

Sharp Restricted Isometry Property Bounds for Low-Rank Matrix Recovery Problems with Corrupted Measurements

Ziye Ma^{*1}, Yingjie Bi^{*2}, Javad Lavaei², Somayeh Sojoudi¹³

¹Department of Electrical Engineering and Computer Science

²Department of Industrial Engineering and Operations Research

³Department of Mechanical Engineering

University of California, Berkeley

{ziyema,yingjiebi,lavaei,sojoudi}@berkeley.edu

Abstract

In this paper, we study a general low-rank matrix recovery problem with linear measurements corrupted by some noise. The objective is to understand under what conditions on the restricted isometry property (RIP) of the problem local search methods can find the ground truth with a small error. By analyzing the landscape of the non-convex problem, we first propose a global guarantee on the maximum distance between an arbitrary local minimizer and the ground truth under the assumption that the RIP constant is smaller than $1/2$. We show that this distance shrinks to zero as the intensity of the noise reduces. Our new guarantee is sharp in terms of the RIP constant and is much stronger than the existing results. We then present a local guarantee for problems with an arbitrary RIP constant, which states that any local minimizer is either considerably close to the ground truth or far away from it. Next, we prove the strict saddle property, which guarantees the global convergence of the perturbed gradient descent method in polynomial time. The developed results demonstrate how the noise intensity and the RIP constant of the problem affect the landscape of the problem.

Introduction

Low-rank matrix recovery problems arise in various applications, such as matrix completion (Candès and Recht 2009; Recht, Fazel, and Parrilo 2010), phase synchronization/retrieval (Singer 2011; Boumal 2016; Shechtman et al. 2015), robust PCA (Ge, Jin, and Zheng 2017), and several others (Chen and Chi 2018; Chi, Lu, and Chen 2019). In this paper, we study a class of low-rank matrix recovery problems, where the goal is to recover a symmetric and positive semidefinite ground truth matrix M^* with $\text{rank}(M^*) = r$ from certain linear measurements corrupted by noise. This problem can be formulated as the following optimization problem:

$$\begin{aligned} \min_{M \in \mathbb{R}^{n \times n}} \quad & \frac{1}{2} \|\mathcal{A}(M) - b + w\|^2 \\ \text{s. t.} \quad & \text{rank}(M) \leq r, \quad M \succeq 0. \end{aligned} \quad (1)$$

Here, $\mathcal{A} : \mathbb{R}^{n \times n} \rightarrow \mathbb{R}^m$ is a linear operator whose action on a matrix M is given by

$$\mathcal{A}(M) = [\langle A_1, M \rangle, \dots, \langle A_m, M \rangle]^T,$$

where $A_1, \dots, A_m \in \mathbb{R}^{n \times n}$ are called sensing matrices. In addition, $b = \mathcal{A}(M^*)$ represents the perfect measurement on the ground truth M^* and w comes from an arbitrary probability distribution. Note that only the noisy measurement $b - w$ is available to the user, and indeed b is unknown. In other words, from a problem-solving perspective, the random variable w is hidden to the user, and it is explicitly modeled here only for the sake of analysis.

The modeling of the noise in this problem is critical for many practical applications in which the influence of the noise cannot be ignored. For example, state estimation is a major data analytic problem for the operation of power grids, which can be modeled as matrix sensing (Jin et al. 2019b). In this problem, each measurement comes from a physical device, and the noise in our formulation models not only the sensory noise but also mismatch between the true model of the system and the one used by a power operator, changes in the measurements due to cyber-attacks, mechanical faults, etc. In other words, by a noisy problem we mean a learning problem where a part of the data is wrong for various reasons, and as with this case, it is impossible to assume that the measurements are noise-free.

Although it may be possible to solve the problem (1) based on convex relaxations (Candès and Recht 2009; Recht, Fazel, and Parrilo 2010; Candès and Tao 2010), the computational complexity associated with solving a semidefinite program presents a major challenge for large-scale problems. A more scalable approach is to use the Burer-Monteiro factorization (Burer and Monteiro 2003) by expressing M as XX^T with $X \in \mathbb{R}^{n \times r}$, which leads to the following equivalent formulation of the original problem (1):

$$\min_{X \in \mathbb{R}^{n \times r}} f(X) = \frac{1}{2} \|\mathcal{A}(XX^T) - b + w\|^2. \quad (2)$$

The unconstrained problem (2) is often solved by local search methods such as gradient descent. Since the objective function $f(X)$ in (2) is non-convex, local search algorithms may converge to a local minimizer, leading to a suboptimal or plainly wrong solution. Hence, it is desirable to provide

^{*}These authors contributed equally.

guarantees on the maximum distance between these local minimizers and the ground truth M^* . This problem will be addressed in this paper.

Related Works

The special noiseless case of the problem (2) can be obtained by setting $w = 0$. In this case, any solution Z with $ZZ^T = M^*$ is a global minimizer of the problem (2). Many previous researches such as Ge, Jin, and Zheng (2017); Bhojanapalli, Neyshabur, and Srebro (2016); Park et al. (2017); Zhang et al. (2018b); Zhu et al. (2018); Zhang, Sojoudi, and Lavaei (2019); Zhang and Zhang (2020); Bi and Lavaei (2020); Ha, Liu, and Barber (2020); Zhu et al. (2021); Zhang (2021) focused on proving that the problem has no spurious (non-global) local minimizers under the assumption of restricted isometry property (RIP). Moreover, as demonstrated in previous works such as Ge, Jin, and Zheng (2017), the developed techniques under the RIP condition can be adopted to show that other low-rank matrix recovery problems, such as matrix completion under incoherence condition and robust PCA problem, also have benign landscape. The RIP condition is equivalent to the restricted strongly convex and smooth property used in Wang, Zhang, and Gu (2017); Park et al. (2018); Zhu et al. (2021), and its formal definition is given below.

Definition 1. The linear operator $\mathcal{A}(\cdot) : \mathbb{R}^{n \times n} \rightarrow \mathbb{R}^m$ is said to satisfy the δ -RIP $_{2r}$ property for some constant $\delta \in [0, 1)$ if the inequality

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

holds for all $M \in \mathbb{R}^{n \times n}$ with $\text{rank}(M) \leq 2r$.

In the recent paper by Zhang (2021), the author developed a sharp bound on the absence of spurious local minima for the noiseless case of problem (2), which says that the problem has no spurious local minima if the measurement operator $\mathcal{A}$ satisfies the δ -RIP $_{2r}$ property with $\delta < 1/2$. This result is tight since there is a known counterexample (Zhang et al. 2018a) having spurious local minima under $\delta = 1/2$.

