

Optimal tuning-free convex relaxation for noisy matrix completion

Yuepeng Yang* Cong Ma*

July 12, 2022

Abstract

This paper is concerned with noisy matrix completion—the problem of recovering a low-rank matrix from partial and noisy entries. Under uniform sampling and incoherence assumptions, we prove that a tuning-free square-root matrix completion estimator (**square-root MC**) achieves optimal statistical performance for solving the noisy matrix completion problem. Similar to the square-root Lasso estimator in high-dimensional linear regression, **square-root MC** does not rely on the knowledge of the size of the noise. While solving **square-root MC** is a convex program, our statistical analysis of **square-root MC** hinges on its intimate connections to a nonconvex rank-constrained estimator.

1 Introduction

Low-rank matrix completion [CR09, KMO10] aims to reconstruct a low-rank data matrix from its partially observed entries. This problem finds numerous applications in collaborative filtering [RS05], causal inference [ABD⁺21], sensor network localization [BLWY06], etc.

In this paper, we focus on the *noisy* matrix completion problem, in which the revealed entries are further corrupted by random noise. Mathematically, let $\mathbf{L}^* \in \mathbb{R}^{n \times n}$ be a rank- r matrix of interest, and $\mathbf{E} \in \mathbb{R}^{n \times n}$ denote the noise matrix. We observe a subset of entries

$$M_{ij} = L_{ij}^* + E_{ij}, \quad \text{for } (i, j) \in \Omega, \quad (1)$$

where $\Omega \subseteq \{1, 2, \dots, n\} \times \{1, 2, \dots, n\}$ represents the index set of the observations. The goal of noisy matrix completion is to recover the underlying low-rank matrix \mathbf{L}^* given the observation $\mathbf{M} = [M_{ij}]$.

Arguably, one of the most natural approaches to solving noisy matrix completion is the following nuclear norm regularized least-squares estimator [CP10, CCF⁺20]:

$$\min_{\mathbf{L} \in \mathbb{R}^{n \times n}} \sum_{(i,j) \in \Omega} (L_{ij} - M_{ij})^2 + \lambda \|\mathbf{L}\|_*, \quad (2)$$

where $\|\mathbf{L}\|_*$ denotes the nuclear norm (i.e., sum of singular values) of the matrix \mathbf{L} , and $\lambda > 0$ is a tuning parameter. Here, the least-squares loss $\sum_{(i,j) \in \Omega} (L_{ij} - M_{ij})^2$ measures the fidelity of the estimate \mathbf{L} to the observation \mathbf{M} , while the nuclear norm penalty $\lambda \|\mathbf{L}\|_*$ encourages the low-rank property of the solution. In a recent work [CCF⁺20], it has been shown that with properly chosen regularization parameter λ , the nuclear norm regularized least-squares estimator (2) achieves optimal statistical performance in terms of estimating the low-rank matrix \mathbf{L}^* . However, this optimal choice depends on the noise size, which is often *unknown* in practice. This begs the question:

Can we develop an estimator for noisy matrix completion that does not rely on the unknown noise size (a.k.a., tuning-free), and at the same time achieves optimal statistical performance?

Motivated by the success of the square-root Lasso estimator [BCW11] for sparse recovery problems, we consider in this paper the following square-root matrix completion estimator (dubbed **square-root MC**):

*Department of Statistics, University of Chicago; Email: {yuepengyang, congma}@uchicago.edu

$$\min_{\mathbf{L} \in \mathbb{R}^{n \times n}} \sqrt{\sum_{(i,j) \in \Omega} (L_{ij} - M_{ij})^2 + \lambda \|\mathbf{L}\|_*}. \quad (3)$$

A notable difference from the vanilla least-squares estimator (2) is that **square-root MC** (3) aims at minimizing the regularized ℓ_2 error instead of the regularized *squared* ℓ_2 error.

Our contributions. The main result of this paper (cf. Theorem 1) shows that **square-root MC** (3) with a noise-size-oblivious choice $\lambda \asymp 1/\sqrt{n}$ (e.g., $\lambda = 32/\sqrt{n}$) achieves the optimal error guarantees for recovering the low-rank matrix \mathbf{L}^* over a wide range of noise sizes. Such guarantees are on par with those established for the vanilla least-squares estimator (2) with a choice of λ depending the noise size [CCF⁺20]. Clearly, the tuning-free property and statistical optimality of **square-root MC** together answer our motivating question in the affirmative.

To put our contributions into context, we would like to immediately point out two relevant pieces of prior work, while deferring other related ones to Section 4. First and foremost, a variant of the **square-root MC** estimator has been proposed and studied by Klopp [Klo14], in which an extra element-wise max norm constraint is added to the problem (3). In the same paper, it was shown that **square-root MC** achieves optimal statistical performance when the size of the noise is sufficiently large compared to the entries of the low-rank matrix. However, when the noise size is relatively small, the upper bound proved therein fails to uncover the optimal performance of the **square-root MC** estimator. In particular, it falls short of uncovering the exact recovery property when there is no noise, i.e., when $\mathbf{E} = \mathbf{0}$. More recently, Zhang et al. [ZYW21] focuses on a closely related noisy robust PCA problem [CLMW11, CFMY21] and studies a similar tuning-free estimator. Their results, however, even in the full observation setting (i.e., $\Omega = \{1, 2, \dots, n\} \times \{1, 2, \dots, n\}$), has a poor dependence on the problem dimension n , which is far from optimality. Detailed comparisons between our results and those in the papers [Klo14, ZYW21] can be found in Section 2.

Notation. For a vector \mathbf{v} , we use $\|\mathbf{v}\|_2$ to denote its Euclidean norm. For a matrix \mathbf{M} , we use $\|\mathbf{M}\|, \|\mathbf{M}\|_{\text{F}},$ and $\|\mathbf{M}\|_{\infty}$ to denote its spectral norm, Frobenius norm, and the elementwise ℓ_{∞} norm. In addition, $\|\mathbf{M}\|_{2,\infty}$ denotes the largest ℓ_2 norm of the rows. We also use $\sigma_j(\mathbf{M})$ to denote the j -th largest singular value of \mathbf{M} .

Additionally, the standard notation $f(n) = O(g(n))$ or $f(n) \lesssim g(n)$ means that there exists a constant $c > 0$ such that $|f(n)| \leq c|g(n)|$, $f(n) \gtrsim g(n)$ means that there exists a constant $c > 0$ such that $|f(n)| \geq c|g(n)|$. Also, $f(n) \gg g(n)$ means that there exists some large enough constant $c > 0$ such that $|f(n)| \geq c|g(n)|$. Similarly, $f(n) \ll g(n)$ means that there exists some sufficiently small constant $c > 0$ such that $|f(n)| \leq c|g(n)|$.

2 Main results

We start with introducing the model assumptions for noisy matrix completion. The first assumption is on the observation pattern.

Assumption 1. *Each index (i, j) belongs to the set Ω independently with probability p .*

The next assumption is concerned with the noise matrix.

Assumption 2. *The noise matrix $\mathbf{E} = [E_{ij}]$ is composed of i.i.d. zero-mean sub-Gaussian random variables with variance σ^2 and sub-Gaussian norm $O(\sigma)$, i.e., $\|E_{i,j}\|_{\psi_2} = O(\sigma)$; see Definition 5.7 in the article [Ver10].*

In the end, we turn to the assumptions on the groundtruth matrix \mathbf{L}^* . Let $\sigma_{\min}, \sigma_{\max}$ be the smallest and largest singular values of \mathbf{L}^* , respectively, and let $\kappa := \sigma_{\max}/\sigma_{\min}$ be its condition number. We require the matrix \mathbf{L}^* to be μ -incoherent defined in the following way.

Assumption 3. *The rank- r matrix \mathbf{L}^* with SVD $\mathbf{L}^* = \mathbf{U}^* \mathbf{\Sigma}^* \mathbf{V}^{*\top}$ is μ -incoherent in the sense that*

$$\|\mathbf{U}^*\|_{2,\infty} \leq \sqrt{\frac{\mu}{n}} \|\mathbf{U}^*\|_{\text{F}} = \sqrt{\frac{\mu r}{n}}, \quad \text{and} \quad \|\mathbf{V}^*\|_{2,\infty} \leq \sqrt{\frac{\mu}{n}} \|\mathbf{V}^*\|_{\text{F}} = \sqrt{\frac{\mu r}{n}}.$$

Now we are in position to state our main results regarding the **square-root MC** estimator, with the proof deferred to Section 3.

Theorem 1. *Suppose that Assumptions 1-3 hold. In addition, assume that the sample size and the noise level satisfy*

$$n^2 p \geq C_{\text{sample}} \kappa^4 \mu^2 r^2 n \log^3 n, \quad \text{and} \quad \frac{\sigma}{\sigma_{\min}} \sqrt{\frac{n}{p}} \leq \frac{C_{\text{noise}}}{\sqrt{\kappa^4 \mu r \log n}}$$

for some sufficient large (resp. small) constant $C_{\text{sample}} > 0$ (resp. $C_{\text{noise}} > 0$). Set $\lambda = C_\lambda / \sqrt{n}$ for the **square-root MC** estimator (3), where C_λ is some large constant (e.g., 32). With probability at least $1 - O(n^{-3})$, any solution \mathbf{L}_{cvx} to the **square-root MC** problem (3) obeys

$$\|\mathbf{L}_{\text{cvx}} - \mathbf{L}^*\|_{\text{F}} \leq C_{\text{F}} \kappa \frac{\sigma}{\sigma_{\min}} \sqrt{\frac{n}{p}} \|\mathbf{L}^*\|_{\text{F}}; \quad (4\text{a})$$

$$\|\mathbf{L}_{\text{cvx}} - \mathbf{L}^*\|_{\infty} \leq C_{\infty} \sqrt{\kappa^3 \mu r} \frac{\sigma}{\sigma_{\min}} \sqrt{\frac{n \log n}{p}} \|\mathbf{L}^*\|_{\infty}; \quad (4\text{b})$$

$$\|\mathbf{L}_{\text{cvx}} - \mathbf{L}^*\| \leq C_{\text{op}} \frac{\sigma}{\sigma_{\min}} \sqrt{\frac{n}{p}} \|\mathbf{L}^*\|. \quad (4\text{c})$$

Here $C_{\text{F}}, C_{\infty}, C_{\text{op}} > 0$ are three universal constants.

Several remarks on Theorem 1 are in order.

Minimax-optimal ℓ_{F} estimation error. When the condition number κ is of a constant order, the **square-root MC** estimator enjoys minimax-optimal ℓ_{F} estimation error [NW12, CCF+20]. In contrast, the upper bound in the paper [Klo14] reads $\|\mathbf{L}_{\text{cvx}} - \mathbf{L}^*\|_{\text{F}} \lesssim \max\{\sigma, \|\mathbf{L}^*\|_{\infty}\} \sqrt{n \log n / p}$, which is only statistically optimal when $\sigma \gtrsim \|\mathbf{L}^*\|_{\infty}$. In addition, translating the bound in the paper [ZYW21] from robust PCA to the matrix completion setting, one obtains $\|\mathbf{L}_{\text{cvx}} - \mathbf{L}^*\|_{\text{F}} \lesssim \sigma n^2$, which has a much worse (and hence sub-optimal) dependence on the problem dimension n .

Tuning-free property. More importantly, the optimal performance of **square-root MC** is achieved in a completely tuning-free fashion. The regularization parameter λ can be set to be $32/\sqrt{n}$, that does not depend on the noise variance σ^2 , the observation probability p , nor the true rank r of the matrix \mathbf{L}^* . This is in stark contrast to the vanilla nuclear norm regularized least-squares estimator (2) in which λ is set to be on the order of $\sigma \sqrt{np}$ (cf. [CCF+20]).

Entrywise error guarantees. Also, our main results provide upper bounds on the entrywise estimation error (cf. bound (4b)). Compared to the ℓ_{F} estimation error (4a), it can be seen that the **square-root MC** estimator is uniformly good in the sense that there is no spiky entry estimate with large estimation error.

To corroborate our main results, we perform numerical experiments on noisy matrix completion with simulated data. We fix the rank r to be 5 throughout the experiment. For each problem dimension n , we generate two $n \times r$ random orthonormal matrices as \mathbf{X}^* and \mathbf{Y}^* and take $\mathbf{L}^* := \mathbf{X}^* \mathbf{Y}^{*\top}$ as the rank- r $n \times n$ groundtruth matrix. The entrywise noise is taken to be Gaussian with variance σ^2 . For all the experiments, we set $\lambda = 4/\sqrt{n}$ in **square-root MC**, and report the average results over 20 Monte-Carlo simulations. Figure 1 reports the relative error of the **square-root MC** estimator in Frobenius, spectral, and infinity norms. More specifically, Figure 1(a) fixes $n = 500$, $p = 0.5$, and varies σ ; Figure 1(b) fixes $\sigma = 10^{-4}$, $p = 0.5$, and varies n ; Figure 1(c) fixes $\sigma = 10^{-4}$, $n = 2000$, and varies p . Overall, the plots showcase a linear relationship between the performance and the noise size σ , the problem dimension \sqrt{n} , and the observation probability p . This is consistent with the $O(\sigma \sqrt{n/p})$ scaling proved in Theorem 1.




Figure 1: (a) Relative estimation error of \mathbf{L}_{cvx} vs. noise size σ on a log-log scale, where we fix $n = 500, r = 5, p = 0.5$; (b) Relative estimation error of \mathbf{L}_{cvx} vs. problem size \sqrt{n} , where we fix $r = 5, \sigma = 10^{-4}, p = 0.5$; (c) Relative estimation error of \mathbf{L}_{cvx} vs. observation probability p on a log-log scale, where we fix $n = 2000, r = 5, \sigma = 10^{-4}$. For all three plots, $\lambda = 4/\sqrt{n}$ and each point represents the average of 20 independent trials.

3 Outline of the proof

In this section, we provide the key steps for proving our main result, i.e., Theorem 1. The proof follows the general strategy of bridging convex and nonconvex solutions, first appeared in the paper [CCF⁺20], with several important modifications to handle the non-smooth ℓ_F norm (as opposed to the smooth squared ℓ_F norm).

