

Bridging Convex and Nonconvex Optimization in Robust PCA: Noise, Outliers, and Missing Data

Yuxin Chen* Jianqing Fan† Cong Ma† Yuling Yan†

January 2020

Abstract

This paper delivers improved theoretical guarantees for the convex programming approach in low-rank matrix estimation, in the presence of (1) random noise, (2) gross sparse outliers, and (3) missing data. This problem, often dubbed as *robust principal component analysis (robust PCA)*, finds applications in various domains. Despite the wide applicability of convex relaxation, the available statistical support (particularly the stability analysis vis-à-vis random noise) remains highly suboptimal, which we strengthen in this paper. When the unknown matrix is well-conditioned, incoherent, and of constant rank, we demonstrate that a principled convex program achieves near-optimal statistical accuracy, in terms of both the Euclidean loss and the ℓ_∞ loss. All of this happens even when nearly a constant fraction of observations are corrupted by outliers with arbitrary magnitudes. The key analysis idea lies in bridging the convex program in use and an auxiliary nonconvex optimization algorithm, and hence the title of this paper.

Keywords: robust principal component analysis, nonconvex optimization, convex relaxation, ℓ_∞ guarantees, leave-one-out analysis

Contents

1	Introduction	2
1.1	A principled convex programming approach	3
1.2	Theory-practice gaps under random noise	3
1.3	Models, assumptions and notations	5
1.4	Main results	6
1.5	Random signs of outliers?	9
1.6	A peek at our technical approach	10
2	Prior art	12
3	Architecture of the proof	13
3.1	Crude estimation error bounds for convex relaxation	13
3.2	Approximate stationary points of the nonconvex formulation	14
3.3	Constructing an approximate stationary point via nonconvex algorithms	15
3.4	Proof of Theorem 2	16
4	Discussion	18
A	An equivalent probabilistic model of Ω^* used throughout the proof	18

Author names are sorted alphabetically.

*Department of Electrical Engineering, Princeton University, Princeton, NJ 08544, USA; Email: yuxin.chen@princeton.edu.

†Department of Operations Research and Financial Engineering, Princeton University, Princeton, NJ 08544, USA; Email: [jqfan, congma, yulingy@princeton.edu](mailto:{jqfan, congma, yulingy}@princeton.edu).

B Preliminaries	19
B.1 A few preliminary facts	19
B.2 Proof of Lemma 4	20
C Proof of Lemma 2	21
D Crude error bounds (Proof of Theorem 3)	21
D.1 Proof of Lemma 6	25
D.2 Proof of Lemma 7	26
E Equivalence between convex and nonconvex solutions (Proof of Theorem 4)	27
E.1 Preliminary facts	27
E.2 Proof of Theorem 4	28
E.2.1 Step 1: showing that $\mathcal{P}_{\Omega_{\text{obs}}}(\Delta_L + \Delta_S) \approx \mathbf{0}$	28
E.2.2 Step 2: showing that $\mathcal{P}_{T^\perp}(\Delta_L) \approx \mathbf{0}$ and $\mathcal{P}_{(\Omega^*)^c}(\Delta_S) \approx \mathbf{0}$	29
E.2.3 Step 3: controlling the size of Δ_S (and hence that of Δ_L)	30
E.2.4 Proof of Claim 1	31
F Analysis of the nonconvex procedure (Proof of Theorem 5)	32
F.1 Leave-one-out analysis	33
F.2 Key lemmas for establishing the induction hypotheses	35
F.3 Proof of Lemma 10	37
F.4 Proof of Lemma 17	39
F.5 Proof of Lemma 18	41
F.6 Proof of Lemma 19	45
F.7 Proof of Lemma 20	46

1 Introduction

A diverse array of science and engineering applications (e.g. video surveillance, joint shape matching, graph clustering, learning graphical models) involves estimation of low-rank matrices [CLC19, CLMW11, CGH14, JCSX11, CPW12, DR16]. The imperfectness of data acquisition processes, however, presents several common yet critical challenges: (1) random noise: which reflects the uncertainty of the environment and/or the measurement processes; (2) outliers: which represent a sort of corruption that exhibits abnormal behavior; and (3) incomplete data, namely, we might only get to observe a fraction of the matrix entries. Low-rank matrix estimation algorithms aimed at addressing these challenges have been extensively studied under the umbrella of *robust principal component analysis (Robust PCA)* [CSPW11, CLMW11], a terminology popularized by the seminal work [CLMW11].

To formulate the above-mentioned problem in a more precise manner, imagine that we seek to estimate an unknown low-rank matrix $\mathbf{L}^* \in \mathbb{R}^{n \times n}$.¹ What we can obtain is a collection of partially observed and corrupted entries as follows

$$M_{ij} = L_{ij}^* + S_{ij}^* + E_{ij}, \quad (i, j) \in \Omega_{\text{obs}}, \quad (1.1)$$

where $\mathbf{S}^* = [S_{ij}^*]_{1 \leq i, j \leq n}$ is a matrix consisting of outliers, $\mathbf{E} = [E_{ij}]_{1 \leq i, j \leq n}$ represents the random noise, and we only observe entries over an index subset $\Omega_{\text{obs}} \subseteq [n] \times [n]$ with $[n] := \{1, 2, \dots, n\}$. The current paper assumes that \mathbf{S}^* is a relatively sparse matrix whose non-zero entries might have arbitrary magnitudes. This assumption has been commonly adopted in prior work to model gross outliers, while enabling reliable disentanglement of the outlier component and the low-rank component [CSPW11, CLMW11, CJSC13, Li13]. In addition, we suppose that the entries $\{E_{ij}\}$ are independent zero-mean sub-Gaussian random variables, as commonly assumed in the statistics literature to model a large family of random noise. The aim is to reliably estimate \mathbf{L}^* given the grossly corrupted and possibly incomplete data (1.1). Ideally, this task should be accomplished without knowing the locations and magnitudes of the outliers \mathbf{S}^* .

¹To avoid cluttered notation, this paper works with square matrices of size n by n . Our results and analysis can be extended to accommodate rectangular matrices.

1.1 A principled convex programming approach

Focusing on the noiseless case with $\mathbf{E} = \mathbf{0}$, the papers [CSPW11, CLMW11] delivered a positive and somewhat surprising message: both the low-rank component \mathbf{L}^* and the sparse component \mathbf{S}^* can be efficiently recovered with absolutely no error by means of a principled convex program

$$\underset{\mathbf{L}, \mathbf{S} \in \mathbb{R}^{n \times n}}{\text{minimize}} \quad \|\mathbf{L}\|_* + \tau \|\mathbf{S}\|_1 \quad \text{subject to} \quad \mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{L} + \mathbf{S} - \mathbf{M}) = \mathbf{0}, \quad (1.2)$$

provided that certain “separation” and “incoherence” conditions on $(\mathbf{L}^*, \mathbf{S}^*, \Omega_{\text{obs}})$ hold² and that the regularization parameter τ is properly chosen. Here, $\|\mathbf{L}\|_*$ denotes the nuclear norm (i.e. the sum of the singular values) of \mathbf{L} , $\|\mathbf{S}\|_1 = \sum_{i,j} |S_{ij}|$ denotes the usual entrywise ℓ_1 norm, and $\mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{M})$ denotes the Euclidean projection of a matrix \mathbf{M} onto the subspace of matrices supported on Ω_{obs} . Given that the nuclear norm $\|\cdot\|_*$ (resp. the ℓ_1 norm $\|\cdot\|_1$) is the convex relaxation of the rank function $\text{rank}(\cdot)$ (resp. the ℓ_0 counting norm $\|\cdot\|_0$), the rationale behind (1.2) is rather clear: it seeks a decomposition (\mathbf{L}, \mathbf{S}) of \mathbf{M} by promoting the low-rank structure of \mathbf{L} as well as the sparsity structure of \mathbf{S} .

Moving on to the more realistic noisy setting, a natural strategy is to solve the following regularized least-squares problem

$$\underset{\mathbf{L}, \mathbf{S} \in \mathbb{R}^{n \times n}}{\text{minimize}} \quad \frac{1}{2} \|\mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{L} + \mathbf{S} - \mathbf{M})\|_{\text{F}}^2 + \lambda \|\mathbf{L}\|_* + \tau \|\mathbf{S}\|_1. \quad (1.3)$$

With the regularization parameters $\lambda, \tau > 0$ properly chosen, one hopes to strike a balance between enhancing the goodness of fit (by enforcing $\mathbf{L} + \mathbf{S} - \mathbf{M}$ to be small) and promoting the desired low-complexity structures (by regularizing both the nuclear norm of \mathbf{L} and the ℓ_1 norm of \mathbf{S}). A natural and important question comes into our mind:

Where does the algorithm (1.3) stand in terms of its statistical performance vis-à-vis random noise, sparse outliers and missing data?

Unfortunately, however simple this program (1.3) might seem, the existing theoretical support remains far from satisfactory, as we shall discuss momentarily.

1.2 Theory-practice gaps under random noise

To assess the tightness of prior statistical guarantees for (1.3), we find it convenient to first look at a simple setting where (i) \mathbf{E} consists of independent Gaussian components, namely, $E_{ij} \sim \mathcal{N}(0, \sigma^2)$, and (ii) there is no missing data. This simple scenario is sufficient to illustrate the sub-optimality of prior theory.

Prior statistical guarantees. The paper [ZLW⁺10] was the first to derive a sort of statistical performance guarantees for the above convex program. Under mild conditions, [ZLW⁺10] demonstrated that any minimizer $(\hat{\mathbf{L}}, \hat{\mathbf{S}})$ of (1.3) achieves³

$$\|\hat{\mathbf{L}} - \mathbf{L}^*\|_{\text{F}} = O(n \|\mathbf{E}\|_{\text{F}}) = O(\sigma n^2) \quad (1.4)$$

with high probability, where we have substituted in the well-known high-probability bound $\|\mathbf{E}\|_{\text{F}} = O(\sigma n)$ under i.i.d. Gaussian noise. While this theory corroborates the potential stability of convex relaxation against both additive noise and sparse outliers, it remains unclear whether the estimation error bound (1.4) reflects the true performance of the convex program in use. In what follows, we shall compare it with an oracle error bound and collect some numerical evidence.

²Clearly, if the low-rank matrix \mathbf{L}^* is also sparse, one cannot possibly separate \mathbf{S}^* from \mathbf{L}^* . The same holds true if the matrix \mathbf{S}^* is simultaneously sparse and low-rank.

³Mathematically, [ZLW⁺10] investigated an equivalent constrained form of (1.3) and developed an upper bound on the corresponding estimation error.




Figure 1: (a) Euclidean estimation errors of (1.3) and (1.5) vs. the problem size \sqrt{n} , where we fix $r = 5, \sigma = 10^{-3}$; (b) Euclidean estimation errors of (1.3) and (1.5) vs. the noise level σ in a log-log plot, where we fix $n = 1000, r = 5$. For both plots, we take $\lambda = 5\sigma\sqrt{n}$ and $\tau = 2\sigma\sqrt{\log n}$. The results are averaged over 50 independent trials.

Comparisons with an oracle bound. Consider an idealistic scenario where an oracle informs us of the outlier matrix \mathbf{S}^* . With the assistance of this oracle, the task of estimating \mathbf{L}^* reduces to a low-rank matrix denoising problem [DG14]. By fixing \mathbf{S} to be \mathbf{S}^* in (1.3), we arrive at a simplified convex program

$$\underset{\mathbf{L} \in \mathbb{R}^{n \times n}}{\text{minimize}} \quad \frac{1}{2} \|\mathbf{L} - (\mathbf{L}^* + \mathbf{E})\|_F^2 + \lambda \|\mathbf{L}\|_* . \quad (1.5)$$

It is known that (e.g. [DG14, CCF⁺19]): under mild conditions and with a properly chosen λ , the estimation error of (1.5) satisfies

$$\|\hat{\mathbf{L}} - \mathbf{L}^*\|_F = O(\sigma\sqrt{nr}) , \quad (1.6)$$

where we abuse the notation and denote by $\hat{\mathbf{L}}$ the minimizer of (1.5). The large gap between the above two bounds (1.4) and (1.6) is self-evident; in particular, if $r = O(1)$, the gap between these two bounds can be as large as an order of $n^{1.5}$.

A numerical example without oracles. One might naturally wonder whether the discrepancy between the two bounds (1.4) and (1.6) stems from the magical oracle information (i.e. \mathbf{S}^*) which (1.3) does not have the luxury to know. To demonstrate that this is not the case, we conduct some numerical experiments to assess the importance of such oracle information. Generate $\mathbf{L}^* = \mathbf{X}^*\mathbf{Y}^{*\top}$, where $\mathbf{X}^*, \mathbf{Y}^* \in \mathbb{R}^{n \times r}$ are random orthonormal matrices. Each entry of \mathbf{S}^* is generated independently from a mixed distribution: with probability 1/10, the entry is drawn from $\mathcal{N}(0, 10)$; otherwise, it is set to be zero. In other words, approximately 10% of the entries in \mathbf{L}^* are corrupted by large outliers. Throughout the experiments, we set $\lambda = 5\sigma\sqrt{n}$ and $\tau = 2\sigma\sqrt{\log n}$ with σ the standard deviation of each noise entry $\{E_{ij}\}$. Figure 1(a) fixes $r = 5, \sigma = 10^{-3}$ and examines the dependency of the Euclidean error $\|\hat{\mathbf{L}} - \mathbf{L}^*\|_F$ on the size \sqrt{n} of the matrix \mathbf{L}^* . Similarly, Figure 1(b) fixes $r = 5, n = 1000$ and displays the estimation error $\|\hat{\mathbf{L}} - \mathbf{L}^*\|_F$ as the noise size σ varies in a log-log plot. As can be seen from Figure 1, the performance of the oracle-aided estimator (1.5) matches the theoretical prediction (1.6), namely, the numerical estimation error $\|\hat{\mathbf{L}} - \mathbf{L}^*\|_F$ is proportional to both \sqrt{n} and σ . Perhaps more intriguingly, even without the help of the oracle, the original estimator (1.3) performs quite well and behaves qualitatively similarly. In comparison with the bound (1.4) derived in the prior work [ZLW⁺10], our numerical experiments suggest that the convex estimator (1.3) might perform much better than previously predicted.

All in all, there seems to be a large gap between the practical performance of (1.3) and the existing theoretical support. This calls for a new theory that better explains practice, which we pursue in the current paper. We remark in passing that statistical guarantees have been developed in [ANW12, KLT17] for other convex

estimators (i.e. the ones that are different from the convex estimator (1.3) considered herein). We shall compare our results with theirs later in Section 1.4.

1.3 Models, assumptions and notations

As it turns out, the appealing empirical performance of the convex program (1.3) in the presence of both sparse outliers and zero-mean random noise can be justified in theory. Towards this end, we need to introduce several notations and model assumptions that will be used throughout. Let $\mathbf{U}^* \boldsymbol{\Sigma}^* \mathbf{V}^{*\top}$ be the singular value decomposition (SVD) of the unknown rank- r matrix $\mathbf{L}^* \in \mathbb{R}^{n \times n}$, where $\mathbf{U}^*, \mathbf{V}^* \in \mathbb{R}^{n \times r}$ consist of orthonormal columns and $\boldsymbol{\Sigma}^* = \text{diag}\{\sigma_1^*, \dots, \sigma_r^*\}$ is a diagonal matrix. Here, we let

$$\sigma_{\max} := \sigma_1^* \geq \sigma_2^* \geq \dots \geq \sigma_r^* =: \sigma_{\min} \quad \text{and} \quad \kappa := \sigma_{\max} / \sigma_{\min}$$

represent the singular values and the condition number of \mathbf{L}^* , respectively. We denote by Ω^* the support set of \mathbf{S}^* , that is,

$$\Omega^* := \{(i, j) \in \Omega_{\text{obs}} : S_{ij}^* \neq 0\}. \quad (1.7)$$

With these notations in place, we list below our key model assumptions.

Assumption 1 (Incoherence). *The low-rank matrix \mathbf{L}^* with SVD $\mathbf{L}^* = \mathbf{U}^* \boldsymbol{\Sigma}^* \mathbf{V}^{*\top}$ is assumed to be μ -incoherent in the sense that*

$$\|\mathbf{U}^*\|_{2,\infty} \leq \sqrt{\frac{\mu}{n}} \|\mathbf{U}^*\|_{\text{F}} = \sqrt{\frac{\mu r}{n}} \quad \text{and} \quad \|\mathbf{V}^*\|_{2,\infty} \leq \sqrt{\frac{\mu}{n}} \|\mathbf{V}^*\|_{\text{F}} = \sqrt{\frac{\mu r}{n}}. \quad (1.8)$$

Here, $\|\mathbf{U}\|_{2,\infty}$ denotes the largest ℓ_2 norm of all rows of a matrix \mathbf{U} .

Assumption 2 (Random sampling). *Each entry is observed independently with probability p , namely,*

$$\mathbb{P}\{(i, j) \in \Omega_{\text{obs}}\} = p. \quad (1.9)$$

Assumption 3 (Random locations of outliers). *Each observed entry is independently corrupted by an outlier with probability ρ_s , namely,*

$$\mathbb{P}\{(i, j) \in \Omega^* \mid (i, j) \in \Omega_{\text{obs}}\} = \rho_s, \quad (1.10)$$

where $\Omega^* \subseteq \Omega_{\text{obs}}$ is the support of the outlier matrix \mathbf{S}^* .

Assumption 4 (Random signs of outliers). *The signs of the nonzero entries of \mathbf{S}^* are i.i.d. symmetric Bernoulli random variables (independent from the locations), namely,*

$$\text{sign}(S_{ij}^*) \stackrel{\text{ind.}}{=} \begin{cases} 1, & \text{with probability } 1/2, \\ -1, & \text{else,} \end{cases} \quad \text{for all } (i, j) \in \Omega^*. \quad (1.11)$$

Assumption 5 (Random noise). *The noise matrix $\mathbf{E} = [E_{ij}]_{1 \leq i, j \leq n}$ is composed of independent symmetric⁴ zero-mean sub-Gaussian random variables with sub-Gaussian norm at most $\sigma > 0$, i.e. $\|E_{ij}\|_{\psi_2} \leq \sigma$ (see [Ver12, Definition 5.7] for precise definitions).*

We take a moment to expand on our model assumptions. Assumption 1 is standard in the low-rank matrix recovery literature [CR09, CLMW11, Che15, CLC19]. If μ is small, then this assumption specifies that the singular spaces of \mathbf{L}^* is not sparse in the standard basis, thus ensuring that \mathbf{L}^* is not simultaneously low-rank and sparse. Assumption 3 requires the sparsity pattern of the outliers \mathbf{S}^* to be random, which precludes it from being simultaneously sparse and low-rank. In essence, Assumptions 1 and 3 taken together serve as a sort of separation condition on $(\mathbf{L}^*, \mathbf{S}^*)$, which plays a crucial role in guaranteeing exact recovery in the noiseless case (i.e. $\mathbf{E} = \mathbf{0}$); see [CLMW11] for more discussions on these conditions. Assumption 4

⁴In fact, we only require E_{ij} to be symmetric for all $(i, j) \in \Omega^*$.

requires the signs of the outliers to be random, which has also been made in [ZLW⁺10, WL17].⁵ We shall discuss in detail the crucial role of this random sign assumption (as opposed to deterministic sign patterns) in Section 1.5.

Finally, we introduce a few notations that are useful throughout. Denote by $f(n) \lesssim g(n)$ or $f(n) = O(g(n))$ the condition $|f(n)| \leq Cg(n)$ for some constant $C > 0$ when n is sufficiently large; we use $f(n) \gtrsim g(n)$ to denote $f(n) \geq C|g(n)|$ for some constant $C > 0$ when n is sufficiently large; we also use $f(n) \asymp g(n)$ to indicate that $f(n) \lesssim g(n)$ and $f(n) \gtrsim g(n)$ hold simultaneously. The notation $f(n) \gg g(n)$ (resp. $f(n) \ll g(n)$) means that there exists a sufficiently large (resp. small) constant $c_1 > 0$ (resp. $c_2 > 0$) such that $f(n) \geq c_1g(n)$ (resp. $f(n) \leq c_2g(n)$). For any subspace T , we denote by $\mathcal{P}_T(\mathbf{M})$ the Euclidean projection of a matrix \mathbf{M} onto the subspace T , and let $\mathcal{P}_{T^\perp}(\mathbf{M}) := \mathbf{M} - \mathcal{P}_T(\mathbf{M})$. For any index set Ω , we denote by $\mathcal{P}_\Omega(\mathbf{M})$ the Euclidean projection of a matrix \mathbf{M} onto the subspace of matrices supported on Ω , and define $\mathcal{P}_{\Omega^c}(\mathbf{M}) := \mathbf{M} - \mathcal{P}_\Omega(\mathbf{M})$. For any matrix \mathbf{M} , we let $\|\mathbf{M}\|$, $\|\mathbf{M}\|_F$, $\|\mathbf{M}\|_*$, $\|\mathbf{M}\|_1$ and $\|\mathbf{M}\|_\infty$ denote its spectral norm, Frobenius norm, nuclear norm, entrywise ℓ_1 norm, and entrywise ℓ_∞ norm, respectively.

1.4 Main results

Armed with the above model assumptions, we are positioned to present our improved statistical guarantees for convex relaxation (1.3) in the random noise setting. As we shall elucidate in Sections 1.6 and Sections 3, our theory is established by exploiting an intriguing and intimate connection between convex relaxation and nonconvex optimization, and hence the title of this paper.

For the sake of simplicity, we shall start by presenting our statistical guarantees when the rank r , the condition number κ and the incoherence parameter μ of \mathbf{L}^* are all bounded by some constants. Despite its simplicity, this setting subsumes as special cases a wide array of fundamentally important applications, including angular and phase synchronization [Sin11] in computational biology, joint shape mapping problem [HG13, CGH14] in computer vision, and so on. All of these problems involve estimating a very well-conditioned matrix \mathbf{L}^* with a small rank.