For the noisy problem, the relation $X^*X^{*T} = M^*$ is unlikely to be satisfied, where X^* denotes a global minimizer of problem (2). However, in this situation, X^*X^{*T} should be close to the ground truth M^* if the noise w is small. As a generalization of the above-mentioned results for the noiseless problem, it is natural to study whether all local minimizers, including the global minimizers, are close to the ground truth M^* under the RIP assumption. One such result is presented in Bhojanapalli, Neyshabur, and Srebro (2016) and given below.

Theorem 1 (from Theorem 3.1 in Bhojanapalli, Neyshabur, and Srebro (2016)). *Suppose that $w \sim \mathcal{N}(0, \sigma_w^2 I_m)$ and $\mathcal{A}(\cdot)$ has the δ -RIP $_{4r}$ property with $\delta < 1/10$. Then, with probability at least $1 - 10/n^2$, any local minimizer $\hat{X}$ of problem (2) satisfies the inequality*

$$\|\hat{X}\hat{X}^T - M^*\|_F \leq 20\sqrt{\frac{\log(n)}{m}}\sigma_w.$$

Theorem 31 in Ge, Jin, and Zheng (2017) further improves the above result by replacing the δ -RIP $_{4r}$ property with the δ -RIP $_{2r}$ property. Li et al. (2020) studies a similar noisy low-rank matrix recovery problem with l_1 norm.

As compared above, there is an evident gap between the state-of-the-art results for the noiseless and noisy problems. The result for the noiseless problem only requires the RIP constant $\delta < 1/2$, but the one for the noisy problem requires $\delta < 1/10$ no matter how small the noise is. This gap will be addressed in this paper by showing that a major generalization of Theorem 1 holds for the noisy problem under the same RIP assumption as in the sharp bound for the noiseless problem.

Notations

In this paper, I_n refers to the identity matrix of size $n \times n$. The notation $M \succeq 0$ means that M is a symmetric and positive semidefinite matrix. $\sigma_i(M)$ denotes the i -th largest singular value of a matrix M , and $\lambda_i(M)$ denotes the i -th largest eigenvalue of M . $\|v\|$ denotes the Euclidean norm of a vector v , while $\|M\|_F$ and $\|M\|_2$ denote the Frobenius norm and induced l_2 norm of a matrix M , respectively. $\langle A, B \rangle$ is defined to be $\text{tr}(A^T B)$ for two matrices A and B of the same size. The Kronecker product between A and B is denoted as $A \otimes B$. For a matrix M , $\text{vec}(M)$ is the usual vectorization operation by stacking the columns of the matrix M into a vector. For a vector $v \in \mathbb{R}^{n^2}$, $\text{mat}(v)$ converts v to a square matrix and $\text{mat}_S(v)$ converts v to a symmetric matrix, i.e., $\text{mat}(v) = M$ and $\text{mat}_S(v) = (M + M^T)/2$, where $M \in \mathbb{R}^{n \times n}$ is the unique matrix satisfying $v = \text{vec}(M)$. Finally, $\mathcal{N}(\mu, \Sigma)$ refers to the multivariate Gaussian distribution with mean μ and covariance Σ .

Main Results

We first present the global guarantee on the local minimizers of the problem (2). To simplify the notation, we use a matrix representation of the measurement operator $\mathcal{A}$ as follows:

$$\mathbf{A} = [\text{vec}(A_1), \text{vec}(A_2), \dots, \text{vec}(A_m)]^T \in \mathbb{R}^{m \times n^2}.$$

Then, $\mathbf{A} \text{vec}(M) = \mathcal{A}(M)$ for every matrix $M \in \mathbb{R}^{n \times n}$.

Theorem 2. *Assume that the linear operator $\mathcal{A}$ satisfies the δ -RIP $_{2r}$ property with $\delta < 1/2$. For every $\epsilon > 0$, with probability at least $\mathbb{P}(\|\mathbf{A}^T w\| \leq \epsilon)$, either of the following two inequalities*

$$(1 - \delta)\|\hat{X}\hat{X}^T - M^*\|_F^2 \leq \epsilon\sqrt{r}\|\hat{X}\hat{X}^T - M^*\|_F + 4\epsilon\sqrt{r}\|M^*\|_F, \quad (3a)$$

$$\frac{2(1 - 2\delta)}{3(1 + \delta)}\|\hat{X}\hat{X}^T - M^*\|_F \leq 2\epsilon\sqrt{r} + 2\sqrt{2\epsilon(1 + \delta)}(\|\hat{X}\hat{X}^T - M^*\|_F^{1/2} + \|M^*\|_F^{1/2}) \quad (3b)$$

holds for every arbitrary local minimizer $\hat{X} \in \mathbb{R}^{n \times r}$ of problem (2).

Note that two upper bounds on the distance $\|\hat{X}\hat{X}^T - M^*\|_F$ can be obtained for any local minimizer $\hat{X}$ by solving the two quadratic-like inequalities (3a) and (3b), and the

larger bound needs to be used because only one of the two inequalities is guaranteed to hold. The explicit upper bound solved from Theorem 2 is given in Appendix A. The reason for the existence of two inequalities in Theorem 2 is the split of its proof into two cases. The first case is when the r -th smallest singular value of $\hat{X}$ is small, and the second case is the opposite, which are respectively handled by Lemma 2 and Lemma 3.

The reason for the existence of two inequalities in Theorem 2 is the split of its proof into two cases. The first case is when the r -th smallest singular value of $\hat{X}$ is small, and the second case is the opposite, which are respectively handled by Lemma 2 and Lemma 3.

Theorem 2 is a major extension of the existing sharp result stating that the noiseless problem has no spurious local minima under the same assumption of the δ -RIP_{2r} property with $\delta < 1/2$. The reason is that in the case when the noise w is equal to zero, one can choose an arbitrarily small ϵ in Theorem 2 to conclude from the inequalities (3a) and (3b) that $\hat{X}\hat{X}^T = M^*$ for every local minimizer $\hat{X}$. Moreover, when the RIP constant δ further decreases from $1/2$, the upper bound on $\|\hat{X}\hat{X}^T - M^*\|_F$ will also decrease, which means that a local minimizer found by local search methods will be closer to the ground truth M^* . This suggests that the RIP condition is able to not only guarantee the absence of spurious local minima as shown in the previous literature but also mitigate the influence of the noise in the measurements.