A central object in our analysis is the following nonconvex optimization problem

$$\min_{\mathbf{X}, \mathbf{Y} \in \mathbb{R}^{n \times r}, \theta > 0} f(\mathbf{X}, \mathbf{Y}, \theta) := \frac{1}{2} \left(\frac{\|\mathcal{P}_\Omega(\mathbf{X}\mathbf{Y}^\top - \mathbf{M})\|_F^2}{\theta} + \theta \right) + \frac{\lambda}{2} (\|\mathbf{X}\|_F^2 + \|\mathbf{Y}\|_F^2), \quad (5)$$

which is closely related to the original convex square-root MC formulation (3). To see this, first, for any rank- r matrix \mathbf{Z} , one has

$$\|\mathbf{Z}\|_* = \inf_{\mathbf{X}, \mathbf{Y} \in \mathbb{R}^{n \times r}: \mathbf{X}\mathbf{Y}^\top = \mathbf{Z}} \frac{1}{2} (\|\mathbf{X}\|_F^2 + \|\mathbf{Y}\|_F^2).$$

Second and more importantly, we have that for any matrix $\mathbf{Z} = \mathbf{X}\mathbf{Y}^\top$,

$$\|\mathcal{P}_\Omega(\mathbf{Z} - \mathbf{M})\|_F = \inf_{\theta > 0} \frac{1}{2} \left(\frac{\|\mathcal{P}_\Omega(\mathbf{X}\mathbf{Y}^\top - \mathbf{M})\|_F^2}{\theta} + \theta \right).$$

It turns out that the (approximate) solution to the nonconvex optimization problem (5) serves as an extremely tight approximation to the square-root MC estimator, which facilitates the statistical analysis of the latter.

In sum, our proof involves two main steps:

1. We first show—via an explicit construction—that an approximate stationary point \mathbf{L}_{ncvx} of the nonconvex problem (5) exists and is also close to the groundtruth matrix \mathbf{L}^* .
2. We then establish that such an approximate stationary point \mathbf{L}_{ncvx} is extremely close to the solution \mathbf{L}_{cvx} to the convex problem (3).

Combining the two key steps via triangle inequality finishes the proof.

Step 1: Nonconvex optimization. The nonconvex optimization problem (5) has two groups of decision variables, i.e., (\mathbf{X}, \mathbf{Y}) and θ . Also note that given a fixed pair (\mathbf{X}, \mathbf{Y}) , the optimal choice of θ is simply given by $\theta = \|\mathcal{P}_\Omega(\mathbf{X}\mathbf{Y}^\top - \mathbf{M})\|_F$. Therefore it is natural to consider an alternating minimization method to construct an approximate stationary point of the nonconvex program (5); see Algorithm 1. Given a current iterate $(\mathbf{X}_t, \mathbf{Y}_t, \theta_t)$, the algorithm first runs one step of gradient descent on (\mathbf{X}, \mathbf{Y}) while fixing θ_t . It then

Algorithm 1 Gradient descent on the nonconvex formulation of square root matrix completion

Input: initialization $\mathbf{X}_0 = \mathbf{X}^*$, $\mathbf{Y}_0 = \mathbf{Y}^*$, $\theta_0 = \|\mathcal{P}_\Omega(\mathbf{X}^*\mathbf{Y}^{*\top} - \mathbf{M})\|_F$, step size $\eta \asymp \sigma/(\sqrt{p}\kappa^3\sigma_{\max})$, and total number of iterations $t_0 = n^{18}$.

Gradient updates: for $t = 0, 1, \dots, t_0 - 1$ do

$$\mathbf{X}_{t+1} = \mathbf{X}_t - \eta \nabla_{\mathbf{X}} f(\mathbf{X}_t, \mathbf{Y}_t, \theta_t) = \mathbf{X}_t - \eta \left(\frac{1}{\theta_t} \mathcal{P}_\Omega(\mathbf{X}_t \mathbf{Y}_t^\top - \mathbf{M}) \mathbf{Y}_t + \lambda \mathbf{X}_t \right); \quad (6a)$$

$$\mathbf{Y}_{t+1} = \mathbf{Y}_t - \eta \nabla_{\mathbf{Y}} f(\mathbf{X}_t, \mathbf{Y}_t, \theta_t) = \mathbf{Y}_t - \eta \left(\frac{1}{\theta_t} [\mathcal{P}_\Omega(\mathbf{X}_t \mathbf{Y}_t^\top - \mathbf{M})]^\top \mathbf{X}_t + \lambda \mathbf{Y}_t \right); \quad (6b)$$

$$\theta_{t+1} = \|\mathcal{P}_\Omega(\mathbf{X}_{t+1} \mathbf{Y}_{t+1}^\top - \mathbf{M})\|_F. \quad (6c)$$

Define

$$t^* := \arg \min_{0 \leq t \leq t_0} \|\nabla_{\mathbf{X}, \mathbf{Y}} f(\mathbf{X}_t, \mathbf{Y}_t, \theta_t)\|_F,$$

where

$$\nabla_{\mathbf{X}, \mathbf{Y}} f(\mathbf{X}_t, \mathbf{Y}_t, \theta_t) = \begin{bmatrix} \frac{1}{\theta_t} \mathcal{P}_\Omega(\mathbf{X}_t \mathbf{Y}_t^\top - \mathbf{M}) \mathbf{Y}_t + \lambda \mathbf{X}_t \\ \frac{1}{\theta_t} [\mathcal{P}_\Omega(\mathbf{X}_t \mathbf{Y}_t^\top - \mathbf{M})]^\top \mathbf{X}_t + \lambda \mathbf{Y}_t \end{bmatrix}.$$

Output: $\mathbf{L}_{\text{ncvx}} := \mathbf{X}_{t^*} \mathbf{Y}_{t^*}^\top$, $\mathbf{X}_{\text{ncvx}} := \mathbf{X}_{t^*}$, and $\mathbf{Y}_{\text{ncvx}} := \mathbf{Y}_{t^*}$.

updates $\theta_{t+1} = \|\mathcal{P}_\Omega(\mathbf{X}_{t+1} \mathbf{Y}_{t+1}^\top - \mathbf{M})\|_F$ to be the optimal choice given the new iterate $(\mathbf{X}_{t+1}, \mathbf{Y}_{t+1})$. In the end, Algorithm 1 returns the point \mathbf{L}_{ncvx} with the smallest gradient among the iterates as an approximate stationary point.

The following lemma ensures that \mathbf{L}_{ncvx} is an approximate stationary point of the nonconvex problem and more importantly is close to the groundtruth matrix \mathbf{L}^* . The proof is deferred to Section 3.1.

Lemma 1. *Instate the assumptions of Theorem 1. With probability at least $1 - O(n^{-3})$, one has*

$$\|\mathbf{L}_{\text{ncvx}} - \mathbf{L}^*\|_F \leq 3\kappa C_F \left(\frac{\sigma}{\sigma_{\min}} \sqrt{\frac{n}{p}} \right) \|\mathbf{L}^*\|_F, \quad (7a)$$

$$\|\mathbf{L}_{\text{ncvx}} - \mathbf{L}^*\|_\infty \leq 3\sqrt{\kappa^3 \mu r} C_\infty \left(\frac{\sigma}{\sigma_{\min}} \sqrt{\frac{n \log n}{p}} \right) \|\mathbf{L}^*\|_\infty, \quad (7b)$$

$$\|\mathbf{L}_{\text{ncvx}} - \mathbf{L}^*\| \leq 3C_{\text{op}} \left(\frac{\sigma}{\sigma_{\min}} \sqrt{\frac{n}{p}} \right) \|\mathbf{L}^*\|, \quad (7c)$$

where $C_F, C_\infty, C_{\text{op}}$ are three universal positive constants.

Step 2: Bridging convex and nonconvex solutions. It remains to show that \mathbf{L}_{ncvx} is extremely close to the convex solution \mathbf{L}_{cvx} , which is provided in the following lemma.

Lemma 2. *Instate the assumptions of Theorem 1. With probability exceeding $1 - O(n^{-3})$, one has*

$$\|\mathbf{L}_{\text{ncvx}} - \mathbf{L}_{\text{cvx}}\|_F \leq \frac{1}{n^5} \frac{\lambda \sigma}{\sigma_{\min}} \|\mathbf{L}^*\|_F.$$

See Section 3.2 for the proof of this lemma.

We remark in passing that the polynomial factor n^{-5} in Lemma 2 is arbitrarily chosen, and the exponent 5 can be replaced with any large constant. The essence is that the difference between \mathbf{L}_{ncvx} and \mathbf{L}_{cvx} is orderwise much smaller compared to the estimation error of \mathbf{L}_{ncvx} itself. Such proximity between \mathbf{L}_{ncvx} and \mathbf{L}_{cvx} is verified empirically in Figure 2.

Now we are ready to combine the previous two steps and finish the proof of Theorem 1.




Figure 2: Relative Frobenius estimation error of convex and nonconvex solutions and their distance. The parameters are chosen as: $n = 200, r = 5, p = 0.5$ while σ varies from 10^{-5} to 10^{-3} .

Proof of Theorem 1. Combine Lemmas 1-2 with the triangle inequality to arrive at

$$\begin{aligned}
\|\mathbf{L}_{\text{cvx}} - \mathbf{L}^*\|_{\text{F}} &\leq \|\mathbf{L}_{\text{ncvx}} - \mathbf{L}_{\text{cvx}}\|_{\text{F}} + \|\mathbf{L}_{\text{ncvx}} - \mathbf{L}^*\|_{\text{F}} \\
&\leq \left[\frac{1}{n^5} \frac{\lambda\sigma}{\sigma_{\min}} + 3\kappa C_{\text{F}} \left(\frac{\sigma}{\sigma_{\min}} \sqrt{\frac{n}{p}} \right) \right] \|\mathbf{L}^*\|_{\text{F}} \\
&\leq 4\kappa C_{\text{F}} \left(\frac{\sigma}{\sigma_{\min}} \sqrt{\frac{n}{p}} \right) \|\mathbf{L}^*\|_{\text{F}},
\end{aligned}$$

where the last relation uses the facts that $\lambda \asymp 1/\sqrt{n}$ and that $p \gtrsim 1/\sqrt{n}$. Redefine $4C_{\text{F}}$ to be C_{F} to complete the proof of the bound (4a). The other two bounds on the operator norm and the ℓ_{∞} norm follow from similar arguments. We omit here for brevity.

3.1 Proof of Lemma 1

Since Algorithm 1 operates in the space of low-rank factors, we start with establishing guarantees for the stacked low-rank factor $\mathbf{F}_t := \begin{bmatrix} \mathbf{X}_t \\ \mathbf{Y}_t \end{bmatrix} \in \mathbb{R}^{2n \times r}$, and then translate the guarantees to the matrix space $\mathbf{L}_t = \mathbf{X}_t \mathbf{Y}_t^{\top}$. Special care is needed as the decomposition $\mathbf{L} = \mathbf{X} \mathbf{Y}^{\top}$ is not unique in (\mathbf{X}, \mathbf{Y}) , and hence we need to account for the rotational ambiguity in (\mathbf{X}, \mathbf{Y}) . To this end, for each $t \geq 0$, we define the optimal rotation matrix to be

$$\mathbf{H}_t := \operatorname{argmin}_{\mathbf{R} \in \mathcal{O}^{r \times r}} \|\mathbf{X}_t \mathbf{R} - \mathbf{X}^*\|_{\text{F}}^2 + \|\mathbf{Y}_t \mathbf{R} - \mathbf{Y}^*\|_{\text{F}}^2. \quad (8)$$

Introducing leave-one-out sequences. In order to control the $\ell_{2,\infty}$ error of \mathbf{F}_t (and hence ℓ_{∞} error of \mathbf{L}_t), we construct $2n$ leave-one-out auxiliary sequences $\{\mathbf{F}_t^{(l)}\}_{1 \leq l \leq 2n, t \geq 0}$. The hope is that $\{\mathbf{F}_t^{(l)}\}_{1 \leq l \leq 2n, t \geq 0}$ serves as a good approximation to the original sequence $\{\mathbf{F}_t\}_{t \geq 0}$, while at the same time is more amenable to statistical analysis.

To formally construct such leave-one-out sequences, we first define $2n$ auxiliary loss functions. For each $1 \leq l \leq n$, define

$$f^{(l)}(\mathbf{X}, \mathbf{Y}, \theta) = \frac{1}{2} \left(\frac{\|\mathcal{P}_{\Omega_{-l}}(\mathbf{L} - \mathbf{M})\|_{\text{F}}^2 + p \|\mathcal{P}_{l,\cdot}(\mathbf{L} - \mathbf{M})\|_{\text{F}}^2}{\theta} + \theta \right) + \frac{\lambda}{2} (\|\mathbf{X}\|_{\text{F}}^2 + \|\mathbf{Y}\|_{\text{F}}^2)$$

where

$$[\mathcal{P}_{\Omega_{-l}}(\mathbf{B})]_{ij} = \begin{cases} B_{ij}, & \text{if } (i, j) \in \Omega \text{ and } i \neq l \\ 0, & \text{otherwise} \end{cases}, \quad \text{and} \quad [\mathcal{P}_{l,\cdot}(\mathbf{B})]_{ij} = \begin{cases} B_{ij}, & \text{if } i = l \\ 0, & \text{otherwise} \end{cases}.$$

Algorithm 2 Gradient descent generating the leave-one-out sequences

Initialization: $\mathbf{X}_0^{(l)} = \mathbf{X}^*$, $\mathbf{Y}_0^{(l)} = \mathbf{Y}^*$, $\theta_0^{(l)} = \|\mathcal{P}_\Omega(\mathbf{X}^* \mathbf{Y}^{*\top} - \mathbf{M})\|_F$, step size $\eta \asymp \sigma / (\sqrt{p} \kappa^3 \sigma_{\max})$, and total number of iterations $t_0 = n^{18}$.

Gradient updates: for $t = 0, 1, \dots, t_0 - 1$ do

$$\mathbf{X}_{t+1}^{(l)} = \mathbf{X}_t^{(l)} - \eta \nabla_{\mathbf{X}} f^{(l)}(\mathbf{X}_t^{(l)}, \mathbf{Y}_t^{(l)}, \theta_t); \quad (9a)$$

$$\mathbf{Y}_{t+1}^{(l)} = \mathbf{Y}_t^{(l)} - \eta \nabla_{\mathbf{Y}} f^{(l)}(\mathbf{X}_t^{(l)}, \mathbf{Y}_t^{(l)}, \theta_t); \quad (9b)$$

$$\theta_{t+1} = \|\mathcal{P}_\Omega(\mathbf{X}_{t+1} \mathbf{Y}_{t+1}^\top - \mathbf{M})\|_F. \quad (9c)$$

Similarly, for each $n + 1 \leq l \leq 2n$, we define

$$f^{(l)}(\mathbf{X}, \mathbf{Y}, \theta) = \frac{1}{2} \left(\frac{\|\mathcal{P}_{\Omega, -(l-n)}(\mathbf{L} - \mathbf{M})\|_F^2 + p \|\mathcal{P}_{l, \cdot}(\mathbf{L} - \mathbf{M})\|_F^2}{\theta} + \theta \right) + \frac{\lambda}{2} (\|\mathbf{X}\|_F^2 + \|\mathbf{Y}\|_F^2)$$

where

$$[\mathcal{P}_{\Omega, -(l-n)}(\mathbf{B})]_{ij} = \begin{cases} B_{ij}, & \text{if } (i, j) \in \Omega \text{ and } j \neq l - n \\ 0, & \text{otherwise} \end{cases} \quad \text{and} \quad [\mathcal{P}_{l, \cdot}(\mathbf{B})]_{ij} = \begin{cases} B_{ij}, & \text{if } j = l - n \\ 0, & \text{otherwise} \end{cases}.$$

With these notations in place, Algorithm 2 details the way we construct the leave-one-out sequences.

