2\langle \nabla f_s(UU^T) - \nabla f_s(VV^T), WW^T \rangle. \end{aligned}$$

There are four terms in the above expression. The first term can be upper bounded as

$$\begin{aligned} \mathcal{A}_1 &:= [\nabla^2 f_s(UU^T) - \nabla^2 f_s(VV^T)](UW^T + WU^T, UW^T + WU^T) \\ &\leq \rho_2 \|UU^T - VV^T\|_F \|UW^T + WU^T\|_F^2 \\ &\leq 4\rho_2 \|UU^T - VV^T\|_F \|U\|_F^2 \|W\|_F^2 \\ &\leq 8\rho_2 r^{3/4} R^{3/2} \|U - V\|_F \|W\|_F^2. \end{aligned}$$

Similarly, the sum of the second and third terms can be bounded as

$$\begin{aligned}
\mathcal{A}_2 &:= [\nabla^2 f_s(VV^T)](UW^T + WU^T, UW^T + WU^T) \\
&\quad - [\nabla^2 f_s(VV^T)](VW^T + WV^T, VW^T + WV^T) \\
&= [\nabla^2 f_s(VV^T)](UW^T + WU^T, (U - V)W^T + W(U - V)^T) \\
&\quad + [\nabla^2 f_s(VV^T)]((U - V)W^T + W(U - V)^T, VW^T + WV^T) \\
&\leq (1 + 2\delta)(\|UW^T + WU^T\|_F + \|VW^T + WV^T\|_F)\|(U - V)W^T + W(U - V)^T\|_F \\
&\leq 4(1 + 2\delta)(\|U\|_F + \|V\|_F)\|U - V\|_F\|W\|_F^2 \\
&\leq 8\rho_1 r^{1/4} R^{1/2} \|U - V\|_F \|W\|_F^2,
\end{aligned}$$

where Lemma 8 is applied in the second step. Moreover, we can upper bound the last term as

$$\begin{aligned}
\mathcal{A}_3 &:= 2\langle \nabla f_s(UU^T) - \nabla f_s(VV^T), WW^T \rangle \\
&\leq 2\rho_1 \|UU^T - VV^T\|_F \|W\|_F^2 \\
&\leq 4\rho_1 r^{1/4} R^{1/2} \|U - V\|_F \|W\|_F^2.
\end{aligned}$$

Therefore,

$$\begin{aligned}
[\nabla^2 g_s(U)](W, W) - [\nabla^2 g_s(V)](W, W) &= \mathcal{A}_1 + \mathcal{A}_2 + \mathcal{A}_3 \\
&\leq 4\rho_1 r^{1/4} R^{1/2} (2r^{1/2} R\rho_2/\rho_1 + 3) \|U - V\|_F \|W\|_F^2,
\end{aligned}$$

which implies that the Hessian $\nabla^2 g_s$ has the desired Lipschitz property. \square

Next, we verify some facts about the augmented ground truth \tilde{M}^* for the asymmetric problem, which will be useful in the transformation from asymmetric problems to symmetric problems.

Lemma 10. *The augmented ground truth \tilde{M}^* defined in (8) is independent of the balanced factorization of the ground truth M^* . Furthermore,*

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

Proof. By expanding all the terms, it can be checked that the inequality

$$\|U_1 U_1^T - U_2 U_2^T\|_F^2 + \|V_1 V_1^T - V_2 V_2^T\|_F^2 \leq 2\|U_1 V_1^T - U_2 V_2^T\|_F^2$$

holds for all $U_1, U_2 \in \mathbb{R}^{n \times r}$ and $V_1, V_2 \in \mathbb{R}^{m \times r}$ with $U_1^T U_1 = V_1^T V_1$ and $U_2^T U_2 = V_2^T V_2$ (see Appendix F in Zhu et al. (2018)). Then, if (U_1, V_1) and (U_2, V_2) are two balanced factorizations of the ground truth M^* , we must have

$$U_1 U_1^T = U_2 U_2^T, \quad V_1 V_1^T = V_2 V_2^T$$

and thus \tilde{M}^* is unique.

Assume that (U^*, V^*) is a balanced factorization of M^* , the remaining equalities follow from the fact that

$$\begin{aligned}
\sigma_i(M^*)^2 &= \sigma_i(U^* V^{*T} V^* U^{*T}) = \sigma_i(U^* U^{*T} U^* U^{*T}) \\
&= \sigma_i(U^* U^{*T})^2 = \sigma_i(U^{*T} U^*)^2 \\
&= \frac{1}{4} \sigma_i(U^{*T} U^* + V^{*T} V^*)^2 = \frac{1}{4} \sigma_i(\tilde{M}^*)^2
\end{aligned}$$

for all $i \in \{1, \dots, r\}$. \square

In the following, we will show that the gradient and the Hessian of the objective function g_a in the transformed asymmetric problem (7) satisfies the same Lipschitz property as in Lemma 9. This means that those proofs in the remainder of this paper that depend on the Lipschitz property of g_s can be applied to both the symmetric problem (4) and the transformed asymmetric problem (7).

Lemma 11. *The gradient ∇g_a and the Hessian $\nabla^2 g_a$ in the transformed asymmetric problem (7) satisfy the same Lipschitz property as in Lemma 9.*

Proof. Consider arbitrary low-rank matrices $N, N', K \in \mathbb{R}^{(n+m) \times (n+m)}$ written in block forms in the same way as in (6), with $\text{rank}(N) \leq r$, $\text{rank}(N') \leq r$ and $\text{rank}(K) \leq 2r$. First, we will prove that the gradient ∇F and the Hessian $\nabla^2 F$ of the transformed function F are still ρ_1 -restricted Lipschitz continuous and ρ_2 -restricted Lipschitz continuous, respectively. Given the gradient

$$\nabla F(N) = \frac{1}{2} \begin{bmatrix} 0 & \nabla f_a(N_{12}) \\ (\nabla f_a(N_{21}^T))^T & 0 \end{bmatrix} + \frac{\phi}{2} \begin{bmatrix} N_{11} & -N_{12} \\ -N_{21} & N_{22} \end{bmatrix},$$