Theorem 1. *Suppose that Assumptions 1-5 hold, and that $r, \kappa, \mu = O(1)$. Take $\lambda = C_\lambda \sigma \sqrt{np}$ and $\tau = C_\tau \sigma \sqrt{\log n}$ in (1.3) for some large enough constants $C_\lambda, C_\tau > 0$. Define*

$$\delta_n := \frac{\sigma}{\sigma_{\min}} \sqrt{\frac{n}{p}}. \quad (1.12)$$

Further assume that

$$n^2 p \geq C_{\text{sample}} n \log^6 n, \quad \delta_n \leq \frac{c_{\text{noise}}}{\sqrt{\log n}} \quad \text{and} \quad \rho_s \leq \frac{c_{\text{outlier}}}{\log n} \quad (1.13)$$

for some sufficiently large constant $C_{\text{sample}} > 0$ and some sufficiently small constants $c_{\text{noise}}, c_{\text{outlier}} > 0$. Then with probability exceeding $1 - O(n^{-3})$, the following holds:

1. Any minimizer $(\mathbf{L}_{\text{cvx}}, \mathbf{S}_{\text{cvx}})$ of the convex program (1.3) obeys

$$\|\mathbf{L}_{\text{cvx}} - \mathbf{L}^*\|_F \lesssim \delta_n \|\mathbf{L}^*\|_F, \quad (1.14a)$$

$$\|\mathbf{L}_{\text{cvx}} - \mathbf{L}^*\|_\infty \lesssim \delta_n \sqrt{\log n} \|\mathbf{L}^*\|_\infty, \quad (1.14b)$$

$$\|\mathbf{L}_{\text{cvx}} - \mathbf{L}^*\| \lesssim \delta_n \|\mathbf{L}^*\|. \quad (1.14c)$$

2. Letting $\mathbf{L}_{\text{cvx},r} := \arg \min_{\mathbf{L}: \text{rank}(\mathbf{L}) \leq r} \|\mathbf{L} - \mathbf{L}_{\text{cvx}}\|_F$ be the best rank- r approximation of \mathbf{L}_{cvx} , we have

$$\|\mathbf{L}_{\text{cvx},r} - \mathbf{L}_{\text{cvx}}\|_F \leq \frac{1}{n^5} \delta_n \|\mathbf{L}^*\|_F, \quad (1.15)$$

and the statistical guarantees (1.14) hold unchanged if \mathbf{L}_{cvx} is replaced by $\mathbf{L}_{\text{cvx},r}$.

⁵Note that while the theorems in [ZLW⁺10, WL17] do not make explicit this random sign assumption, the proofs therein do rely on this assumption to guarantee the existence of certain approximate dual certificates.

Table 1: Comparison of our statistical guarantee and prior theory when $\kappa, \mu, r = O(1)$.

	Euclidean estimation error	Accounting for missing data
[ZLW ⁺ 10]	σn^2	no
[ANW12]	$\sigma\sqrt{n} \max\{1, \sqrt{n\rho_s} \log n\} + \ \mathbf{L}^*\ _\infty n\sqrt{\rho_s}$	no
[WL17]	$\sigma n^{1.5}$	yes ($p \gtrsim 1$)
[KLT17]	$\max\{\sigma, \ \mathbf{L}^*\ _\infty, \ \mathbf{S}^*\ _\infty\} \sqrt{(n \log n)/p} \max\{1, \sqrt{np\rho_s}\}$	yes ($p \gtrsim (\text{poly log } n)/n$)
This paper	$\sigma\sqrt{n/p}$	yes ($p \gtrsim (\text{poly log } n)/n$)

Before we embark on interpreting our statistical guarantees, let us first parse the required conditions (1.13) in Theorem 1.

- *Missing data.* Theorem 1 accommodates the case where a dominant fraction of entries are unobserved (more precisely, the sample size can be as low as an order of $n \text{ poly log } n$). This is an appealing result since, even when there is no noise and no outlier (i.e. $\mathbf{E} = \mathbf{0}$ and $\rho_s = 0$), the minimal sample size required for exact matrix completion is at least on the order of $n \log n$ [CT10]. In comparison, prior theory on robust PCA with both sparse outliers and dense additive noise is either based on full observations [ZLW⁺10, ANW12], or assumes the sampling rate p exceeds some universal constant [WL17]. In other words, these prior results require the number of observed entries to exceed the order of n^2 . The only exception is [KLT17], which also allows a significant amount of missing data (i.e. $p \gtrsim (\text{poly log } n)/n$).
- *Noise levels.* The noise condition, namely $\sigma\sqrt{n \log n/p} \lesssim \sigma_{\min}$, accommodates a wide range of noise levels. To see this, it is straightforward to check that this noise condition is equivalent to

$$\sigma \lesssim \sqrt{\frac{np}{\log n}} \|\mathbf{L}^*\|_\infty$$

as long as $r, \mu, \kappa \asymp 1$. In other words, the entrywise noise level σ is allowed to be significantly larger than the maximum magnitude of the entries in the low-rank matrix \mathbf{L}^* , as long as $p \gg (\log n)/n$.

- *Tolerable fraction of outliers.* The above theorem assumes that no more than a fraction $\rho_s \lesssim 1/\log n$ of observations are corrupted by outliers. In words, our theory allows *nearly* a constant proportion (up to a logarithmic order) of the entries of \mathbf{L}^* to be corrupted with arbitrary magnitudes.

Next, we move on to the interpretation of our statistical guarantees.

- *Near-optimal statistical guarantees.* Our first result (1.14a) gives an Euclidean estimation error bound of (1.3)

$$\|\mathbf{L}_{\text{cvx}} - \mathbf{L}^*\|_F \lesssim \sigma\sqrt{\frac{n}{p}}. \quad (1.16)$$

This cannot be improved even when an oracle has informed us of the outliers \mathbf{S}^* and the tangent space of \mathbf{L}^* ; see [CP10, Section III.B]. We remark that under similar model assumptions, the paper [WL17] derived an estimation error bound for the constrained version of the convex program (1.3), which reads

$$\|\mathbf{L}_{\text{cvx}} - \mathbf{L}^*\|_F \lesssim \sigma n^{1.5}. \quad (1.17)$$

Clearly, this is sub-optimal compared to our results. Nevertheless, it is worth pointing out that the bound therein accommodates arbitrary noise matrix \mathbf{E} (e.g. deterministic, adversary), and here in (1.17) we specialize their result to the random noise setting, namely the noise \mathbf{E} obeys Assumption 5. In addition, under the full observation (i.e. $p = 1$) setting, the paper [ANW12] derived an estimation error bound for a convex program similar to (1.3), but with an additional constraint regularizing the spikiness




Figure 2: The relative estimation error of \mathbf{L}_{cvx} measured by both $\|\cdot\|_\infty$ (i.e. $\|\mathbf{L}_{\text{cvx}} - \mathbf{L}^*\|_\infty / \|\mathbf{L}^*\|_\infty$) and $\|\cdot\|$ (i.e. $\|\mathbf{L}_{\text{cvx}} - \mathbf{L}^*\| / \|\mathbf{L}^*\|$) vs. the standard deviation σ of the noise in a log-log plot. The results are reported for $n = 1000$, $r = 5$, $p = 0.2$, $\rho_s = 0.1$, $\lambda = 5\sigma\sqrt{n\rho_s}$, $\tau = 2\sigma\sqrt{\log n}$, and are averaged over 50 independent trials. In addition, the data generating process is similar to that in Figure 1.

of the low-rank component. When $\{E_{ij}\}$ are i.i.d. drawn from $\mathcal{N}(0, \sigma^2)$ and when there is no missing data (i.e. $p = 1$), the Euclidean estimation error bound achievable by their estimator $\mathbf{L}_{\text{cvx}}^{\text{ANW}}$ reads

$$\|\mathbf{L}_{\text{cvx}}^{\text{ANW}} - \mathbf{L}^*\|_{\text{F}} \lesssim \sigma\sqrt{n} \max\left\{1, \sqrt{n\rho_s \log n}\right\} + \|\mathbf{L}^*\|_\infty n\sqrt{\rho_s}, \quad (1.18)$$

which is sub-optimal compared to our results. In particular, (i) the bound (1.18) does not vanish even as the noise level decreases to zero, and (ii) it becomes looser as ρ_s grows (e.g. if $\rho_s \asymp 1/\log n$, the bound (1.18) is $O(\sqrt{n})$ larger than our bound). Moreover, the work [ANW12] did not account for missing data. Similar to [ANW12], the paper [KLT17] derived an estimation error bound for a constrained convex program, with a new constraint regularizing the spikiness of the sparse outliers. Their Euclidean estimation error bound reads

$$\|\mathbf{L}_{\text{cvx}}^{\text{KLT}} - \mathbf{L}^*\|_{\text{F}} \lesssim \max\{\sigma, \|\mathbf{L}^*\|_\infty, \|\mathbf{S}^*\|_\infty\} \sqrt{\frac{n \log n}{p}} \max\{1, \sqrt{n\rho_s}\}, \quad (1.19)$$

which is also sub-optimal compared to our results. In particular, (1) their error bound degrades as the magnitude $\|\mathbf{S}^*\|_\infty$ of the outlier increases; (2) when there is no missing data (i.e. $p = 1$), their bound might be off by a factor as large as $O(\sqrt{n})$. See Table 1 for a summary of our results vs. prior statistical guarantees.

- *Entrywise and spectral norm error control.* Moving beyond Euclidean estimation errors, our theory also provides statistical guarantees measured by two other important metrics: the entrywise ℓ_∞ norm (cf. (1.14b)) and the spectral norm (cf. (1.14c)). In particular, our entrywise error bound (1.14b) in reads

$$\|\mathbf{L}_{\text{cvx}} - \mathbf{L}^*\|_\infty \lesssim \sigma \sqrt{\frac{\log n}{np}} \quad (1.20)$$

as long as $r, \kappa, \mu \asymp 1$, which is about $O(n)$ times smaller than the Euclidean loss (1.16) modulo some logarithmic factor. This uncovers an appealing “delocalization” behavior of the estimation errors, namely, the estimation errors of \mathbf{L}^* are fairly spread out across all entries. This can also be viewed as an “implicit regularization” phenomenon: the convex program automatically controls the spikiness of the low-rank solution, without the need of explicitly regularizing it (e.g. adding a constraint $\|\mathbf{L}\|_\infty \leq \alpha$ as adopted in the prior work [ANW12, KLT17]). See Figure 2 for the numerical evidence for the relative entrywise and spectral norm error of \mathbf{L}_{cvx} .

- *Approximate low-rank structure of the convex estimator \mathbf{L}_{cvx} .* Last but not least, Theorem 1 ensures that the convex estimate \mathbf{L}_{cvx} is nearly rank- r , so that a rank- r approximation of \mathbf{L}_{cvx} is extremely accurate.

In other words, the convex program automatically adapts to the true rank of \mathbf{L}^* without having any prior knowledge about r . As we shall see shortly, this is a crucial observation underlying the intimate connection between convex relaxation and a certain nonconvex approach.

Moving beyond the setting with $r, \kappa, \mu \asymp 1$, we have developed theoretical guarantees that allow r, κ, μ to grow with the problem dimension n . The result is this.

Theorem 2. *Suppose that Assumptions 1-5 hold. Take $\lambda = C_\lambda \sigma \sqrt{np}$ and $\tau = C_\tau \sigma \sqrt{\log n}$ in (1.3) for some large enough constants $C_\lambda, C_\tau > 0$. Define $\delta_n = (\sigma/\sigma_{\min}) \cdot \sqrt{n/p}$, and further assume that*

$$n^2 p \geq C_{\text{sample}} \kappa^4 \mu^2 r^2 n \log^6 n, \quad \delta_n \leq \frac{c_{\text{noise}}}{\sqrt{\kappa^4 \mu r \log n}}, \quad \text{and} \quad \rho_s \leq \frac{c_{\text{outlier}}}{\kappa^3 \mu r \log n} \quad (1.21)$$

for some sufficiently large constant $C_{\text{sample}} > 0$ and some sufficiently small constants $c_{\text{noise}}, c_{\text{outlier}} > 0$. Then with probability exceeding $1 - O(n^{-3})$, the following holds:

1. Any minimizer $(\mathbf{L}_{\text{cvx}}, \mathbf{S}_{\text{cvx}})$ of the convex program (1.3) obeys

$$\|\mathbf{L}_{\text{cvx}} - \mathbf{L}^*\|_{\text{F}} \lesssim \delta_n \kappa \|\mathbf{L}^*\|_{\text{F}}, \quad (1.22a)$$

$$\|\mathbf{L}_{\text{cvx}} - \mathbf{L}^*\|_{\infty} \lesssim \delta_n \sqrt{\kappa^3 \mu r \log n} \|\mathbf{L}^*\|_{\infty}, \quad (1.22b)$$

$$\|\mathbf{L}_{\text{cvx}} - \mathbf{L}^*\| \lesssim \delta_n \|\mathbf{L}^*\|. \quad (1.22c)$$

2. Letting $\mathbf{L}_{\text{cvx},r} := \arg \min_{\mathbf{L}: \text{rank}(\mathbf{L}) \leq r} \|\mathbf{L} - \mathbf{L}_{\text{cvx}}\|_{\text{F}}$ be the best rank- r approximation of \mathbf{L}_{cvx} , we have

$$\|\mathbf{L}_{\text{cvx},r} - \mathbf{L}_{\text{cvx}}\|_{\text{F}} \leq \frac{1}{n^5} \delta_n \|\mathbf{L}^*\|_{\text{F}}, \quad (1.23)$$

and the statistical guarantees (1.22) hold unchanged if \mathbf{L}_{cvx} is replaced by $\mathbf{L}_{\text{cvx},r}$.

Similar to Theorem 1, our general theory (i.e. Theorem 2) provides the estimation error of the convex estimator \mathbf{L}_{cvx} in three different norms (i.e. the Euclidean, entrywise and operator norms), and reveals the near low-rankness of the convex estimator (cf. (1.23)) as well as the implicit regularization phenomenon (cf. (1.22b)).

Finally, we make note of several aspects of our general theory that call for further improvement. For instance, when there is no missing data, the rank r of the unknown matrix \mathbf{L}^* needs to satisfy $r \lesssim \sqrt{n}$. On the positive side, our result allows r to grow with the problem dimension n . However, prior results in the noiseless case [CLMW11, Li13] allow r to grow almost linearly with n . This unsatisfactory aspect arises from the suboptimal analysis (in terms of the dependency on r) of a tightly related nonconvex estimation algorithm (to be elaborated on later), which, to the best of our knowledge, has not been resolved in the nonconvex low-rank matrix recovery literature [MWCC17, CLL19]. See Section 2 for more discussions about this point. Moreover, when $\mathbf{E} = \mathbf{0}$, it is known that ρ_s can be as large as a constant even when r grows with n [Li13, CJSC13] — a case not covered by our current theory for noisy case.

1.5 Random signs of outliers?

The careful reader might wonder whether it is possible to remove the random sign assumption on \mathbf{S}^* (namely, Assumption 4) without compromising our statistical guarantees. After all, the results of [CSPW11, CLMW11, Li13] derived for the noise-free case do not rely on such a random sign assumption at all.⁶ Unfortunately, removal of such a condition might be problematic in general, as illustrated by the following example.

An example with non-random signs. Suppose that (i) each non-zero entry of \mathbf{S}^* obeys $S_{ij}^* = \sigma$, (ii) $\rho_s = c_0/\log n$ for some sufficiently small constant $c_0 > 0$, and (iii) there is no missing data (i.e. $p = 1$). In such a scenario, the data matrix can be decomposed as

$$\mathbf{M} = \mathbf{L}^* + \mathbf{S}^* + \mathbf{E} = \underbrace{\mathbf{L}^* + \mathbb{E}[\mathbf{S}^*]}_{=: \tilde{\mathbf{L}}^*} + \underbrace{\mathbf{S}^* - \mathbb{E}[\mathbf{S}^*]}_{=: \tilde{\mathbf{E}}} + \mathbf{E}.$$

⁶Notably, in the noisy setting, prior theory [ZLW⁺10, WL17] also implicitly assumes this random sign condition, while [ANW12, KLT17] do not require this condition.

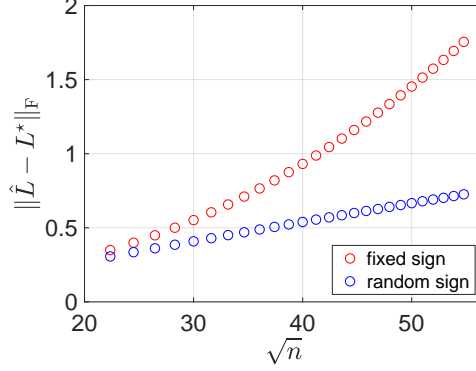


Figure 3: The Euclidean estimation error of (1.3) vs. \sqrt{n} under two different sign patterns of \mathbf{S}^* . The results are reported for $r = 5$, $p = 1$, and $\sigma = 10^{-3}$, with $\lambda = 5\sigma\sqrt{n\bar{p}}$ and $\tau = 2\sigma\sqrt{\log n}$ and are averaged over 50 independent trials. For the random sign setting, the nonzero entries of \mathbf{S}^* are independently generated from $\mathcal{N}(0, 10)$. For the fixed sign setting, each nonzero entry of \mathbf{S}^* is independently generated following the same distribution as $|z|$, where $z \sim \mathcal{N}(0, 10)$.

Two observations are worth noting: (1) given that $\mathbb{E}[\mathbf{S}^*] = \rho_s\sigma\mathbf{1}\mathbf{1}^\top$ with $\mathbf{1}$ the all-one vector, the rank of the matrix $\tilde{\mathbf{L}}^* = \mathbf{L}^* + \mathbb{E}[\mathbf{S}^*]$ is at most $r + 1$; (2) $\tilde{\mathbf{E}}$ is a zero-mean random matrix consisting of independent entries with sub-Gaussian norm $O(\sigma)$. In other words, the decomposition $\mathbf{M} = \tilde{\mathbf{L}}^* + \tilde{\mathbf{E}}$ corresponds to a case with random noise but no outliers. Consequently, we can invoke Theorem 1 to conclude that (assuming $r = O(1)$ and $\tilde{\mathbf{L}}^*$ is incoherent with condition number $O(1)$): any minimizer $(\mathbf{L}_{\text{cvx}}, \mathbf{S}_{\text{cvx}})$ of (1.3) obeys

$$\|\mathbf{L}_{\text{cvx}} - \mathbf{L}^* - \rho_s\sigma\mathbf{1}\mathbf{1}^\top\|_{\text{F}} = \|\mathbf{L}_{\text{cvx}} - \tilde{\mathbf{L}}^*\|_{\text{F}} \lesssim \sigma\sqrt{n}$$

with high probability. This, however, leads to a lower bound on the estimation error

$$\begin{aligned} \|\mathbf{L}_{\text{cvx}} - \mathbf{L}^*\|_{\text{F}} &\geq \|\rho_s\sigma\mathbf{1}\mathbf{1}^\top\|_{\text{F}} - \|\mathbf{L}_{\text{cvx}} - \mathbf{L}^* - \rho_s\sigma\mathbf{1}\mathbf{1}^\top\|_{\text{F}} = \sigma(\rho_s n - O(\sqrt{n})) \\ &= (1 - o(1)) \frac{c_0\sigma n}{\log n}, \end{aligned}$$

which can be $O(\sqrt{n}/\log n)$ times larger than the desired estimation error $O(\sigma\sqrt{n})$. This issue has also been observed in numerical experiments; see Figure 3.

The take-away message is this: when the entries of \mathbf{S}^* are of non-random signs, it might sometimes be possible to decompose \mathbf{S}^* into (1) a low-rank bias component with a large Euclidean norm, and (2) a random fluctuation component whose typical size does not exceed that of \mathbf{E} . If this is the case, then the convex program (1.3) might mistakenly treat the bias component as a part of the low-rank matrix \mathbf{L}^* , thus dramatically hampering its estimation accuracy.

1.6 A peek at our technical approach

Before delving into the proof details, we immediately highlight our key technical ideas and novelties.

Connections between convex and nonconvex optimization. Instead of directly analyzing the convex program (1.3), we turn attention to a seemingly different, but in fact closely related, nonconvex program

$$\underset{\mathbf{X}, \mathbf{Y} \in \mathbb{R}^{n \times r}, \mathbf{S} \in \mathbb{R}^{n \times n}}{\text{minimize}} \quad \frac{1}{2} \|\mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{X}\mathbf{Y}^\top + \mathbf{S} - \mathbf{M})\|_{\text{F}}^2 + \frac{\lambda}{2} (\|\mathbf{X}\|_{\text{F}}^2 + \|\mathbf{Y}\|_{\text{F}}^2) + \tau \|\mathbf{S}\|_1. \quad (1.24)$$

This idea is inspired by an interesting numerical finding (cf. Figure 4) that the solution to the convex program (1.3), and an estimate obtained by attempting to solve the nonconvex formulation (1.24), are

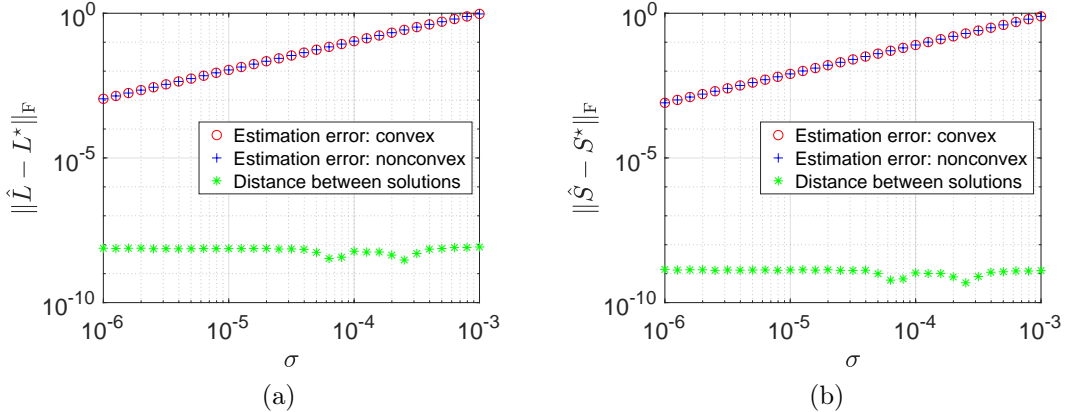


Figure 4: (a) The relative estimation errors of both \mathbf{L}_{cvx} (the convex estimator (1.3)) and \mathbf{L}_{ncvx} (the estimate returned by the nonconvex approach tailored to (1.24)) and the relative distance between them vs. the standard deviation σ of the noise. (b) The relative estimation errors of both \mathbf{S}_{cvx} (the convex estimator in (1.3)) and \mathbf{S}_{ncvx} (the estimate returned by the nonconvex approach tailored to (1.24)) and the relative distance between them vs. the standard deviation σ of the noise. The results are reported for $n = 1000$, $r = 5$, $p = 0.2$, $\rho_s = 0.1$, $\lambda = 5\sigma\sqrt{np}$, $\tau = 2\sigma\sqrt{\log n}$ and are averaged over 50 independent trials.

exceedingly close in our experiments. If such an intimate connection can be formalized, then it suffices to analyze the statistical performance of the nonconvex approach instead.⁷ Fortunately, recent advances in nonconvex low-rank factorization (see [CLC19] for an overview) provide powerful tools for analyzing nonconvex low-rank estimation, allowing us to derive the desired statistical guarantees that can then be transferred to the convex approach. Of course, this is merely a high-level picture of our proof strategy, and we defer the details to Section 3.