Similar constructions have been deployed in the papers [CCF+20] and [CFMY21]. However, it is worth pointing out that the sequence $\{\theta_t\}$ is produced according to the original sequence, instead of the leave-one-out sequence. This change is tailored to the analysis of the **square-root MC** estimator as it aligns better with the original loss function f , while allowing us to reuse several keys results in the paper [CCF+20].

Properties of the iterates. As planned, we aim to show that the leave-one-out iterates $\{\mathbf{F}_t^{(l)}\}_{1 \leq l \leq 2n, t \geq 0}$ stay extremely close to the original iterates $\{\mathbf{F}_t\}_{t \geq 0}$, and that $\{\mathbf{F}_t\}_{t \geq 0}$ is close to the groundtruth factor \mathbf{F}^* . Such properties are collected in the following lemma.

Lemma 3. *With probability at least $1 - O(n^{-3})$, the following statements hold for all iterations $0 \leq t \leq t_0$:*

$$\|\mathbf{F}_t \mathbf{H}_t - \mathbf{F}^*\|_F \leq C_F \frac{\sigma}{\sigma_{\min}} \sqrt{\frac{n}{p}} \|\mathbf{X}^*\|_F, \quad (10a)$$

$$\|\mathbf{F}_t \mathbf{H}_t - \mathbf{F}^*\| \leq C_{\text{op}} \frac{\sigma}{\sigma_{\min}} \sqrt{\frac{n}{p}} \|\mathbf{X}^*\|, \quad (10b)$$

$$\max_{1 \leq l \leq 2n} \|\mathbf{F}_t \mathbf{H}_t - \mathbf{F}_t^{(l)} \mathbf{R}_t^{(l)}\|_F \leq C_3 \frac{\sigma}{\sigma_{\min}} \sqrt{\frac{n \log n}{p}} \|\mathbf{F}^*\|_{2, \infty}, \quad (10c)$$

$$\max_{1 \leq l \leq 2n} \|(\mathbf{F}_t^{(l)} \mathbf{H}_t^{(l)} - \mathbf{F}^*)_{l, \cdot}\|_F \leq C_4 \kappa \frac{\sigma}{\sigma_{\min}} \sqrt{\frac{n \log n}{p}} \|\mathbf{F}^*\|_{2, \infty}, \quad (10d)$$

$$\|\mathbf{F}_t \mathbf{H}_t - \mathbf{F}^*\|_{2, \infty} \leq C_\infty \kappa \frac{\sigma}{\sigma_{\min}} \sqrt{\frac{n \log n}{p}} \|\mathbf{F}^*\|_{2, \infty}, \quad (10e)$$

for some positive constants $C_F, C_{\text{op}}, C_3, C_4, C_\infty$. Here $\mathbf{H}_t^{(l)}$ and $\mathbf{R}_t^{(l)}$ are defined as

$$\begin{aligned} \mathbf{H}_t^{(l)} &:= \operatorname{argmin}_{\mathbf{R} \in \mathcal{O}^{r \times r}} \|\mathbf{F}_t^{(l)} \mathbf{R} - \mathbf{F}^*\|_F \\ \mathbf{R}_t^{(l)} &:= \operatorname{argmin}_{\mathbf{R} \in \mathcal{O}^{r \times r}} \|\mathbf{F}_t^{(l)} \mathbf{R} - \mathbf{F}_t \mathbf{H}_t\|_F. \end{aligned}$$

Furthermore the output $(\mathbf{X}_{t^*}, \mathbf{Y}_{t^*})$ has small gradient:

$$\|\nabla_{\mathbf{X}, \mathbf{Y}} f(\mathbf{X}_{t^*}, \mathbf{Y}_{t^*}, \theta_{t^*})\|_{\text{F}} \leq C_{\text{grad}} \frac{1}{n^8} \sqrt{\frac{\sigma_{\max}}{p}}. \quad (11)$$

See Section A for the proof of this lemma.

Now we are ready to prove Lemma 1 based on the results presented in Lemma 3.

Proof of Lemma 1. By the triangle inequality, one has

$$\begin{aligned} \|\mathbf{X}_{t^*} \mathbf{Y}_{t^*}^\top - \mathbf{L}^*\| &\leq \|\mathbf{X}_{t^*} \mathbf{Y}_{t^*}^\top - \mathbf{X}_{t^*} \mathbf{Y}^{*\top}\| + \|\mathbf{X}_{t^*} \mathbf{Y}^{*\top} - \mathbf{L}^*\| \\ &\leq \|\mathbf{Y}_{t^*} - \mathbf{Y}^*\| \|\mathbf{X}_{t^*}\| + \|\mathbf{X}_{t^*} - \mathbf{X}^*\| \|\mathbf{Y}^*\|. \end{aligned}$$

Use relation (10b) to obtain

$$\|\mathbf{X}_{t^*} \mathbf{Y}_{t^*}^\top - \mathbf{L}^*\| \leq 3C_{\text{op}} \frac{\sigma}{\sigma_{\min}} \sqrt{\frac{n}{p}} \|\mathbf{X}^*\| \|\mathbf{X}^*\| = 3C_{\text{op}} \frac{\sigma}{\sigma_{\min}} \sqrt{\frac{n}{p}} \|\mathbf{L}^*\|.$$

The first inequality uses $\|\mathbf{X}_{t^*}\| \leq 2\|\mathbf{X}^*\|$, which is a direct consequence of (10b) and the last line uses $\|\mathbf{L}^*\| = \sigma_{\max} = \|\mathbf{X}^*\|^2$. Similarly we have

$$\|\mathbf{X}_{t^*} \mathbf{Y}_{t^*}^\top - \mathbf{L}^*\|_{\text{F}} \leq 3C_{\text{F}} \frac{\sigma}{\sigma_{\min}} \sqrt{\frac{n}{p}} \|\mathbf{X}^*\|_{\text{F}} \|\mathbf{X}^*\| \stackrel{(i)}{\leq} 3\kappa C_{\text{op}} \frac{\sigma}{\sigma_{\min}} \sqrt{\frac{n}{p}} \|\mathbf{L}^*\|_{\text{F}}$$

and

$$\|\mathbf{X}_{t^*} \mathbf{Y}_{t^*}^\top - \mathbf{L}^*\|_{\infty} \leq 3C_{\infty} \frac{\sigma}{\sigma_{\min}} \sqrt{\frac{n \log n}{p}} \|\mathbf{F}^*\|_{2, \infty} \|\mathbf{F}^*\|_{2, \infty} \stackrel{(ii)}{\leq} 3\sqrt{\kappa^3 \mu r} C_{\infty} \frac{\sigma}{\sigma_{\min}} \sqrt{\frac{n}{p}} \|\mathbf{L}^*\|_{\infty}.$$

Here step (i) uses the fact $\|\mathbf{X}^*\|_{\text{F}} \|\mathbf{X}^*\| \leq \kappa \|\mathbf{L}^*\|_{\text{F}}$, whereas in step (ii) we use $\|\mathbf{F}^*\|_{2, \infty} \|\mathbf{F}^*\|_{2, \infty} \leq \sqrt{\kappa^3 \mu r} \|\mathbf{L}^*\|_{\infty}$.

3.2 Proof of Lemma 2

Before embarking on the main proof, we state a few useful properties of the noise matrix \mathbf{E} and the nonconvex solution \mathbf{L}_{ncvx} . These properties allow us to establish the proximity between the approximate stationary point \mathbf{L}_{ncvx} and the convex solution \mathbf{L}_{cvx} .

The first property is concerned with the size of the regularization parameter, which appeared as Lemma 3 in the paper [CCF⁺20].

Lemma 4. *Suppose that $n^2 p \geq C n \log^2 n$ for some sufficiently large constant $C > 0$. Take $\lambda = C_{\lambda} / \sqrt{n}$ for some constant C_{λ} . Then with probability at least $1 - O(n^{-10})$, one has*

$$\|\mathcal{P}_{\Omega}(\mathbf{E})\| \leq \frac{\lambda}{16} n p^{1/2} \sigma. \quad (12)$$

The next property is on the injectivity of \mathcal{P}_{Ω} in the tangent space T at \mathbf{L}_{ncvx} . More precisely, letting $\mathbf{U} \mathbf{\Sigma} \mathbf{V}^\top$ be the SVD of \mathbf{L}_{ncvx} , we define the tangent space T at \mathbf{L}_{ncvx} as

$$T = \{\mathbf{U} \mathbf{A}^\top + \mathbf{B} \mathbf{V}^\top \mid \mathbf{A}, \mathbf{B} \in \mathbb{R}^{n \times r}\}.$$

Lemma 5. *Instate the assumptions of Theorem 1. With probability exceeding $1 - O(n^{-3})$, for all $\mathbf{H} \in T$*

$$p^{-1/2} \|\mathcal{P}_{\Omega}(\mathbf{H})\|_{\text{F}} \geq C_{\text{inj}} \|\mathbf{H}\|_{\text{F}}, \quad \text{where } C_{\text{inj}} = (32\kappa)^{-1/2}. \quad (13)$$

Proof. This is an easy consequence of Lemma 4 in the paper [CCF⁺20] and the relation (10e). \square

Last but not least, the lemma collects several interesting properties of the nonconvex solution \mathbf{L}_{ncvx} , as well as its low-rank factors $\mathbf{X}_{\text{ncvx}}, \mathbf{Y}_{\text{ncvx}}$.

Lemma 6. *The approximate stationary point \mathbf{L}_{ncvx} satisfies*

$$\sqrt{\sigma_{\min}/2} \leq \sigma_{\min}(\mathbf{X}_{\text{ncvx}}) \leq \sigma_{\max}(\mathbf{X}_{\text{ncvx}}) \leq \sqrt{2\sigma_{\max}}; \quad (14a)$$

$$\sqrt{\sigma_{\min}/2} \leq \sigma_{\min}(\mathbf{Y}_{\text{ncvx}}) \leq \sigma_{\max}(\mathbf{Y}_{\text{ncvx}}) \leq \sqrt{2\sigma_{\max}}; \quad (14b)$$

$$\frac{1}{2}np^{1/2}\sigma \leq \|\mathcal{P}_{\Omega}(\mathbf{L}_{\text{ncvx}} - \mathbf{M})\|_{\text{F}} \leq 2np^{1/2}\sigma; \quad (14c)$$

$$\|\mathcal{P}_{\Omega}(\mathbf{X}\mathbf{Y}^{\top} - \mathbf{L}^*) - p(\mathbf{X}\mathbf{Y}^{\top} - \mathbf{L}^*)\| \leq \frac{\lambda}{16}np^{1/2}\sigma. \quad (14d)$$

See Section A.4 for the proof of this lemma.

For notational simplicity, we define

$$g(\mathbf{X}, \mathbf{Y}) := f(\mathbf{X}, \mathbf{Y}, \|\mathcal{P}_{\Omega}(\mathbf{X}\mathbf{Y}^{\top} - \mathbf{M})\|_{\text{F}}) = \|\mathcal{P}_{\Omega}(\mathbf{X}\mathbf{Y}^{\top} - \mathbf{M})\|_{\text{F}} + \frac{\lambda}{2}(\|\mathbf{X}\|_{\text{F}}^2 + \|\mathbf{Y}\|_{\text{F}}^2).$$

In other words, $g(\mathbf{X}, \mathbf{Y})$ is the minimal value of $f(\mathbf{X}, \mathbf{Y}, \theta)$ when (\mathbf{X}, \mathbf{Y}) is fixed.

Now we are ready to present the key lemma of this section, which relates the difference between \mathbf{L}_{ncvx} and \mathbf{L}_{cvx} to the size of the gradient $\nabla g(\mathbf{X}_{\text{ncvx}}, \mathbf{Y}_{\text{ncvx}})$. The proof is deferred to Section B.

Lemma 7. *Suppose that $(\mathbf{X}_{\text{ncvx}}, \mathbf{Y}_{\text{ncvx}})$ has small gradient in the sense that*

$$\|\nabla g(\mathbf{X}_{\text{ncvx}}, \mathbf{Y}_{\text{ncvx}})\|_{\text{F}} \leq \frac{\sqrt{\sigma_{\min}}}{280\kappa} \max \left\{ C_{\text{inj}}\sqrt{p}, \frac{1}{2}\lambda^2 n\sigma \right\}. \quad (15)$$

Then on the event that Lemmas 4-6 hold, any minimizer \mathbf{L}_{cvx} of the convex program (3) satisfies

$$\|\mathbf{L}_{\text{ncvx}} - \mathbf{L}_{\text{cvx}}\|_{\text{F}} \leq \frac{\lambda\kappa^2}{\sqrt{p\sigma_{\min}}}n\sigma\|\nabla g(\mathbf{X}_{\text{ncvx}}, \mathbf{Y}_{\text{ncvx}})\|_{\text{F}}.$$

Remark 1. Observe that if $\|\nabla g(\mathbf{X}_{\text{ncvx}}, \mathbf{Y}_{\text{ncvx}})\|_{\text{F}} = 0$, i.e., if \mathbf{L}_{ncvx} is an exact stationary point of the nonconvex square-root MC problem, \mathbf{L}_{ncvx} is also a solution to the convex problem (3).

With the help of Lemma 7, we can prove Lemma 2 now.

Proof of Lemma 2. First, Lemma 3 tells us that the nonconvex solution $(\mathbf{X}_{t^*}, \mathbf{Y}_{t^*})$ satisfies the bound (15) on the size of the gradient. This together with Lemmas 4 to 6 allows us to invoke Lemma 7 to obtain

$$\|\mathbf{L}_{\text{ncvx}} - \mathbf{L}_{\text{cvx}}\|_{\text{F}} \lesssim \frac{\lambda\kappa^2}{\sqrt{p\sigma_{\min}}}n\sigma\|\nabla g(\mathbf{X}_{\text{ncvx}}, \mathbf{Y}_{\text{ncvx}})\|_{\text{F}} \lesssim \frac{1}{n^5} \frac{\lambda\sigma}{\sigma_{\min}}\|\mathbf{L}^*\|_{\text{F}},$$

where the last inequality uses the gradient upper bound (15), $\|\mathbf{L}^*\|_{\text{F}} \geq \|\mathbf{L}^*\| \geq \sigma_{\max} = \kappa\sigma_{\min}$, and the fact that the sample size assumption $n^2p \geq C_{\text{sample}}\kappa^4\mu^2r^2n\log^3n$ implies $np \gtrsim 1$ and $\kappa \lesssim n$.