Compared with the existing results such as Theorem 1, our new result has two advantages. First, by improving the RIP constant from $1/10$ to $1/2$, one can apply the results on the location of spurious local minima to a much broader class of problems, which can often help reduce the number of measurements. For example, in the case when the measurements are given by random Gaussian matrices, it is proven in Candès and Recht (2009) that to achieve the δ -RIP_{2r} property the minimum number of measurements needed is in the order of $O(1/\delta^2)$. By improving the RIP constant in the bound, we can significantly reduce the number of measurements while still keeping the benign landscape. In applications such as learning for energy networks, there is a fundamental limit on the number of measurements that can be collected due to the physics of the problem (Jin et al. 2021). Finding a better bound on RIP helps with addressing the issues with the number of measurements needed to reliably solve the problem. Second, Theorem 1 is just about the probability of having all spurious solutions in a fixed ball around the ground truth of radius $O(\sigma_w)$ instead of balls of arbitrary radii, and this fixed ball could be a large one depending on whether the noise level σ_w is fixed or scales with the problem. On the other hand, in Theorem 2, we consider the probability $\mathbb{P}(\|\mathbf{A}^T w\| \leq \epsilon)$ for any arbitrary value of ϵ . By having a flexible ϵ , our work not only improves the RIP constant but also allows computing the probability of having all spurious solutions in any given ball.

In the special case of rank $r = 1$, the conditions (3a) and (3b) in Theorem 2 can be substituted with a simpler condition as shown below. Its proof is highly similar to that of

Lemma 3, and obviated in this paper for succinctness.

Theorem 3. Consider the case $r = 1$ and assume that the linear operator $\mathcal{A}$ satisfies the δ -RIP₂ property with $\delta < 1/2$. For every $\epsilon > 0$, with probability at least $\mathbb{P}(\|\mathbf{A}^T w\| \leq \epsilon)$, every arbitrary local minimizer $\hat{X} \in \mathbb{R}^{n \times r}$ of problem (2) satisfies

$$\|\hat{X}\hat{X}^T - M^*\|_F \leq \frac{3(1 + \sqrt{2})\epsilon(1 + \delta)}{1 - 2\delta}. \quad (4)$$

In the case when the RIP constant δ is not less than $1/2$, it is not possible to achieve a global guarantee similar to Theorem 2 or Theorem 3 since it is known that the problem may have a spurious solution even in the noiseless case. Instead, we turn to local guarantees by showing that every arbitrary local minimizer $\hat{X}$ of problem (2) is either close to the ground truth M^* or far away from it in terms of the distance $\|\hat{X}\hat{X}^T - M^*\|_F$.

Theorem 4. Assume that the linear operator $\mathcal{A}$ satisfies the δ -RIP_{2r} property for some $\delta \in [0, 1)$. Consider arbitrary constants $\epsilon > 0$ and $\tau \in (0, 2(\sqrt{2} - 1))$ such that

$$\delta < \sqrt{1 - \frac{3 + 2\sqrt{2}}{4}\tau^2}.$$

Every arbitrary local minimizer $\hat{X} \in \mathbb{R}^{n \times r}$ of problem (2) satisfying

$$\|\hat{X}\hat{X}^T - M^*\|_F \leq \tau\lambda_r(M^*) \quad (5)$$

will also satisfy

$$\|\hat{X}\hat{X}^T - M^*\|_F \leq \frac{\epsilon(1 + \delta)C(\tau, M^*)}{\sqrt{1 - \frac{3 + 2\sqrt{2}}{4}\tau^2} - \delta} \quad (6)$$

with probability at least $\mathbb{P}(\|\mathbf{A}^T w\| \leq \epsilon)$, where

$$C(\tau, M^*) = \sqrt{\frac{2(\lambda_1(M^*) + \tau\lambda_r(M^*))}{(1 - \tau)\lambda_r(M^*)}}.$$

The upper bounds in (5) and (6) define an outer ball and an inner ball centered at the ground truth M^* . Theorem 4 states that there is no local minimizer in the ring between the two balls. As ϵ approaches zero, the inner ball shrinks to the ground truth. This theorem shows that bad local minimizers are located outside the outer ball. Note that the problem could be highly non-convex when δ is close to 1, while this theorem shows a benign landscape in a local neighborhood of the solution. Furthermore, all the theorems in this section are applicable to arbitrary noise models since they make no explicit use of the probability distribution of the noise. The only required information is the probability $\mathbb{P}(\|\mathbf{A}^T w\| \leq \epsilon)$, which can be computed or bounded when the probability distribution of the noise is given as illustrated in Section .

The results presented above are all about the location of the local minimizers. They do not directly lead to the global convergence of local search methods with a fast convergence rate. To provide performance guarantees for local search methods, the next theorem establishes a stronger property

for the landscape of the noisy problem that is usually called the strict saddle property in the literature, which essentially says that all approximate second-order critical points are close to the ground truth.

Theorem 5. Assume that the linear operator $\mathcal{A}$ satisfies the δ -RIP_{2r} property with $\delta < 1/2$. For every $\epsilon > 0$ and $\kappa \geq 0$, with probability at least $\mathbb{P}(\|\mathbf{A}^T w\| \leq \epsilon)$, either of the following two inequalities

$$(1 - \delta) \|\hat{X} \hat{X}^T - M^*\|_F^2 \leq \epsilon \sqrt{r} \|\hat{X} \hat{X}^T - M^*\|_F + \frac{r^{1/4} \kappa}{2} \|\hat{X} \hat{X}^T - M^*\|_F^{1/2} + \frac{r^{1/4} \kappa}{2} \|M^*\|_F^{1/2} \quad (7a)$$

$$+ (4\sqrt{r}\epsilon + \frac{5\sqrt{r}\kappa}{2}) \|M^*\|_F$$

$$\frac{2(1 - 2\delta)}{3(1 + \delta)} \|\hat{X} \hat{X}^T - M^*\|_F \leq (2\epsilon + \kappa) \sqrt{r} + \sqrt{2\kappa(1 + \delta)} + 2\sqrt{2\epsilon(1 + \delta)} (\|\hat{X} \hat{X}^T - M^*\|_F^{1/2} + \|M^*\|_F^{1/2}) \quad (7b)$$

holds for every matrix $\hat{X} \in \mathbb{R}^{n \times r}$ satisfying

$$\|\nabla f(\hat{X})\| \leq \kappa, \quad \nabla^2 f(\hat{X}) \succeq -\kappa I_{nr}.$$

Note that this property is not proven in the literature for $\delta < 1/2$ even in the noiseless case, and thus our result generalizes the existing ones even in this scenario. On the other hand, it is proven by Jin et al. (2017) that the perturbed gradient descent method with an arbitrary initialization will find a solution $\hat{X}$ satisfying the requirements in Theorem 5 with a high probability in $O(\text{poly}(1/\kappa))$ number of iterations. By Theorem 5, $\hat{X} \hat{X}^T$ will be close to the ground truth if ϵ and κ are chosen to be relatively small.