we have

$$\begin{aligned} & \|\nabla F(N) - \nabla F(N')\|_F \\ & \leq \frac{1}{2} \sqrt{\|\nabla f_a(N_{12}) - \nabla f_a(N'_{12})\|_F^2 + \|\nabla f_a(N_{21}^T) - \nabla f_a(N'_{21}^T)\|_F^2} + \frac{\phi}{2} \|N - N'\|_F \\ & \leq \frac{\rho_1}{2} \sqrt{\|N_{12} - N'_{12}\|_F^2 + \|N_{21} - N'_{21}\|_F^2} + \frac{\phi}{2} \|N - N'\|_F \\ & \leq \frac{1}{2}(\rho_1 + \phi) \|N - N'\|_F \leq \rho_1 \|N - N'\|_F, \end{aligned}$$

in which the second inequality is due to the ρ_1 -restricted Lipschitz continuity of ∇f_a , while the last inequality follows from the choice $\phi = (1 - \delta)/2$ in Assumption 4 and $\rho_1 \geq 1 + 2\delta$ in Assumption 2. Moreover, since

$$\begin{aligned} [\nabla^2 F(N)](K, K) &= \frac{1}{2}([\nabla^2 f_a(N_{12})](K_{12}, K_{12}) + [\nabla^2 f_a(N_{21}^T)](K_{21}^T, K_{21}^T)) \\ & \quad + \frac{\phi}{2}(\|K_{11}\|_F^2 + \|K_{22}\|_F^2 - \|K_{12}\|_F^2 - \|K_{21}\|_F^2), \end{aligned}$$

it is clear that $\nabla^2 F$ is $\rho_2/2$ -restricted Lipschitz continuous as the second term in the above equation is independent of N . Next, we can repeat the argument in Lemma 9 with the function f_s replaced with $4F/(1 + \delta)$, noting that the latter function satisfies the $2\delta/(1 + \delta)$ -RIP $_{2r}$ property as proven in Theorem 12 of Zhang et al. (2021). \square

Using the Lipschitz properties proven in Lemma 9, we will show that the objective value decreases at each iteration of the gradient descent algorithm with a sufficiently small step size η . Although the following lemma is stated for the symmetric problem (4), a similar result holds for the transformed asymmetric problem (7).

Lemma 12. *Given a matrix $X \in \mathbb{R}^{n \times r}$ satisfying*

$$\|XX^T - M^*\|_F \leq R,$$

let $X' = X - \eta \nabla g_s(X)$ be the result of a one-step gradient descent applied to the symmetric problem (4) with the step size η satisfying

$$1/\eta \geq 12\rho_1 r^{1/2}(R + D)$$

Then, $g_s(X') \leq g_s(X) - \eta \|\nabla g_s(X)\|_F^2/2$.

Proof. The assumption on η implies that $\eta\rho_1 R \leq 1/12$. Define $\tilde{X}(t) = X - t\eta \nabla g_s(X)$ for $t \in [0, 1]$. We have $\tilde{X}(1) = X'$ and

$$\begin{aligned} \|\tilde{X}(t)\tilde{X}(t)^T - M^*\|_F &\leq 2t\eta \|\nabla g_s(X)X^T\|_F + t^2\eta^2 \|\nabla g_s(X)\nabla g_s(X)^T\|_F + \|XX^T - M^*\|_F \\ &\leq 4t\eta \|\nabla f_s(XX^T)XX^T\|_F + 4t^2\eta^2 \|\nabla f_s(XX^T)XX^T \nabla f_s(XX^T)^T\|_F + R \\ &\leq 4\eta\rho_1 \|XX^T - M^*\|_F \|XX^T\|_F + 4\eta^2\rho_1^2 \|XX^T - M^*\|_F^2 \|XX^T\|_F + R \\ &\leq 4\eta\rho_1 R(R + D)(1 + \eta\rho_1 R) + R \leq \frac{3}{2}(R + D). \end{aligned}$$

By the Lipschitz property of the function g_s proven in Lemma 9 and the assumption on η , we have

$$\|\nabla g_s(\tilde{X}(t)) - \nabla g_s(X)\|_F \leq 12\rho_1 r^{1/2}(R + D) \|\tilde{X}(t) - X\|_F \leq t \|\nabla g_s(X)\|_F.$$

Now, one can write

$$\begin{aligned}
g_s(X') - g_s(X) &= \int_0^1 \langle \nabla g_s(\tilde{X}(t)), X' - X \rangle dt \\
&= -\eta \|\nabla g_s(X)\|_F^2 + \eta \int_0^1 \langle \nabla g_s(X) - \nabla g_s(\tilde{X}(t)), \nabla g_s(X) \rangle dt \\
&\leq -\eta \|\nabla g_s(X)\|_F^2 + \frac{\eta}{2} \|\nabla g_s(X)\|_F^2.
\end{aligned}$$

As a result, $g_s(X') \leq g_s(X) - \eta \|\nabla g_s(X)\|_F^2/2$. \square

B Proofs for Section 3

First, we need to introduce some notations that will be used throughout this section and next section. For every $X \in \mathbb{R}^{n \times r}$, define

$$\mathbf{e} := \text{vec}(XX^T - M^*)$$

and let $\mathbf{X} \in \mathbb{R}^{n^2 \times nr}$ be the matrix satisfying

$$\mathbf{X} \text{vec} U = \text{vec}(XU^T + UX^T), \quad \forall U \in \mathbb{R}^{n \times r}.$$

The following lemma borrowed from Bhojanapalli et al. (2016b) will be useful.

Lemma 13. *Let $X, Z \in \mathbb{R}^{n \times r}$ be two arbitrary matrices such that $X^T Z$ is symmetric and positive semidefinite. It holds that*

$$\sigma_r(ZZ^T) \|X - Z\|_F^2 \leq \frac{1}{2(\sqrt{2} - 1)} \|XX^T - ZZ^T\|_F^2.$$

Proof of Theorem 1. Define

$$q_1 = \sqrt{1 - \frac{\tilde{C}^2}{2(\sqrt{2} - 1)\sigma_r(M^*)}}, \quad q_2 = \frac{\sqrt{2}\mu'}{\sigma_r(M^*)^{1/2} - \tilde{C}}. \quad (16)$$