4 Prior art

Matrix completion. Convex relaxation has been extensively studied for the matrix completion problem both in the noiseless setting [CR09, CT10, Gro11, Rec11, Che15], and the noisy case [CP10, NW12, KLT11, Klo14, CCF+20]. In the noiseless setting, convex relaxation achieves exact recovery as soon as the number of observed entries n^2p exceeds $nr\log^2n$ —roughly the degrees of freedom of a rank- r matrix, which is information theoretically optimal. When it comes to the noisy setting, Candès and Plan [CP10] focuses on arbitrary noise (e.g., noise could be deterministic and adversarial), and proves that convex relaxation is stable w.r.t. the noise size. The theoretical guarantees for convex relaxation are strengthened by Chen et al. [CCF+20] in the stochastic noise case, which is the same setting we study in the current paper. Such a discrepancy between stochastic and deterministic noise for convex relaxation is also documented in [KS21].

Pioneered by the work [KMO10, KMO09], nonconvex optimization has gained a lot of attentions during the past decade for solving matrix completion owing to its computational efficiency. Efficient computational and statistical guarantees have been provided for manifold optimization [KMO10, KMO09], gradient descent [MWCC18, CLL20], projected gradient descent [CW15, ZL16], alternating minimization [JNS13, Har14], scaled gradient descent [TMC21], singular value projection [DC20], etc. See the recent surveys [DR16, CLC19] for more related work on matrix completion.

Tuning-free methods. A variety of tuning-free methods have been proposed to tackle high-dimensional linear regression. The seminal work [BCW11] proposes the square-root Lasso estimator which does not rely on knowing the size of the noise and is also statistically optimal. [SZ12] proposes an equivalent method named scaled sparse linear regression, which originates from the concomitant scale estimation [Hub11, Owe07]. [LM15] proposes TREX, a method similar to square-root Lasso and is completely parameter-free. [WPB⁺20] borrows ideas from non-parametric statistics and proposes Rank Lasso, whose optimal choice of tuning parameter can be simulated easily in the case with unknown variance of the noise. See [WW19] for a survey on the selection of tuning-parameters for high-dimensional regression and [GHV12] for a survey on regression with unknown variance of noise.

Bridging convex and nonconvex optimization. The connections between convex and nonconvex optimization has been extensively used in a recent line of work. Chen et al. [CCF⁺20] uses this to prove the optimality of the vanilla least-squares estimator for noisy matrix completion; Later, the papers [CFMY21, CFWY21, WF22] extend the technique to the robust PCA problem, the blind deconvolution problem, and matrix completion with heavy-tailed noise.

Leave-one-out analysis. Leave-one-out analysis is powerful in decoupling statistical dependence and obtain element-wise performance guarantees. It has been successfully applied to high-dimensional regression [EKBB⁺13, EK18], phase synchornization [ZB18], ranking [CFMW19, CGZ22], matrix completion [MWCC18, CLL20, AFWZ20], reinforcement learning [PW20], high-dimensional inference [CFMY19, YCF21] to name a few. Interested readers are referred to a recent overview [CCF⁺21] for detailed discussions.

5 Discussions

Focusing on the noisy matrix completion problem, this paper shows that a tuning-free estimator—**square-root MC** achieves optimal statistical performance. This opens up several interesting avenues for future research. Below, we list a few of them.

- *Extensions to robust PCA.* While our work focuses on matrix completion, a natural extension is to further consider partial observations with outliers, i.e., robust PCA. As mentioned, Zhang et al. [ZYW21] has studied this problem (with full observation) and provides an error guarantee of order $O(\sigma n^2)$, which is sub-optimal in its dependency on the problem dimension. By contrast, a vanilla least-squares estimator with noise-size-dependent choice of λ has been shown to be optimal [CFMY21]. It remains to be seen whether one can devise an optimal tuning-free method for robust PCA with noise and missing data.
- *Inference for **square-root MC** estimator.* The current paper discusses solely the estimation performance of the tuning-free estimator. As statistical inference for matrix completion is equally important, one wishes to develop inferential procedures around the **square-root MC** estimator as that has been done in the paper [CFMY19] for the vanilla least-squares estimator.
- *Robustness to non-uniform design.* In high-dimensional linear regression, optimal tuning-free methods have been developed to be adaptive to both the unknown noise size and the design matrix. In the matrix completion setting, the design is governed by the sampling pattern, which is assumed to be uniform in the current paper. It is of great interest to develop robust and tuning-free approaches for noisy matrix completion with non-uniform sampling that improved over the max-norm constrained estimator in [Klo14].

References

- [ABD⁺21] Susan Athey, Mohsen Bayati, Nikolay Doudchenko, Guido Imbens, and Khashayar Khosravi. Matrix completion methods for causal panel data models. *Journal of the American Statistical Association*, 116(536):1716–1730, 2021.

- [AFWZ20] Emmanuel Abbe, Jianqing Fan, Kaizheng Wang, and Yiqiao Zhong. Entrywise eigenvector analysis of random matrices with low expected rank. *Annals of statistics*, 48(3):1452, 2020.
- [BCW11] Alexandre Belloni, Victor Chernozhukov, and Lie Wang. Square-root lasso: pivotal recovery of sparse signals via conic programming. *Biometrika*, 98(4):791–806, 2011.
- [BLWY06] Pratik Biswas, Tzu-Chen Lian, Ta-Chung Wang, and Yinyu Ye. Semidefinite programming based algorithms for sensor network localization. *ACM Transactions on Sensor Networks (TOSN)*, 2(2):188–220, 2006.
- [CCF⁺20] Yuxin Chen, Yuejie Chi, Jianqing Fan, Cong Ma, and Yuling Yan. Noisy matrix completion: Understanding statistical guarantees for convex relaxation via nonconvex optimization. *SIAM journal on optimization*, 30(4):3098–3121, 2020.
- [CCF⁺21] Yuxin Chen, Yuejie Chi, Jianqing Fan, Cong Ma, et al. Spectral methods for data science: A statistical perspective. *Foundations and Trends® in Machine Learning*, 14(5):566–806, 2021.
- [CFMW19] Yuxin Chen, Jianqing Fan, Cong Ma, and Kaizheng Wang. Spectral method and regularized mle are both optimal for top-k ranking. *Annals of statistics*, 47(4):2204, 2019.
- [CFMY19] Yuxin Chen, Jianqing Fan, Cong Ma, and Yuling Yan. Inference and uncertainty quantification for noisy matrix completion. *Proceedings of the National Academy of Sciences*, 116(46):22931–22937, 2019.
- [CFMY21] Yuxin Chen, Jianqing Fan, Cong Ma, and Yuling Yan. Bridging convex and nonconvex optimization in robust pca: Noise, outliers and missing data. *The Annals of Statistics*, 49(5):2948–2971, 2021.
- [CFWY21] Yuxin Chen, Jianqing Fan, Bingyan Wang, and Yuling Yan. Convex and nonconvex optimization are both minimax-optimal for noisy blind deconvolution under random designs. *Journal of the American Statistical Association*, pages 1–11, 2021.
- [CGZ22] Pinhan Chen, Chao Gao, and Anderson Y Zhang. Partial recovery for top-k ranking: Optimality of mle and suboptimality of the spectral method. *The Annals of Statistics*, 50(3):1618–1652, 2022.
- [Che15] Yudong Chen. Incoherence-optimal matrix completion. *IEEE Transactions on Information Theory*, 61(5):2909–2923, 2015.
- [CLC19] Yuejie Chi, Yue M Lu, and Yuxin Chen. Nonconvex optimization meets low-rank matrix factorization: An overview. *IEEE Transactions on Signal Processing*, 67(20):5239–5269, 2019.
- [CLL20] Ji Chen, Dekai Liu, and Xiaodong Li. Nonconvex rectangular matrix completion via gradient descent without $\ell_{2,\infty}$ regularization. *IEEE Transactions on Information Theory*, 66(9):5806–5841, 2020.
- [CLMW11] Emmanuel J Candès, Xiaodong Li, Yi Ma, and John Wright. Robust principal component analysis? *Journal of the ACM (JACM)*, 58(3):1–37, 2011.
- [CP10] Emmanuel J Candès and Yaniv Plan. Matrix completion with noise. *Proceedings of the IEEE*, 98(6):925–936, 2010.
- [CR09] Emmanuel J Candès and Benjamin Recht. Exact matrix completion via convex optimization. *Foundations of Computational mathematics*, 9(6):717–772, 2009.
- [CT10] Emmanuel J Candès and Terence Tao. The power of convex relaxation: Near-optimal matrix completion. *IEEE Transactions on Information Theory*, 56(5):2053–2080, 2010.
- [CW15] Yudong Chen and Martin J Wainwright. Fast low-rank estimation by projected gradient descent: General statistical and algorithmic guarantees. *arXiv preprint arXiv:1509.03025*, 2015.

- [DC20] Lijun Ding and Yudong Chen. Leave-one-out approach for matrix completion: Primal and dual analysis. *IEEE Transactions on Information Theory*, 66(11):7274–7301, 2020.
- [DR16] Mark A Davenport and Justin Romberg. An overview of low-rank matrix recovery from incomplete observations. *IEEE Journal of Selected Topics in Signal Processing*, 10(4):608–622, 2016.
- [EK18] Noureddine El Karoui. On the impact of predictor geometry on the performance on high-dimensional ridge-regularized generalized robust regression estimators. *Probability Theory and Related Fields*, 170(1):95–175, 2018.
- [EKBB⁺13] Noureddine El Karoui, Derek Bean, Peter J Bickel, Chinghway Lim, and Bin Yu. On robust regression with high-dimensional predictors. *Proceedings of the National Academy of Sciences*, 110(36):14557–14562, 2013.
- [GHV12] Christophe Giraud, Sylvie Huet, and Nicolas Verzelen. High-dimensional regression with unknown variance. *Statistical Science*, 27(4):500–518, 2012.
- [Gro11] David Gross. Recovering low-rank matrices from few coefficients in any basis. *IEEE Transactions on Information Theory*, 57(3):1548–1566, 2011.
- [Har14] Moritz Hardt. Understanding alternating minimization for matrix completion. In *2014 IEEE 55th Annual Symposium on Foundations of Computer Science*, pages 651–660. IEEE, 2014.
- [Hub11] Peter J Huber. Robust statistics. In *International encyclopedia of statistical science*, pages 1248–1251. Springer, 2011.
- [JNS13] Prateek Jain, Praneeth Netrapalli, and Sujay Sanghavi. Low-rank matrix completion using alternating minimization. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*, pages 665–674, 2013.
- [Klo14] Olga Klopp. Noisy low-rank matrix completion with general sampling distribution. *Bernoulli*, 20(1):282–303, 2014.
- [KLT11] Vladimir Koltchinskii, Karim Lounici, and Alexandre B Tsybakov. Nuclear-norm penalization and optimal rates for noisy low-rank matrix completion. *The Annals of Statistics*, 39(5):2302–2329, 2011.
- [KMO09] Raghunandan Keshavan, Andrea Montanari, and Sewoong Oh. Matrix completion from noisy entries. *Advances in neural information processing systems*, 22, 2009.
- [KMO10] Raghunandan H Keshavan, Andrea Montanari, and Sewoong Oh. Matrix completion from a few entries. *IEEE transactions on information theory*, 56(6):2980–2998, 2010.
- [KS21] Felix Krahermer and Dominik Stöger. On the convex geometry of blind deconvolution and matrix completion. *Communications on Pure and Applied Mathematics*, 74(4):790–832, 2021.
- [LM15] Johannes Lederer and Christian Müller. Don’t fall for tuning parameters: tuning-free variable selection in high dimensions with the trex. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 29, 2015.
- [MWCC18] Cong Ma, Kaizheng Wang, Yuejie Chi, and Yuxin Chen. Implicit regularization in nonconvex statistical estimation: Gradient descent converges linearly for phase retrieval and matrix completion. In *International Conference on Machine Learning*, pages 3345–3354. PMLR, 2018.
- [NW12] Sahand Negahban and Martin J Wainwright. Restricted strong convexity and weighted matrix completion: Optimal bounds with noise. *The Journal of Machine Learning Research*, 13(1):1665–1697, 2012.

- [Owe07] Art B Owen. A robust hybrid of lasso and ridge regression. *Contemporary Mathematics*, 443(7):59–72, 2007.
- [PW20] Ashwin Pananjady and Martin J Wainwright. Instance-dependent ℓ_∞ -bounds for policy evaluation in tabular reinforcement learning. *IEEE Transactions on Information Theory*, 67(1):566–585, 2020.
- [Rec11] Benjamin Recht. A simpler approach to matrix completion. *Journal of Machine Learning Research*, 12(12), 2011.
- [RS05] Jasson DM Rennie and Nathan Srebro. Fast maximum margin matrix factorization for collaborative prediction. In *Proceedings of the 22nd international conference on Machine learning*, pages 713–719, 2005.
- [SZ12] Tingni Sun and Cun-Hui Zhang. Scaled sparse linear regression. *Biometrika*, 99(4):879–898, 2012.
- [TMC21] Tian Tong, Cong Ma, and Yuejie Chi. Accelerating ill-conditioned low-rank matrix estimation via scaled gradient descent. *J. Mach. Learn. Res.*, 22:150–1, 2021.
- [Ver10] Roman Vershynin. Introduction to the non-asymptotic analysis of random matrices. *arXiv preprint arXiv:1011.3027*, 2010.
- [WF22] Bingyan Wang and Jianqing Fan. Robust matrix completion with heavy-tailed noise. *arXiv preprint arXiv:2206.04276*, 2022.
- [WPB⁺20] Lan Wang, Bo Peng, Jelena Bradic, Runze Li, and Yunan Wu. A tuning-free robust and efficient approach to high-dimensional regression. *Journal of the American Statistical Association*, 115(532):1700–1714, 2020.
- [WW19] Yunan Wu and Lan Wang. A survey of tuning parameter selection for high-dimensional regression. *arXiv preprint arXiv:1908.03669*, 2019.
- [YCF21] Yuling Yan, Yuxin Chen, and Jianqing Fan. Inference for heteroskedastic pca with missing data. *arXiv preprint arXiv:2107.12365*, 2021.
- [ZB18] Yiqiao Zhong and Nicolas Boumal. Near-optimal bounds for phase synchronization. *SIAM Journal on Optimization*, 28(2):989–1016, 2018.
- [ZL16] Qinqing Zheng and John Lafferty. Convergence analysis for rectangular matrix completion using burer-monteiro factorization and gradient descent. *arXiv preprint arXiv:1605.07051*, 2016.
- [ZYW21] Junhui Zhang, Jingkai Yan, and John Wright. Square root principal component pursuit: Tuning-free noisy robust matrix recovery. *Advances in Neural Information Processing Systems*, 34, 2021.