It is worth emphasizing that our key idea — that is, bridging convex and nonconvex optimization — is drastically different from previous technical approaches for analyzing convex estimators (e.g. (1.3)). As it turns out, these prior approaches, which include constructing dual certificates and/or exploiting restricted strong convexity, have their own deficiencies in analyzing (1.3) and fall short of explaining the effectiveness of (1.3) in the random noise setting. For instance, constructing dual certificates in the noisy case is notoriously challenging given that we do not have closed-form expressions for the primal solutions (so that it is difficult to invoke the powerful dual construction strategies like the golfing scheme [Gro11] developed for the noiseless case). If we directly utilize the dual certificates constructed for the noiseless case, we would end up with an overly conservative bound like (1.4), which is exactly why the results in [ZLW⁺10] are sub-optimal. On the other hand, while it is viable to show certain strong convexity of (1.3) when restricted to some highly local sets and directions, it is unclear how (1.3) forces its solution to stay within the desired set and follow the desired directions, without adding further (and often unnecessary) constraints to (1.3).

Nonconvex low-rank estimation with nonsmooth loss functions. A similar connection between convex and nonconvex optimization has been pointed out by [CCF⁺19] in understanding the power of convex relaxation for noisy matrix completion. Due to the absence of sparse outliers in the noisy matrix completion problem, the nonconvex loss function considered therein is smooth in nature, thus simplifying both the algorithmic and theoretical development. By contrast, the nonsmoothness inherent in (1.24) makes it particularly challenging to achieve the two desiderata mentioned above, namely, connecting the convex and nonconvex solutions and establishing the optimality of the nonconvex solution. To address this issue, we develop an alternating minimization scheme — which alternates between gradient updates on (\mathbf{X}, \mathbf{Y}) and minimization of \mathbf{S} — aimed at minimizing the nonsmooth nonconvex loss function (1.24); see Algorithm 1

⁷On the surface, the convex program (1.3) and the nonconvex one (1.24) are closely related: the convex solution $(\mathbf{L}_{\text{cvx}}, \mathbf{S}_{\text{cvx}})$ coincides with that of the nonconvex program (1.24) if \mathbf{L}_{cvx} is rank- r . This is an immediate consequence of the algebraic identity $\|\mathbf{Z}\|_* = \inf_{\mathbf{X}, \mathbf{Y} \in \mathbb{R}^{n \times r}: \mathbf{X}\mathbf{Y}^\top = \mathbf{Z}} (\|\mathbf{X}\|_F^2 + \|\mathbf{Y}\|_F^2)$ [SS05, MHT10]. However, it is difficult to know *a priori* the rank of the convex solution. Hence such a connection does not prove useful in establishing the statistical properties of the convex estimator.

for details. As it turns out, such a simple algorithm allows us to track the proximity of the convex and nonconvex solutions and establish the optimality of the nonconvex solution all at once.

2 Prior art

Principal component analysis (PCA) [Pea01, Jol11, FSZZ18] is one of the most widely used statistical methods for dimension reduction in data analysis. However, PCA is known to be quite sensitive to adversarial outliers — even a single corrupted data point can make PCA completely off. This motivated the investigation of robust PCA, which aims at making PCA robust to gross adversarial outliers. As formulated in [CLMW11, CSPW11], this is closely related to the problem of disentangling a low-rank matrix \mathbf{L}^* and a sparse outlier matrix \mathbf{S}^* (with unknown locations and magnitudes) from a superposition of them. Consequently, robust PCA can be viewed as an outlier-robust extension of the low-rank matrix estimation/completion tasks [CR09, KMO10, CLC19]. In a similar vein, robust PCA has also been extensively studied in the context of structured covariance estimation under approximate factor models [FFL08, FLM13, FWZ18, FWZ19], where the population covariance of certain random sample vectors is a mixture of a low-rank matrix and a sparse matrix, corresponding to the factor component and the idiosyncratic component, respectively.

Focusing on the convex relaxation approach, [CSPW11, CLMW11] started by considering the noiseless case with no missing data (i.e. $\mathbf{E} = \mathbf{0}$ and $p = 1$) and demonstrated that, under mild conditions, convex relaxation succeeds in exactly decomposing both \mathbf{L}^* and \mathbf{S}^* from the data matrix $\mathbf{L}^* + \mathbf{S}^*$. More specifically, [CSPW11] adopted a deterministic model without assuming any probabilistic structure on the outlier matrix \mathbf{S}^* . As shown in [CSPW11] and several subsequent work [CJSC13, HKZ11], convex relaxation is guaranteed to work as long as the fraction of outliers in each row/column does not exceed $O(1/r)$. In contrast, [CLMW11] proposed a random model by assuming that \mathbf{S}^* has random support (cf. Assumption 3); under this model, exact recovery is guaranteed even if a constant fraction of the entries of \mathbf{S}^* are nonzero with arbitrary magnitudes. Following the random location model proposed in [CLMW11], the paper [GWL+10] showed that, in the absence of noise, convex programming can provably tolerate a dominant fraction of outliers, provided that the signs of the nonzero entries of \mathbf{S}^* are randomly generated (cf. Assumption 4). Later, the papers [CJSC13, Li13] extended these results to the case when most entries of the matrix are unseen; even in the presence of highly incomplete data, convex relaxation still succeeds when a constant proportion of the observed entries are arbitrarily corrupted. It is worth noting that the results of [CJSC13] accommodated both models proposed in [CSPW11] and [CLMW11], while the results of [Li13] focused on the latter model.

The literature on robust PCA with not only sparse outliers but also dense noise — namely, when the measurements take the form $\mathbf{M} = \mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{L}^* + \mathbf{S}^* + \mathbf{E})$ — is relatively scarce. [ZLW+10, ANW12] were among the first to present a general theory for robust PCA with dense noise, which was further extended in [WL17, KLT17]. As we mentioned before, the first three [ZLW+10, ANW12, WL17] accommodated arbitrary noise with the last one [KLT17] focusing on the random noise. As we have discussed in Section 1.4, the statistical guarantees provided in these papers are highly suboptimal when it comes to the random noise setting considered herein. The paper [CC14] extended the robust PCA results to the case where the truth is not only low-rank but also of Hankel structure. The results therein, however, suffered from the same sub-optimality issue.

Moving beyond convex relaxation methods, another line of work proposed nonconvex approaches for robust PCA [NNS+14, YPCC16, CGJ17, CCD+19, LMCC18, CCW19], largely motivated by the recent success of nonconvex methods in low-rank matrix factorization [CLC19, KMO10, CLS15, SL16, CC17, CW15, ZCL16, CC18, JNS13, NJS13, MWCC17, CCFM19, WGE17, WCCL16, CW18, ZL16, CDDD19]. Following the deterministic model of [CSPW11], the paper [NNS+14] proposed an alternating projection / minimization scheme to seek a low-rank and sparse decomposition of the observed data matrix. In the noiseless setting, i.e. $\mathbf{E} = \mathbf{0}$, this alternating minimization scheme provably disentangles the low-rank and sparse matrix from their superposition under mild conditions. In addition, [NNS+14] extended their result to the arbitrary noise case where the size of the noise is extremely small, namely, $\|\mathbf{E}\|_{\infty} \ll \sigma_{\min}/n$. When the noise $\{E_{ij}\} \sim \mathcal{N}(0, \sigma^2)$, this is equivalent to the condition $\sigma \ll \sigma_{\min}/(n\sqrt{\log n})$. Comparing this with our noise condition $\sigma \ll \sigma_{\min}/(\sqrt{n}\log n)$ (cf. (1.13)) when $r, \mu, \kappa \asymp 1$, one sees that our theoretical guarantees cover a wider range of noise levels. Similarly, [YPCC16] applied regularized gradient descent on a smooth nonconvex loss function which enjoys provable convergence guarantees to $(\mathbf{L}^*, \mathbf{S}^*)$ under the noiseless and partial obser-

vation setting. A recent paper [CCD⁺19] considered the nonsmooth nonconvex formulation for robust PCA and established rigorously the convergence of subgradient-type methods in the rank-1 setting, i.e. $r = 1$. However, the extension to more general rank remains out of reach.

It is worth noting that noisy matrix completion problem [CP10, CCF⁺19] is subsumed as a special case by the model studied in this paper (namely, it is a special case with $\mathbf{S}^* = \mathbf{0}$). Statistical optimality under the random noise setting (cf. Assumption 5) — including the convex relaxation approach [CCF⁺19, NW12, KLT11, Klo14] and the nonconvex approach [MWCC17, CLL19] — has been extensively studied. Focusing on arbitrary deterministic noise, [CP10] established the stability of the convex approach, whose resulting estimation error bound is similar to the one established for robust PCA with noise in [ZLW⁺10]) (see (1.4)). The paper [KS19] later confirmed that the estimation error bound established in [CP10] is the best one can hope for in the arbitrary noise setting for matrix completion, although it might be highly suboptimal if we restrict attention to random noise.

Finally, there is also a large literature considering robust PCA under different settings and/or from different perspectives. For instance, the computational efficiency in solving the convex optimization problem (1.3) and its variants has been studied in the optimization literature (e.g. [TY11, GMS13, SWZ14, MA18]). The problem has also been investigated under a streaming / online setting [GQV14, QV10, FXY13, ZLGV16, QVLH14, VN18]. These are beyond the scope of the current paper.

3 Architecture of the proof

In this section, we give an outline for proving Theorem 2. The proof of Theorem 1 follows immediately as it is a special case of Theorem 2.

The main ingredient of the proof lies in establishing an intimate link between convex and nonconvex optimization. Unless otherwise noted, we shall set the regularization parameters as

$$\lambda = C_\lambda \sigma \sqrt{np} \quad \text{and} \quad \tau = C_\tau \sigma \sqrt{\log n} \quad (3.1)$$

throughout. In addition, the soft thresholding operator at level τ is defined such that

$$\mathcal{S}_\tau(x) := \text{sign}(x) \max(|x| - \tau, 0) = \begin{cases} x - \tau, & \text{if } x > \tau, \\ x + \tau, & \text{if } x < -\tau, \\ 0, & \text{otherwise.} \end{cases} \quad (3.2)$$

For any matrix \mathbf{X} , the matrix $\mathcal{S}_\tau(\mathbf{X})$ is obtained by applying the soft thresholding operator $\mathcal{S}_\tau(\cdot)$ to each entry of \mathbf{X} separately. Additionally, we define the true low-rank factors as follows

$$\mathbf{X}^* := \mathbf{U}^* (\boldsymbol{\Sigma}^*)^{1/2} \quad \text{and} \quad \mathbf{Y}^* := \mathbf{V}^* (\boldsymbol{\Sigma}^*)^{1/2}, \quad (3.3)$$

where $\mathbf{U}^* \boldsymbol{\Sigma}^* \mathbf{V}^{*\top}$ is the SVD of the true low-rank matrix \mathbf{L}^* .

3.1 Crude estimation error bounds for convex relaxation

We start by delivering a crude upper bound on the Euclidean estimation error, built upon the (approximate) duality certificate previously constructed in [CJSC13]. The proof is postponed to Appendix D.

Theorem 3. *Consider any given $\lambda > 0$ and set $\tau \asymp \lambda \sqrt{(\log n)/np}$. Suppose that Assumptions 1-4 hold, and that*

$$n^2 p \geq C \mu^2 r^2 n \log^6 n \quad \text{and} \quad \rho_s \leq c$$

hold for some sufficiently large (resp. small) constant $C > 0$ (resp. $c > 0$). Then with probability at least $1 - O(n^{-10})$, any minimizer $(\mathbf{L}_{\text{cvx}}, \mathbf{S}_{\text{cvx}})$ of the convex program (1.3) satisfies

$$\|\mathbf{L}_{\text{cvx}} - \mathbf{L}^*\|_{\text{F}}^2 + \|\mathbf{S}_{\text{cvx}} - \mathbf{S}^*\|_{\text{F}}^2 \lesssim \lambda^2 n^5 \log n + \frac{n}{\lambda^2} \|\mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{E})\|_{\text{F}}^4. \quad (3.4)$$

It is worth noting that the above theorem holds true for an arbitrary noise matrix \mathbf{E} . When specialized to the case with independent sub-Gaussian noise, this crude bound admits a simpler expression as follows.

Corollary 1. Take $\lambda = C_\lambda \sigma \sqrt{np}$ and $\tau = C_\tau \sigma \sqrt{\log n}$ for some universal constant $C_\lambda, C_\tau > 0$. Under the assumptions of Theorem 3 and Assumption 5, we have — with probability exceeding $1 - O(n^{-10})$ — that

$$\|\mathbf{L}_{\text{cvx}} - \mathbf{L}^*\|_{\text{F}} \lesssim \sigma n^3 \sqrt{\log n} \quad \text{and} \quad \|\mathbf{S}_{\text{cvx}} - \mathbf{S}^*\|_{\text{F}} \lesssim \sigma n^3 \sqrt{\log n}. \quad (3.5)$$

Proof. This corollary follows immediately by combining Theorem 3 and Lemma 1 below. \square

Lemma 1. Suppose that Assumption 5 holds and that $n^2 p > C_1 n \log^2 n$ for some sufficiently large constant $C_1 > 0$. Then with probability exceeding $1 - O(n^{-10})$, one has

$$\|\mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{E})\| \lesssim \sigma \sqrt{np} \quad \text{and} \quad \|\mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{E})\|_{\text{F}} \lesssim \sigma n \sqrt{p}.$$

While the above results often lose a polynomial factor in n vis-à-vis the optimal error bound, it serves as an important starting point that paves the way for subsequent analytical refinement.

3.2 Approximate stationary points of the nonconvex formulation

Instead of analyzing the convex estimator directly, we take a detour by considering the following nonconvex optimization problem

$$\underset{\mathbf{X}, \mathbf{Y} \in \mathbb{R}^{n \times r}, \mathbf{S} \in \mathbb{R}^{n \times n}}{\text{minimize}} \quad F(\mathbf{X}, \mathbf{Y}, \mathbf{S}) := \underbrace{\frac{1}{2p} \|\mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{X}\mathbf{Y}^\top + \mathbf{S} - \mathbf{M})\|_{\text{F}}^2 + \frac{\lambda}{2p} \|\mathbf{X}\|_{\text{F}}^2 + \frac{\lambda}{2p} \|\mathbf{Y}\|_{\text{F}}^2 + \frac{\tau}{p} \|\mathbf{S}\|_1}_{=: f(\mathbf{X}, \mathbf{Y}; \mathbf{S})}. \quad (3.6)$$

Here, $f(\mathbf{X}, \mathbf{Y}; \mathbf{S})$ is a function of \mathbf{X} and \mathbf{Y} with \mathbf{S} frozen, which contains the smooth component of the loss function $F(\mathbf{X}, \mathbf{Y}, \mathbf{S})$. As it turns out, the solution to convex relaxation (1.3) is exceedingly close to an estimate obtained by a nonconvex algorithm aimed at solving (3.6). This fundamental connection between the two algorithmic paradigms provides a powerful framework that allows us to understand convex relaxation by studying nonconvex optimization.

In what follows, we set out to develop the above-mentioned intimate connection. Before proceeding, we first state the following conditions concerned with the interplay between the noise size, the estimation accuracy, and the regularization parameters.

Condition 1. The regularization parameters λ and $\tau \asymp \lambda \sqrt{(\log n)/np}$ satisfy

- $\|\mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{E})\| < \lambda/16$ and $\|\mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{E})\|_{\infty} \leq \tau/4$;
- $\|\mathbf{S} - \mathbf{S}^*\| < \lambda/16$ and $\|\mathbf{X}\mathbf{Y}^\top - \mathbf{L}^*\|_{\infty} \leq \tau/4$;
- $\|\mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{X}\mathbf{Y}^\top - \mathbf{L}^*) - p(\mathbf{X}\mathbf{Y}^\top - \mathbf{L}^*)\| < \lambda/8$.

As an interpretation, the above condition says that: (1) the regularization parameters are not too small compared to the size of the noise, so as to ensure that we enforce a sufficiently large degree of regularization; (2) the estimate represented by the point $(\mathbf{X}\mathbf{Y}^\top, \mathbf{S})$ is sufficiently close to the truth. At this point, whether this condition is meaningful or not remains far from clear; we shall return to justify its feasibility shortly.

In addition, we need another condition concerning the injectivity of \mathcal{P}_{Ω^*} w.r.t. a certain tangent space. Again, the validity of this condition will be discussed momentarily.

Condition 2 (Injectivity). Let T be the tangent space of the set of rank- r matrices at the point $\mathbf{X}\mathbf{Y}^\top$. Assume that there exist constants $c_{\text{inj}} > 0$ such that for all $\mathbf{H} \in T$, one has

$$p^{-1} \|\mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{H})\|_{\text{F}}^2 \geq \frac{c_{\text{inj}}}{\kappa} \|\mathbf{H}\|_{\text{F}}^2 \quad \text{and} \quad p^{-1} \|\mathcal{P}_{\Omega^*}(\mathbf{H})\|_{\text{F}}^2 \leq \frac{c_{\text{inj}}}{4\kappa} \|\mathbf{H}\|_{\text{F}}^2.$$

With the above conditions in place, we are ready to make precise the intimate link between convex relaxation and a candidate nonconvex solution. The proof is deferred to Appendix E.

Theorem 4. Suppose that $n \geq \kappa$ and $\rho_s \leq c/\kappa$ for some sufficiently small constant $c > 0$. Assume that there exists a triple $(\mathbf{X}, \mathbf{Y}, \mathbf{S})$ such that

$$\|\nabla f(\mathbf{X}, \mathbf{Y}; \mathbf{S})\|_{\text{F}} \leq \frac{1}{n^{20}} \frac{\lambda}{p} \sqrt{\sigma_{\min}}, \quad \text{and} \quad \mathbf{S} = \mathcal{P}_{\Omega_{\text{obs}}}(\mathcal{S}_{\tau}(\mathbf{M} - \mathbf{X}\mathbf{Y}^{\top})). \quad (3.7)$$

Further, assume that any singular value of \mathbf{X} and \mathbf{Y} lies in $[\sqrt{\sigma_{\min}/2}, \sqrt{2\sigma_{\max}}]$. If the solution $(\mathbf{L}_{\text{cvx}}, \mathbf{S}_{\text{cvx}})$ to the convex program (1.3) admits the following crude error bound

$$\|\mathbf{L}_{\text{cvx}} - \mathbf{L}^*\|_{\text{F}} \lesssim \sigma n^4, \quad (3.8)$$

then under Conditions 1-2 we have

$$\|\mathbf{X}\mathbf{Y}^{\top} - \mathbf{L}_{\text{cvx}}\|_{\text{F}} \lesssim \frac{\sigma}{n^5} \quad \text{and} \quad \|\mathbf{S} - \mathbf{S}_{\text{cvx}}\|_{\text{F}} \lesssim \frac{\sigma}{n^5}.$$

This theorem is a deterministic result, focusing on some sort of ‘‘approximate stationary points’’ of $F(\mathbf{X}, \mathbf{Y}, \mathbf{S})$. To interpret this, observe that in view of (3.7), one has $\nabla f(\mathbf{X}, \mathbf{Y}; \mathbf{S}) \approx \mathbf{0}$, and \mathbf{S} minimizes $F(\mathbf{X}, \mathbf{Y}, \cdot)$ for any fixed \mathbf{X} and \mathbf{Y} . If one can identify such an approximate stationary point that is sufficiently close to the truth (so that it satisfies Condition 1), then under mild conditions our theory asserts that

$$\mathbf{X}\mathbf{Y}^{\top} \approx \mathbf{L}_{\text{cvx}} \quad \text{and} \quad \mathbf{S} \approx \mathbf{S}_{\text{cvx}}.$$

This would in turn formalize the intimate relation between the solution to convex relaxation and an approximate stationary point of the nonconvex formulation.

The careful reader might immediately remark that this theorem does not say anything explicit about the minimizer of the nonconvex optimization problem (3.6); rather, it only pays attention to a special class of approximate stationary points of the nonconvex formulation. This arises mainly due to a technical consideration: it seems more difficult to analyze the nonconvex optimizer directly than to study certain approximate stationary points. Fortunately, our theorem indicates that any approximate stationary point obeying the above conditions serves as an extremely tight approximation of the convex estimate, and, therefore, it suffices to identify and analyze any such points.

3.3 Constructing an approximate stationary point via nonconvex algorithms

By virtue of Theorem 4, the key to understanding convex relaxation is to construct an approximate stationary point of the nonconvex problem (3.6) that enjoys desired statistical properties. For this purpose, we resort to the following iterative algorithm (Algorithm 1) to solve the nonconvex program (3.6).