Table 1 briefly summarizes our result compared with the existing literature.

Proofs of Main Results

Before presenting the proofs, we first compute the gradient and the Hessian of the objective function $f(\hat{X})$ of the problem (2):

$$\nabla f(\hat{X}) = \hat{\mathbf{X}}^T \mathbf{A}^T (\mathbf{A} \mathbf{e} + w),$$

$$\nabla^2 f(\hat{X}) = 2I_r \otimes \text{mat}_S(\mathbf{A}^T (\mathbf{A} \mathbf{e} + w)) + \hat{\mathbf{X}}^T \mathbf{A}^T \mathbf{A} \hat{\mathbf{X}},$$

where

$$\mathbf{e} = \text{vec}(\hat{X} \hat{X}^T - M^*),$$

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

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

The first step in the proofs is to derive necessary conditions for a matrix $\hat{X} \in \mathbb{R}^{n \times r}$ to be an approximate second-order critical point, which depend on the linear operator $\mathcal{A}$, the noise $w \in \mathbb{R}^m$, the solution $\hat{X}$, and the parameter κ characterizing how close $\hat{X}$ is to a true second-order critical point.

Lemma 1. Given $\kappa \geq 0$, assume that $\hat{X} \in \mathbb{R}^{n \times r}$ satisfies

$$\|\nabla f(\hat{X})\| \leq \kappa, \quad \nabla^2 f(\hat{X}) \succeq -\kappa I_{nr}.$$

Then, it must satisfy the following inequalities:

$$\|\hat{\mathbf{X}}^T \mathbf{H} \mathbf{e}\| \leq 2\|\hat{X}\|_2 \|\mathbf{A}^T w\| + \kappa, \quad (8a)$$

$$2I_r \otimes \text{mat}_S(\mathbf{H} \mathbf{e}) + \hat{\mathbf{X}}^T \mathbf{H} \hat{\mathbf{X}} \succeq -(2\|\mathbf{A}^T w\| + \kappa) I_{nr}, \quad (8b)$$

where $\mathbf{H} = \mathbf{A}^T \mathbf{A}$.

Proof. To obtain condition (8a), notice that $\|\nabla f(\hat{X})\| \leq \kappa$ implies that

$$\|\hat{\mathbf{X}}^T \mathbf{H} \mathbf{e}\| \leq \|\hat{\mathbf{X}}^T \mathbf{A}^T w\| + \kappa \leq \|\hat{X}\|_2 \|\mathbf{A}^T w\| + \kappa \leq 2\|\hat{X}\|_2 \|\mathbf{A}^T w\| + \kappa,$$

in which the last inequality is due to

$$\|\hat{\mathbf{X}} \text{vec}(U)\| = \|\hat{X} U^T + U \hat{X}^T\|_F \leq 2\|\hat{X}\|_2 \|U\|_F,$$

for every $U \in \mathbb{R}^{n \times r}$. Similarly, $\nabla^2 f(\hat{X}) \succeq -\kappa I_{nr}$ implies that

$$2I_r \otimes \text{mat}_S(\mathbf{H} \mathbf{e}) + \hat{\mathbf{X}}^T \mathbf{H} \hat{\mathbf{X}} \succeq -2I_r \otimes \text{mat}_S(\mathbf{A}^T w) - \kappa I_{nr}.$$

On the other hand, the eigenvalues of $I_r \otimes \text{mat}_S(\mathbf{A}^T w)$ are the same as those of $\text{mat}_S(\mathbf{A}^T w)$, and each eigenvalue $\lambda_i(\text{mat}_S(\mathbf{A}^T w))$ of the latter matrix further satisfies

$$|\lambda_i(\text{mat}_S(\mathbf{A}^T w))| \leq \|\text{mat}_S(\mathbf{A}^T w)\|_F \leq \|\mathbf{A}^T w\|,$$

which proves condition (8b). $\square$

If $\hat{X}$ is a local minimizer of the problem (2), Lemma 1 shows that $\hat{X}$ satisfies the inequalities (8a) and (8b) with $\kappa = 0$. Similarly, Theorem 2 can also be regarded as a special case of Theorem 5 with $\kappa = 0$. The proofs of these two theorems consist of inspecting two cases. The following lemma deals with the first case in which $\hat{X}$ is an approximate second-order critical point with $\sigma_r(\hat{X})$ being close to zero.

Lemma 2. Given $\hat{X} \in \mathbb{R}^{n \times r}$ and arbitrary constants $\epsilon > 0$ and $\kappa \geq 0$, the inequalities

$$\sigma_r(\hat{X}) \leq \sqrt{\frac{\epsilon + \kappa}{1 + \delta}}, \quad \|\nabla f(\hat{X})\| \leq \kappa, \quad \nabla^2 f(\hat{X}) \succeq -\kappa I_{nr}$$

and $\|\mathbf{A}^T w\| \leq \epsilon$ will together imply the inequality (7a).

Proof. Let $G = \text{mat}_S(\mathbf{H} \mathbf{e})$ and $u \in \mathbb{R}^n$ be a unit eigenvector of G corresponding to its minimum eigenvalue, i.e.,

$$\|u\| = 1, \quad Gu = \lambda_{\min}(G)u.$$

In addition, let $v \in \mathbb{R}^r$ be a singular vector of $\hat{X}$ such that

$$\|v\| = 1, \quad \|\hat{X} v\| = \sigma_r(\hat{X}).$$

Let $\mathbf{U} = \text{vec}(uv^T)$. Then, $\|\mathbf{U}\| \leq 1$ and (8b) implies that

$$\begin{aligned} -2\epsilon - \kappa &\leq 2\mathbf{U}^T (I_r \otimes \text{mat}_S(\mathbf{H} \mathbf{e})) \mathbf{U} + \mathbf{U}^T \hat{\mathbf{X}}^T \mathbf{H} \hat{\mathbf{X}} \mathbf{U} \\ &\leq 2\text{tr}(vu^T G uv^T) + (1 + \delta) \|\hat{X} v u^T + u v^T \hat{X}^T\|_F^2 \\ &\leq 2\lambda_{\min}(G) + 4(1 + \delta) \sigma_r(\hat{X})^2 \\ &\leq 2\lambda_{\min}(G) + 4\epsilon + 4\kappa. \end{aligned} \quad (9)$$

Paper	Noise	RIP Assumption	Rank	Convergence
Bhojanapalli, Neyshabur, and Srebro (2016)	Isotropic Gaussian	$\delta < 1/10$	rank r	N/A
Zhang, Sojoudi, and Lavaei (2019)	Noiseless	$\delta < 1/2$	rank 1	N/A
Zhang (2021)	Noiseless	$\delta < 1/2$	rank r	N/A
Ours	Finite Variance	$\delta < 1/2$	rank r	Polynomial

Table 1: Comparison between our result and the existing literature.