A Proof of Lemma 3

We prove Lemma 3 via induction. Since all the algorithms start from the groundtruth, it is trivial to see that the hypotheses (10) hold for $t = 0$. We also record two important properties of the iterates at $t = 0$, namely,

$$\frac{1}{2}np^{1/2}\sigma \leq \theta_t \leq 2np^{1/2}\sigma \quad (16)$$

and

$$\|\mathbf{X}_t^\top \mathbf{X}_t - \mathbf{Y}_t^\top \mathbf{Y}_t\|_F \leq C_B \kappa \eta \frac{\sigma}{\sigma_{\min}} \sqrt{\frac{n}{p}} \sqrt{r} \sigma_{\max}^2, \quad (17)$$

where $C_B > 0$ is a universal constant. Note that at $t = 0$, we have $\theta_0 = \|\mathcal{P}_\Omega(\mathbf{E})\|_F$, which concentrates sharply around $np^{1/2}\sigma$ under the noise assumption and uniform sampling.

Now suppose the hypotheses (10), (16), and (17) hold for the t -th iterates. We aim to show that the same set of hypotheses continue to hold for the $(t+1)$ -th iterates. Sections A.1 and A.2 are devoted to this induction step. In addition, we prove the last claim (11) in Section A.3. In Section A.4 we prove Lemma 6 which is a consequence of (10) and (16).

A.1 Induction on hypotheses (10) and (17)

Define

$$\tilde{\lambda}_t := \lambda\theta_t, \quad \text{and} \quad \tilde{\eta}_t := \eta/\theta_t.$$

We make a key observation that the t -th iterations of Algorithm 1 and 2 are exactly the same as the t -th iterations of Algorithm 1 (vanilla gradient descent) and 2 (construction of the leave-one-out sequence) in the paper [CCF+20] with the parameters $\tilde{\lambda}_t$ and $\tilde{\eta}_t$. Moreover, given the induction hypothesis (16) one has $\frac{1}{2}n\sqrt{\bar{\rho}}\sigma \leq \theta_{t-1} \leq 2n\sqrt{\bar{\rho}}\sigma$. Combine this with our choice of $\lambda = C_\lambda n^{-1/2}$ to see that

$$\tilde{\lambda}_t \asymp \sigma\sqrt{np}, \quad \text{and} \quad \tilde{\eta}_t \asymp 1/(np\kappa^3\sigma_{\max}).$$

which are consistent with the choice of λ and η in [CCF+20]. These allow us to invoke Lemmas 10-15 in [CCF+20] to prove that claims (10) and (17) hold for the $(t+1)$ -th iterates.

A.2 Induction on hypotheses (16)

In this section, we aim to show that the claim (16) holds for the $(t+1)$ -th iterates.

Observe that

$$\mathcal{P}_\Omega(\mathbf{X}_{t+1}\mathbf{Y}_{t+1}^\top - \mathbf{M}) = \mathcal{P}_\Omega(\mathbf{X}_{t+1}\mathbf{Y}_{t+1}^\top - \mathbf{L}^*) - \mathcal{P}_\Omega(\mathbf{E}).$$

Similar to the proof of Lemma 1, using the incoherence assumption $\|\mathbf{F}^*\|_{2,\infty} = \max\{\|\mathbf{X}^*\|_{2,\infty}, \|\mathbf{Y}^*\|_{2,\infty}\} \leq \sqrt{\mu r \sigma_{\max}/n}$, we have

$$\begin{aligned} \|\mathbf{X}_{t+1}\mathbf{Y}_{t+1}^\top - \mathbf{L}^*\|_\infty &\leq 3C_\infty \frac{\sigma}{\sigma_{\min}} \sqrt{\frac{n \log n}{p}} \|\mathbf{F}^*\|_{2,\infty} \|\mathbf{F}^*\|_{2,\infty} \\ &\leq 3C_\infty \frac{\mu r \sigma_{\max}}{n} \frac{\sigma}{\sigma_{\min}} \sqrt{\frac{n \log n}{p}}. \end{aligned}$$

Then $\|\mathcal{P}_\Omega(\mathbf{X}_{t+1}\mathbf{Y}_{t+1}^\top - \mathbf{L}^*)\|_{\text{F}} \lesssim n\sqrt{p} \|\mathbf{X}_{t+1}\mathbf{Y}_{t+1}^\top - \mathbf{L}^*\|_\infty \lesssim \kappa\mu r \sigma \sqrt{n \log n}$. As the sample size satisfies $n^2 p \gg \kappa^4 \mu^2 r^2 n \log^3 n$, we have $\|\mathcal{P}_\Omega(\mathbf{X}_{t+1}\mathbf{Y}_{t+1}^\top - \mathbf{L}^*)\|_{\text{F}} \ll n\sqrt{p}\sigma$. As mentioned before, $\|\mathcal{P}_\Omega(\mathbf{E})\|_{\text{F}}$ sharply concentrates around $np^{1/2}\sigma$. Therefore by the triangle inequality, we have

$$\frac{1}{2} \sigma n \sqrt{p} \leq \|\mathcal{P}_\Omega(\mathbf{X}_t \mathbf{Y}_t^\top - \mathbf{M})\|_{\text{F}} \leq 2 \sigma n \sqrt{p}$$

for large enough n .

A.3 Proof of bound (11)

Suppose for the moment that

$$f(\mathbf{X}_t, \mathbf{Y}_t, \theta_t) \leq f(\mathbf{X}_{t-1}, \mathbf{Y}_{t-1}, \theta_{t-1}) - \frac{\eta}{2} \|\nabla_{\mathbf{X}, \mathbf{Y}} f(\mathbf{X}_{t-1}, \mathbf{Y}_{t-1}, \theta_{t-1})\|_{\text{F}}^2 \quad (18)$$

holds for all $t \geq 1$. Then a telescoping argument would yield the conclusion that

$$\begin{aligned} f(\mathbf{X}_0, \mathbf{Y}_0, \theta_0) - f(\mathbf{X}_{t_0}, \mathbf{Y}_{t_0}, \theta_{t_0}) &\geq \frac{\eta}{2} \sum_{t=0}^{t_0-1} \|\nabla_{\mathbf{X}, \mathbf{Y}} f(\mathbf{X}_t, \mathbf{Y}_t, \theta_t)\|_{\text{F}}^2 \\ &\geq \frac{\eta t_0}{2} \min_{0 \leq t < t_0} \|\nabla_{\mathbf{X}, \mathbf{Y}} f(\mathbf{X}_t, \mathbf{Y}_t, \theta_t)\|_{\text{F}}^2. \end{aligned}$$

Expanding the left hand side, we see that it is upper bounded by

$$\begin{aligned} f(\mathbf{X}_0, \mathbf{Y}_0, \theta_0) - f(\mathbf{X}_{t_0}, \mathbf{Y}_{t_0}, \theta_{t_0}) &= \|\mathcal{P}_\Omega(\mathbf{E})\|_F - \|\mathcal{P}_\Omega(\mathbf{X}_{t_0} \mathbf{Y}_{t_0}^\top - \mathbf{M})\|_F + \frac{\lambda}{2} (\|\mathbf{X}^*\|_F^2 - \|\mathbf{X}_{t_0}\|_F^2 + \|\mathbf{Y}^*\|_F^2 - \|\mathbf{Y}_{t_0}\|_F^2) \\ &\leq \|\mathcal{P}_\Omega(\mathbf{E})\|_F + \frac{\lambda}{2} (\|\mathbf{X}^*\|_F^2 - \|\mathbf{X}_{t_0} \mathbf{H}_{t_0}\|_F^2 + \|\mathbf{Y}^*\|_F^2 - \|\mathbf{Y}_{t_0} \mathbf{H}_{t_0}\|_F^2), \end{aligned}$$

where the last line uses the nonnegativity of norms and the invariance of Frobenius norm under rotation. In view of the properties (10) and the noise size assumption $\frac{\sigma}{\sigma_{\min}} \sqrt{\frac{n}{p}} \ll 1$, we have

$$\|\mathbf{X}^* - \mathbf{X}_{t_0} \mathbf{H}_{t_0}\|_F \lesssim \frac{\sigma}{\sigma_{\min}} \sqrt{\frac{n}{p}} \|\mathbf{X}^*\|_F, \quad \text{and} \quad \|\mathbf{X}_{t_0}\|_F = \|\mathbf{X}_{t_0} \mathbf{H}_{t_0}\|_F \leq 2 \|\mathbf{X}^*\|_F.$$

Then,

$$\begin{aligned} \left| \|\mathbf{X}^*\|_F^2 - \|\mathbf{X}_{t_0} \mathbf{H}_{t_0}\|_F^2 \right| &\leq \left| \|\mathbf{X}^*\|_F - \|\mathbf{X}_{t_0} \mathbf{H}_{t_0}\|_F \right| (\|\mathbf{X}^*\|_F + \|\mathbf{X}_{t_0} \mathbf{H}_{t_0}\|_F) \\ &\lesssim \frac{\sigma}{\sigma_{\min}} \sqrt{\frac{n}{p}} \|\mathbf{X}^*\|_F \|\mathbf{X}^*\|_F \\ &\leq \sigma r \kappa \sqrt{\frac{n}{p}}, \end{aligned} \tag{19}$$

where the last line uses the fact that $\|\mathbf{X}^*\|_F \leq \sqrt{r \sigma_{\max}}$. Similarly, we have $\left| \|\mathbf{Y}^*\|_F^2 - \|\mathbf{Y}_{t_0} \mathbf{H}_{t_0}\|_F^2 \right| \lesssim \sigma r \kappa \sqrt{\frac{n}{p}}$. These combined with the fact that $\|\mathcal{P}_\Omega(\mathbf{E})\|_F \lesssim n \sqrt{p} \sigma$ implies, as $t_0 = n^{18}, \eta \asymp \sigma / (\sqrt{p} \kappa^3 \sigma_{\max})$, and $\lambda \asymp 1/\sqrt{n}$,

$$\begin{aligned} \min_{0 \leq t < t_0} \|\nabla_{\mathbf{X}, \mathbf{Y}} f(\mathbf{X}_t, \mathbf{Y}_t, \theta_t)\|_F &\leq \left[\frac{f(\mathbf{X}_0, \mathbf{Y}_0, \theta_0) - f(\mathbf{X}_{t_0}, \mathbf{Y}_{t_0}, \theta_{t_0})}{\eta t_0 / 2} \right]^{1/2} \\ &\lesssim \left[\frac{\sigma_{\max}}{n^{18} \sqrt{p} \sigma} \left(\frac{n \sigma}{\sqrt{p}} \right) \right]^{1/2} \\ &\lesssim \frac{1}{n^8} \sqrt{\frac{\sigma_{\max}}{p}}. \end{aligned}$$

To simplify the expression we use $\kappa \lesssim n$ and $r \lesssim \sqrt{n}$ which are consequences of the sample size assumption $n^2 \geq n^2 p \gg \kappa^4 \mu^2 r^2 n \log n$.

Proof of bound (18). Define $h(\mathbf{X}, \mathbf{Y}) := \theta_t [f(\mathbf{X}, \mathbf{Y}, \theta_t) - \theta_t/2]$. Then $h(\mathbf{X}, \mathbf{Y})$ matches the form of the objective function in Lemma 16 of the paper [CCF+20]. Then Lemma 16 therein tells us that

$$h(\mathbf{X}_{t+1}, \mathbf{Y}_{t+1}) \leq h(\mathbf{X}_t, \mathbf{Y}_t) - \frac{\tilde{\eta}_t}{2} \|\nabla h(\mathbf{X}_t, \mathbf{Y}_t)\|_F^2,$$

where we recall $\tilde{\eta}_t = \eta/\theta_t$. Rewriting the bound in terms of f yields

$$f(\mathbf{X}_{t+1}, \mathbf{Y}_{t+1}, \theta_t) \leq f(\mathbf{X}_t, \mathbf{Y}_t, \theta_t) - \frac{\eta}{2} \|\nabla_{\mathbf{X}, \mathbf{Y}} f(\mathbf{X}_t, \mathbf{Y}_t, \theta_t)\|_F^2. \tag{20}$$

In addition, by the optimality of θ_{t+1} , one has

$$f(\mathbf{X}_{t+1}, \mathbf{Y}_{t+1}, \theta_{t+1}) \leq f(\mathbf{X}_{t+1}, \mathbf{Y}_{t+1}, \theta_t). \tag{21}$$

Combining equations (20) and (21) completes the proof.

A.4 Proof of Lemma 6

By Lemma 3, we know that \mathbf{X}_{ncvx} satisfies

$$\|\mathbf{X}_{\text{ncvx}} - \mathbf{X}^*\| \leq C_{\text{op}} \left(\frac{\sigma}{\sigma_{\min}} \sqrt{\frac{n}{p}} \right) \|\mathbf{X}^*\| \ll \sqrt{\sigma_{\min}},$$

where the last relation arises from the noise level assumption $\frac{\sigma}{\sigma_{\min}} \sqrt{\frac{n}{p}} \ll 1/\sqrt{\kappa^4 \mu r \log n}$. Therefore we can apply Weyl's inequality to obtain

$$\begin{aligned} \sigma_{\max}(\mathbf{X}_{\text{ncvx}}) &\leq \sqrt{\sigma_{\max}} + \|\mathbf{X}_{\text{ncvx}} - \mathbf{X}^*\| \leq \sqrt{2\sigma_{\max}}; \\ \sigma_{\min}(\mathbf{X}_{\text{ncvx}}) &\geq \sqrt{\sigma_{\min}} - \|\mathbf{X}_{\text{ncvx}} - \mathbf{X}^*\| \geq \sqrt{\sigma_{\min}/2} \end{aligned}$$

for large enough n . These hold similarly for the singular values of \mathbf{Y}_{ncvx} .

On the other hand, the relations (14c) come directly from (16), and (14d) follows from Lemma 4 in [CCF+20].

B Proof of Lemma 7

To simplify the notation, we denote $\theta := \|\mathcal{P}_{\Omega}(\mathbf{L}_{\text{ncvx}} - \mathbf{M})\|_{\text{F}}$, and $\Delta := \mathbf{L}_{\text{cvx}} - \mathbf{L}_{\text{ncvx}}$ throughout this section. In view of Lemma 6, we know that $\theta \neq 0$, and hence θ^{-1} is well defined.

Recall that $\mathbf{U}\Sigma\mathbf{V}^{\top}$ is the SVD for \mathbf{L}_{ncvx} , and T is the tangent space at \mathbf{L}_{ncvx} . The following lemma is useful in controlling the size of Δ .