Algorithm 1 Alternating minimization method for solving the nonconvex problem (3.6).

Suitable initialization: $\mathbf{X}^0 = \mathbf{X}^*, \mathbf{Y}^0 = \mathbf{Y}^*, \mathbf{S}^0 = \mathbf{S}^*$.

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

$$\mathbf{X}^{t+1} = \mathbf{X}^t - \eta \nabla_{\mathbf{X}} f(\mathbf{X}^t, \mathbf{Y}^t; \mathbf{S}^t) = \mathbf{X}^t - \frac{\eta}{p} [\mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{X}^t \mathbf{Y}^{t\top} + \mathbf{S}^t - \mathbf{M}) \mathbf{Y}^t + \lambda \mathbf{X}^t]; \quad (3.9a)$$

$$\mathbf{Y}^{t+1} = \mathbf{Y}^t - \eta \nabla_{\mathbf{Y}} f(\mathbf{X}^t, \mathbf{Y}^t; \mathbf{S}^t) = \mathbf{Y}^t - \frac{\eta}{p} \left\{ [\mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{X}^t \mathbf{Y}^{t\top} + \mathbf{S}^t - \mathbf{M})]^{\top} \mathbf{X}^t + \lambda \mathbf{Y}^t \right\}; \quad (3.9b)$$

$$\mathbf{S}^{t+1} = \mathcal{S}_{\tau}[\mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{M} - \mathbf{X}^{t+1} \mathbf{Y}^{t+1\top})]. \quad (3.9c)$$

In a nutshell, Algorithm 1 alternates between one iteration of gradient updates (w.r.t. the decision matrices \mathbf{X} and \mathbf{Y}) and optimization of the non-smooth problem w.r.t. \mathbf{S} (with \mathbf{X} and \mathbf{Y} frozen).⁸ For the sake of simplicity, we initialize this algorithm from the ground truth $(\mathbf{X}^*, \mathbf{Y}^*, \mathbf{S}^*)$, but our analysis framework might be extended to accommodate other more practical initialization (e.g. the one obtained by a spectral method).

⁸Note that for any given \mathbf{X} and \mathbf{Y} , the solution to minimize $F(\mathbf{X}, \mathbf{Y}, \mathbf{S})$ is given precisely by $\mathcal{S}_{\tau}(\mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{M} - \mathbf{X}\mathbf{Y}^{\top}))$.

The following theorem makes precise the statistical guarantees of the above nonconvex optimization algorithm; the proof is deferred to Appendix F. Here and throughout, we define

$$\mathbf{H}^t := \arg \min_{\mathbf{R} \in \mathcal{O}^{r \times r}} (\|\mathbf{X}^t \mathbf{R} - \mathbf{X}^*\|_{\text{F}}^2 + \|\mathbf{Y}^t \mathbf{R} - \mathbf{Y}^*\|_{\text{F}}^2)^{1/2}, \quad (3.10)$$

where $\mathcal{O}^{r \times r}$ denotes the set of $r \times r$ orthonormal matrices.

Theorem 5. *Instate the assumptions of Theorem 2 and recall that $\delta_n = \sigma/\sigma_{\min} \cdot \sqrt{n/p}$ therein. Take $t_0 = n^{47}$ and $\eta \asymp 1/(n\kappa^3\sigma_{\max})$ in Algorithm 1. With probability at least $1 - O(n^{-3})$, the iterates $\{(\mathbf{X}^t, \mathbf{Y}^t, \mathbf{S}^t)\}_{0 \leq t \leq t_0}$ of Algorithm 1 satisfy*

$$\max \left\{ \|\mathbf{X}^t \mathbf{H}^t - \mathbf{X}^*\|_{\text{F}}, \|\mathbf{Y}^t \mathbf{H}^t - \mathbf{Y}^*\|_{\text{F}} \right\} \lesssim \delta_n \|\mathbf{X}^*\|_{\text{F}}, \quad (3.11a)$$

$$\max \left\{ \|\mathbf{X}^t \mathbf{H}^t - \mathbf{X}^*\|, \|\mathbf{Y}^t \mathbf{H}^t - \mathbf{Y}^*\| \right\} \lesssim \delta_n \|\mathbf{X}^*\|, \quad (3.11b)$$

$$\max \left\{ \|\mathbf{X}^t \mathbf{H}^t - \mathbf{X}^*\|_{2,\infty}, \|\mathbf{Y}^t \mathbf{H}^t - \mathbf{Y}^*\|_{2,\infty} \right\} \lesssim \kappa \sqrt{\log n} \delta_n \max \left\{ \|\mathbf{X}^*\|_{2,\infty}, \|\mathbf{Y}^*\|_{2,\infty} \right\}, \quad (3.11c)$$

$$\|\mathbf{S}^t - \mathbf{S}^*\| \lesssim \sigma \sqrt{np}. \quad (3.11d)$$

In addition, with probability at least $1 - O(n^{-3})$, one has

$$\min_{0 \leq t < t_0} \|\nabla f(\mathbf{X}^t, \mathbf{Y}^t; \mathbf{S}^t)\|_{\text{F}} \leq \frac{1}{n^{20}} \frac{\lambda}{p} \sqrt{\sigma_{\min}}. \quad (3.12)$$

In short, the bounds (3.11a)-(3.11c) reveal that the entire sequence $\{\mathbf{X}^t, \mathbf{Y}^t\}_{0 \leq t \leq t_0}$ stays sufficiently close to the truth (measured by $\|\cdot\|_{\text{F}}$, $\|\cdot\|$, and more importantly, $\|\cdot\|_{2,\infty}$), the inequality (3.11d) demonstrates the goodness of fit of $\{\mathbf{S}^t\}_{0 \leq t \leq t_0}$ in terms of the spectral norm accuracy, whereas the last bound (3.12) indicates that there is at least one point in the sequence $\{\mathbf{X}^t, \mathbf{Y}^t, \mathbf{S}^t\}_{0 \leq t \leq t_0}$ that can serve as an approximate stationary point of the nonconvex formulation.

We shall also gather a few immediate consequences of Theorem 5 as follows, which contain basic properties that will be useful throughout.

Corollary 2. *Instate the assumptions of Theorem 5. Suppose that the sample size obeys $n^2 p \gg \kappa^4 \mu^2 r^2 n \log^4 n$, the noise satisfies $\delta_n \ll 1/\sqrt{\kappa^4 \mu r \log n}$, the outlier fraction satisfies $\rho_s \ll 1/(\kappa^3 \mu r \log n)$. With probability at least $1 - O(n^{-3})$, the iterates of Algorithm 1 satisfy*

$$\|\mathbf{X}^t \mathbf{Y}^{t\top} - \mathbf{L}^*\|_{\text{F}} \lesssim \kappa \delta_n \|\mathbf{L}^*\|_{\text{F}}, \quad (3.13a)$$

$$\|\mathbf{X}^t \mathbf{Y}^{t\top} - \mathbf{L}^*\|_{\infty} \lesssim \sqrt{\kappa^3 \mu r \log n} \delta_n \|\mathbf{L}^*\|_{\infty}, \quad (3.13b)$$

$$\|\mathbf{X}^t \mathbf{Y}^{t\top} - \mathbf{L}^*\| \lesssim \delta_n \|\mathbf{L}^*\| \quad (3.13c)$$

simultaneously for all $t \leq t_0$.

Proof. See [CCF⁺19, Appendix D.12]. □

3.4 Proof of Theorem 2

Define

$$t_* := \arg \min_{0 \leq t < t_0} \|\nabla f(\mathbf{X}^t, \mathbf{Y}^t; \mathbf{S}^t)\|_{\text{F}}; \quad (3.14)$$

$$(\mathbf{X}_{\text{ncvx}}, \mathbf{Y}_{\text{ncvx}}, \mathbf{S}_{\text{ncvx}}) := (\mathbf{X}^{t_*} \mathbf{H}^{t_*}, \mathbf{Y}^{t_*} \mathbf{H}^{t_*}, \mathbf{S}^{t_*}). \quad (3.15)$$

Theorem 5 and Corollary 2 have established appealing statistical performance of the nonconvex solution $(\mathbf{X}_{\text{ncvx}}, \mathbf{Y}_{\text{ncvx}}, \mathbf{S}_{\text{ncvx}})$. To transfer this desired statistical property to that of $(\mathbf{L}_{\text{cvx}}, \mathbf{S}_{\text{cvx}})$, it remains to show that the nonconvex estimator $(\mathbf{X}_{\text{ncvx}} \mathbf{Y}_{\text{ncvx}}^\top, \mathbf{S}_{\text{ncvx}})$ is extremely close to the convex estimator $(\mathbf{L}_{\text{cvx}}, \mathbf{S}_{\text{cvx}})$. Towards this end, we intend to invoke Theorem 4; therefore, it boils down to verifying the conditions therein.

1. The small gradient condition (cf. (3.7)) holds automatically under (3.12).
2. By virtue of the spectral norm bound (3.11b), one has

$$\|\mathbf{X}_{\text{ncvx}} - \mathbf{X}^*\| = \|\mathbf{X}^{t_*} \mathbf{H}^{t_*} - \mathbf{X}^*\| \lesssim \frac{\sigma}{\sigma_{\min}} \sqrt{\frac{n}{p}} \|\mathbf{L}^*\| \leq \frac{\sqrt{\sigma_{\min}}}{10},$$

as long as $\sigma \sqrt{\kappa n/p} \ll \sigma_{\min}$. This together with the Weyl inequality verifies the constraints on the singular values of $(\mathbf{X}_{\text{ncvx}}, \mathbf{Y}_{\text{ncvx}})$.

3. The crude error bounds are valid in view of Theorem 3.
4. Regarding Condition 1 and Condition 2, Lemma 1 and standard inequalities about sub-Gaussian random variables imply that $\|\mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{E})\| < \lambda/16$ and $\|\mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{E})\|_{\infty} \leq \tau/4$. In addition, the bounds (3.11d) and (3.13b) ensure the second assumption $\|\mathbf{S}_{\text{ncvx}} - \mathbf{S}^*\| \leq \lambda/16$ and $\|\mathbf{X}\mathbf{Y}^{\top} - \mathbf{L}^*\|_{\infty} \leq \tau/4$ in Condition 1. We are left with the last assumption in Condition 1 and Condition 2, which are guaranteed to hold in view of the following lemma (see Appendix C for the proof).

Lemma 2. *Instate the notations and assumptions of Theorem 2. Then with probability exceeding $1 - O(n^{-10})$, we have*

$$\|\mathcal{P}_{\Omega}(\mathbf{X}\mathbf{Y}^{\top} - \mathbf{M}^*) - p(\mathbf{X}\mathbf{Y}^{\top} - \mathbf{M}^*)\| < \lambda/8, \quad (3.16a)$$

$$\frac{1}{p} \|\mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{H})\|_{\mathbb{F}}^2 \geq \frac{1}{32\kappa} \|\mathbf{H}\|_{\mathbb{F}}^2, \quad \forall \mathbf{H} \in T, \quad (3.16b)$$

$$p^{-1} \|\mathcal{P}_{\Omega^*}(\mathbf{H})\|_{\mathbb{F}}^2 \leq \frac{1}{128\kappa} \|\mathbf{H}\|_{\mathbb{F}}^2, \quad \forall \mathbf{H} \in T \quad (3.16c)$$

simultaneously for all (\mathbf{X}, \mathbf{Y}) obeying

$$\|\mathbf{X} - \mathbf{X}^*\|_{2,\infty} \leq C_{\infty} \kappa \frac{\sigma}{\sigma_{\min}} \sqrt{\frac{n \log n}{p}} \max \left\{ \|\mathbf{X}^*\|_{2,\infty}, \|\mathbf{Y}^*\|_{2,\infty} \right\}; \quad (3.17a)$$

$$\|\mathbf{Y} - \mathbf{Y}^*\|_{2,\infty} \leq C_{\infty} \kappa \frac{\sigma}{\sigma_{\min}} \sqrt{\frac{n \log n}{p}} \max \left\{ \|\mathbf{X}^*\|_{2,\infty}, \|\mathbf{Y}^*\|_{2,\infty} \right\}. \quad (3.17b)$$

Here, T denotes the tangent space of the set of rank- r matrices at the point $\mathbf{X}\mathbf{Y}^{\top}$, and $C_{\infty} > 0$ is an absolute constant.

Armed with the above conditions, we can readily invoke Theorem 4 to reach

$$\|\mathbf{X}_{\text{ncvx}} \mathbf{Y}_{\text{ncvx}}^{\top} - \mathbf{L}_{\text{cvx}}\|_{\mathbb{F}} \lesssim \frac{\sigma}{n^5} \quad \text{and} \quad \|\mathbf{S}_{\text{ncvx}} - \mathbf{S}_{\text{cvx}}\|_{\mathbb{F}} \lesssim \frac{\sigma}{n^5}$$

with high probability. This taken collectively with Corollary 2 gives

$$\begin{aligned} \|\mathbf{L}_{\text{cvx}} - \mathbf{L}^*\|_{\mathbb{F}} &\leq \|\mathbf{X}_{\text{ncvx}} \mathbf{Y}_{\text{ncvx}}^{\top} - \mathbf{L}_{\text{cvx}}\|_{\mathbb{F}} + \|\mathbf{X}_{\text{ncvx}} \mathbf{Y}_{\text{ncvx}}^{\top} - \mathbf{L}^*\|_{\mathbb{F}} \\ &\lesssim \frac{\sigma}{n^5} + \kappa \frac{\sigma}{\sigma_{\min}} \sqrt{\frac{n}{p}} \|\mathbf{L}^*\|_{\mathbb{F}} \\ &\asymp \kappa \frac{\sigma}{\sigma_{\min}} \sqrt{\frac{n}{p}} \|\mathbf{L}^*\|_{\mathbb{F}}. \end{aligned}$$

Similar arguments lead to the advertised high-probability bounds

$$\begin{aligned} \|\mathbf{L}_{\text{cvx}} - \mathbf{L}^*\|_{\infty} &\lesssim \sqrt{\kappa^3 \mu r} \frac{\sigma}{\sigma_{\min}} \sqrt{\frac{n \log n}{p}} \|\mathbf{L}^*\|_{\infty}, \\ \|\mathbf{L}_{\text{cvx}} - \mathbf{L}^*\| &\lesssim \frac{\sigma}{\sigma_{\min}} \sqrt{\frac{n}{p}} \|\mathbf{L}^*\|. \end{aligned}$$

Finally, given that $\mathbf{X}_{\text{ncvx}}\mathbf{Y}_{\text{ncvx}}^\top$ is a rank- r matrix, the rank- r approximation $\mathbf{L}_{\text{cvx},r} := \arg \min_{\mathbf{Z}:\text{rank}(\mathbf{Z})\leq r} \|\mathbf{Z} - \mathbf{L}_{\text{cvx}}\|_{\text{F}}$ of \mathbf{L}_{cvx} necessarily satisfies

$$\|\mathbf{L}_{\text{cvx},r} - \mathbf{L}_{\text{cvx}}\|_{\text{F}} \leq \|\mathbf{X}_{\text{ncvx}}\mathbf{Y}_{\text{ncvx}}^\top - \mathbf{L}_{\text{cvx}}\|_{\text{F}} \lesssim \frac{\sigma}{n^5} \leq \frac{1}{n^5} \cdot \delta_n \|\mathbf{L}^*\|,$$

which establishes (1.23). In view of the triangle inequality, the properties (1.22) hold unchanged if \mathbf{L}_{cvx} is replaced by $\mathbf{L}_{\text{cvx},r}$.

4 Discussion

This paper investigates the unreasonable effectiveness of convex programming in estimating an unknown low-rank matrix from grossly corrupted data. We develop an improved theory that confirms the optimality of convex relaxation in the presence of random noise, gross sparse outliers, and missing data. In particular, our results significantly improve upon the prior statistical guarantees [ZLW⁺10] under random noise, while further allowing for missing data. Our theoretical analysis is built upon an appealing connection between convex and nonconvex optimization, which has not been established previously.

Having said this, our current work leaves open several important issues that call for further investigation. To begin with, the conditions (1.21) stated in the main theorem are likely suboptimal in terms of the dependency on both the rank r and the condition number κ . For example, we shall keep in mind that in the noise-free setting, the sample size can be as low as $O(nr \text{poly} \log n)$ and the tolerable outlier fraction can be as large as a constant [Li13, CJSC13], both of which exhibit more favorable scalings w.r.t. r and κ compared to our current condition (1.21). Moving forward, our analysis ideas suggest a possible route for analyzing convex relaxation for other structured estimation problems under both random noise and outliers, including but not limited to sparse PCA (the case with a simultaneously low-rank and sparse matrix) [CMW13], low-rank Hankel matrix estimation (the case involving a low-rank Hankel matrix) [CC14], and blind deconvolution (the case that aims to recover a low-rank matrix from structured Fourier measurements) [ARR14]. Last but not least, we would like to point out that it is possible to design a similar debiasing procedure as in [CFMY19] for correcting the bias in the convex estimator, which further allows uncertainty quantification and statistical inference on the unknown low-rank matrix of interest.

Acknowledgements

Y. Chen is supported in part by the AFOSR YIP award FA9550-19-1-0030, by the ONR grant N00014-19-1-2120, by the ARO grant W911NF-18-1-0303, by the NSF grants CCF-1907661 and IIS-1900140, and by the Princeton SEAS innovation award. J. Fan is supported in part by the NSF grants DMS-1662139 and DMS-1712591, the ONR grant N00014-19-1-2120, and the NIH grant 2R01-GM072611-14.

A An equivalent probabilistic model of Ω^* used throughout the proof

In this section, we introduce an equivalent probabilistic model of Ω^* , which is more amenable to analysis and which shall be assumed throughout the proof.

- **The original model.** Recall from Assumption 3 the way we generate Ω^* : (1) sample Ω_{obs} from the i.i.d. Bernoulli model with parameter p ; (2) for each $(i, j) \in \Omega_{\text{obs}}$, let $(i, j) \in \Omega^*$ independently with probability ρ_s .
- **An equivalent model.** We introduce an equivalent model: (1) sample Ω_{obs} from the i.i.d. Bernoulli model with parameter p ; (2) generate an augmented index set $\Omega_{\text{aug}} \subseteq \Omega_{\text{obs}}$ such that: for each $(i, j) \in \Omega_{\text{obs}}$, we $(i, j) \in \Omega_{\text{aug}}$ independently with probability ρ_{aug} ; (3) for any $(i, j) \in \Omega_{\text{aug}}$, include (i, j) in Ω^* independently with probability ρ_s/ρ_{aug} .

It is straightforward to verify that the two models for Ω^* are equivalent as long as $\rho_s \leq \rho_{\text{aug}} \leq 1$. Two important remarks are in order. First, by construction, we have $\Omega^* \subseteq \Omega_{\text{aug}}$. Second, the choice of ρ_{aug} can vary as needed as long as $\rho_s \leq \rho_{\text{aug}} \leq 1$.

The introduction of this augmented index set Ω_{aug} comes in handy when we would like to control the size $\|\mathcal{P}_{\Omega^*}(\mathbf{A})\|_{\text{F}}$ for some matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$. The first inclusion property $\Omega^* \subseteq \Omega_{\text{aug}}$ allows us to upper bound $\|\mathcal{P}_{\Omega^*}(\mathbf{A})\|_{\text{F}}$ by $\|\mathcal{P}_{\Omega_{\text{aug}}}(\mathbf{A})\|_{\text{F}}$, and the freedom to choose ρ_{aug} allows us to leverage stronger concentration results, which might not hold for the smaller ρ_s . See Corollary 3 in the next section for an example.

B Preliminaries

B.1 A few preliminary facts

This subsection collects several results that are useful throughout the proof. To begin with, the incoherence assumption (cf. Assumption 1) asserts that

$$\|\mathbf{X}^*\|_{2,\infty} \leq \sqrt{\mu r/n} \|\mathbf{X}^*\| \quad \text{and} \quad \|\mathbf{Y}^*\|_{2,\infty} \leq \sqrt{\mu r/n} \|\mathbf{Y}^*\|. \quad (\text{B.1})$$

This is because

$$\|\mathbf{X}^*\|_{2,\infty} = \|\mathbf{U}^*(\boldsymbol{\Sigma}^*)^{1/2}\|_{2,\infty} \leq \|\mathbf{U}^*\|_{2,\infty} \|(\boldsymbol{\Sigma}^*)^{1/2}\| \leq \sqrt{\mu r/n} \|\mathbf{X}^*\|,$$

where the first inequality comes from the elementary inequality $\|\mathbf{AB}\|_{2,\infty} \leq \|\mathbf{A}\|_{2,\infty} \|\mathbf{B}\|$, and the last inequality is a consequence of the incoherence assumption as well as the fact that $\|(\boldsymbol{\Sigma}^*)^{1/2}\| = \|\mathbf{X}^*\|$.

The next lemma is extensively used in the low-rank matrix completion literature.

Lemma 3. *Suppose that each (i, j) is included in $\Omega_0 \subseteq [n] \times [n]$ independently with probability ρ_0 . Then with probability exceeding $1 - O(n^{-10})$, one has*

$$\|\mathcal{P}_{T^*} - \rho_0^{-1} \mathcal{P}_{T^*} \mathcal{P}_{\Omega_0} \mathcal{P}_{T^*}\| \leq \frac{1}{2}, \quad (\text{B.2})$$

provided that $n^2 \rho_0 \gg \mu r n \log n$. Here, T^* denotes the tangent space of the set of rank- r matrices at the point $\mathbf{L}^* = \mathbf{X}^* \mathbf{Y}^{*\top}$.

Proof. See [CR09, Theorem 4.1] □

In fact, the bound (B.2) uncovers certain near-isometry of the operator $\rho_0^{-1} \mathcal{P}_{\Omega_0}(\cdot)$ when restricted to the tangent space T^* . This property is formalized in the following fact.