The assumption (9) on \tilde{C} implies that $\delta < q_1$, and thus one can always find a sufficiently small $\mu' > 0$ such that

$$\frac{1 - \delta}{1 + \delta} > \frac{1 - q_1 + q_2}{1 + q_1}. \quad (17)$$

We choose $\mu = \mu'^2/(1 + \delta)$. Assume on the contrary that

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

at a particular matrix X in the region (10). Obviously, $XX^T \neq M^*$. It results from (11) that

$$\frac{1}{2} \|\nabla g_s(X)\|_F^2 < \mu(f_s(XX^T) - f_s(M^*)) \leq \frac{\mu(1 + \delta)}{2} \|XX^T - M^*\|_F^2,$$

and thus

$$\|\nabla g_s(X)\|_F \leq \mu' \|XX^T - M^*\|_F.$$

For every $U \in \mathbb{R}^{n \times r}$, the above inequality implies that

$$\langle \nabla f_s(XX^T), XU^T + UX^T \rangle = \langle \nabla g_s(X), U \rangle \leq \|\nabla g_s(X)\|_F \|U\|_F \leq \mu' \|XX^T - M^*\|_F \|U\|_F.$$

Define an auxiliary function $h : \mathbb{R}^{n \times n} \rightarrow \mathbb{R}$ by letting

$$h(M) = \langle \nabla f_s(M), XU^T + UX^T \rangle, \quad \forall M \in \mathbb{R}^{n \times n}.$$

The mean value theorem implies that

$$\begin{aligned}
\mu' \|XX^T - M^*\|_F \|U\|_F &\geq h(XX^T) - h(M^*) \\
&= \int_0^1 \langle \nabla h((1-t)XX^T + tM^*), XX^T - M^* \rangle dt \\
&= \int_0^1 [\nabla^2 f_s((1-t)XX^T + tM^*)](XX^T - M^*, XU^T + UX^T) dt \\
&= \mathbf{e}^T \mathbf{H} \mathbf{X} \mathbf{U},
\end{aligned}$$

where $\mathbf{U} = \text{vec}(U)$ and $\mathbf{H} \in \mathbb{R}^{n^2 \times n^2}$ is the symmetric matrix satisfying

$$(\text{vec}(K))^T \mathbf{H} \text{vec}(L) = \int_0^1 [\nabla^2 f_s((1-t)XX^T + tM^*)](K, L) dt$$

for all $K, L \in \mathbb{R}^{n \times n}$. After choosing $\mathbf{U} = \mathbf{X}^T \mathbf{H} \mathbf{e}$, we obtain

$$\|\mathbf{X}^T \mathbf{H} \mathbf{e}\| \leq \mu' \|\mathbf{e}\|.$$

Since $\nabla^2 f_s((1-t)XX^T + tM^*)$ satisfies the δ -RIP $_{2r}$ property for all $t \in [0, 1]$, \mathbf{H} also satisfies the δ -RIP $_{2r}$. Therefore, if we define $\delta_f^*(X, \mu')$ to be the optimal value of the optimization problem

$$\begin{aligned}
&\min_{\delta, \mathbf{H}} \delta \\
&\text{s. t. } \|\mathbf{X}^T \mathbf{H} \mathbf{e}\| \leq \mu' \|\mathbf{e}\|, \\
&\quad \mathbf{H} \text{ is symmetric and satisfies } \delta\text{-RIP}_{2r},
\end{aligned} \tag{18}$$

then the above argument shows that $\delta_f^*(X, \mu') \leq \delta$. However, Lemma 14 (to be stated next) shows that

$$\frac{1 - \delta}{1 + \delta} \leq \frac{1 - \delta_f^*(X, \mu')}{1 + \delta_f^*(X, \mu')} \leq \frac{1 - q_1 + q_2}{1 + q_1},$$

which contradicts the inequality (17). \square

Lemma 14. *If $X \in \mathbb{R}^{n \times r}$ is a matrix in the region (10) such that $XX^T \neq M^*$, then the optimal value $\delta_f^*(X, \mu')$ of the optimization problem (18) satisfies*

$$\frac{1 - \delta_f^*(X, \mu')}{1 + \delta_f^*(X, \mu')} \leq \frac{1 - q_1 + q_2}{1 + q_1},$$

where q_1 and q_2 are defined in (16).

Proof. Let $Z \in \mathcal{Z}$ be a global minimizer such that $ZZ^T = M^*$. The fact that X is in the region (10) implies that $\|X - Z\|_F \leq \tilde{C}$. Without loss of generality, it can be assumed that $X^T Z$ is symmetric and positive semidefinite. If this is not the case, then we use the singular value decomposition $X^T Z = PDQ^T$ in which $P, Q \in \mathbb{R}^{n \times n}$ are orthogonal and $D \in \mathbb{R}^{n \times n}$ is diagonal. By defining $R = QP^T$, the matrix ZR becomes another global minimizer and

$$X^T(ZR) = PDQ^T QP^T = PDP^T \succeq 0,$$

implying that we can continue the following argument with ZR instead of Z .

By an argument similar to Lemma 14 in Bi and Lavaei (2020), the optimal value of the problem (18) is equal to that of the problem

$$\begin{aligned}
&\min_{\delta, \mathbf{H}} \delta \\
&\text{s. t. } \|\mathbf{X}^T \mathbf{H} \mathbf{e}\| \leq \mu' \|\mathbf{e}\|, \\
&\quad (1 - \delta)I_{n^2} \preceq \mathbf{H} \preceq (1 + \delta)I_{n^2}.
\end{aligned} \tag{19}$$

Furthermore, the norm constraint $\|\mathbf{X}^T \mathbf{H} \mathbf{e}\| \leq \mu' \|\mathbf{e}\|$ in (19) is equivalent to

$$\begin{bmatrix} I_{nr} & \mathbf{X}^T \mathbf{H} \mathbf{e} \\ (\mathbf{X}^T \mathbf{H} \mathbf{e})^T & \mu'^2 \|\mathbf{e}\|^2 \end{bmatrix} \succeq 0.$$

Now, define $\eta_f^*(X, \mu')$ to be the optimal value of the following optimization problem:

$$\begin{aligned} & \max_{\eta, \mathbf{H}} \quad \eta \\ & \text{s. t.} \quad \begin{bmatrix} I_{nr} & \mathbf{X}^T \mathbf{H} \mathbf{e} \\ (\mathbf{X}^T \mathbf{H} \mathbf{e})^T & \mu'^2 \|\mathbf{e}\|^2 \end{bmatrix} \succeq 0, \\ & \quad \eta I_{n^2} \preceq \mathbf{H} \preceq I_{n^2}. \end{aligned} \quad (20)$$

Given any feasible solution (δ, \mathbf{H}) to (19),

$$\left(\frac{1-\delta}{1+\delta}, \frac{1}{1+\delta} \mathbf{H} \right)$$

is a feasible solution to the above problem (20). Therefore,

$$\eta_f^*(X, \mu') \geq \frac{1 - \delta_f^*(X, \mu')}{1 + \delta_f^*(X, \mu')}. \quad (21)$$