Lemma 8. *Under the notations and assumptions of Lemma 7, we have*

$$\frac{1}{\theta} \mathcal{P}_{\Omega}(\mathbf{L}_{\text{ncvx}} - \mathbf{M}) = -\lambda(\mathbf{U}\mathbf{V}^{\top} + \mathbf{R}), \quad (22)$$

where \mathbf{R} is a residual matrix such that

$$\|\mathcal{P}_T(\mathbf{R})\|_{\text{F}} \leq 70\kappa\sigma_{\min}^{-1/2} \|\nabla g(\mathbf{X}, \mathbf{Y})\|_{\text{F}}, \quad \text{and} \quad \|\mathcal{P}_{T^{\perp}}(\mathbf{R})\| < 1/2.$$

See Section B.1 for the proof.

We decompose the proof into three steps. In Step 1, we show that the difference matrix Δ mainly lies in the tangent space T . In Step 2, the previous fact is leveraged to show an upper bound on $\mathcal{P}_{\Omega}(\Delta)$. In the last step (Step 3), we connect the previous steps with the injectivity property (cf. Lemma 5) to reach the desired conclusion.

Step 1: showing that Δ lies primarily in the tangent space T . By the optimality of \mathbf{L}_{cvx} , we have

$$0 \geq \|\mathcal{P}_{\Omega}(\mathbf{L}_{\text{cvx}} - \mathbf{M})\|_{\text{F}} - \|\mathcal{P}_{\Omega}(\mathbf{L}_{\text{ncvx}} - \mathbf{M})\|_{\text{F}} + \lambda(\|\mathbf{L}_{\text{cvx}}\|_* - \|\mathbf{L}_{\text{ncvx}}\|_*). \quad (23)$$

Use the convexity of $\|\cdot\|_{\text{F}}$ and $\|\cdot\|_*$ and the decomposition $\mathbf{L}_{\text{ncvx}} = \mathbf{U}\Sigma\mathbf{V}^{\top}$ to see that

$$0 \geq \left\langle \frac{1}{\theta} \mathcal{P}_{\Omega}(\mathbf{L}_{\text{ncvx}} - \mathbf{M}), \Delta \right\rangle + \lambda \langle \mathbf{U}\mathbf{V}^{\top} + \mathbf{W}_0, \Delta \rangle$$

holds for any $\mathbf{W}_0 \in T^{\perp}$ with $\|\mathbf{W}_0\| \leq 1$. Apply Lemma 8 to further obtain

$$0 \geq -\lambda \langle \mathbf{R}, \Delta \rangle + \lambda \langle \mathbf{W}_0, \Delta \rangle.$$

In particular, one can choose $\mathbf{W}_0 \in T^{\perp}$ such that $\|\mathcal{P}_{T^{\perp}}(\Delta)\|_* = \langle \mathbf{W}_0, \Delta \rangle$, which yields the inequality

$$\begin{aligned} 0 &\geq \lambda \|\mathcal{P}_{T^{\perp}}(\Delta)\|_* - \lambda \langle \mathbf{R}, \Delta \rangle \\ &= \lambda \|\mathcal{P}_{T^{\perp}}(\Delta)\|_* - \lambda \langle \mathcal{P}_T(\mathbf{R}), \Delta \rangle - \lambda \langle \mathcal{P}_{T^{\perp}}(\mathbf{R}), \Delta \rangle \\ &\geq \lambda \|\mathcal{P}_{T^{\perp}}(\Delta)\|_* - \lambda \|\mathcal{P}_T(\mathbf{R})\|_{\text{F}} \|\mathcal{P}_T(\Delta)\|_{\text{F}} - \lambda \|\mathcal{P}_{T^{\perp}}(\mathbf{R})\| \|\mathcal{P}_{T^{\perp}}(\Delta)\|_*. \end{aligned}$$

Here the last line arises from Holder's inequality.

Again, by Lemma 8, we have the bounds $\|\mathcal{P}_T(\mathbf{R})\|_F \leq 70\kappa\sigma_{\min}^{-1/2}\|\nabla g(\mathbf{X}, \mathbf{Y})\|_F$ and $\|\mathcal{P}_{T^\perp}(\mathbf{R})\| < 1/2$, which allow us to further arrive at

$$0 \geq \frac{\lambda}{2}\|\mathcal{P}_{T^\perp}(\mathbf{\Delta})\|_* - 70\lambda\kappa\sigma_{\min}^{-1/2}\|\nabla g(\mathbf{X}, \mathbf{Y})\|_F\|\mathcal{P}_T(\mathbf{\Delta})\|_F.$$

This further implies

$$\|\mathcal{P}_{T^\perp}(\mathbf{\Delta})\|_F \leq \|\mathcal{P}_{T^\perp}(\mathbf{\Delta})\|_* \leq 140\kappa\sigma_{\min}^{-1/2}\|\nabla g(\mathbf{X}, \mathbf{Y})\|_F\|\mathcal{P}_T(\mathbf{\Delta})\|_F. \quad (24)$$

As an immediate consequence, under the assumed upper bound (15) for $\|\nabla g(\mathbf{X}, \mathbf{Y})\|_F$, we have $140\kappa\sigma_{\min}^{-1/2}\|\nabla g(\mathbf{X}, \mathbf{Y})\|_F \leq 1$, and hence

$$\|\mathbf{\Delta}\|_F \leq \|\mathcal{P}_{T^\perp}(\mathbf{\Delta})\|_F + \|\mathcal{P}_T(\mathbf{\Delta})\|_F \leq 2\|\mathcal{P}_T(\mathbf{\Delta})\|_F \quad (25)$$

Step 2: bounding $\|\mathcal{P}_\Omega(\mathbf{\Delta})\|_F^2$. We start with presenting an identity involving $\|\mathcal{P}_\Omega(\mathbf{\Delta})\|_F^2$:

$$\begin{aligned} \|\mathcal{P}_\Omega(\mathbf{\Delta})\|_F^2 &= (\|\mathcal{P}_\Omega(\mathbf{L}_{\text{cvx}} - \mathbf{M})\|_F - \|\mathcal{P}_\Omega(\mathbf{L}_{\text{ncvx}} - \mathbf{M})\|_F) (\|\mathcal{P}_\Omega(\mathbf{L}_{\text{cvx}} - \mathbf{M})\|_F + \|\mathcal{P}_\Omega(\mathbf{L}_{\text{ncvx}} - \mathbf{M})\|_F) \\ &\quad - 2\langle \mathbf{\Delta}, \mathcal{P}_\Omega(\mathbf{L}_{\text{ncvx}} - \mathbf{M}) \rangle \\ &= (\|\mathcal{P}_\Omega(\mathbf{L}_{\text{cvx}} - \mathbf{M})\|_F - \|\mathcal{P}_\Omega(\mathbf{L}_{\text{ncvx}} - \mathbf{M})\|_F)^2 \\ &\quad + 2\|\mathcal{P}_\Omega(\mathbf{L}_{\text{ncvx}} - \mathbf{M})\|_F \cdot \left(\|\mathcal{P}_\Omega(\mathbf{L}_{\text{cvx}} - \mathbf{M})\|_F - \|\mathcal{P}_\Omega(\mathbf{L}_{\text{ncvx}} - \mathbf{M})\|_F - \left\langle \frac{1}{\theta}\mathcal{P}_\Omega(\mathbf{L}_{\text{ncvx}} - \mathbf{M}), \mathbf{\Delta} \right\rangle \right). \end{aligned} \quad (26)$$

Lemma 8 and Equation (23) tell us that

$$\begin{aligned} &\|\mathcal{P}_\Omega(\mathbf{L}_{\text{cvx}} - \mathbf{M})\|_F - \|\mathcal{P}_\Omega(\mathbf{L}_{\text{ncvx}} - \mathbf{M})\|_F - \left\langle \frac{1}{\theta}\mathcal{P}_\Omega(\mathbf{L}_{\text{ncvx}} - \mathbf{M}), \mathbf{\Delta} \right\rangle \\ &\leq \lambda\|\mathbf{L}_{\text{ncvx}}\|_* - \lambda\|\mathbf{L}_{\text{cvx}}\|_* + \lambda\langle \mathbf{UV}^\top + \mathbf{R}, \mathbf{\Delta} \rangle. \end{aligned}$$

By convexity of $\|\cdot\|_*$, this further simplifies to

$$\begin{aligned} &\lambda\|\mathbf{L}_{\text{ncvx}}\|_* - \lambda\|\mathbf{L}_{\text{cvx}}\|_* + \lambda\langle \mathbf{UV}^\top + \mathbf{R}, \mathbf{\Delta} \rangle \\ &\leq -\lambda\langle \mathbf{UV}^\top + \mathbf{W}, \mathbf{\Delta} \rangle + \lambda\langle \mathbf{UV}^\top + \mathbf{R}, \mathbf{\Delta} \rangle \\ &= \lambda\langle \mathbf{\Delta}, \mathbf{R} - \mathbf{W} \rangle, \end{aligned} \quad (27)$$

for any $\mathbf{W} \in T^\perp$ with $\|\mathbf{W}\| \leq 1$. Combine Equation (26) and (27) to reach

$$\|\mathcal{P}_\Omega(\mathbf{\Delta})\|_F^2 \leq \underbrace{(\|\mathcal{P}_\Omega(\mathbf{L}_{\text{cvx}} - \mathbf{M})\|_F - \|\mathcal{P}_\Omega(\mathbf{L}_{\text{ncvx}} - \mathbf{M})\|_F)^2}_{=: \alpha_1} + \underbrace{2\lambda\|\mathcal{P}_\Omega(\mathbf{L}_{\text{ncvx}} - \mathbf{M})\|_F |\langle \mathbf{\Delta}, \mathbf{R} - \mathbf{W} \rangle|}_{=: \alpha_2}.$$

We prove in the end of this section that the two terms α_1 and α_2 obey

$$\alpha_1 \leq \lambda^2(\sqrt{r} + 140\kappa\sigma_{\min}^{-1/2}\|\nabla g(\mathbf{X}, \mathbf{Y})\|_F)^2\|\mathcal{P}_T(\mathbf{\Delta})\|_F^2; \quad (28a)$$

$$\alpha_2 \leq 560\lambda\kappa\sigma_{\min}^{-1/2}\theta\|\nabla g(\mathbf{X}, \mathbf{Y})\|_F\|\mathcal{P}_T(\mathbf{\Delta})\|_F, \quad (28b)$$

which yields the upper bound on $\|\mathcal{P}_\Omega(\mathbf{\Delta})\|_F^2$ in terms of $\|\mathcal{P}_T(\mathbf{\Delta})\|_F$:

$$\begin{aligned} \|\mathcal{P}_\Omega(\mathbf{\Delta})\|_F^2 &\leq \lambda^2(\sqrt{r} + 140\kappa\sigma_{\min}^{-1/2}\|\nabla g(\mathbf{X}, \mathbf{Y})\|_F)^2\|\mathcal{P}_T(\mathbf{\Delta})\|_F^2 \\ &\quad + 560\lambda\kappa\sigma_{\min}^{-1/2}\theta\|\nabla g(\mathbf{X}, \mathbf{Y})\|_F\|\mathcal{P}_T(\mathbf{\Delta})\|_F. \end{aligned}$$

Step 3: final calculations. Using the decomposition $\mathcal{P}_\Omega(\Delta) = \mathcal{P}_\Omega\mathcal{P}_T(\Delta) + \mathcal{P}_\Omega\mathcal{P}_{T^\perp}(\Delta)$, we obtain

$$\begin{aligned}\|\mathcal{P}_\Omega(\Delta)\|_F &= \|\mathcal{P}_\Omega\mathcal{P}_T(\Delta) + \mathcal{P}_\Omega\mathcal{P}_{T^\perp}(\Delta)\|_F \\ &\geq \|\mathcal{P}_\Omega\mathcal{P}_T(\Delta)\|_F - \|\mathcal{P}_\Omega\mathcal{P}_{T^\perp}(\Delta)\|_F.\end{aligned}$$

Together with Lemma 5 and Equation 24, we have

$$\begin{aligned}\|\mathcal{P}_\Omega(\Delta)\|_F &\geq (\sqrt{p}C_{\text{inj}} - 140\kappa\sigma_{\min}^{-1/2}\|\nabla g(\mathbf{X}, \mathbf{Y})\|_F)\|\mathcal{P}_T(\Delta)\|_F \\ &\geq \frac{\sqrt{p}}{2}C_{\text{inj}}\|\mathcal{P}_T(\Delta)\|_F.\end{aligned}$$

where the last line uses (15). As a result, we arrive at the sandwich formula

$$\begin{aligned}\frac{1}{4}pC_{\text{inj}}^2\|\mathcal{P}_T(\Delta)\|_F^2 &\leq \|\mathcal{P}_\Omega(\Delta)\|_F^2 \\ &\leq \lambda^2(\sqrt{r} + 140\kappa\sigma_{\min}^{-1/2}\|\nabla g(\mathbf{X}, \mathbf{Y})\|_F)^2\|\mathcal{P}_T(\Delta)\|_F^2 \\ &\quad + 560\lambda\kappa\sigma_{\min}^{-1/2}\theta\|\nabla g(\mathbf{X}, \mathbf{Y})\|_F\|\mathcal{P}_T(\Delta)\|_F,\end{aligned}$$

which further implies

$$\begin{aligned}\left\{\frac{pC_{\text{inj}}^2}{4} - \lambda^2(\sqrt{r} + 140\kappa\sigma_{\min}^{-1/2}\|\nabla g(\mathbf{X}, \mathbf{Y})\|_F)^2\right\}\|\mathcal{P}_T(\Delta)\|_F^2 \\ \leq 560\lambda\kappa\sigma_{\min}^{-1/2}\theta\|\nabla g(\mathbf{X}, \mathbf{Y})\|_F\|\mathcal{P}_T(\Delta)\|_F.\end{aligned}$$

Reorganize and substitute in (15) to see that for large enough n ,

$$\frac{pC_{\text{inj}}^2}{4} - \lambda^2(\sqrt{r} + 140\kappa\sigma_{\min}^{-1/2}\|\nabla g(\mathbf{X}, \mathbf{Y})\|_F)^2 \geq \frac{pC_{\text{inj}}^2}{8}.$$

Combine the above two relations to reach

$$\frac{pC_{\text{inj}}^2}{8}\|\mathcal{P}_T(\Delta)\|_F^2 \leq 560\lambda\kappa\sigma_{\min}^{-1/2}\theta\|\nabla g(\mathbf{X}, \mathbf{Y})\|_F\|\mathcal{P}_T(\Delta)\|_F,$$

which together with $C_{\text{inj}} = (32\kappa)^{-1/2}$ and (14c) implies

$$\|\mathcal{P}_T(\Delta)\|_F \lesssim \frac{\lambda\kappa^2}{\sqrt{p\sigma_{\min}}}n\sigma\|\nabla g(\mathbf{X}, \mathbf{Y})\|_F.$$

Use (25), we obtain the bound on $\|\Delta\|_F$,

$$\|\Delta\|_F \leq 2\|\mathcal{P}_T(\Delta)\|_F \lesssim \frac{\lambda\kappa^2}{\sqrt{p\sigma_{\min}}}n\sigma\|\nabla g(\mathbf{X}, \mathbf{Y})\|_F.$$

Proof of the bound (28a). For α_1 we consider the cases when $\|\mathcal{P}_\Omega(\mathbf{L}_{\text{cvx}} - \mathbf{M})\|_F - \|\mathcal{P}_\Omega(\mathbf{L}_{\text{ncvx}} - \mathbf{M})\|_F$ is positive and non-positive separately.