Fact 1. *Suppose that $\|\mathcal{P}_{T^*} - \rho_0^{-1} \mathcal{P}_{T^*} \mathcal{P}_{\Omega_0} \mathcal{P}_{T^*}\| \leq 1/2$. Then one has*

$$\frac{1}{2} \|\mathbf{H}\|_{\text{F}}^2 \leq \frac{1}{\rho_0} \|\mathcal{P}_{\Omega_0}(\mathbf{H})\|_{\text{F}}^2 \leq \frac{3}{2} \|\mathbf{H}\|_{\text{F}}^2, \quad \text{for all } \mathbf{H} \in T^*.$$

Proof. The proof has actually been documented in the literature. For completeness, we present the proof for the lower bound here; the upper bound follows from a very similar argument. For any $\mathbf{H} \in \mathbb{R}^{n \times n}$, one has

$$\begin{aligned} \|\mathcal{P}_{\Omega_0} \mathcal{P}_{T^*}(\mathbf{H})\|_{\text{F}}^2 &= \langle \mathcal{P}_{\Omega_0} \mathcal{P}_{T^*}(\mathbf{H}), \mathcal{P}_{\Omega_0} \mathcal{P}_{T^*}(\mathbf{H}) \rangle = \langle \mathcal{P}_{T^*}(\mathbf{H}), \mathcal{P}_{T^*} \mathcal{P}_{\Omega_0} \mathcal{P}_{T^*}(\mathbf{H}) \rangle \\ &= \rho_0 \|\mathcal{P}_{T^*}(\mathbf{H})\|_{\text{F}}^2 - \rho_0 \langle \mathcal{P}_{T^*}(\mathbf{H}), (\mathcal{P}_{T^*} - \rho_0^{-1} \mathcal{P}_{T^*} \mathcal{P}_{\Omega_0} \mathcal{P}_{T^*})(\mathcal{P}_{T^*}(\mathbf{H})) \rangle \\ &\geq \rho_0 \|\mathcal{P}_{T^*}(\mathbf{H})\|_{\text{F}}^2 - \rho_0 \|\mathcal{P}_{T^*} - \rho_0^{-1} \mathcal{P}_{T^*} \mathcal{P}_{\Omega_0} \mathcal{P}_{T^*}\| \|\mathcal{P}_{T^*}(\mathbf{H})\|_{\text{F}}^2 \\ &\geq \frac{\rho_0}{2} \|\mathcal{P}_{T^*}(\mathbf{H})\|_{\text{F}}^2. \end{aligned}$$

Here, the penultimate inequality relies on the elementary fact that $\langle \mathbf{A}, \mathbf{B} \rangle \leq \|\mathbf{A}\|_{\text{F}} \|\mathbf{B}\|_{\text{F}}$, and the last step follows from the assumption $\|\mathcal{P}_{T^*} - \rho_0^{-1} \mathcal{P}_{T^*} \mathcal{P}_{\Omega_0} \mathcal{P}_{T^*}\| \leq 1/2$. □

The following corollary is an immediate consequence of Lemma 3 and Fact 1.

Corollary 3. *Suppose that $\rho_s \leq \rho_{\text{aug}} \leq 1/12$ and that $n^2 p \rho_{\text{aug}} \gg \mu r n \log n$. Then with probability at least $1 - O(n^{-10})$, we have*

$$\|\mathcal{P}_{\Omega^*} \mathcal{P}_{T^*}\|^2 \leq p/8.$$

Proof. Recall the auxiliary index set Ω_{aug} introduced in Appendix A. Since $\Omega^* \subseteq \Omega_{\text{aug}}$, we have for any $\mathbf{H} \in \mathbb{R}^{n \times n}$

$$\|\mathcal{P}_{\Omega^*} \mathcal{P}_{T^*}(\mathbf{H})\|_{\text{F}}^2 \leq \|\mathcal{P}_{\Omega_{\text{aug}}} \mathcal{P}_{T^*}(\mathbf{H})\|_{\text{F}}^2 \leq \frac{3p\rho_{\text{aug}}}{2} \|\mathcal{P}_{T^*}(\mathbf{H})\|_{\text{F}}^2 \leq \frac{3p\rho_{\text{aug}}}{2} \|\mathbf{H}\|_{\text{F}}^2.$$

Here, the second inequality arises from Lemma 3 and Fact 1 (by taking $\Omega_0 = \Omega_{\text{aug}}$ and $\rho_0 = p\rho_{\text{aug}}$). The proof is complete by recognizing the assumption $\rho_{\text{aug}} \leq 1/12$. \square

As it turns out, the near-isometry property of $\rho_0^{-1} \mathcal{P}_{\Omega_0}(\cdot)$ can be strengthened to a uniform version (uniform over a large collection of tangent spaces), as shown in the lemma below.

Lemma 4. *Suppose that each (i, j) is included in $\Omega_0 \subseteq [n] \times [n]$ independently with probability ρ_0 , and that $n^2 \rho_0 \gg \mu r n \log n$. Then with probability at least $1 - O(n^{-10})$,*

$$\frac{1}{32\kappa} \|\mathbf{H}\|_{\text{F}}^2 \leq \frac{1}{\rho_0} \|\mathcal{P}_{\Omega_0}(\mathbf{H})\|_{\text{F}}^2 \leq 40\kappa \|\mathbf{H}\|_{\text{F}}^2, \quad \text{for all } \mathbf{H} \in T$$

holds simultaneously for all (\mathbf{X}, \mathbf{Y}) obeying

$$\max \left\{ \|\mathbf{X} - \mathbf{X}^*\|_{2,\infty}, \|\mathbf{Y} - \mathbf{Y}^*\|_{2,\infty} \right\} \leq \frac{c}{\kappa\sqrt{n}} \|\mathbf{X}^*\|.$$

Here, $c > 0$ is some sufficiently small constant, and T denotes the tangent space of the set of rank- r matrices at the point $\mathbf{X}\mathbf{Y}^\top$.

Proof. See Appendix B.2. \square

In the end, we recall a useful lemma which relates the operator norm to the $\ell_{2,\infty}$ norm of a matrix.

Lemma 5. *Suppose that each (i, j) is included in $\Omega_0 \subseteq [n] \times [n]$ independently with probability ρ_0 , and that $n^2 \rho_0 \gg n \log n$. Then there exists some absolute constant $C > 0$ such that with probability at least $1 - O(n^{-10})$,*

$$\|\mathcal{P}_{\Omega_0}(\mathbf{A}\mathbf{B}^\top) - \rho_0 \mathbf{A}\mathbf{B}^\top\| \leq C\sqrt{n\rho_0} \|\mathbf{A}\|_{2,\infty} \|\mathbf{B}\|_{2,\infty}$$

holds simultaneously for all \mathbf{A} and \mathbf{B} .

Proof. See [CLL19, Lemmas 4.2 and 4.3]. \square

B.2 Proof of Lemma 4

The lower bound has been established in [CCF⁺19, Lemma 7], and hence we focus on the upper bound. We start by expressing $\mathbf{H} \in T$ as $\mathbf{H} = \mathbf{X}\mathbf{A}^\top + \mathbf{B}\mathbf{Y}^\top$, where $\mathbf{A}, \mathbf{B} \in \mathbb{R}^{n \times r}$ are chosen to be

$$(\mathbf{A}, \mathbf{B}) := \arg \min_{(\tilde{\mathbf{A}}, \tilde{\mathbf{B}}): \mathbf{H} = \mathbf{X}\tilde{\mathbf{A}}^\top + \tilde{\mathbf{B}}\mathbf{Y}^\top} \left\{ \|\tilde{\mathbf{A}}\|_{\text{F}}^2/2 + \|\tilde{\mathbf{B}}\|_{\text{F}}^2/2 \right\}.$$

The optimality condition of (\mathbf{A}, \mathbf{B}) requires

$$\mathbf{X}^\top \mathbf{B} = \mathbf{A}^\top \mathbf{Y}; \tag{B.3}$$

see [CCF⁺19, Section C.3.1] for the justification of this identity. The proof then consists of two steps:

1. Showing that $\|\mathbf{H}\|_{\text{F}}^2$ is bounded from below, namely,

$$\|\mathbf{H}\|_{\text{F}}^2 \geq \frac{49}{100} \sigma_{\min} \left(\|\mathbf{A}\|_{\text{F}}^2 + \|\mathbf{B}\|_{\text{F}}^2 \right).$$

To see this, we can invoke the bound on α_2 stated in [CCF⁺19, Appendix C.3.1] to yield

$$\begin{aligned} \|\mathbf{H}\|_{\text{F}}^2 &= \|\mathbf{X}\mathbf{A}^\top + \mathbf{B}\mathbf{Y}^\top\|_{\text{F}}^2 \geq \frac{1}{2} \left(\|\mathbf{X}^* \mathbf{A}^\top\|_{\text{F}}^2 + \|\mathbf{B}\mathbf{Y}^*\|_{\text{F}}^2 \right) - \frac{1}{100} \sigma_{\min} \left(\|\mathbf{A}\|_{\text{F}}^2 + \|\mathbf{B}\|_{\text{F}}^2 \right) \\ &\geq \frac{1}{2} \sigma_{\min} \left(\|\mathbf{A}\|_{\text{F}}^2 + \|\mathbf{B}\|_{\text{F}}^2 \right) - \frac{1}{100} \sigma_{\min} \left(\|\mathbf{A}\|_{\text{F}}^2 + \|\mathbf{B}\|_{\text{F}}^2 \right) \geq \frac{49}{100} \sigma_{\min} \left(\|\mathbf{A}\|_{\text{F}}^2 + \|\mathbf{B}\|_{\text{F}}^2 \right). \end{aligned}$$

2. Showing that $\|\mathcal{P}_{\Omega^*}(\mathbf{H})\|_{\mathbb{F}}^2$ is bounded from above, namely,

$$\frac{1}{2\rho_0} \|\mathcal{P}_{\Omega^*}(\mathbf{H})\|_{\mathbb{F}}^2 \leq 9\sigma_{\max} \left(\|\mathbf{A}\|_{\mathbb{F}}^2 + \|\mathbf{B}\|_{\mathbb{F}}^2 \right).$$

To this end, one starts with the following decomposition

$$\frac{1}{2\rho_0} \|\mathcal{P}_{\Omega_0}(\mathbf{H})\|_{\mathbb{F}}^2 = \frac{1}{2} \|\mathbf{H}\|_{\mathbb{F}}^2 + \frac{1}{2\rho_0} \|\mathcal{P}_{\Omega_0}(\mathbf{H})\|_{\mathbb{F}}^2 - \frac{1}{2} \|\mathbf{H}\|_{\mathbb{F}}^2. \quad (\text{B.4})$$

Apply [CCF⁺19, Equation (83)] to obtain

$$\frac{1}{2} \|\mathbf{X}\mathbf{A}^\top + \mathbf{B}\mathbf{Y}^\top\|_{\mathbb{F}}^2 \leq 8\sigma_{\max} \left(\|\mathbf{A}\|_{\mathbb{F}}^2 + \|\mathbf{B}\|_{\mathbb{F}}^2 \right).$$

In addition, the bound on α_1 stated in [CCF⁺19, Appendix C.3.1] tells us that

$$\begin{aligned} \frac{1}{2\rho_0} \|\mathcal{P}_{\Omega_0}(\mathbf{H})\|_{\mathbb{F}}^2 - \frac{1}{2} \|\mathbf{H}\|_{\mathbb{F}}^2 &= \frac{1}{2\rho_0} \|\mathcal{P}_{\Omega_0}(\mathbf{X}\mathbf{A}^\top + \mathbf{B}\mathbf{Y}^\top)\|_{\mathbb{F}}^2 - \frac{1}{2} \|\mathbf{X}\mathbf{A}^\top + \mathbf{B}\mathbf{Y}^\top\|_{\mathbb{F}}^2 \\ &\leq \frac{1}{32} \left(\|\mathbf{X}^*\mathbf{A}^\top\|_{\mathbb{F}}^2 + \|\mathbf{B}\mathbf{Y}^*\|_{\mathbb{F}}^2 \right) + \frac{1}{25} \sigma_{\min} \left(\|\mathbf{A}\|_{\mathbb{F}}^2 + \|\mathbf{B}\|_{\mathbb{F}}^2 \right) \\ &\leq \left(\frac{1}{32} \sigma_{\max} + \frac{1}{25} \sigma_{\min} \right) \left(\|\mathbf{A}\|_{\mathbb{F}}^2 + \|\mathbf{B}\|_{\mathbb{F}}^2 \right). \end{aligned}$$

Substitution into (B.4) gives

$$\begin{aligned} \frac{1}{2\rho_0} \|\mathcal{P}_{\Omega_0}(\mathbf{H})\|_{\mathbb{F}}^2 &\leq 8\sigma_{\max} \left(\|\mathbf{A}\|_{\mathbb{F}}^2 + \|\mathbf{B}\|_{\mathbb{F}}^2 \right) + \left(\frac{1}{32} \sigma_{\max} + \frac{1}{25} \sigma_{\min} \right) \left(\|\mathbf{A}\|_{\mathbb{F}}^2 + \|\mathbf{B}\|_{\mathbb{F}}^2 \right) \\ &\leq 9\sigma_{\max} \left(\|\mathbf{A}\|_{\mathbb{F}}^2 + \|\mathbf{B}\|_{\mathbb{F}}^2 \right). \end{aligned}$$

Putting the above two bounds together, we conclude that

$$\frac{1}{2\rho_0} \|\mathcal{P}_{\Omega_0}(\mathbf{H})\|_{\mathbb{F}}^2 \leq 9\sigma_{\max} \left(\|\mathbf{A}\|_{\mathbb{F}}^2 + \|\mathbf{B}\|_{\mathbb{F}}^2 \right) \leq \frac{900}{49} \kappa \cdot \frac{49}{100} \sigma_{\min} \left(\|\mathbf{A}\|_{\mathbb{F}}^2 + \|\mathbf{B}\|_{\mathbb{F}}^2 \right) \leq 20\kappa \|\mathbf{H}\|_{\mathbb{F}}^2$$

as claimed.

C Proof of Lemma 2

With Lemma 4 in place, we can immediately justify Lemma 2.

To begin with, the first two parts (3.16a) and (3.16b) are the same as [CCF⁺19, Lemma 4]. Hence, it suffices to verify the last one (3.16c). Recall from Appendix A that $\Omega^* \subseteq \Omega_{\text{aug}}$, where Ω_{aug} is randomly sampled such that each (i, j) is included in Ω_{aug} independently with probability $p\rho_{\text{aug}}$. Applying Lemma 4 on Ω_{aug} finishes the proof, with the proviso that $\rho_{\text{aug}} \asymp 1/\kappa^2$ and $\rho_s \leq \rho_{\text{aug}}$.

D Crude error bounds (Proof of Theorem 3)

This section is devoted to establishing our crude statistical error bounds on $\|\mathbf{L}_{\text{cvx}} - \mathbf{L}^*\|_{\mathbb{F}}$ and $\|\mathbf{S}_{\text{cvx}} - \mathbf{S}^*\|_{\mathbb{F}}$. Without loss of generality, we only consider the case when $\tau = \lambda\sqrt{\frac{\log n}{np}}$. The proof works for general choices $\tau \asymp \lambda\sqrt{\frac{\log n}{np}}$ with slight modification. To simplify the notation hereafter, we denote

$$\begin{aligned} \mathbf{\Lambda}_L &:= \mathbf{L}_{\text{cvx}} - \mathbf{L}^*, & \text{and} & & \mathbf{\Lambda}_S &:= \mathbf{S}_{\text{cvx}} - \mathbf{S}^*, \\ \mathbf{\Lambda}^+ &:= (\mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{\Lambda}_L) + \mathbf{\Lambda}_S)/2, & \text{and} & & \mathbf{\Lambda}^- &:= (\mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{\Lambda}_L) - \mathbf{\Lambda}_S)/2, \end{aligned}$$

which immediately imply

$$\mathbf{\Lambda}_L = \mathbf{\Lambda}^+ + \mathbf{\Lambda}^- + \mathcal{P}_{\Omega_{\text{obs}}^c}(\mathbf{\Lambda}_L), \quad \text{and} \quad \mathbf{\Lambda}_S = \mathbf{\Lambda}^+ - \mathbf{\Lambda}^-.$$

These in turn allow us to decompose $\|\mathbf{\Lambda}_L\|_{\mathbb{F}}^2 + \|\mathbf{\Lambda}_S\|_{\mathbb{F}}^2$ as follows

$$\begin{aligned} \|\mathbf{\Lambda}_L\|_{\mathbb{F}}^2 + \|\mathbf{\Lambda}_S\|_{\mathbb{F}}^2 &= \|\mathbf{\Lambda}^+ + \mathbf{\Lambda}^- + \mathcal{P}_{\Omega_{\text{obs}}^c}(\mathbf{\Lambda}_L)\|_{\mathbb{F}}^2 + \|\mathbf{\Lambda}^+ - \mathbf{\Lambda}^-\|_{\mathbb{F}}^2 \\ &= \|\mathbf{\Lambda}^+\|_{\mathbb{F}}^2 + \|\mathbf{\Lambda}^- + \mathcal{P}_{\Omega_{\text{obs}}^c}(\mathbf{\Lambda}_L)\|_{\mathbb{F}}^2 + 2\langle \mathbf{\Lambda}^+, \mathbf{\Lambda}^- + \mathcal{P}_{\Omega_{\text{obs}}^c}(\mathbf{\Lambda}_L) \rangle + \|\mathbf{\Lambda}^+\|_{\mathbb{F}}^2 + \|\mathbf{\Lambda}^-\|_{\mathbb{F}}^2 - 2\langle \mathbf{\Lambda}^+, \mathbf{\Lambda}^- \rangle \\ &= 2\|\mathbf{\Lambda}^+\|_{\mathbb{F}}^2 + \|\mathbf{\Lambda}^- + \mathcal{P}_{\Omega_{\text{obs}}^c}(\mathbf{\Lambda}_L)\|_{\mathbb{F}}^2 + \|\mathbf{\Lambda}^-\|_{\mathbb{F}}^2 + 2\langle \mathbf{\Lambda}^+, \mathcal{P}_{\Omega_{\text{obs}}^c}(\mathbf{\Lambda}_L) \rangle. \end{aligned} \quad (\text{D.1})$$

Since $(\mathbf{L}_{\text{cvx}}, \mathbf{S}_{\text{cvx}})$ is the minimizer of (1.3), it is self-evident that \mathbf{S}_{cvx} must be supported on Ω_{obs} . Then by construction, $\mathbf{\Lambda}_S, \mathbf{\Lambda}^+$ and $\mathbf{\Lambda}^-$ are all necessarily supported on Ω_{obs} , thus indicating that

$$\langle \mathbf{\Lambda}^+, \mathcal{P}_{\Omega_{\text{obs}}^c}(\mathbf{\Lambda}_L) \rangle = 0.$$

Making use of this relation, we can continue the derivation (D.1) above to obtain

$$\begin{aligned} \|\mathbf{\Lambda}_L\|_{\mathbb{F}}^2 + \|\mathbf{\Lambda}_S\|_{\mathbb{F}}^2 &= 2\|\mathbf{\Lambda}^+\|_{\mathbb{F}}^2 + \|\mathbf{\Lambda}^- + \mathcal{P}_{\Omega_{\text{obs}}^c}(\mathbf{\Lambda}_L)\|_{\mathbb{F}}^2 + \|\mathbf{\Lambda}^-\|_{\mathbb{F}}^2 \\ &= \underbrace{2\|\mathbf{\Lambda}^+\|_{\mathbb{F}}^2}_{=:\alpha_1} + \underbrace{\|\mathcal{P}_{T^*}(\mathbf{\Lambda}^- + \mathcal{P}_{\Omega_{\text{obs}}^c}(\mathbf{\Lambda}_L))\|_{\mathbb{F}}^2 + \|\mathcal{P}_{\Omega^*}(\mathbf{\Lambda}^-)\|_{\mathbb{F}}^2}_{=:\alpha_2} + \underbrace{\|\mathcal{P}_{T^{\perp}}(\mathbf{\Lambda}^- + \mathcal{P}_{\Omega_{\text{obs}}^c}(\mathbf{\Lambda}_L))\|_{\mathbb{F}}^2 + \|\mathcal{P}_{\Omega^{*c}}(\mathbf{\Lambda}^-)\|_{\mathbb{F}}^2}_{=:\alpha_3}. \end{aligned}$$

In the sequel, we shall control the three terms α_1, α_2 and α_3 separately.