To prove the desired inequality, it is sufficient to upper bound $\eta_f^*(X, \mu')$ by finding a feasible solution to the dual problem of (20) given below:

$$\begin{aligned} & \min_{U_1, U_2, G, \lambda, y} \quad \text{tr}(U_2) + \mu'^2 \|\mathbf{e}\|^2 \lambda + \text{tr}(G), \\ & \text{s. t.} \quad \text{tr}(U_1) = 1, \\ & \quad (\mathbf{X}y)\mathbf{e}^T + \mathbf{e}(\mathbf{X}y)^T = U_1 - U_2, \\ & \quad \begin{bmatrix} G & -y \\ -y^T & \lambda \end{bmatrix} \succeq 0, \\ & \quad U_1 \succeq 0, \quad U_2 \succeq 0. \end{aligned} \quad (22)$$

As shown in the first part of the proof of (Bi and Lavaei, 2020, Lemma 19), there exists a nonzero vector $y \in \mathbb{R}^{nr}$ such that

$$\|\mathbf{X}y\|^2 \geq 2\sigma_r(\mathbf{X}\mathbf{X}^T)\|y\|^2 \quad (23)$$

and

$$\|\mathbf{e} - \mathbf{X}y\| \leq \|X - Z\|_F^2.$$

By Lemma 13, we have

$$\frac{\|\mathbf{e} - \mathbf{X}y\|}{\|\mathbf{e}\|} \leq \frac{\|X - Z\|_F^2}{\|\mathbf{X}\mathbf{X}^T - M^*\|_F} \leq \sqrt{\frac{1}{2(\sqrt{2}-1)\sigma_r(M^*)}} \tilde{C} < 1.$$

If θ is the angle between \mathbf{e} and $\mathbf{X}y$, then the above inequality implies that $\theta < \pi/2$ and

$$\sin \theta \leq \frac{\|\mathbf{e} - \mathbf{X}y\|}{\|\mathbf{e}\|} \leq \sqrt{\frac{1}{2(\sqrt{2}-1)\sigma_r(M^*)}}.$$

Therefore,

$$\cos \theta \geq q_1. \quad (24)$$

On the other hand, the Wielandt–Hoffman theorem implies that

$$|\sigma_r(\mathbf{X}\mathbf{X}^T)^{1/2} - \sigma_r(M^*)^{1/2}| = |\sigma_r(X) - \sigma_r(Z)| \leq \|X - Z\|_F \leq \tilde{C}.$$

Combining the above inequality and (23) gives

$$\|y\| \leq \frac{\|\mathbf{X}y\|}{\sqrt{2}(\sigma_r(M^*)^{1/2} - \tilde{C})}. \quad (25)$$

Let

$$M = (\mathbf{X}y)\mathbf{e}^T + \mathbf{e}(\mathbf{X}y)^T,$$

with y given above, and decompose M as

$$M = [M]_+ - [M]_-$$

such that $[M]_+ \succeq 0$ and $[M]_- \succeq 0$. By Lemma 14 in Zhang et al. (2019), we have

$$\begin{aligned}\text{tr}([M]_+) &= \|\mathbf{e}\| \|\mathbf{X}y\| (1 + \cos \theta), \\ \text{tr}([M]_-) &= \|\mathbf{e}\| \|\mathbf{X}y\| (1 - \cos \theta).\end{aligned}$$

Again, θ is the angle between \mathbf{e} and $\mathbf{X}y$. Then,

$$\begin{aligned}U_1^* &= \frac{[M]_+}{\text{tr}([M]_+)}, \quad U_2^* = \frac{[M]_-}{\text{tr}([M]_+)}, \\ G^* &= \frac{1}{\lambda^*} y^* y^{*T}, \quad \lambda^* = \frac{\|y^*\|}{\mu' \|\mathbf{e}\|}, \quad y^* = \frac{y}{\text{tr}([M]_+)}\end{aligned}$$

form a feasible solution to the dual problem (22) with the objective value

$$\frac{\text{tr}([M]_-) + 2\mu' \|\mathbf{e}\| \|y\|}{\text{tr}([M]_+)} = \frac{1 - \cos \theta + 2\mu' \|y\| / \|\mathbf{X}y\|}{1 + \cos \theta}.$$

The inequalities (24) and (25) imply that

$$\eta_f^*(X, \eta) \leq \frac{1 - q_1 + q_2}{1 + q_1}.$$

The proof is completed by the above inequality and (21). \square

Proof of Theorem 2. Define

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

Then, it follows from Theorem 1 that there exists a constant $\mu > 0$ such that the PL inequality

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

is satisfied in the region

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

By Lemma 13, in order to prove that a matrix X belongs to \mathcal{D} , it suffices to show that

$$\|X X^T - M^*\|_F \leq \sqrt{2(\sqrt{2} - 1)\sigma_r(M^*)}^{1/2} \tilde{C} = \sqrt{\frac{1 + \delta}{1 - \delta}} \|X_0 X_0^T - M^*\|_F. \quad (26)$$

Next, we prove by induction that X_t satisfies (26) and $g_s(X_t) \leq g_s(X_{t-1})$ at each step of the iteration. Obviously, (26) holds for X_0 . At step t , by Lemma 12, the induction assumption

$$\|X_{t-1} X_{t-1}^T - M^*\|_F \leq \sqrt{\frac{1 + \delta}{1 - \delta}} \|X_0 X_0^T - M^*\|_F$$

and our choice of the step size η imply that $g_s(X_t) \leq g_s(X_{t-1}) \leq \dots \leq g_s(X_0)$. Then, the inequality (12) immediately implies that X_t satisfies (26).

Finally, since X_t is guaranteed to be contained in a region satisfying the PL inequality for all t , we can apply Theorem 1 in Karimi et al. (2016) to obtain

$$g_s(X_t) - f_s(M^*) \leq (1 - \mu\eta)^t (g_s(X_0) - f_s(M^*)).$$

Now, (13) follows from the above inequality and (11). \square

After the transformation from asymmetric problems to symmetric problems, the proof of Theorem 3 is similar to that of Theorem 2, and thus it is omitted here.