Case of $\|\mathcal{P}_\Omega(\mathbf{L}_{\text{cvx}} - \mathbf{M})\|_F - \|\mathcal{P}_\Omega(\mathbf{L}_{\text{ncvx}} - \mathbf{M})\|_F \leq 0$. By convexity of $\|\cdot\|_F$,

$$0 \geq \|\mathcal{P}_\Omega(\mathbf{L}_{\text{cvx}} - \mathbf{M})\|_F - \|\mathcal{P}_\Omega(\mathbf{L}_{\text{ncvx}} - \mathbf{M})\|_F > \left\langle \frac{1}{\theta}\mathcal{P}(\mathbf{L}_{\text{ncvx}} - \mathbf{M}), \Delta \right\rangle.$$

Using the representation in Lemma 8, the last term can be written as $\lambda\langle \mathbf{UV}^\top + \mathbf{R}, \Delta \rangle$. Splitting the parts into T and T^\perp , we have

$$\begin{aligned}(\|\mathcal{P}_\Omega(\mathbf{L}_{\text{cvx}} - \mathbf{M})\|_F - \|\mathcal{P}_\Omega(\mathbf{L}_{\text{ncvx}} - \mathbf{M})\|_F)^2 \\ \leq \lambda^2\langle \mathbf{UV}^\top + \mathbf{R}, \Delta \rangle^2 \\ \leq \lambda^2(\|\mathbf{UV}^\top\|_F\|\mathcal{P}_T(\Delta)\|_F + \|\mathcal{P}_T(\mathbf{R})\|_F\|\mathcal{P}_T(\Delta)\|_F + \|\mathcal{P}_{T^\perp}(\mathbf{R})\|_F\|\mathcal{P}_{T^\perp}(\Delta)\|_F)^2.\end{aligned}$$

Together with Equation (24) and Lemma 8, we arrive at

$$(\|\mathcal{P}_\Omega(\mathbf{L}_{\text{cvx}} - \mathbf{M})\|_F - \|\mathcal{P}_\Omega(\mathbf{L}_{\text{ncvx}} - \mathbf{M})\|_F)^2 \leq \lambda^2(\sqrt{r} + 140\kappa\sigma_{\min}^{-1/2}\|\nabla g(\mathbf{X}, \mathbf{Y})\|_F)^2\|\mathcal{P}_T(\Delta)\|_F^2.$$

Case of $\|\mathcal{P}_\Omega(\mathbf{L}_{\text{cvx}} - \mathbf{M})\|_F - \|\mathcal{P}_\Omega(\mathbf{L}_{\text{ncvx}} - \mathbf{M})\|_F > 0$. By optimality of \mathbf{L}_{cvx} and convexity of $\|\cdot\|_*$,

$$0 < \|\mathcal{P}_\Omega(\mathbf{L}_{\text{cvx}} - \mathbf{M})\|_F - \|\mathcal{P}_\Omega(\mathbf{L}_{\text{ncvx}} - \mathbf{M})\|_F \leq -\lambda(\|\mathbf{L}_{\text{cvx}}\|_* - \|\mathbf{L}_{\text{ncvx}}\|_*) \leq -\lambda \langle \mathbf{UV}^\top, \mathbf{\Delta} \rangle.$$

Then similar to the case of $\|\mathcal{P}_\Omega(\mathbf{L}_{\text{cvx}} - \mathbf{M})\|_F - \|\mathcal{P}_\Omega(\mathbf{L}_{\text{ncvx}} - \mathbf{M})\|_F \leq 0$,

$$(\|\mathcal{P}_\Omega(\mathbf{L}_{\text{cvx}} - \mathbf{M})\|_F - \|\mathcal{P}_\Omega(\mathbf{L}_{\text{ncvx}} - \mathbf{M})\|_F)^2 \leq \lambda^2 r \|\mathcal{P}_T(\mathbf{\Delta})\|_F^2.$$

Combining the two cases yields (28a).

Proof of the bound (28b). For α_2 , we can split the parts into T and T^\perp similar to the proof for (28a). Using Equation (24) and Lemma 8, we have

$$\begin{aligned} 2\theta \cdot \lambda \langle \mathbf{\Delta}, \mathbf{R} - \mathbf{W} \rangle &\leq 2\lambda\theta (|\langle \mathbf{\Delta}, \mathbf{R} \rangle| + |\langle \mathbf{\Delta}, \mathbf{W} \rangle|) \\ &\leq 2\lambda\theta (\|\mathcal{P}_T(\mathbf{R})\|_F \|\mathcal{P}_T(\mathbf{\Delta})\|_F + (\|\mathcal{P}_{T^\perp}(\mathbf{R})\| + \|\mathcal{P}_{T^\perp}(\mathbf{W})\|) \|\mathcal{P}_{T^\perp}(\mathbf{\Delta})\|_*) \\ &\leq 560\lambda\kappa\sigma_{\min}^{-1/2}\theta \|\nabla g(\mathbf{X}, \mathbf{Y})\|_F \|\mathcal{P}_T(\mathbf{\Delta})\|_F. \end{aligned} \quad (29)$$

B.1 Proof of Lemma 8

The proof relies on the following representation of the low-rank factors \mathbf{X}, \mathbf{Y} of the nonconvex solution \mathbf{L}_{ncvx} .

Lemma 9. *Under the assumptions and notations of Lemma 7, there exists an invertible matrix $\mathbf{Q} \in \mathbb{R}^{r \times r}$ such that $\mathbf{X} = \mathbf{U}\mathbf{\Sigma}^{1/2}\mathbf{Q}$, $\mathbf{Y} = \mathbf{V}\mathbf{\Sigma}^{1/2}\mathbf{Q}^{-\top}$, $\|\mathbf{Q}\| \leq 2$ and*

$$\left\| \mathbf{\Sigma}^{1/2}\mathbf{Q}\mathbf{Q}^\top\mathbf{\Sigma}^{-1/2} - \mathbf{I}_r \right\| \leq \frac{32\kappa}{\sqrt{\sigma_{\min}}} \|\nabla g(\mathbf{X}, \mathbf{Y})\|_F \leq 1/3. \quad (30)$$

where $\mathbf{U}\mathbf{Q}\mathbf{\Sigma}\mathbf{Q}^\top\mathbf{V}$ is the SVD of \mathbf{Q} .

See Section B.2 for the proof.

Denote the partial gradients of $g(\mathbf{X}, \mathbf{Y})$ as $\mathbf{B}_1, \mathbf{B}_2$, i.e.,

$$\mathbf{B}_1 := \nabla_{\mathbf{X}} g(\mathbf{X}, \mathbf{Y}) = \frac{1}{\theta} \mathcal{P}_\Omega(\mathbf{X}\mathbf{Y}^\top - \mathbf{M})\mathbf{Y} + \lambda\mathbf{X}; \quad (31)$$

$$\mathbf{B}_2 := \nabla_{\mathbf{Y}} g(\mathbf{X}, \mathbf{Y}) = \frac{1}{\theta} \mathcal{P}_\Omega(\mathbf{X}\mathbf{Y}^\top - \mathbf{M})^\top\mathbf{X} + \lambda\mathbf{Y}, \quad (32)$$

where we recall $\theta = \|\mathcal{P}_\Omega(\mathbf{X}\mathbf{Y}^\top - \mathbf{M})\|_F$. By definition, we know that $\max\{\|\mathbf{B}_1\|_F, \|\mathbf{B}_2\|_F\} \leq \|\nabla g(\mathbf{X}, \mathbf{Y})\|_F$.

Let \mathbf{R} be the matrix that is defined by equation (22). We now control its component in T and T^\perp separately.

Part 1: Bounding $\|\mathcal{P}_T(\mathbf{R})\|_F$. By the definition of the projection operator \mathcal{P}_T , we have

$$\begin{aligned} \|\mathcal{P}_T(\mathbf{R})\|_F &= \|\mathbf{U}\mathbf{U}^\top\mathbf{R}(\mathbf{I} - \mathbf{V}\mathbf{V}^\top) + \mathbf{R}\mathbf{V}\mathbf{V}^\top\|_F \\ &\leq \|\mathbf{U}\mathbf{U}^\top\mathbf{R}(\mathbf{I} - \mathbf{V}\mathbf{V}^\top)\|_F + \|\mathbf{R}\mathbf{V}\mathbf{V}^\top\|_F \\ &\leq \|\mathbf{U}^\top\mathbf{R}\|_F + \|\mathbf{R}\mathbf{V}\|_F. \end{aligned}$$

For the term $\mathbf{R}\mathbf{V}$, we use the definitions of \mathbf{B}_1 and \mathbf{R} to see that

$$\lambda\mathbf{U}\mathbf{V}^\top\mathbf{Y} + \lambda\mathbf{R}\mathbf{Y} = \lambda\mathbf{X} - \mathbf{B}_1,$$

which together with the representations in Lemma 9 implies

$$\mathbf{R}\mathbf{V} = \mathbf{U}\mathbf{\Sigma}^{1/2}(\mathbf{Q}\mathbf{Q}^\top - \mathbf{I}_r)\mathbf{\Sigma}^{-1/2} - \mathbf{B}_1\mathbf{Q}^\top\mathbf{\Sigma}^{-1/2}.$$

In view of the relation (30), we have

$$\begin{aligned}\|\mathbf{R}\mathbf{V}\|_{\mathbb{F}} &\leq \|\Sigma^{1/2}(\mathbf{Q}\mathbf{Q}^{\top} - \mathbf{I}_r)\Sigma^{-1/2}\|_{\mathbb{F}} + \|\Sigma^{-1/2}\| \|\mathbf{Q}\| \|\mathbf{B}_1\|_{\mathbb{F}} \\ &\leq \frac{32\kappa}{\sqrt{\sigma_{\min}}} \|\nabla g(\mathbf{X}, \mathbf{Y})\|_{\mathbb{F}} + 2\sqrt{\frac{2}{\sigma_{\min}}} \|\nabla g(\mathbf{X}, \mathbf{Y})\|_{\mathbb{F}} \\ &\leq \frac{35\kappa}{\sqrt{\sigma_{\min}}} \|\nabla g(\mathbf{X}, \mathbf{Y})\|_{\mathbb{F}},\end{aligned}$$

where we have used the fact that $\|\Sigma^{-1}\| \leq \sigma_{\min}/2$. Similarly we can establish that $\|\mathbf{U}^{\top}\mathbf{R}\|_{\mathbb{F}} \leq \frac{35\kappa}{\sqrt{\sigma_{\min}}} \|\nabla g(\mathbf{X}, \mathbf{Y})\|_{\mathbb{F}}$. Combine the two inequalities to arrive at

$$\|\mathcal{P}_T(\mathbf{R})\|_{\mathbb{F}} \leq \frac{70\kappa}{\sqrt{\sigma_{\min}}} \|\nabla g(\mathbf{X}, \mathbf{Y})\|_{\mathbb{F}}.$$

Part 2: Bounding $\|\mathcal{P}_{T^{\perp}}(\mathbf{R})\|$. For any matrix \mathbf{A} , define $\mathcal{P}_{\Omega}^{\text{debias}}(\mathbf{A}) := \mathcal{P}_{\Omega}(\mathbf{A}) - p\mathbf{A}$. We can rewrite the identities (31) and (32) as

$$\begin{aligned}\frac{1}{\theta} [p\mathbf{L}^* + \mathcal{P}_{\Omega}(\mathbf{E}) - \mathcal{P}_{\Omega}^{\text{debias}}(\mathbf{X}\mathbf{Y}^{\top} - \mathbf{L}^*)] \mathbf{Y} &= \frac{p}{\theta} \mathbf{X}\mathbf{Y}^{\top} \mathbf{Y} + \lambda \mathbf{X} - \mathbf{B}_1; \\ \frac{1}{\theta} [p\mathbf{L}^* + \mathcal{P}_{\Omega}(\mathbf{E}) - \mathcal{P}_{\Omega}^{\text{debias}}(\mathbf{X}\mathbf{Y}^{\top} - \mathbf{L}^*)]^{\top} \mathbf{X} &= \frac{p}{\theta} \mathbf{Y}\mathbf{X}^{\top} \mathbf{X} + \lambda \mathbf{Y} - \mathbf{B}_2.\end{aligned}$$

Again, using the representations in Lemma 9, we have the following two identities

$$\frac{1}{\theta} [p\mathbf{L}^* + \mathcal{P}_{\Omega}(\mathbf{E}) - \mathcal{P}_{\Omega}^{\text{debias}}(\mathbf{X}\mathbf{Y}^{\top} - \mathbf{L}^*)] \mathbf{V} = \frac{1}{\theta} p\mathbf{U}\Sigma + \lambda\mathbf{U}\Sigma^{1/2}\mathbf{Q}\mathbf{Q}^{\top}\Sigma^{-1/2} - \mathbf{B}_1\mathbf{Q}^{\top}\Sigma^{-1/2}; \quad (33a)$$

$$\frac{1}{\theta} [p\mathbf{L}^* + \mathcal{P}_{\Omega}(\mathbf{E}) - \mathcal{P}_{\Omega}^{\text{debias}}(\mathbf{X}\mathbf{Y}^{\top} - \mathbf{L}^*)]^{\top} \mathbf{U} = \frac{1}{\theta} p\mathbf{V}\Sigma + \lambda\mathbf{V}\Sigma^{1/2}\mathbf{Q}^{-\top}\mathbf{Q}^{-1}\Sigma^{-1/2} - \mathbf{B}_2\mathbf{Q}^{-1}\Sigma^{-1/2}. \quad (33b)$$

These two equations motivate us to define a matrix $\tilde{\mathbf{R}}$ using

$$\frac{1}{\theta} [p\mathbf{L}^* + \mathcal{P}_{\Omega}(\mathbf{E}) - \mathcal{P}_{\Omega}^{\text{debias}}(\mathbf{X}\mathbf{Y}^{\top} - \mathbf{L}^*)] = \frac{1}{\theta} p\mathbf{U}\Sigma\mathbf{V}^{\top} + \lambda\mathbf{U}\Sigma^{1/2}\mathbf{Q}\mathbf{Q}^{\top}\Sigma^{-1/2}\mathbf{V}^{\top} + \lambda\tilde{\mathbf{R}}, \quad (34)$$

where $\tilde{\mathbf{R}}$ obeys $\mathcal{P}_{T^{\perp}}(\mathbf{R}) = \mathcal{P}_{T^{\perp}}(\tilde{\mathbf{R}})$. To see this, we use the definition of \mathbf{R} to write

$$\mathcal{P}_{T^{\perp}}(\mathbf{R}) = -\frac{1}{\lambda}\mathcal{P}_{T^{\perp}}(\theta^{-1}\mathcal{P}_{\Omega}(\mathbf{X}\mathbf{Y}^{\top} - \mathbf{M})) = -\frac{1}{\lambda\theta}\mathcal{P}_{T^{\perp}}[\mathcal{P}_{\Omega}(\mathbf{X}\mathbf{Y}^{\top} - \mathbf{L}^*) - \mathcal{P}_{\Omega}(\mathbf{E})]. \quad (35)$$

Since $\mathcal{P}_{T^{\perp}}(\mathbf{X}\mathbf{Y}^{\top}) = 0$, by definition of $\tilde{\mathbf{R}}$, we obtain

$$\mathcal{P}_{T^{\perp}}(\mathbf{R}) = \frac{1}{\lambda\theta}\mathcal{P}_{T^{\perp}}[p(\mathbf{L}^* - \mathbf{X}\mathbf{Y}^{\top}) + \mathcal{P}_{\Omega}(\mathbf{E}) - \mathcal{P}_{\Omega}(\mathbf{X}\mathbf{Y}^{\top} - \mathbf{L}^*)] = \mathcal{P}_{T^{\perp}}(\tilde{\mathbf{R}}).$$

Therefore from now on, we concentrate on bounding $\|\mathcal{P}_{T^{\perp}}(\tilde{\mathbf{R}})\|$.