On the other hand,

$$\begin{aligned}
(1 - \delta) \|\hat{X} \hat{X}^T - M^*\|_F^2 &\leq \mathbf{e}^T \mathbf{H} \mathbf{e} \\
&= \text{vec}(\hat{X} \hat{X}^T)^T \mathbf{H} \mathbf{e} - \text{vec}(M^*)^T \mathbf{H} \mathbf{e} \\
&= \frac{1}{2} \text{vec}(\hat{X})^T \hat{\mathbf{X}}^T \mathbf{H} \mathbf{e} - \langle M^*, \text{mat}_S(\mathbf{H} \mathbf{e}) \rangle \\
&\leq \frac{1}{2} \|\hat{X}\|_2 \|\hat{\mathbf{X}}^T \mathbf{H} \mathbf{e}\| + \left(3\epsilon + \frac{5\kappa}{2}\right) \text{tr}(M^*) \\
&\leq \epsilon \|\hat{X}\|_2^2 + \frac{\kappa}{2} \|\hat{X}\|_2 + \left(3\epsilon + \frac{5\kappa}{2}\right) \text{tr}(M^*),
\end{aligned}$$

in which the second last inequality is due to (9) and the last inequality is due to (8a). Furthermore, the right-hand side of the above inequality can be relaxed as

$$\begin{aligned}
&\epsilon \|\hat{X}\|_2^2 + \frac{\kappa}{2} \|\hat{X}\|_2 + \left(3\epsilon + \frac{5\kappa}{2}\right) \text{tr}(M^*) \leq \epsilon \sqrt{r} \|\hat{X} \hat{X}^T\|_F \\
&\quad + \frac{r^{1/4} \kappa}{2} \|\hat{X} \hat{X}^T\|_F^{1/2} + \left(3\epsilon + \frac{5\kappa}{2}\right) \sqrt{r} \|M^*\|_F \\
&\leq \epsilon \sqrt{r} \|\hat{X} \hat{X}^T - M^*\|_F + \frac{r^{1/4} \kappa}{2} \|\hat{X} \hat{X}^T - M^*\|_F^{1/2} \\
&\quad + \left(4\sqrt{r}\epsilon + \frac{5\sqrt{r}\kappa}{2}\right) \|M^*\|_F + \frac{r^{1/4} \kappa}{2} \|M^*\|_F^{1/2},
\end{aligned}$$

which leads to the inequality (7a). $\square$

The remaining case with

$$\sigma_r(\hat{X}) > \sqrt{\frac{\epsilon + \kappa}{1 + \delta}}$$

will be handled in the following lemma using a different method.

Lemma 3. Assume that the linear operator $\mathcal{A}$ satisfies the δ -RIP $_{2r}$ property with $\delta < 1/2$. Given $\hat{X} \in \mathbb{R}^{n \times r}$ and arbitrary constants $\epsilon > 0$ and $\kappa \geq 0$, the inequalities

$$\sigma_r(\hat{X}) > \sqrt{\frac{\epsilon + \kappa}{1 + \delta}}, \quad \|\nabla f(\hat{X})\| \leq \kappa, \quad \nabla^2 f(\hat{X}) \succeq -\kappa I_{nr}$$

and $\|\mathbf{A}^T w\| \leq \epsilon$ will together imply the inequality (7b).

The proofs of both Lemma 3 and the local guarantee in Theorem 4 generalize the proof of the absence of spurious local minima for the noiseless problem in Zhang and Zhang (2020); Zhang (2021). Our innovation here is to develop

new techniques to analyze approximate optimality conditions for the solutions because unlike the noiseless problem the local minimizers of the noisy one are only approximate second-order critical points of the distance function $\|\mathcal{A}(X X^T) - b\|^2$. For a fixed solution $\hat{X}$ and noise w , one can find an operator $\hat{\mathcal{A}}$ satisfying the δ -RIP $_{2r}$ property with the smallest possible δ such that $\hat{X}$ and $\hat{\mathcal{A}}$ satisfy the necessary conditions stated in Lemma 1. Let $\delta^*(\hat{X})$ be the RIP constant of the found measurement operator $\hat{\mathcal{A}}$ in the worst-case scenario. Then, if $\hat{X}$ in Lemma 3 is a solution of the current problem with the linear operator $\mathcal{A}$ satisfying the δ -RIP $_{2r}$ property, it holds that $\delta \geq \delta^*(\hat{X})$, which can further lead to an upper bound on the distance $\|\hat{X} \hat{X}^T - M^*\|_F$.

To compute $\delta^*(\hat{X})$ defined above, let $q = \mathbf{A}^T w$ and solve the following optimization problem whose optimal value is $\delta^*(\hat{X})$:

$$\begin{aligned}
&\min_{\delta, \hat{\mathbf{H}}} \delta \\
&\text{s.t.} \quad \|\hat{\mathbf{X}}^T \hat{\mathbf{H}} \mathbf{e}\| \leq 2\|\hat{X}\|_2 \|q\| + \kappa, \\
&\quad 2I_r \otimes \text{mat}_S(\hat{\mathbf{H}} \mathbf{e}) + \hat{\mathbf{X}}^T \hat{\mathbf{H}} \hat{\mathbf{X}} \succeq -(2\|q\| + \kappa) I_{nr}, \\
&\quad \hat{\mathbf{H}} \text{ is symmetric and satisfies the } \delta\text{-RIP}_{2r} \text{ property.}
\end{aligned} \tag{10}$$

Note that a matrix $\hat{\mathbf{H}} \in \mathbb{R}^{n^2 \times n^2}$ is said to satisfy the δ -RIP $_{2r}$ property if

$$(1 - \delta) \|\mathbf{U}\|^2 \leq \mathbf{U}^T \hat{\mathbf{H}} \mathbf{U} \leq (1 + \delta) \|\mathbf{U}\|^2$$

holds for every matrix $U \in \mathbb{R}^{n \times n}$ with $\text{rank}(U) \leq 2r$ and $\mathbf{U} = \text{vec}(U)$. Obviously, for a linear operator $\hat{\mathcal{A}}$, $\hat{\mathbf{H}} = \hat{\mathcal{A}}^T \hat{\mathcal{A}}$ satisfies the δ -RIP $_{2r}$ property if and only if $\hat{\mathcal{A}}$ satisfies the δ -RIP $_{2r}$ property.