C Proofs for Section 4

We first present the perturbed gradient descent algorithm with local improvement adapted from the general algorithm in Jin et al. (2017) for solving the symmetric problem (4), which can be also used to solve (5) after the transformation from asymmetric problems to symmetric problems. In Algorithm 1, X_0 is the initial point and $1 - \gamma$ is the success probability of the algorithm, while η and w are respectively the step size and perturbation size which are further determined by the parameter c . Furthermore, the parameter κ determines at what time the corresponding point is sufficiently close to the ground truth so that it belongs to the local convergence region and thus perturbations are no longer necessary in future iterations. After the first loop ends, the current matrix X_t will satisfy (14). The choice of the parameters c and κ will be given in the proof of Theorem 5, but they can also be selected empirically.

Algorithm 1 Perturbed Gradient Descent Method With Local Improvement

```

 $R \leftarrow 3D(1 + \delta)/(1 - \delta)$ 
 $\ell_1 \leftarrow 8\rho_1 r^{1/2}R, \quad \ell_2 \leftarrow 4\rho_1 r^{1/4}R^{1/2}(2r^{1/2}R\rho_2/\rho_1 + 3)$ 
 $\hat{\epsilon} \leftarrow \min\{\kappa, \kappa^2/\ell_2\}, \quad \Delta \leftarrow 2(1 + \delta)D^2$ 
 $\chi \leftarrow 3 \max\{\log((nr\ell_1\Delta)/(c\hat{\epsilon}^2\gamma)), 4\}, \quad \eta \leftarrow c/\ell_1, \quad w \leftarrow \sqrt{c\hat{\epsilon}}/(\chi^2\ell_1)$ 
 $g_{\text{thres}} \leftarrow \sqrt{c\hat{\epsilon}}/\chi^2, \quad f_{\text{thres}} \leftarrow c\sqrt{\hat{\epsilon}^3}/\ell_2/\chi^3, \quad t_{\text{thres}} \leftarrow \chi\ell_1/(c^2\sqrt{\ell_2\hat{\epsilon}})$ 
 $t \leftarrow 0, \quad t_{\text{noise}} \leftarrow -t_{\text{thres}} - 1$ 
loop
  if  $\|\nabla g_s(X_t)\|_F \leq g_{\text{thres}}$  and  $t - t_{\text{noise}} > t_{\text{thres}}$  then
     $\tilde{X}_t \leftarrow X_t, \quad t_{\text{noise}} \leftarrow t$ 
     $X_t \leftarrow X_t + W$ , where  $W$  is drawn uniformly from the ball with radius  $w$ 
  end if
  if  $t - t_{\text{noise}} = t_{\text{thres}}$  and  $g_s(X_t) - g_s(\tilde{X}_{t_{\text{noise}}}) > -f_{\text{thres}}$  then
     $X_t \leftarrow \tilde{X}_{t_{\text{noise}}}$ 
    break
  end if
   $X_{t+1} \leftarrow X_t - \eta\nabla g_s(X_t), \quad t \leftarrow t + 1$ 
end loop
loop
   $X_{t+1} \leftarrow X_t - \eta\nabla g_s(X_t), \quad t \leftarrow t + 1$ 
end loop

```

The following lemma will be useful in the proof of Lemma 4, which can be obtained by combining Lemma 6 and Lemma 7 in Zhang et al. (2021).

Lemma 15. *For any $C > 0$, there exist some $\kappa > 0$ and $\zeta > 0$ such that for each $X \in \mathbb{R}^{n \times r}$ the two inequalities in (14) together with $\sigma_r(X) \leq \zeta$ will imply $\|XX^T - ZZ^T\|_F < C$.*

Proof of Lemma 4. Let ζ be the constant given by Lemma 15. We only need to consider all $X \in \mathbb{R}^{n \times r}$ satisfying $\sigma_r(X) > \zeta$, since the opposite case can be directly handled by applying Lemma 15.

In the linear measurement case, the gradient and the Hessian of the objective function g_s are given in the vectorized form by

$$\begin{aligned} \nabla g(X) &= \mathbf{X}^T \mathbf{H} \mathbf{e}, \\ \nabla^2 g(X) &= 2I_r \otimes \text{mat}_S(\mathbf{H} \mathbf{e}) + \mathbf{X}^T \mathbf{H} \mathbf{X}, \end{aligned}$$

where \mathbf{e} and \mathbf{X} are defined at the beginning of Appendix B, while $\mathbf{H} \in \mathbb{R}^{n^2 \times n^2}$ is the matrix satisfying

$$(\text{vec}(K))^T \mathbf{H} \text{vec}(L) = \langle \mathcal{A}(K), \mathcal{A}(L) \rangle, \quad \forall K, L \in \mathbb{R}^{n \times n}.$$

Hence, if X satisfies the approximate first-order and second-order necessary optimality conditions (14), we must have $\delta \geq \delta^*(X, \kappa)$, where $\delta^*(X, \kappa)$ is the optimal value of the following optimization

problem:

$$\begin{aligned}
& \min_{\delta, \mathbf{H}} \quad \delta \\
& \text{s. t.} \quad \|\mathbf{X}^T \mathbf{H} \mathbf{e}\| \leq \kappa, \\
& \quad \quad 2I_r \otimes \text{mat}_S(\mathbf{H} \mathbf{e}) + \mathbf{X}^T \mathbf{H} \mathbf{X} \succeq -\kappa I_{nr}, \\
& \quad \quad \mathbf{H} \text{ is symmetric and satisfies } \delta\text{-RIP}_{2r}.
\end{aligned} \tag{27}$$

On the other hand, both the assumption $\delta < 1/2$ and Lemma 16 imply that

$$\frac{1}{3} < \frac{1-\delta}{1+\delta} \leq \frac{1-\delta^*(X, \kappa)}{1+\delta^*(X, \kappa)} \leq \frac{1}{3} + \Gamma \frac{\kappa}{\|\mathbf{e}\|},$$

for some constant Γ , which further implies that

$$\kappa \geq \frac{\|\mathbf{e}\|}{\Gamma} \left(\frac{1-\delta}{1+\delta} - \frac{1}{3} \right).$$

The strict saddle property can then be proved by choosing a sufficiently small κ . \square

Lemma 16. *Given a constant $\zeta > 0$, if $X \in \mathbb{R}^{n \times r}$ is a matrix satisfying $XX^T \neq M^*$ and $\sigma_r(X) > \zeta$, then the optimal value $\delta^*(X, \kappa)$ of the optimization problem (27) satisfies*

$$\frac{1-\delta^*(X, \kappa)}{1+\delta^*(X, \kappa)} \leq \frac{1}{3} + \Gamma \frac{\kappa}{\|\mathbf{e}\|},$$

where $\Gamma = \sqrt{r} + \sqrt{2}/\zeta$.