To this end, we rewrite equation (34) as

$$\frac{1}{\theta} [p\mathbf{L}^* + \mathcal{P}_{\Omega}(\mathbf{E}) - \mathcal{P}_{\Omega}^{\text{debias}}(\mathbf{X}\mathbf{Y}^{\top} - \mathbf{L}^*)] - \lambda\mathcal{P}_T(\tilde{\mathbf{R}}) = \frac{1}{\theta} p\mathbf{U}\Sigma\mathbf{V}^{\top} + \lambda\mathbf{U}\Sigma^{1/2}\mathbf{Q}\mathbf{Q}^{\top}\Sigma^{-1/2}\mathbf{V}^{\top} + \lambda\mathcal{P}_{T^{\perp}}(\tilde{\mathbf{R}}).$$

Suppose that

$$\|\mathcal{P}_T(\tilde{\mathbf{R}})\| \leq \frac{\lambda}{4}\theta,$$

which together with Lemma 4 and Lemma 6 implies that

$$\frac{1}{\theta} \left\| \mathcal{P}_{\Omega}(\mathbf{E}) - \mathcal{P}_{\Omega}^{\text{debias}}(\mathbf{X}\mathbf{Y}^{\top} - \mathbf{L}^*) - \lambda\mathcal{P}_T(\tilde{\mathbf{R}}) \right\| \leq \lambda/8 + \lambda/8 + \lambda/4 = \lambda/2.$$

By Weyl's inequality and the fact that \mathbf{L}^* is of rank r , for each $i = r + 1, \dots, n$, one has

$$\sigma_i \left(\frac{1}{\theta} p \mathbf{U} \boldsymbol{\Sigma} \mathbf{V}^\top + \lambda \mathbf{U} \boldsymbol{\Sigma}^{1/2} \mathbf{Q} \mathbf{Q}^\top \boldsymbol{\Sigma}^{-1/2} \mathbf{V}^\top + \lambda \mathcal{P}_{T^\perp}(\tilde{\mathbf{R}}) \right) \quad (36)$$

$$\begin{aligned} &\leq \frac{1}{\theta} \left\| \mathcal{P}_\Omega(\mathbf{E}) - \mathcal{P}_\Omega^{\text{debias}}(\mathbf{X} \mathbf{Y}^\top - \mathbf{L}^*) - \lambda \mathcal{P}_T(\tilde{\mathbf{R}}) \right\| \\ &\leq \lambda/2. \end{aligned} \quad (37)$$

At the same time, for each $i = 1, \dots, r$, we have

$$\begin{aligned} &\sigma_i \left(\frac{1}{\theta} p \mathbf{U} \boldsymbol{\Sigma} \mathbf{V}^\top + \lambda \mathbf{U} \boldsymbol{\Sigma}^{1/2} \mathbf{Q} \mathbf{Q}^\top \boldsymbol{\Sigma}^{-1/2} \mathbf{V}^\top \right) \\ &\geq \sigma_r \left[\mathbf{U} \left(\frac{1}{\theta} p \boldsymbol{\Sigma} + \lambda \mathbf{I}_r + \lambda (\boldsymbol{\Sigma}^{1/2} \mathbf{Q} \mathbf{Q}^\top \boldsymbol{\Sigma}^{-1/2} - \mathbf{I}_r) \right) \mathbf{V}^\top \right] \\ &\geq \sigma_r \left(\frac{1}{\theta} p \boldsymbol{\Sigma} + \lambda \mathbf{I}_r \right) - \lambda \left\| \boldsymbol{\Sigma}^{1/2} \mathbf{Q} \mathbf{Q}^\top \boldsymbol{\Sigma}^{-1/2} - \mathbf{I}_r \right\| \\ &\geq \lambda - \lambda/3 > \lambda/2, \end{aligned} \quad (38)$$

where the last line uses the claim (30). As a result, the singular values of $\lambda \mathcal{P}_{T^\perp}(\tilde{\mathbf{R}})$ must fall below $\lambda/2$, i.e.,

$$\|\mathcal{P}_{T^\perp}(\mathbf{R})\| = \|\mathcal{P}_{T^\perp}(\tilde{\mathbf{R}})\| < \lambda/2.$$

We are left with controlling $\|\mathcal{P}_T(\tilde{\mathbf{R}})\|$. Similar to bounding $\|\mathcal{P}_T(\mathbf{R})\|$, using (33a) and (33b) we have

$$\begin{aligned} \|\tilde{\mathbf{R}} \mathbf{V}\|_{\text{F}} &= \frac{1}{\lambda} \|\mathbf{B}_1 \mathbf{Q}^\top \boldsymbol{\Sigma}^{-1/2} \mathbf{V}\|_{\text{F}} \\ &\leq \frac{1}{\lambda} \|\mathbf{Q}\| \|\boldsymbol{\Sigma}^{-1/2}\| \|\mathbf{B}_1\|_{\text{F}} \\ &\leq \frac{2}{\lambda \sqrt{\sigma_{\min}/2}} \|\nabla g(\mathbf{X}, \mathbf{Y})\|_{\text{F}} \end{aligned}$$

and

$$\begin{aligned} \|\tilde{\mathbf{R}}^\top \mathbf{U}\|_{\text{F}} &= \|\mathbf{V} (\boldsymbol{\Sigma}^{-1/2} \mathbf{Q} \mathbf{Q}^\top \boldsymbol{\Sigma}^{1/2} - \boldsymbol{\Sigma}^{1/2} \mathbf{Q}^{-\top} \mathbf{Q}^{-1} \boldsymbol{\Sigma}^{-1/2}) - \frac{1}{\lambda} \mathbf{B}_2 \mathbf{Q}^\top \boldsymbol{\Sigma}^{-1/2} \mathbf{U}\|_{\text{F}} \\ &\leq \|\boldsymbol{\Sigma}^{-1/2} (\mathbf{Q} \mathbf{Q}^\top - \mathbf{I}_r) \boldsymbol{\Sigma}^{1/2}\|_{\text{F}} + \|\boldsymbol{\Sigma}^{1/2} (\mathbf{Q}^{-\top} \mathbf{Q}^{-1} - \mathbf{I}_r) \boldsymbol{\Sigma}^{-1/2}\|_{\text{F}} + \frac{1}{\lambda} \|\mathbf{B}_2 \mathbf{Q}^\top \boldsymbol{\Sigma}^{-1/2} \mathbf{U}\|_{\text{F}} \\ &\leq \frac{64\kappa}{\sqrt{\sigma_{\min}}} \|\nabla g(\mathbf{X}, \mathbf{Y})\|_{\text{F}} + \frac{2}{\lambda \sqrt{\sigma_{\min}/2}} \|\nabla g(\mathbf{X}, \mathbf{Y})\|_{\text{F}}. \end{aligned}$$

Combining the two bounds we have

$$\begin{aligned} \|\mathcal{P}_T(\tilde{\mathbf{R}})\| &\leq \|\mathcal{P}_T(\tilde{\mathbf{R}})\|_{\text{F}} \leq \|\tilde{\mathbf{R}}^\top \mathbf{U}\|_{\text{F}} + \|\tilde{\mathbf{R}} \mathbf{V}\|_{\text{F}} \\ &\leq \frac{64\kappa + 8/\lambda}{\sqrt{\sigma_{\min}}} \|\nabla g(\mathbf{X}, \mathbf{Y})\|_{\text{F}} \\ &\leq \frac{\lambda\theta}{4}, \end{aligned}$$

where the last line comes from equation (15) and Lemma 6.

B.2 Proof of Lemma 9

Reuse the definitions of $\mathbf{B}_1, \mathbf{B}_2$ in equations (31) and (32). We can then write

$$\mathbf{X}^\top \mathbf{X} - \mathbf{Y}^\top \mathbf{Y} = \frac{1}{\lambda} \left[\mathbf{X}^\top \left(\mathbf{B}_1 - \frac{1}{\theta} \mathcal{P}(\mathbf{X} \mathbf{Y}^\top - \mathbf{M}) \mathbf{Y} \right) - \left(\mathbf{B}_2 - \frac{1}{\theta} \mathcal{P}(\mathbf{X} \mathbf{Y}^\top - \mathbf{M})^\top \mathbf{X} \right)^\top \mathbf{Y} \right]$$

$$= \frac{1}{\lambda} (\mathbf{X}^\top \mathbf{B}_1 - \mathbf{B}_2^\top \mathbf{Y}),$$

which further implies

$$\begin{aligned} \|\mathbf{X}^\top \mathbf{X} - \mathbf{Y}^\top \mathbf{Y}\|_{\text{F}} &= \frac{1}{\lambda} \|\mathbf{X}^\top \mathbf{B}_1 - \mathbf{B}_2^\top \mathbf{Y}\|_{\text{F}} \leq \frac{1}{\lambda} (\|\mathbf{X}\| \|\mathbf{B}_1\|_{\text{F}} + \|\mathbf{B}_2\|_{\text{F}} \|\mathbf{Y}\|) \\ &\leq \frac{2\sqrt{2\sigma_{\max}}}{\lambda} \|\nabla g(\mathbf{X}, \mathbf{Y})\|_{\text{F}}. \end{aligned}$$

Here, the last inequality uses the fact that $\max\{\|\mathbf{B}_1\|_{\text{F}}, \|\mathbf{B}_2\|_{\text{F}}\} \leq \|\nabla g(\mathbf{X}, \mathbf{Y})\|_{\text{F}}$, and that $\max\{\|\mathbf{X}\|, \|\mathbf{Y}\|\} \leq \sqrt{2\sigma_{\max}}$.

In addition, since $\min\{\sigma_{\min}(\mathbf{X}), \sigma_{\min}(\mathbf{Y})\} \geq \sqrt{\sigma_{\min}/2}$, we have $\sigma_{\min}(\mathbf{X}\mathbf{Y}^\top) \geq \sigma_{\min}/2$, which together with Lemma 20 in the paper [CCF⁺20] implies the existence of an invertible $\mathbf{Q} \in \mathbb{R}^{r \times r}$ such that $\mathbf{X} = \mathbf{U}\Sigma^{1/2}\mathbf{Q}$, $\mathbf{Y} = \mathbf{V}\Sigma^{1/2}\mathbf{Q}^{-\top}$, and

$$\begin{aligned} \|\Sigma_{\mathbf{Q}} - \Sigma_{\mathbf{Q}}^{-1}\|_{\text{F}} &\leq \frac{2}{\sigma_{\min}} \|\mathbf{X}^\top \mathbf{X} - \mathbf{Y}^\top \mathbf{Y}\|_{\text{F}} \\ &\leq \frac{4\sqrt{2\sigma_{\max}}}{\lambda\sigma_{\min}} \|\nabla g(\mathbf{X}, \mathbf{Y})\|_{\text{F}} = \frac{4\sqrt{2\kappa}}{\lambda\sqrt{\sigma_{\min}}} \|\nabla g(\mathbf{X}, \mathbf{Y})\|_{\text{F}}. \end{aligned}$$

In view of the assumed upper bound $\|\nabla g(\mathbf{X}, \mathbf{Y})\|_{\text{F}} \leq C_{\text{grad}} \frac{1}{n^8} \sqrt{\frac{\sigma_{\max}}{p}}$ and $n^2 p \gg \kappa$, one has

$$\sigma_{\max}(\Sigma_{\mathbf{Q}}) - \sigma_{\max}^{-1}(\Sigma_{\mathbf{Q}}) \leq \|\Sigma_{\mathbf{Q}} - \Sigma_{\mathbf{Q}}^{-1}\|_{\text{F}} \leq C_{\text{grad}} \frac{4\sqrt{2} \cdot \kappa}{\lambda n^8 \sqrt{p}} \leq 1,$$

and hence $\|\mathbf{Q}\| = \|\Sigma_{\mathbf{Q}}\| = \sigma_{\max}(\Sigma_{\mathbf{Q}}) \leq 2$. As a result, we have

$$\begin{aligned} \left\| \Sigma^{1/2} \mathbf{Q} \mathbf{Q}^\top \Sigma^{-1/2} - \mathbf{I}_r \right\| &= \left\| \Sigma^{1/2} (\mathbf{U}_{\mathbf{Q}} \Sigma_{\mathbf{Q}} \Sigma_{\mathbf{Q}} \mathbf{U}_{\mathbf{Q}}^\top - \mathbf{U}_{\mathbf{Q}} \Sigma_{\mathbf{Q}} \Sigma_{\mathbf{Q}}^{-1} \mathbf{U}_{\mathbf{Q}}^\top) \Sigma^{-1/2} \right\| \\ &\leq \|\Sigma^{1/2}\| \|\Sigma^{-1/2}\| \|\mathbf{U}_{\mathbf{Q}}\| \|\mathbf{U}_{\mathbf{Q}}^\top\| \|\Sigma_{\mathbf{Q}}\| \|\Sigma_{\mathbf{Q}} - \Sigma_{\mathbf{Q}}^{-1}\|_{\text{F}} \\ &\leq \frac{32\kappa}{\sqrt{\sigma_{\min}}} \|\nabla g(\mathbf{X}, \mathbf{Y})\|_{\text{F}} \leq 1/3, \end{aligned}$$

where the last inequality again uses the assumed bound (15).