However, since problem (10) is non-convex due to the RIP constraint, we instead solve the following convex reformulation:

$$\begin{aligned}
&\min_{\delta, \hat{\mathbf{H}}} \delta \\
&\text{s.t.} \quad \|\hat{\mathbf{X}}^T \hat{\mathbf{H}} \mathbf{e}\| \leq 2\|\hat{X}\|_2 \|q\| + \kappa, \\
&\quad 2I_r \otimes \text{mat}_S(\hat{\mathbf{H}} \mathbf{e}) + \hat{\mathbf{X}}^T \hat{\mathbf{H}} \hat{\mathbf{X}} \succeq -(2\|q\| + \kappa) I_{nr}, \\
&\quad (1 - \delta) I_{n^2} \preceq \hat{\mathbf{H}} \preceq (1 + \delta) I_{n^2}.
\end{aligned} \tag{11}$$

Lemma 14 in Bi and Lavaei (2020) proves that problem (10) and problem (11) have the same optimal value. The remain-

ing step in the proof of Lemma 3 is to solve the optimization problem (11) for given $\hat{X}$, q and κ . The complete proof of Lemma 3 is lengthy and deferred to Appendix B. Finally, Theorem 2 and Theorem 5 are direct consequences of Lemma 2 and Lemma 3. The proof of Theorem 3 is very similar to that of Lemma 3 and is also given in Appendix B.

Now, we turn to the proof of the local guarantee in Theorem 4. The following existing result will be useful.

Lemma 4 (from Lemma 14 in Zhang, Sojoudi, and Lavaei (2019)). *Given $a, b \in \mathbb{R}^n$, the rank-2 matrix $ab^T + ba^T$ has two possibly nonzero eigenvalues*

$$\|a\|\|b\|(1 + \cos \theta), \quad -\|a\|\|b\|(1 - \cos \theta).$$

Here, θ is the angle between a and b .

Proof of Theorem 4. First, we relax the optimization problem (11) by dropping the constraint related to the second-order necessary optimality condition. This gives rise to the optimization problem

$$\begin{aligned} \min_{\delta, \hat{\mathbf{H}}} \quad & \delta \\ \text{s. t.} \quad & \|\hat{\mathbf{X}}^T \hat{\mathbf{H}} \mathbf{e}\| \leq 2\|\hat{X}\|_2 \|q\|, \\ & (1 - \delta)I_{n^2} \preceq \hat{\mathbf{H}} \preceq (1 + \delta)I_{n^2}. \end{aligned} \quad (12)$$

To further simplify the problem (12), one can replace its decision variable δ with η and introduce the following optimization problem:

$$\begin{aligned} \max_{\eta, \hat{\mathbf{H}}} \quad & \eta \\ \text{s. t.} \quad & \|\hat{\mathbf{X}}^T \hat{\mathbf{H}} \mathbf{e}\| \leq 2\|\hat{X}\|_2 \|q\|, \\ & \eta I_{n^2} \preceq \hat{\mathbf{H}} \preceq I_{n^2}. \end{aligned} \quad (13)$$

Given any feasible solution $(\delta, \hat{\mathbf{H}})$ to (12), the tuple

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

is a feasible solution to problem (13). Therefore, if the optimal value of (12) is denoted as $\delta_f^*(\hat{X})$ and the optimal value of (13) is denoted as $\eta_f^*(\hat{X})$, then it holds that

$$\eta_f^*(\hat{X}) \geq \frac{1 - \delta_f^*(\hat{X})}{1 + \delta_f^*(\hat{X})} \geq \frac{1 - \delta^*(\hat{X})}{1 + \delta^*(\hat{X})} \geq \frac{1 - \delta}{1 + \delta}, \quad (14)$$

in which the last inequality is implied by $\delta \geq \delta^*(\hat{X})$ as shown above. To prove the inequality (6), we need to bound $\eta_f^*(\hat{X})$ from above, which can be achieved by finding a feasible solution to the dual problem of (13) given below:

$$\begin{aligned} \min_{U_1, U_2, G, \lambda, y} \quad & \text{tr}(U_2) + 4\|\hat{X}\|_2^2 \|q\|^2 \lambda + \text{tr}(G) \\ \text{s. t.} \quad & \text{tr}(U_1) = 1, \\ & (\hat{\mathbf{X}}y)\mathbf{e}^T + \mathbf{e}(\hat{\mathbf{X}}y)^T = U_1 - U_2, \\ & \begin{bmatrix} G & -y \\ -y^T & \lambda \end{bmatrix} \succeq 0, \\ & U_1 \succeq 0, \quad U_2 \succeq 0. \end{aligned} \quad (15)$$

For any matrix $\hat{X} \in \mathbb{R}^{n \times r}$ satisfying $\|\hat{X}\hat{X}^T - M^*\|_F \leq \tau\lambda_r(M^*)$, we have $\hat{X} \neq 0$, and it has been shown in the proof of Lemma 19 in Bi and Lavaei (2020) that there exists $y \neq 0$ satisfying the inequalities

$$\|\hat{\mathbf{X}}y\|^2 \geq 2\lambda_r(\hat{X}\hat{X}^T)\|y\|^2, \quad (16a)$$

$$\cos \theta \geq \sqrt{1 - \frac{3 + 2\sqrt{2}}{4}\tau^2}, \quad (16b)$$

where θ is the angle between $\hat{\mathbf{X}}y$ and $\mathbf{e}$. Note that (16b) holds only in the exact-parameterized regime, i.e., the case with $\text{rank}(M^*) = r$, since the derivation of (16b) utilized Lemma 5.4 of Tu et al. (2016), which does not hold when $\text{rank}(M^*) < r$. This is the main reason why our result cannot be directly generalized to the overparameterized case. Now, define