Proof. Let $Z \in \mathcal{Z}$ be a global minimizer such that $ZZ^T = M^*$. In parallel to the proof of Lemma 14, we can introduce a relaxed optimization problem

$$\begin{aligned}
& \max_{\eta, \mathbf{H}} \quad \eta \\
& \text{s. t.} \quad \|\mathbf{X}^T \mathbf{H} \mathbf{e}\| \leq \kappa, \\
& \quad \quad 2I_r \otimes \text{mat}_S(\mathbf{H} \mathbf{e}) + \mathbf{X}^T \mathbf{X} \succeq -\kappa I_{nr}, \\
& \quad \quad \eta I_{n^2} \preceq \mathbf{H} \preceq I_{n^2},
\end{aligned} \tag{28}$$

whose optimal value $\eta^*(X, \kappa)$ satisfies

$$\eta^*(X, \kappa) \geq \frac{1-\delta^*(X, \kappa)}{1+\delta^*(X, \kappa)}.$$

To prove the desired inequality, we need to find an upper bound for $\eta^*(X, \kappa)$, which can be achieved by finding a feasible solution to the dual problem of (28):

$$\begin{aligned}
& \min_{\substack{U_1, U_2, W, \\ G, \lambda, y}} \quad \text{tr}(U_2) + \langle \mathbf{X}^T \mathbf{X}, W \rangle + \kappa \text{tr}(W) + \kappa^2 \lambda + \text{tr}(G) \\
& \text{s. t.} \quad \text{tr}(U_1) = 1, \\
& \quad \quad (\mathbf{X}y - w)\mathbf{e}^T + \mathbf{e}(\mathbf{X}y - w)^T = U_1 - U_2, \\
& \quad \quad \begin{bmatrix} G & -y \\ -y^T & \lambda \end{bmatrix} \succeq 0, \\
& \quad \quad U_1 \succeq 0, \quad U_2 \succeq 0, \quad W = \begin{bmatrix} W_{1,1} & \cdots & W_{r,1}^T \\ \vdots & \ddots & \vdots \\ W_{r,1} & \cdots & W_{r,r} \end{bmatrix} \succeq 0, \\
& \quad \quad w = \sum_{i=1}^r \text{vec}(W_{i,i}).
\end{aligned} \tag{29}$$

Before describing the choice of the dual feasible solution, we need to represent the error vector \mathbf{e} in a different form. Let $\mathcal{P} \in \mathbb{R}^{n \times n}$ be the orthogonal projection matrix onto the range of X , and

$\mathcal{P}_\perp \in \mathbb{R}^{n \times n}$ be the orthogonal projection matrix onto the orthogonal complement of the range of X . Then, Z can be decomposed as $Z = \mathcal{P}Z + \mathcal{P}_\perp Z$, and there exists a matrix $R \in \mathbb{R}^{r \times r}$ such that $\mathcal{P}Z = XR$. Note that

$$ZZ^T = \mathcal{P}ZZ^T\mathcal{P} + \mathcal{P}ZZ^T\mathcal{P}_\perp + \mathcal{P}_\perp ZZ^T\mathcal{P} + \mathcal{P}_\perp ZZ^T\mathcal{P}_\perp.$$

Thus, if we choose

$$\hat{Y} = \frac{1}{2}X - \frac{1}{2}XRR^T - \mathcal{P}_\perp ZR^T, \quad \hat{y} = \text{vec}(\hat{Y}), \quad (30)$$

then it can be verified that

$$\begin{aligned} X\hat{Y}^T + \hat{Y}X^T - \mathcal{P}_\perp ZZ^T\mathcal{P}_\perp &= XX^T - ZZ^T, \\ \langle X\hat{Y}^T + \hat{Y}X^T, \mathcal{P}_\perp ZZ^T\mathcal{P}_\perp \rangle &= 0. \end{aligned}$$

Moreover, we have

$$\begin{aligned} \|X\hat{Y}^T + \hat{Y}X^T\|_F^2 &= 2\text{tr}(X^T X\hat{Y}^T\hat{Y}) + \text{tr}(X^T\hat{Y}X^T\hat{Y}) + \text{tr}(\hat{Y}^T X\hat{Y}^T X) \\ &\geq 2\text{tr}(X^T X\hat{Y}^T\hat{Y}) \geq 2\sigma_r^2(X)\|\hat{Y}\|_F^2, \end{aligned} \quad (31)$$

in which the first inequality is due to

$$\text{tr}(X^T\hat{Y}X^T\hat{Y}) = \frac{1}{4}\text{tr}((X^T X(I_r - RR^T))^2) = \frac{1}{4}\text{tr}((X(I_r - RR^T)X^T)^2) \geq 0.$$

Assume first that $Z_\perp = \mathcal{P}_\perp Z \neq 0$. The other case will be handled at the end of this proof. In the case when $Z_\perp \neq 0$, we also have $X\hat{Y}^T + \hat{Y}X^T \neq 0$. Otherwise, the inequality (31) and the assumption $\sigma_r(X) > 0$ imply that $\hat{Y} = 0$. The orthogonality and the definition of \hat{Y} in (30) then give

$$X - XRR^T = 0, \quad \mathcal{P}_\perp ZR^T = 0.$$

The first equation above implies that R is invertible since X has full column rank, which contradicts $Z_\perp \neq 0$. Now, define the unit vectors

$$\hat{u}_1 = \frac{\mathbf{X}\hat{y}}{\|\mathbf{X}\hat{y}\|}, \quad \hat{u}_2 = \frac{\text{vec}(Z_\perp Z_\perp^T)}{\|Z_\perp Z_\perp^T\|_F}.$$

Then, $\hat{u}_1 \perp \hat{u}_2$ and

$$\mathbf{e} = \|\mathbf{e}\|(\sqrt{1 - \alpha^2}\hat{u}_1 - \alpha\hat{u}_2) \quad (32)$$

with

$$\alpha = \frac{\|Z_\perp Z_\perp^T\|_F}{\|XX^T - ZZ^T\|_F}. \quad (33)$$

We first describe our choices of the dual variables W and y (which will be scaled later). Let

$$X^T X = QSQ^T, \quad Z_\perp Z_\perp^T = PGP^T,$$

with Q, P orthogonal and S, G diagonal, such that $S_{11} = \sigma_r^2(X)$. Fix a constant $\gamma \in [0, 1]$ that is to be determined and define

$$\begin{aligned} V_i &= k^{1/2}G_{ii}^{1/2}PE_{i1}Q^T, \quad \forall i = 1, \dots, r, \\ W &= \sum_{i=1}^r \text{vec}(V_i) \text{vec}(V_i)^T, \quad y = l\hat{y}, \end{aligned}$$

with \hat{y} defined in (30) and

$$k = \frac{\gamma}{\|\mathbf{e}\|\|Z_\perp Z_\perp^T\|_F}, \quad l = \frac{\sqrt{1 - \gamma^2}}{\|\mathbf{e}\|\|\mathbf{X}\hat{y}\|}.$$