Step 1: bounding α_1 . By definition, we have

$$\begin{aligned} \alpha_1 &= 2\|\mathbf{\Lambda}^+\|_{\mathbb{F}}^2 = \frac{1}{2}\|\mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{\Lambda}_L) + \mathbf{\Lambda}_S\|_{\mathbb{F}}^2 = \frac{1}{2}\|\mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{\Lambda}_L + \mathbf{\Lambda}_S)\|_{\mathbb{F}}^2 \\ &= \frac{1}{2}\|\mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{L}_{\text{cvx}} + \mathbf{S}_{\text{cvx}} - \mathbf{M} + \mathbf{M} - \mathbf{L}^* - \mathbf{S}^*)\|_{\mathbb{F}}^2 \\ &\leq \|\mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{L}_{\text{cvx}} + \mathbf{S}_{\text{cvx}} - \mathbf{M})\|_{\mathbb{F}}^2 + \|\mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{L}^* + \mathbf{S}^* - \mathbf{M})\|_{\mathbb{F}}^2 \\ &= \|\mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{L}_{\text{cvx}} + \mathbf{S}_{\text{cvx}} - \mathbf{M})\|_{\mathbb{F}}^2 + \|\mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{E})\|_{\mathbb{F}}^2, \end{aligned} \quad (\text{D.2})$$

where the third identity holds true since $\mathbf{\Lambda}_S = \mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{\Lambda}_S)$, the penultimate relation is due to the elementary inequality $\|\mathbf{A} + \mathbf{B}\|_{\mathbb{F}}^2 \leq 2\|\mathbf{A}\|_{\mathbb{F}}^2 + 2\|\mathbf{B}\|_{\mathbb{F}}^2$, and the last line follows since $\mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{L}^* + \mathbf{S}^* - \mathbf{M}) = \mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{E})$. To upper bound $\|\mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{L}_{\text{cvx}} + \mathbf{S}_{\text{cvx}} - \mathbf{M})\|_{\mathbb{F}}^2$, we leverage the optimality of $(\mathbf{L}_{\text{cvx}}, \mathbf{S}_{\text{cvx}})$ w.r.t. the convex program (1.3) to obtain

$$\begin{aligned} &\frac{1}{2}\|\mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{L}_{\text{cvx}} + \mathbf{S}_{\text{cvx}} - \mathbf{M})\|_{\mathbb{F}}^2 + \lambda\|\mathbf{L}_{\text{cvx}}\|_* + \tau\|\mathbf{S}_{\text{cvx}}\|_1 \\ &\leq \frac{1}{2}\|\mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{L}^* + \mathbf{S}^* - \mathbf{M})\|_{\mathbb{F}}^2 + \lambda\|\mathbf{L}^*\|_* + \tau\|\mathbf{S}^*\|_1. \end{aligned} \quad (\text{D.3})$$

Recognizing again that $\mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{L}^* + \mathbf{S}^* - \mathbf{M}) = \mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{E})$, we can rearrange terms in (D.3) to derive

$$\begin{aligned} \|\mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{L}_{\text{cvx}} + \mathbf{S}_{\text{cvx}} - \mathbf{M})\|_{\mathbb{F}}^2 &\leq \|\mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{E})\|_{\mathbb{F}}^2 + 2\lambda\|\mathbf{L}^*\|_* + 2\tau\|\mathbf{S}^*\|_1 - 2\lambda\|\mathbf{L}_{\text{cvx}}\|_* - 2\tau\|\mathbf{S}_{\text{cvx}}\|_1 \\ &\stackrel{(i)}{\leq} \|\mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{E})\|_{\mathbb{F}}^2 + 2\lambda\|\mathbf{\Lambda}_L\|_* + 2\tau\|\mathbf{\Lambda}_S\|_1 \\ &\stackrel{(ii)}{\leq} \|\mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{E})\|_{\mathbb{F}}^2 + 2\lambda\sqrt{n}\|\mathbf{\Lambda}_L\|_{\mathbb{F}} + 2\tau\sqrt{|\Omega_{\text{obs}}|}\|\mathbf{\Lambda}_S\|_{\mathbb{F}} \\ &\stackrel{(iii)}{\leq} \|\mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{E})\|_{\mathbb{F}}^2 + 2\sqrt{2}\lambda\sqrt{n\log n}(\|\mathbf{\Lambda}_L\|_{\mathbb{F}} + \|\mathbf{\Lambda}_S\|_{\mathbb{F}}), \end{aligned} \quad (\text{D.4})$$

where $|\Omega_{\text{obs}}|$ denotes the cardinality of Ω_{obs} . Here, the relation (i) results from the triangle inequality, the inequality (ii) holds true since $\|\mathbf{A}\|_* \leq \sqrt{n}\|\mathbf{A}\|_{\mathbb{F}}$ for any $\mathbf{A} \in \mathbb{R}^{n \times n}$ and $\|\mathbf{\Lambda}_S\|_1 = \|\mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{\Lambda}_S)\|_1 \leq$

$\sqrt{|\Omega_{\text{obs}}|} \|\mathbf{A}_S\|_F$, and the last line (iii) arises from the fact that $|\Omega_{\text{obs}}| \leq 2n^2p$ with high probability as well as the choice $\tau = \lambda \sqrt{\frac{\log n}{np}}$. Combine (D.2) and (D.4) to reach

$$\begin{aligned} \alpha_1 &\leq 2 \|\mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{E})\|_F^2 + 2\sqrt{2}\lambda\sqrt{n \log n} (\|\mathbf{A}_L\|_F + \|\mathbf{A}_S\|_F) \\ &\leq 2 \|\mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{E})\|_F^2 + 4\lambda\sqrt{n \log n} \sqrt{\|\mathbf{A}_L\|_F^2 + \|\mathbf{A}_S\|_F^2}, \end{aligned} \quad (\text{D.5})$$

where we use the elementary inequality $a + b \leq \sqrt{2} \cdot \sqrt{a^2 + b^2}$.

Step 2: bounding α_2 via α_3 . To relate α_2 to α_3 , the following lemma plays a crucial role, whose proof is deferred to Appendix D.1.

Lemma 6. *Suppose that $\|\mathcal{P}_{\Omega^*} \mathcal{P}_{T^*}\|^2 \leq p/8$ and that $\|\mathcal{P}_{T^*} - p^{-1} \mathcal{P}_{T^*} \mathcal{P}_{\Omega_{\text{obs}}} \mathcal{P}_{T^*}\| \leq 1/2$. Then for any pair (\mathbf{A}, \mathbf{B}) of matrices, we have*

$$\|\mathcal{P}_{T^*}(\mathbf{A})\|_F^2 + \|\mathcal{P}_{\Omega^*}(\mathbf{B})\|_F^2 \leq \frac{4}{p} \|\mathcal{P}_{\Omega_{\text{obs}}}[\mathcal{P}_{T^*}(\mathbf{A}) + \mathcal{P}_{\Omega^*}(\mathbf{B})]\|_F^2. \quad (\text{D.6})$$

Suppose for the moment that the assumptions of Lemma 6 hold. Taking (\mathbf{A}, \mathbf{B}) as $(\mathbf{A}^- + \mathcal{P}_{\Omega_{\text{obs}}^c}(\mathbf{A}_L), -\mathbf{A}^-)$ in Lemma 6 yields

$$\alpha_2 = \|\mathcal{P}_{T^*}(\mathbf{A}^- + \mathcal{P}_{\Omega_{\text{obs}}^c}(\mathbf{A}_L))\|_F^2 + \|\mathcal{P}_{\Omega^*}(\mathbf{A}^-)\|_F^2 \leq \frac{4}{p} \|\mathcal{P}_{\Omega_{\text{obs}}}[\mathcal{P}_{T^*}(\mathbf{A}^- + \mathcal{P}_{\Omega_{\text{obs}}^c}(\mathbf{A}_L)) - \mathcal{P}_{\Omega^*}(\mathbf{A}^-)]\|_F^2.$$

By virtue of the identity

$$\begin{aligned} \mathcal{P}_{T^*}(\mathbf{A}^- + \mathcal{P}_{\Omega_{\text{obs}}^c}(\mathbf{A}_L)) - \mathcal{P}_{\Omega^*}(\mathbf{A}^-) &= \mathbf{A}^- + \mathcal{P}_{\Omega_{\text{obs}}^c}(\mathbf{A}_L) - \mathcal{P}_{T^* \perp}(\mathbf{A}^- + \mathcal{P}_{\Omega_{\text{obs}}^c}(\mathbf{A}_L)) - \mathbf{A}^- + \mathcal{P}_{(\Omega^*)^c}(\mathbf{A}^-) \\ &= \mathcal{P}_{\Omega_{\text{obs}}^c}(\mathbf{A}_L) - \mathcal{P}_{T^* \perp}(\mathbf{A}^- + \mathcal{P}_{\Omega_{\text{obs}}^c}(\mathbf{A}_L)) + \mathcal{P}_{(\Omega^*)^c}(\mathbf{A}^-), \end{aligned}$$

we further obtain

$$\begin{aligned} \alpha_2 &\leq \frac{4}{p} \|\mathcal{P}_{\Omega_{\text{obs}}}[\mathcal{P}_{\Omega_{\text{obs}}^c}(\mathbf{A}_L) - \mathcal{P}_{T^* \perp}(\mathbf{A}^- + \mathcal{P}_{\Omega_{\text{obs}}^c}(\mathbf{A}_L)) + \mathcal{P}_{(\Omega^*)^c}(\mathbf{A}^-)]\|_F^2 \\ &= \frac{4}{p} \|\mathcal{P}_{\Omega_{\text{obs}}}[\mathcal{P}_{T^* \perp}(\mathbf{A}^- + \mathcal{P}_{\Omega_{\text{obs}}^c}(\mathbf{A}_L)) - \mathcal{P}_{(\Omega^*)^c}(\mathbf{A}^-)]\|_F^2 \\ &\leq \frac{4}{p} \|\mathcal{P}_{T^* \perp}(\mathbf{A}^- + \mathcal{P}_{\Omega_{\text{obs}}^c}(\mathbf{A}_L)) - \mathcal{P}_{(\Omega^*)^c}(\mathbf{A}^-)\|_F^2 \\ &\leq \frac{8}{p} \|\mathcal{P}_{T^* \perp}(\mathbf{A}^- + \mathcal{P}_{\Omega_{\text{obs}}^c}(\mathbf{A}_L))\|_F^2 + \frac{8}{p} \|\mathcal{P}_{(\Omega^*)^c}(\mathbf{A}^-)\|_F^2 = \frac{8}{p} \alpha_3. \end{aligned} \quad (\text{D.7})$$

Once again, the derivation has made use of the elementary inequality $\|\mathbf{A} + \mathbf{B}\|_F^2 \leq 2\|\mathbf{A}\|_F^2 + 2\|\mathbf{B}\|_F^2$.

Step 3: bounding α_3 via α_1 and $\|\mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{E})\|_F$. The following lemma proves useful in linking α_3 with α_1 , and we postpone the proof to Appendix D.2.

Lemma 7. *Assume that $n^2p \gg n \log n$, $\rho_s \ll 1$ and $\|\mathcal{P}_{T^*} - p^{-1}(1 - \rho_s)^{-1} \mathcal{P}_{T^*} \mathcal{P}_{\Omega_{\text{obs}} \setminus \Omega^*} \mathcal{P}_{T^*}\| \leq 1/2$. Further assume that there exists a dual certificate $\mathbf{W} \in \mathbb{R}^{n \times n}$ such that*

$$\|\mathcal{P}_{T^*}[\lambda \mathbf{W} + \tau \text{sgn}(\mathbf{S}^*) - \lambda \mathbf{U}^* \mathbf{V}^{*\top}]\|_F \leq \tau / \sqrt{n}, \quad (\text{D.8a})$$

$$\|\mathcal{P}_{T^* \perp}[\lambda \mathbf{W} + \tau \text{sign}(\mathbf{S}^*)]\| < \lambda / 2, \quad (\text{D.8b})$$

$$\mathcal{P}_{(\Omega_{\text{obs}} \setminus \Omega^*)^c}(\mathbf{W}) = \mathbf{0}, \quad (\text{D.8c})$$

$$\|\lambda \mathbf{W}\|_\infty < \tau / 2, \quad (\text{D.8d})$$

where $\text{sign}(\mathbf{S}^*) := [\text{sign}(S_{ij}^*)]_{1 \leq i, j \leq n}$. Then for any $\mathbf{H}_L, \mathbf{H}_S \in \mathbb{R}^{n \times n}$ satisfying $\mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{H}_L) + \mathbf{H}_S = \mathbf{0}$, one has

$$\lambda \|\mathbf{L}^* + \mathbf{H}_L\|_* + \tau \|\mathbf{S}^* + \mathbf{H}_S\|_1 \geq \lambda \|\mathbf{L}^*\|_* + \tau \|\mathbf{S}^*\|_1 + \frac{\lambda}{4} \|\mathcal{P}_{T^* \perp}(\mathbf{H}_L)\|_* + \frac{\tau}{4} \|\mathcal{P}_{\Omega_{\text{obs}} \setminus \Omega^*}(\mathbf{H}_S)\|_1.$$

Again, we assume for the moment that the assumptions in Lemma 7 hold. Setting $\mathbf{H}_L = \mathbf{\Lambda}^- + \mathcal{P}_{\Omega_{\text{obs}}^c}(\mathbf{\Lambda}_L)$ and $\mathbf{H}_S = -\mathbf{\Lambda}^-$ in Lemma 7 gives

$$\begin{aligned} & \lambda \|\mathbf{L}^* + \mathbf{\Lambda}^- + \mathcal{P}_{\Omega_{\text{obs}}^c}(\mathbf{\Lambda}_L)\|_* + \tau \|\mathbf{S}^* - \mathbf{\Lambda}^-\|_1 \\ & \geq \lambda \|\mathbf{L}^*\|_* + \tau \|\mathbf{S}^*\|_1 + \frac{\lambda}{4} \|\mathcal{P}_{T^{\perp}}(\mathbf{\Lambda}^- + \mathcal{P}_{\Omega_{\text{obs}}^c}(\mathbf{\Lambda}_L))\|_* + \frac{\tau}{4} \|\mathcal{P}_{\Omega_{\text{obs}} \setminus \Omega^*}(\mathbf{\Lambda}^-)\|_1. \end{aligned}$$

In addition, recalling the identities $\mathbf{L}_{\text{cvx}} = \mathbf{L}^* + \mathbf{\Lambda}^+ + \mathbf{\Lambda}^- + \mathcal{P}_{\Omega_{\text{obs}}^c}(\mathbf{\Lambda}_L)$ and $\mathbf{S}_{\text{cvx}} = \mathbf{S}^* + \mathbf{\Lambda}^+ - \mathbf{\Lambda}^-$, we can invoke the triangle inequality to obtain

$$\begin{aligned} \lambda \|\mathbf{L}_{\text{cvx}}\|_* + \tau \|\mathbf{S}_{\text{cvx}}\|_1 &= \lambda \|\mathbf{L}^* + \mathbf{\Lambda}^- + \mathbf{\Lambda}^+ + \mathcal{P}_{\Omega_{\text{obs}}^c}(\mathbf{\Lambda}_L)\|_* + \tau \|\mathbf{S}^* - \mathbf{\Lambda}^- + \mathbf{\Lambda}^+\|_1 \\ &\geq \lambda \|\mathbf{L}^* + \mathbf{\Lambda}^- + \mathcal{P}_{\Omega_{\text{obs}}^c}(\mathbf{\Lambda}_L)\|_* + \tau \|\mathbf{S}^* - \mathbf{\Lambda}^-\|_1 - \lambda \|\mathbf{\Lambda}^+\|_* - \tau \|\mathbf{\Lambda}^+\|_1. \end{aligned}$$

Adding the above two inequalities and using the fact $\text{support}(\mathbf{\Lambda}^-) \subseteq \Omega_{\text{obs}}$ lead to

$$\begin{aligned} & \frac{\lambda}{4} \|\mathcal{P}_{T^{\perp}}(\mathbf{\Lambda}^- + \mathcal{P}_{\Omega_{\text{obs}}^c}(\mathbf{\Lambda}_L))\|_* + \frac{\tau}{4} \|\mathcal{P}_{(\Omega^*)^c}(\mathbf{\Lambda}^-)\|_1 = \frac{\lambda}{4} \|\mathcal{P}_{T^{\perp}}(\mathbf{\Lambda}^- + \mathcal{P}_{\Omega_{\text{obs}}^c}(\mathbf{\Lambda}_L))\|_* + \frac{\tau}{4} \|\mathcal{P}_{\Omega_{\text{obs}} \setminus \Omega^*}(\mathbf{\Lambda}^-)\|_1 \\ & \leq \lambda \|\mathbf{L}_{\text{cvx}}\|_* + \tau \|\mathbf{S}_{\text{cvx}}\|_1 + \lambda \|\mathbf{\Lambda}^+\|_* + \tau \|\mathbf{\Lambda}^+\|_1 - \lambda \|\mathbf{L}^*\|_* - \tau \|\mathbf{S}^*\|_1 \\ & \leq \frac{1}{2} \|\mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{L}^* + \mathbf{S}^* - \mathbf{M})\|_{\text{F}}^2 - \frac{1}{2} \|\mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{L}_{\text{cvx}} + \mathbf{S}_{\text{cvx}} - \mathbf{M})\|_{\text{F}}^2 + \lambda \|\mathbf{\Lambda}^+\|_* + \tau \|\mathbf{\Lambda}^+\|_1 \\ & \leq \frac{1}{2} \|\mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{E})\|_{\text{F}}^2 + 4\lambda\sqrt{n} \|\mathbf{\Lambda}^+\|_{\text{F}}. \end{aligned} \tag{D.9}$$

Here, the penultimate line results from the inequality (D.3) and last line follows from the same argument in obtaining (D.4).

We are now ready to establish the upper bound on α_3 . Invoke the elementary inequalities $\|\mathbf{A}\|_{\text{F}} \leq \|\mathbf{A}\|_*$ and $\|\mathbf{A}\|_{\text{F}} \leq \|\mathbf{A}\|_1$ for any $\mathbf{A} \in \mathbb{R}^{n \times n}$ to show that

$$\begin{aligned} \alpha_3 &= \|\mathcal{P}_{T^{\perp}}(\mathbf{\Lambda}^- + \mathcal{P}_{\Omega_{\text{obs}}^c}(\mathbf{\Lambda}_L))\|_{\text{F}}^2 + \|\mathcal{P}_{(\Omega^*)^c}(\mathbf{\Lambda}^-)\|_{\text{F}}^2 \leq \|\mathcal{P}_{T^{\perp}}(\mathbf{\Lambda}^- + \mathcal{P}_{\Omega_{\text{obs}}^c}(\mathbf{\Lambda}_L))\|_*^2 + \|\mathcal{P}_{(\Omega^*)^c}(\mathbf{\Lambda}^-)\|_1^2 \\ &\leq \left(\frac{16}{\lambda^2} + \frac{16}{\tau^2}\right) \left(\frac{\lambda}{4} \|\mathcal{P}_{T^{\perp}}(\mathbf{\Lambda}^- + \mathcal{P}_{\Omega_{\text{obs}}^c}(\mathbf{\Lambda}_L))\|_*^2 + \frac{\tau}{4} \|\mathcal{P}_{(\Omega^*)^c}(\mathbf{\Lambda}^-)\|_1^2\right)^2. \end{aligned}$$

This combined with (D.9) allows us to obtain

$$\alpha_3 \leq \left(\frac{16}{\lambda^2} + \frac{16}{\tau^2}\right) \left(\frac{1}{2} \|\mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{E})\|_{\text{F}}^2 + 4\lambda\sqrt{n} \|\mathbf{\Lambda}^+\|_{\text{F}}\right)^2 \leq 32 \left(\frac{1}{\lambda^2} + \frac{1}{\tau^2}\right) \left(\frac{1}{4} \|\mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{E})\|_{\text{F}}^4 + 16\lambda^2 n \|\mathbf{\Lambda}^+\|_{\text{F}}^2\right),$$

where we have used the elementary inequality $(a+b)^2 \leq 2a^2 + 2b^2$. Recalling that $\tau = \lambda/\sqrt{np/\log n}$ and that $np \geq 1$, we arrive at

$$\alpha_3 \leq \frac{64np}{\lambda^2} \left(\frac{1}{4} \|\mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{E})\|_{\text{F}}^4 + 16\lambda^2 n \|\mathbf{\Lambda}^+\|_{\text{F}}^2\right) = \frac{16np}{\lambda^2} \|\mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{E})\|_{\text{F}}^4 + 2^9 n^2 p \alpha_1, \tag{D.10}$$

where we have identified $2\|\mathbf{\Lambda}^+\|_{\text{F}}^2$ with α_1 .