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

and then decompose M as $M = [M]_+ - [M]_-$ with $[M]_+ \succeq 0$ and $[M]_- \succeq 0$. Then, it is easy to verify that $(U_1^*, U_2^*, G^*, \lambda^*, y^*)$ defined as

$$\begin{aligned} y^* &= \frac{y}{\text{tr}([M]_+)}, \quad U_1^* = \frac{[M]_+}{\text{tr}([M]_+)}, \quad U_2^* = \frac{[M]_-}{\text{tr}([M]_+)}, \\ G^* &= \frac{y^*(y^*)^T}{\lambda^*}, \quad \lambda^* = \frac{\|y^*\|}{2\|\hat{X}\|_2 \|q\|} \end{aligned}$$

forms a feasible solution to the dual problem (15) with the objective value

$$\frac{\text{tr}([M]_-) + 4\|\hat{X}\|_2 \|q\| \|y\|}{\text{tr}([M]_+)}. \quad (17)$$

Furthermore, $\text{rank}(M^*) = r$ implies that $\lambda_r(M^*) > 0$. By the Wielandt–Hoffman theorem,

$$|\lambda_r(\hat{X}\hat{X}^T) - \lambda_r(M^*)| \leq \|\hat{X}\hat{X}^T - M^*\|_F \leq \tau\lambda_r(M^*),$$

$$|\lambda_1(\hat{X}\hat{X}^T) - \lambda_1(M^*)| \leq \|\hat{X}\hat{X}^T - M^*\|_F \leq \tau\lambda_r(M^*).$$

Thus, using the above two inequalities and inequality (16a), we have

$$\begin{aligned} \frac{2\|\hat{X}\|_2 \|y\|}{\|\hat{\mathbf{X}}y\|} &\leq \frac{2\|\hat{X}\|_2}{\sqrt{2\lambda_r(\hat{X}\hat{X}^T)}} \\ &\leq \sqrt{\frac{2(\lambda_1(M^*) + \tau\lambda_r(M^*))}{(1 - \tau)\lambda_r(M^*)}} = C(\tau, M^*). \end{aligned} \quad (18)$$

Next, according to Lemma 4, one can write

$$\text{tr}([M]_+) = \|\hat{\mathbf{X}}y\| \|\mathbf{e}\| (1 + \cos \theta),$$

$$\text{tr}([M]_-) = \|\hat{\mathbf{X}}y\| \|\mathbf{e}\| (1 - \cos \theta).$$


Substituting the above two equations and (18) into the dual objective value (17), one can obtain

$$\eta_f^*(\hat{X}) \leq \frac{1 - \cos \theta + 2C(\tau, M^*)\|q\| \|\mathbf{e}\|}{1 + \cos \theta},$$


which together with (14) implies that

$$\|\mathbf{e}\| \leq (1 + \delta)C(\tau, M^*)\|q\|(\cos \theta - \delta)^{-1}.$$

The inequality (6) can then be proved by combining the above inequality and (16b) under the probabilistic event that $\|q\| \leq \epsilon$. $\square$




(a) The upper bound derived from inequality (3a).




(b) The upper bound derived from inequality (3b).

Figure 1: Comparison of the upper bounds given by Theorem 2 for the distance $\|\hat{X}\hat{X}^T - M^*\|_F$ with $\hat{X}$ being an arbitrary local minimizer.



(a) The upper bound derived from inequality (3a).



(b) The upper bound derived from inequality (3b).

Figure 2: Upper bounds given by Theorem 2 for the distance $\|\hat{X}\hat{X}^T - M^*\|_F$ based on the explicit distribution on the noise w .

Numerical Illustration

In the next, we will empirically study the developed probabilistic guarantees and demonstrate the distance $\|\hat{X}\hat{X}^T - M^*\|_F$ between any local minimizer $\hat{X}$ and the ground truth M^* as well as the value of the RIP constant δ required to be satisfied by the linear operator $\mathcal{A}$.

Before delving into the numerical illustration, note that the probability $\mathbb{P}(\|\mathbf{A}^T w\| \leq \epsilon)$ used in both Theorem 2 and Theorem 4 can be bounded from below by the probability $\mathbb{P}(\|w\| \leq w_0)$ with $w_0 = \epsilon/\|\mathbf{A}\|_2$. The latter probability can be easily estimated when the probability distribution of the noise w is given. As an example, in the simplest case when w is sampled from an isotropic Gaussian distribution, i.e., $w \sim \mathcal{N}(0, \sigma^2 I_m)$, the random variable $\|w/\sigma\|^2$ follows the chi-square distribution and one can apply the Chernoff bound to obtain

$$\begin{aligned} \mathbb{P}(\|w\| \leq w_0) &= 1 - \mathbb{P}\left(\left\|\frac{w}{\sigma}\right\|^2 \geq \frac{w_0^2}{\sigma^2}\right) \\ &\geq 1 - \inf_{0 \leq t < 1/2} (1 - 2t)^{-m/2} e^{-tw_0^2/\sigma^2}. \end{aligned}$$

After solving the minimization problem in the above

equation, we obtain

$$1 - \left(\frac{2m\sigma^2}{w_0^2}\right)^{-m/2} e^{m - \frac{w_0^2}{2\sigma^2}} \leq \mathbb{P}(\|w\| \leq w_0).$$

More generally, if w is a $(\sigma/\sqrt{m})$ -sub-Gaussian vector, then applying Lemma 1 in Jin et al. (2019a) leads to

$$1 - 2e^{-\frac{w_0^2}{16m\sigma^2}} \leq \mathbb{P}(\|w\| \leq w_0).$$

For numerical illustration, assume that $n = 50$, $m = 10$ and $\|\mathbf{A}\|_2 \leq 2$, while the noise w is a $(0.05/\sqrt{m})$ -sub-Gaussian vector. We also assume that the ground truth M^* is of rank 5 with the largest eigenvalue being 1.5 and the smallest eigenvalue being 1.

First, we explore the two inequalities (3a) and (3b) in Theorem 2 to obtain two upper bounds on $\|\hat{X}\hat{X}^T - M^*\|_F$, where $\hat{X}$ denotes any arbitrary (worst) local minimizer. Figure 1 gives the contour plots of the two upper bounds on $\|\hat{X}\hat{X}^T - M^*\|_F$, which hold with the given probability on the y -axis and the given RIP constant δ from 0 to 1/2 on the x -axis. While the final bound on $\|\hat{X}\hat{X}^T - M^*\|_F$ is often determined by the inequality (3b), the inequality (3a) is




Figure 3: Comparison of the maximum RIP constants δ allowed by Theorem 2 and Theorem 4 to guarantee a given maximum distance $\|\hat{X}\hat{X}^T - M^*\|_F$ for an arbitrary local minimizer $\hat{X}$ satisfying (5) with a given probability.

needed theoretically to deal with the case when $\hat{X}$ has a singular value close to 0.