Here, E_{ij} is the elementary matrix of size $n \times r$ with (i, j) -entry being 1. By our construction, $X^T V_i = 0$, which implies that

$$\langle \mathbf{X}^T \mathbf{X}, W \rangle = \sum_{i=1}^r \|XV_i^T + V_i X^T\|_F^2 = 2 \sum_{i=1}^r \text{tr}(X^T XV_i^T V_i) = 2k\sigma_r^2(X) \sum_{i=1}^r G_{ii} = 2\beta\gamma, \quad (34)$$

with

$$\beta = \frac{\sigma_r^2(X) \operatorname{tr}(Z_\perp Z_\perp^T)}{\|X X^T - Z Z^T\|_F \|Z_\perp Z_\perp^T\|_F}. \quad (35)$$

In addition,

$$\operatorname{tr}(W) = \sum_{i=1}^r \|V_i\|_F^2 = k \sum_{i=1}^r G_{ii} = k \operatorname{tr}(Z_\perp Z_\perp^T) \leq \frac{\sqrt{r}}{\|\mathbf{e}\|}, \quad (36)$$

and

$$w = \sum_{i=1}^r \operatorname{vec}(W_{i,i}) = \sum_{i=1}^r V_i V_i^T = k Z_\perp Z_\perp^T.$$

Therefore,

$$\mathbf{X}y - w = \frac{1}{\|\mathbf{e}\|} (\sqrt{1 - \gamma^2} \hat{u}_1 - \gamma \hat{u}_2),$$

which together with (32) implies that

$$\|\mathbf{e}\| \|\mathbf{X}y - w\| = 1, \quad \langle \mathbf{e}, \mathbf{X}y - w \rangle = \gamma \alpha + \sqrt{1 - \gamma^2} \sqrt{1 - \alpha^2} = \psi(\gamma). \quad (37)$$

Next, the inequality (31) and the assumption $\sigma_r(X) > \zeta$ imply that

$$\|y\| \leq \frac{\sqrt{1 - \gamma^2}}{\sqrt{2\zeta} \|\mathbf{e}\|} \leq \frac{1}{\sqrt{2\zeta} \|\mathbf{e}\|}. \quad (38)$$

Define

$$M = (\mathbf{X}y - w) \mathbf{e}^T + \mathbf{e} (\mathbf{X}y - w)^T$$

and decompose

$$M = [M]_+ - [M]_-,$$

in which both $[M]_+ \succeq 0$ and $[M]_- \succeq 0$. Let θ be the angle between \mathbf{e} and $\mathbf{X}y - w$. By Lemma 14 in Zhang et al. (2019), we have

$$\begin{aligned} \operatorname{tr}([M]_+) &= \|\mathbf{e}\| \|\mathbf{X}y - w\| (1 + \cos \theta), \\ \operatorname{tr}([M]_-) &= \|\mathbf{e}\| \|\mathbf{X}y - w\| (1 - \cos \theta). \end{aligned}$$

Now, one can verify that

$$\begin{aligned} U_1^* &= \frac{[M]_+}{\operatorname{tr}([M]_+)}, & U_2^* &= \frac{[M]_-}{\operatorname{tr}([M]_+)}, \\ y^* &= \frac{y}{\operatorname{tr}([M]_+)}, & W^* &= \frac{W}{\operatorname{tr}([M]_+)}, \\ \lambda^* &= \frac{\|y^*\|}{\kappa}, & G^* &= \frac{1}{\lambda^*} y^* y^{*T} \end{aligned}$$

forms a feasible solution to the dual problem (29) whose objective value is equal to

$$\frac{\operatorname{tr}([M]_-) + \langle \mathbf{X}^T \mathbf{X}, W \rangle + \kappa \operatorname{tr}(W) + 2\kappa \|y\|}{\operatorname{tr}([M]_+)}.$$

Putting (34), (36), (37) and (38) into the above equation, we can obtain

$$\eta^*(X, \kappa) \leq \frac{2\beta\gamma + 1 - \psi(\gamma) + (\sqrt{r} + \sqrt{2}/\zeta)\kappa/\|\mathbf{e}\|}{1 + \psi(\gamma)} \leq \frac{2\beta\gamma + 1 - \psi(\gamma)}{1 + \psi(\gamma)} + \Gamma \frac{\kappa}{\|\mathbf{e}\|}.$$

Choosing the best $\gamma \in [0, 1]$ to minimize the far right-side of the above inequality leads to

$$\eta^*(X, \kappa) \leq \eta_0(X) + \Gamma \frac{\kappa}{\|\mathbf{e}\|},$$

with

$$\eta_0(X) = \begin{cases} \frac{1 - \sqrt{1 - \alpha^2}}{1 + \sqrt{1 - \alpha^2}}, & \text{if } \beta \geq \frac{\alpha}{1 + \sqrt{1 - \alpha^2}}, \\ \frac{\beta(\alpha - \beta)}{1 - \beta\alpha}, & \text{if } \beta \leq \frac{\alpha}{1 + \sqrt{1 - \alpha^2}}. \end{cases}$$

Here α and β are defined in (33) and (35), respectively. In the proof of Theorem 1.2 in Zhang (2021), it is shown that $\eta_0(X) \leq 1/3$ for every X with $XX^T \neq ZZ^T$, which gives our desired inequality.