Step 4: putting the above bounds on $\alpha_1, \alpha_2, \alpha_3$ together. Taking the preceding bounds on α_1, α_2 and α_3 collectively yields

$$\begin{aligned} \|\mathbf{\Lambda}_L\|_{\text{F}}^2 + \|\mathbf{\Lambda}_S\|_{\text{F}}^2 &= \alpha_1 + \alpha_2 + \alpha_3 \stackrel{(i)}{\leq} \alpha_1 + \left(1 + \frac{8}{p}\right) \alpha_3 \stackrel{(ii)}{\leq} \alpha_1 + \frac{16}{p} \alpha_3 \\ &\stackrel{(iii)}{\leq} (2^{13} n^2 + 1) \alpha_1 + \frac{2^8 n}{\lambda^2} \|\mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{E})\|_{\text{F}}^4 \\ &\stackrel{(iv)}{\leq} (2^{13} n^2 + 1) \left[2 \|\mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{E})\|_{\text{F}}^2 + 4\lambda\sqrt{n \log n} \sqrt{\|\mathbf{\Lambda}_L\|_{\text{F}}^2 + \|\mathbf{\Lambda}_S\|_{\text{F}}^2}\right] + \frac{2^8 n}{\lambda^2} \|\mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{E})\|_{\text{F}}^4. \end{aligned}$$

Here, the first inequality (i) comes from (D.7), the second (ii) follows from the fact $1 \leq 8/p$, the third relation (iii) is a consequence of (D.10), and the last line (iv) results from (D.5). Note that this forms a quadratic inequality in $\sqrt{\|\mathbf{\Lambda}_L\|_F^2 + \|\mathbf{\Lambda}_S\|_F^2}$. Solving the inequality yields the claimed bound

$$\|\mathbf{\Lambda}_L\|_F^2 + \|\mathbf{\Lambda}_S\|_F^2 \lesssim \lambda^2 n^5 \log n + n^2 \|\mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{E})\|_F^2 + \frac{n}{\lambda^2} \|\mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{E})\|_F^4.$$

Further, the elementary inequality $a^2 + b^2 \geq 2ab$ yields

$$\lambda^2 n^5 + \frac{n}{\lambda^2} \|\mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{E})\|_F^4 \geq 2n^3 \|\mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{E})\|_F^2 \geq n^2 \|\mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{E})\|_F^2,$$

leading to the simplified bound

$$\|\mathbf{\Lambda}_L\|_F^2 + \|\mathbf{\Lambda}_S\|_F^2 \lesssim \lambda^2 n^5 \log n + \frac{n}{\lambda^2} \|\mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{E})\|_F^4.$$

Step 5: checking the conditions in Lemmas 6 and 7. We are left with proving that the conditions in Lemmas 6 and 7 hold with high probability. In view of Lemma 3 and Corollary 3, the conditions $\|\mathcal{P}_{T^*} - p^{-1}\mathcal{P}_{T^*}\mathcal{P}_{\Omega_{\text{obs}}}\mathcal{P}_{T^*}\| \leq 1/2$ and $\|\mathcal{P}_{\Omega^*}\mathcal{P}_{T^*}\|^2 \leq p/8$ hold with high probability, provided that $n^2p \gg \mu rn \log n$ and $\rho_s \leq 1/12$. In addition, Lemma 3 ensures that $\|\mathcal{P}_{T^*} - p^{-1}(1 - \rho_s)^{-1}\mathcal{P}_{T^*}\mathcal{P}_{\Omega_{\text{obs}} \setminus \Omega^*}\mathcal{P}_{T^*}\| \leq 1/2$ holds with high probability, with the proviso that $n^2p(1 - \rho_s) \gg \mu rn \log n$, which holds true under the assumptions $\rho_s \leq 1/12$ and $n^2p \gg \mu rn \log n$. Last but not least, the existence of the dual certificate \mathbf{W} obeying (D.8) is guaranteed with high probability according to [CJSC13, Section III.D], under the conditions $\rho_s \ll 1$ and $n^2p \gg \mu^2 r^2 n \log^6 n$.⁹

D.1 Proof of Lemma 6

Expand $\|\mathcal{P}_{\Omega_{\text{obs}}}(\mathcal{P}_{T^*}(\mathbf{A}) + \mathcal{P}_{\Omega^*}(\mathbf{B}))\|_F^2$ to obtain

$$\begin{aligned} \|\mathcal{P}_{\Omega_{\text{obs}}}[\mathcal{P}_{T^*}(\mathbf{A}) + \mathcal{P}_{\Omega^*}(\mathbf{B})]\|_F^2 &= \|\mathcal{P}_{\Omega_{\text{obs}}}\mathcal{P}_{T^*}(\mathbf{A})\|_F^2 + \|\mathcal{P}_{\Omega^*}(\mathbf{B})\|_F^2 + 2\langle \mathcal{P}_{\Omega_{\text{obs}}}\mathcal{P}_{T^*}(\mathbf{A}), \mathcal{P}_{\Omega^*}(\mathbf{B}) \rangle \\ &\geq \frac{p}{2} \|\mathcal{P}_{T^*}(\mathbf{A})\|_F^2 + \|\mathcal{P}_{\Omega^*}(\mathbf{B})\|_F^2 + 2\langle \mathcal{P}_{\Omega_{\text{obs}}}\mathcal{P}_{T^*}(\mathbf{A}), \mathcal{P}_{\Omega^*}(\mathbf{B}) \rangle. \end{aligned}$$

Here, the equality uses the fact $\Omega^* \subseteq \Omega_{\text{obs}}$, and the inequality holds because of the assumption $\|\mathcal{P}_{T^*} - p^{-1}\mathcal{P}_{T^*}\mathcal{P}_{\Omega_{\text{obs}}}\mathcal{P}_{T^*}\| \leq 1/2$ and Fact 1. Use $\Omega^* \subseteq \Omega_{\text{obs}}$ once again to obtain

$$\begin{aligned} 2\langle \mathcal{P}_{\Omega_{\text{obs}}}\mathcal{P}_{T^*}(\mathbf{A}), \mathcal{P}_{\Omega^*}(\mathbf{B}) \rangle &= 2\langle \mathcal{P}_{\Omega^*}\mathcal{P}_{T^*}\mathcal{P}_{T^*}(\mathbf{A}), \mathcal{P}_{\Omega^*}(\mathbf{B}) \rangle \\ &\geq -2\|\mathcal{P}_{\Omega^*}\mathcal{P}_{T^*}\| \|\mathcal{P}_{T^*}(\mathbf{A})\|_F \|\mathcal{P}_{\Omega^*}(\mathbf{B})\|_F \\ &\geq -2\|\mathcal{P}_{\Omega^*}\mathcal{P}_{T^*}\|^2 \|\mathcal{P}_{T^*}(\mathbf{A})\|_F^2 - \frac{1}{2} \|\mathcal{P}_{\Omega^*}(\mathbf{B})\|_F^2. \end{aligned}$$

Here, the last relation arises from the elementary inequality $ab \leq (a^2 + b^2)/2$ and the fact that $\|\mathcal{P}_{\Omega^*}\mathcal{P}_{T^*}\| \leq 1$. Combine the above two inequalities to obtain

$$\begin{aligned} \|\mathcal{P}_{\Omega_{\text{obs}}}[\mathcal{P}_{T^*}(\mathbf{A}) + \mathcal{P}_{\Omega^*}(\mathbf{B})]\|_F^2 &\geq \left(\frac{p}{2} - 2\|\mathcal{P}_{\Omega^*}\mathcal{P}_{T^*}\|^2\right) \|\mathcal{P}_{T^*}(\mathbf{A})\|_F^2 + \frac{1}{2} \|\mathcal{P}_{\Omega^*}(\mathbf{B})\|_F^2 \\ &\geq \frac{p}{4} \|\mathcal{P}_{T^*}(\mathbf{A})\|_F^2 + \frac{1}{2} \|\mathcal{P}_{\Omega^*}(\mathbf{B})\|_F^2 \\ &\geq \frac{p}{4} \left(\|\mathcal{P}_{T^*}(\mathbf{A})\|_F^2 + \|\mathcal{P}_{\Omega^*}(\mathbf{B})\|_F^2 \right) \end{aligned}$$

as claimed, where we have used the assumption $\|\mathcal{P}_{\Omega^*}\mathcal{P}_{T^*}\|^2 \leq p/8$ in the middle line and the fact $1/2 \geq p/4$ in the last inequality.

⁹Note that [CJSC13, Section III.D] requires $n^2p \gg \max\{\mu, \mu_2\}rn \log^6 n$ under an additional incoherence condition $\|\mathbf{U}^*\mathbf{V}^{*\top}\|_\infty \leq \sqrt{\mu_2 r/n^2}$. While we do not impose this extra condition, it is easily seen that $\|\mathbf{U}^*\mathbf{V}^{*\top}\|_\infty \leq \|\mathbf{U}^*\|_{2,\infty} \|\mathbf{V}^*\|_{2,\infty} \leq \mu r/n$ and hence $\mu_2 \leq \mu^2 r$.

D.2 Proof of Lemma 7

In view of the convexity of the nuclear norm $\|\cdot\|_*$, one has

$$\|\mathbf{L}^* + \mathbf{H}_L\|_* \geq \|\mathbf{L}^*\|_* + \langle \mathbf{U}^* \mathbf{V}^{*\top} + \mathbf{G}_1, \mathbf{H}_L \rangle = \|\mathbf{L}^*\|_* + \langle \mathbf{U}^* \mathbf{V}^{*\top}, \mathbf{H}_L \rangle + \|\mathcal{P}_{T^{\perp}}(\mathbf{H}_L)\|_*.$$

Here, $\mathbf{U}^* \mathbf{V}^{*\top} + \mathbf{G}_1$ is a sub-gradient of $\|\cdot\|_*$ at \mathbf{L}^* . The last identity holds by choosing \mathbf{G}_1 such that $\langle \mathbf{G}_1, \mathbf{H}_L \rangle = \|\mathcal{P}_{T^{\perp}}(\mathbf{H}_L)\|_*$. Similarly, using the assumption $\mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{H}_L) + \mathbf{H}_S = \mathbf{0}$ and the convexity of the ℓ_1 norm $\|\cdot\|_1$, we can obtain

$$\begin{aligned} \|\mathbf{S}^* + \mathbf{H}_S\|_1 &= \|\mathbf{S}^* - \mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{H}_L)\|_1 \\ &\geq \|\mathbf{S}^*\|_1 - \langle \text{sign}(\mathbf{S}^*) + \mathbf{G}_2, \mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{H}_L) \rangle \\ &\stackrel{(i)}{=} \|\mathbf{S}^*\|_1 - \langle \text{sign}(\mathbf{S}^*), \mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{H}_L) \rangle + \|\mathcal{P}_{\Omega_{\text{obs}} \setminus \Omega^*}(\mathbf{H}_L)\|_1 \\ &\stackrel{(ii)}{=} \|\mathbf{S}^*\|_1 - \langle \text{sign}(\mathbf{S}^*), \mathbf{H}_L \rangle + \|\mathcal{P}_{\Omega_{\text{obs}} \setminus \Omega^*}(\mathbf{H}_L)\|_1, \end{aligned}$$

where $\text{sign}(\mathbf{S}^*) := [\text{sign}(S_{ij}^*)]_{1 \leq i, j \leq n}$, and $\text{sign}(\mathbf{S}^*) + \mathbf{G}_2$ is a sub-gradient of $\|\cdot\|_1$ at \mathbf{S}^* . The first equality (i) holds by choosing \mathbf{G}_2 such that $-\langle \mathbf{G}_2, \mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{H}_L) \rangle = \|\mathcal{P}_{\Omega_{\text{obs}} \setminus \Omega^*}(\mathbf{H}_L)\|_1$, and the last relation (ii) arises since $\text{sign}(\mathbf{S}^*)$ is supported on $\Omega^* \subseteq \Omega_{\text{obs}}$. Combine the above two bounds to deduce that

$$\begin{aligned} \Delta &:= \lambda \|\mathbf{L}^* + \mathbf{H}_L\|_* + \tau \|\mathbf{S}^* + \mathbf{H}_S\|_1 - \lambda \|\mathbf{L}^*\|_* - \tau \|\mathbf{S}^*\|_1 \\ &\geq \lambda \langle \mathbf{U}^* \mathbf{V}^{*\top}, \mathbf{H}_L \rangle + \lambda \|\mathcal{P}_{T^{\perp}}(\mathbf{H}_L)\|_* - \tau \langle \text{sgn}(\mathbf{S}^*), \mathbf{H}_L \rangle + \tau \|\mathcal{P}_{\Omega_{\text{obs}} \setminus \Omega^*}(\mathbf{H}_L)\|_1 \\ &= \langle \lambda \mathbf{U}^* \mathbf{V}^{*\top} - \tau \text{sgn}(\mathbf{S}^*), \mathbf{H}_L \rangle + \lambda \|\mathcal{P}_{T^{\perp}}(\mathbf{H}_L)\|_* + \tau \|\mathcal{P}_{\Omega_{\text{obs}} \setminus \Omega^*}(\mathbf{H}_L)\|_1 \\ &= \underbrace{\langle \lambda \mathbf{U}^* \mathbf{V}^{*\top} - \tau \text{sgn}(\mathbf{S}^*) - \lambda \mathbf{W}, \mathbf{H}_L \rangle}_{=: \theta_1} + \underbrace{\langle \lambda \mathbf{W}, \mathbf{H}_L \rangle}_{=: \theta_2} + \lambda \|\mathcal{P}_{T^{\perp}}(\mathbf{H}_L)\|_* + \tau \|\mathcal{P}_{\Omega_{\text{obs}} \setminus \Omega^*}(\mathbf{H}_L)\|_1, \quad (\text{D.11}) \end{aligned}$$

where $\mathbf{W} \in \mathbb{R}^{n \times n}$ is the dual certificate stated in Lemma 7.

In what follows, we shall lower bound the right-hand side of (D.11). To begin with, for θ_1 we have

$$\begin{aligned} \theta_1 &= \langle \mathcal{P}_{T^*} [\lambda \mathbf{U}^* \mathbf{V}^{*\top} - \tau \text{sgn}(\mathbf{S}^*) - \lambda \mathbf{W}], \mathcal{P}_{T^*}(\mathbf{H}_L) \rangle + \langle \mathcal{P}_{T^{\perp}} [\lambda \mathbf{U}^* \mathbf{V}^{*\top} - \tau \text{sgn}(\mathbf{S}^*) - \lambda \mathbf{W}], \mathcal{P}_{T^{\perp}}(\mathbf{H}_L) \rangle \\ &= \langle \mathcal{P}_{T^*} [\lambda \mathbf{U}^* \mathbf{V}^{*\top} - \tau \text{sgn}(\mathbf{S}^*) - \lambda \mathbf{W}], \mathcal{P}_{T^*}(\mathbf{H}_L) \rangle - \langle \mathcal{P}_{T^{\perp}} [\tau \text{sgn}(\mathbf{S}^*) + \lambda \mathbf{W}], \mathcal{P}_{T^{\perp}}(\mathbf{H}_L) \rangle \\ &\geq -\|\mathcal{P}_{T^*} [\lambda \mathbf{U}^* \mathbf{V}^{*\top} - \tau \text{sgn}(\mathbf{S}^*) - \lambda \mathbf{W}]\|_{\text{F}} \|\mathcal{P}_{T^*}(\mathbf{H}_L)\|_{\text{F}} - \|\mathcal{P}_{T^{\perp}} [\tau \text{sgn}(\mathbf{S}^*) + \lambda \mathbf{W}]\| \|\mathcal{P}_{T^{\perp}}(\mathbf{H}_L)\|_* \\ &\geq -\frac{\tau}{\sqrt{n}} \|\mathcal{P}_{T^*}(\mathbf{H}_L)\|_{\text{F}} - \frac{\lambda}{2} \|\mathcal{P}_{T^{\perp}}(\mathbf{H}_L)\|_*. \end{aligned}$$

Here, the penultimate line uses the fact $\mathbf{U}^* \mathbf{V}^{*\top} \in T^*$ and the elementary inequalities $|\langle \mathbf{A}, \mathbf{B} \rangle| \leq \|\mathbf{A}\|_{\text{F}} \|\mathbf{B}\|_{\text{F}}$ and $|\langle \mathbf{A}, \mathbf{B} \rangle| \leq \|\mathbf{A}\| \|\mathbf{B}\|_*$, whereas the last inequality relies on the properties of the dual certificate \mathbf{W} , namely, (D.8a) and (D.8b). Moving on to θ_2 , one has

$$\begin{aligned} \theta_2 &= \langle \lambda \mathcal{P}_{\Omega_{\text{obs}} \setminus \Omega^*}(\mathbf{W}), \mathbf{H}_L \rangle + \langle \lambda \mathcal{P}_{(\Omega_{\text{obs}} \setminus \Omega^*)^c}(\mathbf{W}), \mathbf{H}_L \rangle \\ &= \langle \lambda \mathbf{W}, \mathcal{P}_{\Omega_{\text{obs}} \setminus \Omega^*}(\mathbf{H}_L) \rangle \stackrel{(i)}{\geq} -\|\lambda \mathbf{W}\|_{\infty} \|\mathcal{P}_{\Omega_{\text{obs}} \setminus \Omega^*}(\mathbf{H}_L)\|_1 \stackrel{(ii)}{\geq} -\frac{\tau}{2} \|\mathcal{P}_{\Omega_{\text{obs}} \setminus \Omega^*}(\mathbf{H}_L)\|_1. \end{aligned}$$

Here, the second identity uses the assumption (D.8c), the first inequality (i) uses the elementary inequality $|\langle \mathbf{A}, \mathbf{B} \rangle| \leq \|\mathbf{A}\|_{\infty} \|\mathbf{B}\|_1$, and the last relation (ii) holds because of the assumption (D.8d). Substituting the above two bounds back into (D.11) gives

$$\Delta \geq -\frac{\tau}{\sqrt{n}} \|\mathcal{P}_{T^*}(\mathbf{H}_L)\|_{\text{F}} + \frac{\lambda}{2} \|\mathcal{P}_{T^{\perp}}(\mathbf{H}_L)\|_* + \frac{\tau}{2} \|\mathcal{P}_{\Omega_{\text{obs}} \setminus \Omega^*}(\mathbf{H}_L)\|_1. \quad (\text{D.12})$$

Continuing the lower bound, we have

$$\|\mathcal{P}_{\Omega_{\text{obs}} \setminus \Omega^*}(\mathbf{H}_L)\|_1 \stackrel{(i)}{\geq} \|\mathcal{P}_{\Omega_{\text{obs}} \setminus \Omega^*}(\mathbf{H}_L)\|_{\text{F}} = \|\mathcal{P}_{\Omega_{\text{obs}} \setminus \Omega^*} \mathcal{P}_{T^*}(\mathbf{H}_L) + \mathcal{P}_{\Omega_{\text{obs}} \setminus \Omega^*} \mathcal{P}_{T^{\perp}}(\mathbf{H}_L)\|_{\text{F}}$$

$$\begin{aligned}
&\stackrel{\text{(ii)}}{\geq} \|\mathcal{P}_{\Omega_{\text{obs}} \setminus \Omega^*} \mathcal{P}_{T^*}(\mathbf{H}_L)\|_{\text{F}} - \|\mathcal{P}_{\Omega_{\text{obs}} \setminus \Omega^*} \mathcal{P}_{T^{\perp}}(\mathbf{H}_L)\|_{\text{F}} \\
&\geq \|\mathcal{P}_{\Omega_{\text{obs}} \setminus \Omega^*} \mathcal{P}_{T^*}(\mathbf{H}_L)\|_{\text{F}} - \|\mathcal{P}_{T^{\perp}}(\mathbf{H}_L)\|_{\text{F}},
\end{aligned}$$

where (i) holds because $\|\mathbf{A}\|_1 \geq \|\mathbf{A}\|_{\text{F}}$ for any matrix \mathbf{A} , and (ii) arises from the triangle inequality. Putting the above relation and (D.12) together results in

$$\begin{aligned}
\Delta &\geq -\frac{\tau}{\sqrt{n}} \|\mathcal{P}_{T^*}(\mathbf{H}_L)\|_{\text{F}} + \left(\frac{\lambda}{2} - \frac{\tau}{4}\right) \|\mathcal{P}_{T^{\perp}}(\mathbf{H}_L)\|_* + \frac{\tau}{4} \|\mathcal{P}_{\Omega_{\text{obs}} \setminus \Omega^*} \mathcal{P}_{T^*}(\mathbf{H}_L)\|_{\text{F}} + \frac{\tau}{4} \|\mathcal{P}_{\Omega_{\text{obs}} \setminus \Omega^*}(\mathbf{H}_L)\|_1 \\
&\geq \frac{\tau}{4} \|\mathcal{P}_{\Omega_{\text{obs}} \setminus \Omega^*} \mathcal{P}_{T^*}(\mathbf{H}_L)\|_{\text{F}} - \frac{\tau}{\sqrt{n}} \|\mathcal{P}_{T^*}(\mathbf{H}_L)\|_{\text{F}} + \frac{\lambda}{4} \|\mathcal{P}_{T^{\perp}}(\mathbf{H}_L)\|_* + \frac{\tau}{4} \|\mathcal{P}_{\Omega_{\text{obs}} \setminus \Omega^*}(\mathbf{H}_L)\|_1, \tag{D.13}
\end{aligned}$$

where the last line holds since $\lambda = \tau \sqrt{np/\log n} \geq \tau$ (as long as $np \geq \log n$). Everything then boils down to lower bounding $\|\mathcal{P}_{\Omega_{\text{obs}} \setminus \Omega^*} \mathcal{P}_{T^*}(\mathbf{H}_L)\|_{\text{F}}$. To this end, one can use the assumption $\|\mathcal{P}_{T^*} - p^{-1}(1 - \rho_s)^{-1} \mathcal{P}_{T^*} \mathcal{P}_{\Omega_{\text{obs}} \setminus \Omega^*} \mathcal{P}_{T^*}\| \leq 1/2$ and Fact 1 to obtain

$$\|\mathcal{P}_{\Omega_{\text{obs}} \setminus \Omega^*} \mathcal{P}_{T^*}(\mathbf{H}_L)\|_{\text{F}}^2 \geq \frac{1}{2} p(1 - \rho_s) \|\mathcal{P}_{T^*}(\mathbf{H}_L)\|_{\text{F}}^2. \tag{D.14}$$

Take (D.13) and (D.14) collectively to yield

$$\begin{aligned}
\Delta &\geq \left(\frac{\tau}{4} \sqrt{\frac{1}{2} p(1 - \rho_s)} - \frac{\tau}{\sqrt{n}}\right) \|\mathcal{P}_{T^*}(\mathbf{H}_L)\|_{\text{F}} + \frac{\lambda}{4} \|\mathcal{P}_{T^{\perp}}(\mathbf{H}_L)\|_* + \frac{\tau}{4} \|\mathcal{P}_{\Omega_{\text{obs}} \setminus \Omega^*}(\mathbf{H}_L)\|_1 \\
&\geq \frac{\lambda}{4} \|\mathcal{P}_{T^{\perp}}(\mathbf{H}_L)\|_* + \frac{\tau}{4} \|\mathcal{P}_{\Omega_{\text{obs}} \setminus \Omega^*}(\mathbf{H}_L)\|_1,
\end{aligned}$$

where the last relation is guaranteed by $np \gg 1$ and $\rho_s \ll 1$. Recognizing that $\mathcal{P}_{\Omega_{\text{obs}} \setminus \Omega^*}(\mathbf{H}_L) = -\mathcal{P}_{\Omega_{\text{obs}} \setminus \Omega^*}(\mathbf{H}_S)$ finishes the proof.

E Equivalence between convex and nonconvex solutions (Proof of Theorem 4)

The goal of this section is to establish the intimate connection between the convex and nonconvex solutions (cf. Theorem 4). Before continuing, we remind the readers of the following notations:

- $\mathbf{X}\mathbf{Y}^\top = \mathbf{U}\Sigma\mathbf{V}^\top$: the rank- r singular value decomposition of $\mathbf{X}\mathbf{Y}^\top$;
- T : the tangent space of the set of rank- r matrices at the estimate $\mathbf{X}\mathbf{Y}^\top$.

In addition, we define

$$\Delta_L := \mathbf{L}_{\text{cvx}} - \mathbf{X}\mathbf{Y}^\top, \quad \Delta_S := \mathbf{S}_{\text{cvx}} - \mathbf{S}, \tag{E.1}$$

and denote the support of \mathbf{S} by

$$\Omega := \{(i, j) \mid S_{ij} \neq 0\}. \tag{E.2}$$

E.1 Preliminary facts

We begin with two useful lemmas which demonstrate that the point $(\mathbf{X}\mathbf{Y}^\top, \mathbf{S})$ described in Theorem 4 satisfies approximate optimality conditions w.r.t. the convex program (1.3).