Furthermore, a tighter bound on the success probability can be derived by calculating the exact probability $\mathbb{P}(\|\mathbf{A}^T w\| \leq \epsilon)$ for an explicit distribution on w . Figure 2 is obtained in this way under the same assumptions as those for Figure 1 except that w is isotropic Gaussian with the same parameter in the sub-Gaussian assumption. Compared with Figure 1, the shape is similar, but the bound is tighter.

Next, we illustrate the bounds given by Theorem 2 and Theorem 4. Figure 3 shows the contour plots of the maximum RIP constant δ that is necessary to guarantee that each local minimizer $\hat{X}$ (satisfying the inequality (5) when Theorem 4 is applied) lies within a certain neighborhood of the ground truth (measured via the distance $\|\hat{X}\hat{X}^T - M^*\|_F$ on the x -axis) with a given probability on the y -axis, as implied by the respective global and local guarantees. Figure 3 clearly shows how a smaller RIP constant δ leads to a tighter bound on the distance $\|\hat{X}\hat{X}^T - M^*\|_F$ with a higher probability. In addition, the local guarantee generally requires a looser RIP assumption as it still holds even when $\delta > 1/2$. However, as the parameter τ in Theorem 4 increases, the local bound also degrades quickly, sometimes becoming worse

than the global bound as illustrated in Figure 3(d).

Moreover, in our experiment, we have also tried different values of problem parameters m and n . They all yield similar results and are included in Appendix C for completeness.

Conclusion

In this paper, we develop global and local analyses for the locations of the local minima of the low-rank matrix recovery problem with noisy linear measurements. Unlike the existing results, the probability distribution of the noise is arbitrary and the RIP constant of the problem is free to take any arbitrary value. The developed results encompass the state-of-the-art results on the non-existence of spurious solutions in the noiseless case. Furthermore, we prove the strict saddle property, which guarantees the global convergence of the perturbed gradient descent method in polynomial time. Our analyses show how the value of the RIP constant and the intensity of noise affect the landscape of the non-convex learning problem and the locations of the local minima relative to the ground truth. Future research directions include extending our results into the cases when the matrices are asymmetric, the measurements are nonlinear, or the overparameterized regime in which $\text{rank}(M^*)$ is less than r .

Acknowledgments

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

References

- 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.
- Boumal, N. 2016. Nonconvex phase synchronization. *SIAM Journal on Optimization*, 26(4): 2355–2377.
- Burer, S.; and Monteiro, R. D. 2003. A nonlinear programming algorithm for solving semidefinite programs via low-rank factorization. *Mathematical Programming*, 95(2): 329–357.
- Candès, E. J.; and Recht, B. 2009. Exact matrix completion via convex optimization. *Foundations of Computational Mathematics*, 9(6): 717–772.
- Candès, E. J.; and Tao, T. 2010. The power of convex relaxation: Near-optimal matrix completion. *IEEE Transactions on Information Theory*, 56(5): 2053–2080.
- Chen, Y.; and Chi, Y. 2018. Harnessing structures in big data via guaranteed low-rank matrix estimation: Recent theory and fast algorithms via convex and nonconvex optimization. *IEEE Signal Processing Magazine*, 35(4): 14–31.
- Chi, Y.; Lu, Y. M.; and Chen, Y. 2019. Nonconvex optimization meets low-rank matrix factorization: An overview. *IEEE Transactions on Signal Processing*, 67(20): 5239–5269.
- 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 Journal on Optimization*, 30(4): 2927–2955.
- 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, C.; Netrapalli, P.; Ge, R.; Kakade, S. M.; and Jordan, M. I. 2019a. A short note on concentration inequalities for random vectors with subGaussian norm. ArXiv:1902.03736.
- Jin, M.; Lavaei, J.; Sojoudi, S.; and Baldick, R. 2021. Boundary Defense Against Cyber Threat for Power System State Estimation. *IEEE Transactions on Information Forensics and Security*, 16: 1752–1767.
- Jin, M.; Molybog, I.; Mohammadi-Ghazi, R.; and Lavaei, J. 2019b. Towards robust and scalable power system state estimation. In *2019 IEEE 58th Conference on Decision and Control (CDC)*, 3245–3252. IEEE.
- Li, X.; Zhu, Z.; Man-Cho So, A.; and Vidal, R. 2020. Non-convex robust low-rank matrix recovery. *SIAM Journal on Optimization*, 30(1): 660–686.
- Park, D.; Kyrillidis, A.; Caramanis, C.; and Sanghavi, S. 2018. Finding low-rank solutions via nonconvex matrix factorization, efficiently and provably. *SIAM Journal on Imaging Sciences*, 11(4): 2165–2204.
- Park, D.; Kyrillidis, A.; Carmanis, C.; and Sanghavi, S. 2017. Non-square Matrix Sensing Without Spurious Local Minima via the Burer–Monteiro Approach. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, volume 54 of *Proceedings of Machine Learning Research*, 65–74.
- Recht, B.; Fazel, M.; and Parrilo, P. A. 2010. Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization. *SIAM Review*, 52(3): 471–501.
- Shechtman, Y.; Eldar, Y. C.; Cohen, O.; Chapman, H. N.; Miao, J.; and Segev, M. 2015. Phase retrieval with application to optical imaging: A contemporary overview. *IEEE Signal Processing Magazine*, 32(3): 87–109.
- Singer, A. 2011. Angular synchronization by eigenvectors and semidefinite programming. *Applied and Computational Harmonic Analysis*, 30(1): 20–36.
- 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.
- Wang, L.; Zhang, X.; and Gu, Q. 2017. A unified computational and statistical framework for nonconvex low-rank matrix estimation. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, volume 54 of *Proceedings of Machine Learning Research*, 981–990.
- 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, 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. *Journal of Machine Learning Research*, 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.

Zhu, Z.; Li, Q.; Tang, G.; and Wakin, M. B. 2018. Global Optimality in Low-Rank Matrix Optimization. *IEEE Transactions on Signal Processing*, 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 Transactions on Information Theory*, 67(2): 1308–1331.