Finally, we still need to deal with the case when $\mathcal{P}_\perp Z = 0$. In this case, we know that $X\hat{y} = \mathbf{e}$ with \hat{y} defined in (30). Then, it is easy to check that

$$\begin{aligned} U_1^* &= \frac{\mathbf{e}\mathbf{e}^T}{\|\mathbf{e}\|^2}, & U_2^* &= 0, \\ y^* &= \frac{\hat{y}}{2\|\mathbf{e}\|^2}, & W^* &= 0, \\ \lambda^* &= \frac{\|y^*\|}{\kappa}, & G^* &= \frac{1}{\lambda^*} y^* y^{*T} \end{aligned}$$

forms a feasible solution to the dual problem (29) whose objective value is $2\kappa\|y^*\|$, which is at most $\kappa/(\sqrt{2}\zeta\|\mathbf{e}\|)$ by the inequality (31). \square

Lemma 17. Consider Algorithm 1 for solving the symmetric problem (4). If the initial matrix X_0 satisfies

$$\|X_0 X_0^T\|_F \leq D,$$

the step size η satisfies

$$1/\eta \geq 48\rho_1 r^{1/2} \left(\frac{1+\delta}{1-\delta} D \right),$$

and the perturbation size w satisfies

$$2wr^{1/4} \left(\frac{1+\delta}{1-\delta} \right)^{1/4} \sqrt{3D} + w^2 \leq \sqrt{\frac{1+\delta}{1-\delta}} D,$$

then during the first loop the trajectory X_t is always confined in the region

$$\mathcal{D} = \left\{ X \in \mathbb{R}^{n \times r} \mid \|XX^T - M^*\|_F \leq 3 \left(\frac{1+\delta}{1-\delta} \right) D \right\}.$$

Proof. For convenience, we introduce the set

$$\mathcal{D}_1 = \left\{ X \in \mathbb{R}^{n \times r} \mid \|XX^T - M^*\|_F \leq 2\sqrt{\frac{1+\delta}{1-\delta}} D \right\}.$$

The iteration is initialized at the point $X_0 \in \mathcal{D}_1$. Assume that at some time instance t the current matrix $X_t \in \mathcal{D}_1$, $g_s(X_t) \leq g_s(X_0)$, and some perturbation needs to be added because $\|\nabla g_s(X_t)\|_F$ is small. In this case, a random noise W is generated from the uniform distribution in the ball of radius w . The algorithm saves the original point X_t to \tilde{X}_t and replaces X_t with $X_t + W$. Then, similar to the inequality (15), the old point \tilde{X}_t satisfies

$$\|\tilde{X}_t\|_F \leq r^{1/4} \left(\frac{1+\delta}{1-\delta} \right)^{1/4} \sqrt{3D},$$

and thus the new point X_t satisfies

$$\begin{aligned} \|X_t X_t^T - M^*\|_F &\leq \|\tilde{X}_t \tilde{X}_t^T - M^*\|_F + \|W \tilde{X}_t^T + X_t \tilde{W}^T\|_F + \|W W^T\|_F \\ &\leq 2\sqrt{\frac{1+\delta}{1-\delta}} D + 2wr^{1/4} \left(\frac{1+\delta}{1-\delta} \right)^{1/4} \sqrt{3D} + w^2 \\ &\leq 3\sqrt{\frac{1+\delta}{1-\delta}} D, \end{aligned}$$

by our choice of the parameter w . Due to the design of the perturbed gradient descent algorithm, the perturbation will never be taken in the next t_{thres} number of iterations (t_{thres} is defined in Algorithm 1). As a result, Lemma 12, $X_t \in \mathcal{D}$ and our choice of the step size η imply that $g_s(X_{t+1}) \leq g_s(X_t)$. Hence, the inequality (12) gives

$$\|X_{t+1} X_{t+1}^T - M^*\|_F \leq \sqrt{\frac{1+\delta}{1-\delta}} \|X_t X_t^T - M^*\|_F \leq 3 \left(\frac{1+\delta}{1-\delta} \right) D,$$

which shows that $X_{t+2} \in \mathcal{D}$. Repeating this argument, it can be concluded that $g_s(X_{t+k}) \leq g_s(X_t)$ and $X_{t+k} \in \mathcal{D}$ for all $k = 1, \dots, t_{\text{thres}}$. After $X_{t+t_{\text{thres}}}$ is obtained, the algorithm compares $g_s(X_{t+t_{\text{thres}}})$ with $g_s(\tilde{X}_t)$, and the iteration continues only if $g_s(X_{t+t_{\text{thres}}}) \leq g_s(\tilde{X}_t)$. When this is the case, $g_s(X_{t+t_{\text{thres}}}) \leq g_s(X_0)$, and by the inequality (12) again, we have

$$\|X_{t+t_{\text{thres}}} X_{t+t_{\text{thres}}}^T - M^*\|_F \leq \sqrt{\frac{1+\delta}{1-\delta}} \|X_0 X_0^T - M^*\|_F \leq 2\sqrt{\frac{1+\delta}{1-\delta}} D,$$

and thus $X_{t+t_{\text{thres}}} \in \mathcal{D}_1$. Assume that no perturbation is added at steps $t+t_{\text{thres}}+1, \dots, t+t_{\text{thres}}+l-1$. Then, using a similar argument as above, we can prove that

$$g_s(X_{t+t_{\text{thres}}+k}) \leq g_s(X_{t+t_{\text{thres}}}) \leq g_s(X_0), \quad X_{t+t_{\text{thres}}+k} \in \mathcal{D}_1, \quad \forall k = 1, \dots, l-1.$$

If perturbation needs to be added at step $t+t_{\text{thres}}+l$, we can repeat the above argument with $t+t_{\text{thres}}+l$ instead of t , which leads to the desired result. \square

Proof of Theorem 5. In the first stage of the algorithm, the perturbed gradient descent method is applied. If the parameter c is sufficiently small, then the step size η and the perturbation size w will satisfy the assumptions in Lemma 17. In this case, Lemma 17 implies that the iterations are taken within a region in which ∇g_s and $\nabla^2 g_s$ are Lipschitz continuous. Let κ be the constant given by Lemma 4 such that the approximate second-order necessary optimality conditions (14) will imply that $\|X X^T - Z Z^T\|_F < C$, where

$$C = 2(\sqrt{2}-1)(1-\delta)\sigma_r(M^*)$$

is the radius of the local linear convergence region provided by Theorem 2. Now, Theorem 3 in Jin et al. (2017) shows that with probability $1-\gamma$ the first loop will stop with a solution \tilde{X} satisfying (14), and thus \tilde{X} is within the local convergence region. Note that the number of iterations in this stage is fixed for a given initial matrix X_0 , and that this number is independent of ϵ .

Next, the gradient descent algorithm is run with initialization at the matrix \tilde{X} . Theorem 2 implies that after an additional $O(\log(1/\epsilon))$ number of iterations we find a solution \hat{X} satisfying the accuracy requirement. \square

Theorem 7 can be proved similarly as Theorem 5 using Lemma 17 and Lemma 6.