Lemma 8. *Instate the assumptions in Theorem 4. The triple $(\mathbf{X}, \mathbf{Y}, \mathbf{S})$ as stated in Theorem 4 satisfies*

$$\frac{1}{\lambda} \mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{X}\mathbf{Y}^\top + \mathbf{S} - \mathbf{M}) = -\mathbf{U}\mathbf{V}^\top + \mathbf{R}_1 \tag{E.3}$$

for some matrix $\mathbf{R}_1 \in \mathbb{R}^{n \times n}$ obeying

$$\|\mathcal{P}_T(\mathbf{R}_1)\|_{\text{F}} \lesssim \frac{\kappa p \|\nabla f(\mathbf{X}, \mathbf{Y}; \mathbf{S})\|_{\text{F}}}{\lambda \sqrt{\sigma_{\min}}} \lesssim \frac{1}{n^{19}} \quad \text{and} \quad \|\mathcal{P}_{T^\perp}(\mathbf{R}_1)\| \leq \frac{1}{2}. \tag{E.4}$$

Proof. The proof can be straightforwardly adapted from [CCF⁺19, Claim 2] by replacing \mathbf{E} therein with $\mathbf{E} + \mathbf{S}^* - \mathbf{S}$. We omit it for the sake of brevity. \square

Lemma 9. *The point $(\mathbf{X}\mathbf{Y}^\top, \mathbf{S})$ as stated in Theorem 4 obeys*

$$\frac{1}{\tau} \mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{X}\mathbf{Y}^\top + \mathbf{S} - \mathbf{M}) = -\text{sign}(\mathbf{S}) + \mathbf{R}_2 \quad (\text{E.5})$$

for some matrix $\mathbf{R}_2 \in \mathbb{R}^{n \times n}$, where \mathbf{R}_2 satisfies

$$\mathcal{P}_\Omega(\mathbf{R}_2) = \mathbf{0} \quad \text{and} \quad \|\mathcal{P}_{\Omega^c}(\mathbf{R}_2)\|_\infty \leq 1 \quad (\text{E.6})$$

with Ω defined in (E.2).

Proof. By definition, one has $\mathbf{S} = \mathcal{P}_{\Omega_{\text{obs}}}[\mathcal{S}_\tau(\mathbf{M} - \mathbf{X}\mathbf{Y}^\top)]$. Clearly, this is equivalent to saying that \mathbf{S} is the unique minimizer of the following convex program

$$\mathbf{S} = \arg \min_{\widehat{\mathbf{S}} \in \mathbb{R}^{n \times n}} \frac{1}{2} \left\| \mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{X}\mathbf{Y}^\top + \widehat{\mathbf{S}} - \mathbf{M}) \right\|_{\text{F}}^2 + \frac{\lambda}{2} \|\mathbf{X}\|_{\text{F}}^2 + \frac{\lambda}{2} \|\mathbf{Y}\|_{\text{F}}^2 + \tau \|\widehat{\mathbf{S}}\|_1. \quad (\text{E.7})$$

The claim of this lemma then follows from the optimality condition of this convex program (E.7). \square

Additionally, in view of the crude error bound (3.8) and Condition 1, the matrix $\Delta_{\mathbf{L}}$ (cf. (E.1)) obeys

$$\begin{aligned} \|\Delta_{\mathbf{L}}\|_{\text{F}} &= \|\mathbf{L}_{\text{cvx}} - \mathbf{X}\mathbf{Y}^\top\|_{\text{F}} \leq \|\mathbf{L}_{\text{cvx}} - \mathbf{L}^*\|_{\text{F}} + \|\mathbf{X}\mathbf{Y}^\top - \mathbf{L}^*\|_{\text{F}} \leq \|\mathbf{L}_{\text{cvx}} - \mathbf{L}^*\|_{\text{F}} + n \|\mathbf{X}\mathbf{Y}^\top - \mathbf{L}^*\|_\infty \\ &\lesssim \sigma n^4 + n\tau \asymp \sigma n^4, \end{aligned} \quad (\text{E.8})$$

where we use the elementary inequality $\|\mathbf{A}\|_{\text{F}} \leq n \|\mathbf{A}\|_\infty$ and the fact that $\tau \asymp \sigma \sqrt{\log n}$.

E.2 Proof of Theorem 4

We now present the proof of Theorem 4, which consists of three main steps:

1. Showing that $(\mathbf{X}\mathbf{Y}^\top, \mathbf{S})$ is not far from $(\mathbf{L}_{\text{cvx}}, \mathbf{S}_{\text{cvx}})$ over Ω_{obs} , in the sense that $\mathcal{P}_{\Omega_{\text{obs}}}(\Delta_{\mathbf{L}} + \Delta_{\mathbf{S}}) \approx \mathbf{0}$;
2. Showing that $\Delta_{\mathbf{L}}$ (resp. $\Delta_{\mathbf{S}}$) is extremely small outside the tangent space T (resp. the support Ω^*), and hence most of the energy of $\Delta_{\mathbf{L}}$ (resp. $\Delta_{\mathbf{S}}$) — if it is not vanishingly small — has to reside within T (resp. Ω^*);
3. Showing that $\Delta_{\mathbf{S}} \approx \mathbf{0}$ and $\Delta_{\mathbf{L}} \approx \mathbf{0}$, with the assistance of the preceding two steps.

In what follows, we shall detail each of these steps.

E.2.1 Step 1: showing that $\mathcal{P}_{\Omega_{\text{obs}}}(\Delta_{\mathbf{L}} + \Delta_{\mathbf{S}}) \approx \mathbf{0}$

Since $(\mathbf{L}_{\text{cvx}}, \mathbf{S}_{\text{cvx}})$ is the minimizer of the convex program (1.3), we have

$$\begin{aligned} &\frac{1}{2} \|\mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{L}_{\text{cvx}} + \mathbf{S}_{\text{cvx}} - \mathbf{M})\|_{\text{F}}^2 + \lambda \|\mathbf{L}_{\text{cvx}}\|_* + \tau \|\mathbf{S}_{\text{cvx}}\|_1 \\ &= \frac{1}{2} \|\mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{X}\mathbf{Y}^\top + \Delta_{\mathbf{L}} + \mathbf{S} + \Delta_{\mathbf{S}} - \mathbf{M})\|_{\text{F}}^2 + \lambda \|\mathbf{X}\mathbf{Y}^\top + \Delta_{\mathbf{L}}\|_* + \tau \|\mathbf{S} + \Delta_{\mathbf{S}}\|_1 \\ &\leq \frac{1}{2} \|\mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{X}\mathbf{Y}^\top + \mathbf{S} - \mathbf{M})\|_{\text{F}}^2 + \lambda \|\mathbf{X}\mathbf{Y}^\top\|_* + \tau \|\mathbf{S}\|_1. \end{aligned}$$

Here, the equality arises from the relations $\mathbf{L}_{\text{cvx}} = \mathbf{X}\mathbf{Y}^\top + \Delta_{\mathbf{L}}$ and $\mathbf{S}_{\text{cvx}} = \mathbf{S} + \Delta_{\mathbf{S}}$. Expanding the squares and rearranging terms, we arrive at

$$\begin{aligned} \frac{1}{2} \|\mathcal{P}_{\Omega_{\text{obs}}}(\Delta_{\mathbf{L}} + \Delta_{\mathbf{S}})\|_{\text{F}}^2 &\leq -\langle \mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{X}\mathbf{Y}^\top + \mathbf{S} - \mathbf{M}), \Delta_{\mathbf{L}} + \Delta_{\mathbf{S}} \rangle + \lambda \|\mathbf{X}\mathbf{Y}^\top\|_* - \lambda \|\mathbf{X}\mathbf{Y}^\top + \Delta_{\mathbf{L}}\|_* \\ &\quad + \tau \|\mathbf{S}\|_1 - \tau \|\mathbf{S} + \Delta_{\mathbf{S}}\|_1. \end{aligned} \quad (\text{E.9})$$

In view of the convexity of the nuclear norm $\|\cdot\|_*$ and the ℓ_1 norm $\|\cdot\|_1$, one has

$$\|\mathbf{X}\mathbf{Y}^\top + \Delta_L\|_* \geq \|\mathbf{X}\mathbf{Y}^\top\|_* + \langle \mathbf{U}\mathbf{V}^\top + \mathbf{W}, \Delta_L \rangle \stackrel{(i)}{=} \|\mathbf{X}\mathbf{Y}^\top\|_* + \langle \mathbf{U}\mathbf{V}^\top, \Delta_L \rangle + \|\mathcal{P}_{T^\perp}(\Delta_L)\|_*; \quad (\text{E.10a})$$

$$\|\mathbf{S} + \Delta_S\|_1 \geq \|\mathbf{S}\|_1 + \langle \text{sign}(\mathbf{S}) + \mathbf{G}, \Delta_S \rangle \stackrel{(ii)}{=} \|\mathbf{S}\|_1 + \langle \text{sign}(\mathbf{S}), \Delta_S \rangle + \|\mathcal{P}_{\Omega^c}(\Delta_S)\|_1. \quad (\text{E.10b})$$

Here, $\mathbf{U}\mathbf{V}^\top + \mathbf{W}$ is a sub-gradient of $\|\cdot\|_*$ at $\mathbf{X}\mathbf{Y}^\top$. The identity (i) holds by choosing \mathbf{W} such that $\langle \mathbf{W}, \Delta_L \rangle = \|\mathcal{P}_{T^\perp}(\Delta_L)\|_*$. Similarly, $\text{sign}(\mathbf{S}) + \mathbf{G}$ is a sub-gradient of $\|\cdot\|_1$ at \mathbf{S} and one can choose \mathbf{G} obeying $\langle \mathbf{G}, \Delta_S \rangle = \|\mathcal{P}_{\Omega^c}(\Delta_S)\|_1$ to make (ii) valid. These taken together with (E.9) lead to

$$\begin{aligned} \frac{1}{2} \|\mathcal{P}_{\Omega_{\text{obs}}}(\Delta_L + \Delta_S)\|_{\text{F}}^2 &\leq -\langle \mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{X}\mathbf{Y}^\top + \mathbf{S} - \mathbf{M}), \Delta_L + \Delta_S \rangle - \lambda \langle \mathbf{U}\mathbf{V}^\top, \Delta_L \rangle - \lambda \|\mathcal{P}_{T^\perp}(\Delta_L)\|_* \\ &\quad - \tau \langle \text{sign}(\mathbf{S}), \Delta_S \rangle - \tau \|\mathcal{P}_{\Omega^c}(\Delta_S)\|_1. \end{aligned}$$

Recall the definitions of \mathbf{R}_1 and \mathbf{R}_2 from Lemmas 8 and 9. We can then simplify the above inequality as

$$\frac{1}{2} \|\mathcal{P}_{\Omega_{\text{obs}}}(\Delta_L + \Delta_S)\|_{\text{F}}^2 \leq \underbrace{-\lambda \langle \mathbf{R}_1, \Delta_L \rangle}_{=:\theta_1} - \lambda \|\mathcal{P}_{T^\perp}(\Delta_L)\|_* - \underbrace{\tau \langle \mathbf{R}_2, \Delta_S \rangle - \tau \|\mathcal{P}_{\Omega^c}(\Delta_S)\|_1}_{=:\theta_2}. \quad (\text{E.11})$$

In the sequel, we develop bounds on θ_1 and θ_2 .

1. With regards to θ_1 , one can further decompose it into

$$\begin{aligned} \theta_1 &= -\lambda \langle \mathbf{R}_1, \mathcal{P}_T(\Delta_L) \rangle - \lambda \langle \mathbf{R}_1, \mathcal{P}_{T^\perp}(\Delta_L) \rangle - \lambda \|\mathcal{P}_{T^\perp}(\Delta_L)\|_* \\ &\leq \lambda \|\mathcal{P}_T(\mathbf{R}_1)\|_{\text{F}} \|\mathcal{P}_T(\Delta_L)\|_{\text{F}} - \lambda (1 - \|\mathcal{P}_{T^\perp}(\mathbf{R}_1)\|) \|\mathcal{P}_{T^\perp}(\Delta_L)\|_* \\ &\leq \lambda \|\mathcal{P}_T(\mathbf{R}_1)\|_{\text{F}} \|\mathcal{P}_T(\Delta_L)\|_{\text{F}} - \frac{\lambda}{2} \|\mathcal{P}_{T^\perp}(\Delta_L)\|_*, \end{aligned} \quad (\text{E.12})$$

where the middle line arises from the elementary inequalities $|\langle \mathbf{A}, \mathbf{B} \rangle| \leq \|\mathbf{A}\|_{\text{F}} \|\mathbf{B}\|_{\text{F}}$ and $|\langle \mathbf{A}, \mathbf{B} \rangle| \leq \|\mathbf{A}\| \|\mathbf{B}\|_*$, and the last inequality holds since $\|\mathcal{P}_{T^\perp}(\mathbf{R}_1)\| \leq 1/2$ (see Lemma 8).

2. Similarly, one can decompose θ_2 into

$$\begin{aligned} \theta_2 &= -\tau \langle \mathbf{R}_2, \mathcal{P}_\Omega(\Delta_S) \rangle - \tau \langle \mathbf{R}_2, \mathcal{P}_{\Omega^c}(\Delta_S) \rangle - \tau \|\mathcal{P}_{\Omega^c}(\Delta_S)\|_1 \\ &\leq \tau \langle \mathcal{P}_\Omega(\mathbf{R}_2), \mathcal{P}_\Omega(\Delta_S) \rangle - \tau (1 - \|\mathcal{P}_{\Omega^c}(\mathbf{R}_2)\|_\infty) \|\mathcal{P}_{\Omega^c}(\Delta_S)\|_1 \leq 0. \end{aligned} \quad (\text{E.13})$$

Here, the first inequality comes from the facts that $|\langle \mathbf{A}, \mathbf{B} \rangle| \leq \|\mathbf{A}\|_\infty \|\mathbf{B}\|_1$ and $|\langle \mathbf{A}, \mathbf{B} \rangle| \leq \|\mathbf{A}\|_{\text{F}} \|\mathbf{B}\|_{\text{F}}$, and the second one utilizes the facts that $\mathcal{P}_\Omega(\mathbf{R}_2) = \mathbf{0}$ and $\|\mathcal{P}_{\Omega^c}(\mathbf{R}_2)\|_\infty \leq 1$ (cf. Lemma 9).

Combining (E.11), (E.12) and (E.13) yields

$$\begin{aligned} \frac{1}{2} \|\mathcal{P}_{\Omega_{\text{obs}}}(\Delta_L + \Delta_S)\|_{\text{F}}^2 &\leq \lambda \|\mathcal{P}_T(\mathbf{R}_1)\|_{\text{F}} \|\mathcal{P}_T(\Delta_L)\|_{\text{F}} - \frac{\lambda}{2} \|\mathcal{P}_{T^\perp}(\Delta_L)\|_* \\ &\lesssim \frac{\lambda}{n^{19}} \|\Delta_L\|_{\text{F}} \lesssim \frac{\sigma \sqrt{np}}{n^{19}} \sigma n^4 \lesssim \frac{\sigma^2}{n^{14.5}}, \end{aligned} \quad (\text{E.14})$$

where we make use of the upper bound $\|\mathcal{P}_T(\mathbf{R}_1)\|_{\text{F}} \lesssim n^{-19}$ (cf. Lemma 8), the choice $\lambda \asymp \sigma \sqrt{np}$ as well as the crude error bound $\|\Delta_L\|_{\text{F}} \lesssim \sigma n^4$ (cf. (E.8)). Consequently, we have demonstrated that $\mathcal{P}_{\Omega_{\text{obs}}}(\Delta_L + \Delta_S) \approx \mathbf{0}$ in the sense that

$$\|\mathcal{P}_{\Omega_{\text{obs}}}(\Delta_L + \Delta_S)\|_{\text{F}} \lesssim \frac{\sigma}{n^{7.25}} \leq \frac{\sigma}{n^7}. \quad (\text{E.15})$$

E.2.2 Step 2: showing that $\mathcal{P}_{T^\perp}(\Delta_L) \approx \mathbf{0}$ and $\mathcal{P}_{(\Omega^*)^c}(\Delta_S) \approx \mathbf{0}$

We begin by demonstrating that $\mathcal{P}_{T^\perp}(\Delta_L) \approx \mathbf{0}$. From the inequality (E.14), we have

$$\begin{aligned} \frac{1}{2} \|\mathcal{P}_{T^\perp}(\Delta_L)\|_* &\leq \|\mathcal{P}_T(\mathbf{R}_1)\|_{\text{F}} \|\mathcal{P}_T(\Delta_L)\|_{\text{F}} - \frac{1}{2\lambda} \|\mathcal{P}_{\Omega_{\text{obs}}}(\Delta_L + \Delta_S)\|_{\text{F}}^2 \\ &\leq \|\mathcal{P}_T(\mathbf{R}_1)\|_{\text{F}} \|\mathcal{P}_T(\Delta_L)\|_{\text{F}} \lesssim \frac{1}{n^{19}} \|\Delta_L\|_{\text{F}}, \end{aligned}$$

where the last inequality again results from the estimate $\|\mathcal{P}_T(\mathbf{R}_1)\|_F \lesssim n^{-19}$ given in Lemma 8. Invoking the condition $\|\Delta_L\|_F \lesssim \sigma n^4$ (cf. (E.8)) yields

$$\|\mathcal{P}_{T^\perp}(\Delta_L)\|_F \leq \|\mathcal{P}_{T^\perp}(\Delta_L)\|_* \lesssim \frac{1}{n^{19}} \sigma n^4 \lesssim \frac{\sigma}{n^{15}} \leq \frac{\sigma}{n^{14}}, \quad (\text{E.16})$$

which demonstrates that the energy of Δ_L outside T is extremely small.

We now move on to $\mathcal{P}_{(\Omega^*)^c}(\Delta_S)$. This term obeys

$$\mathcal{P}_{(\Omega^*)^c}(\Delta_S) = \mathcal{P}_{\Omega_{\text{obs}} \setminus \Omega^*}(\Delta_S),$$

where the relation holds since Δ_S is supported on Ω_{obs} . To facilitate the analysis of $\Omega_{\text{obs}} \setminus \Omega^*$, we introduce another index subset

$$\Omega_1 := \{(i, j) \in \Omega_{\text{obs}} : |(\Delta_S)_{ij}| \leq \|\mathcal{P}_{\Omega_{\text{obs}}}(\Delta_L + \Delta_S)\|_\infty\}. \quad (\text{E.17})$$

The usefulness of Ω_1 can be seen through the following claim, whose claim is postponed to the end of this section.

Claim 1. Under Condition 1, we have

$$\Omega_{\text{obs}} \setminus \Omega^* \subseteq \Omega_1.$$

An immediate consequence of Claim 1 is that

$$\begin{aligned} \|\mathcal{P}_{(\Omega^*)^c}(\Delta_S)\|_F &= \|\mathcal{P}_{\Omega_{\text{obs}} \setminus \Omega^*}(\Delta_S)\|_F \leq \|\mathcal{P}_{\Omega_1}(\Delta_S)\|_F \\ &\leq n \|\mathcal{P}_{\Omega_1}(\Delta_S)\|_\infty \leq n \|\mathcal{P}_{\Omega_{\text{obs}}}(\Delta_L + \Delta_S)\|_\infty \\ &\leq n \|\mathcal{P}_{\Omega_{\text{obs}}}(\Delta_L + \Delta_S)\|_F \leq \frac{\sigma}{n^6}, \end{aligned} \quad (\text{E.18})$$

which justifies our assertion that the energy of Δ_S outside Ω^* is extremely small. Here, the last inequality arises from (E.15).

E.2.3 Step 3: controlling the size of Δ_S (and hence that of Δ_L)

In view of (E.15) and the triangle inequality, we have

$$\begin{aligned} \frac{\sigma}{n^7} &\geq \|\mathcal{P}_{\Omega_{\text{obs}}}(\Delta_L + \Delta_S)\|_F \\ &\geq \|\mathcal{P}_{\Omega_{\text{obs}}}\mathcal{P}_T(\Delta_L)\|_F - \|\mathcal{P}_{\Omega_{\text{obs}}}\mathcal{P}_{T^\perp}(\Delta_L)\|_F - \|\mathcal{P}_{\Omega^*}(\Delta_S)\|_F - \|\mathcal{P}_{\Omega_{\text{obs}} \setminus \Omega^*}(\Delta_S)\|_F \\ &\geq \|\mathcal{P}_{\Omega_{\text{obs}}}\mathcal{P}_T(\Delta_L)\|_F - \|\mathcal{P}_{\Omega^*}(\Delta_S)\|_F - \|\mathcal{P}_{T^\perp}(\Delta_L)\|_F - \|\mathcal{P}_{(\Omega^*)^c}(\Delta_S)\|_F \\ &\geq \|\mathcal{P}_{\Omega_{\text{obs}}}\mathcal{P}_T(\Delta_L)\|_F - \|\mathcal{P}_{\Omega^*}(\Delta_S)\|_F - \frac{\sigma}{n^6} - \frac{\sigma}{n^{14}}, \end{aligned} \quad (\text{E.19})$$

where the last step follows from (E.16) and (E.18). By Condition 2, we have

$$\|\mathcal{P}_{\Omega_{\text{obs}}}\mathcal{P}_T(\Delta_L)\|_F \geq \sqrt{\frac{c_{\text{inj}}}{\kappa}} p \|\mathcal{P}_T(\Delta_L)\|_F \quad \text{and} \quad \|\mathcal{P}_{\Omega^*}\mathcal{P}_T(\Delta_L)\|_F \leq \frac{1}{2} \sqrt{\frac{c_{\text{inj}}}{\kappa}} p \|\mathcal{P}_T(\Delta_L)\|_F,$$

given that $\mathcal{P}_T(\Delta_L) \in T$. The latter inequality combined with (E.15) and (E.16) further gives

$$\begin{aligned} \|\mathcal{P}_{\Omega^*}(\Delta_S)\|_F &\leq \|\mathcal{P}_{\Omega^*}(\Delta_S + \Delta_L)\|_F + \|\mathcal{P}_{\Omega^*}(\Delta_L)\|_F \\ &\leq \|\mathcal{P}_{\Omega_{\text{obs}}}(\Delta_S + \Delta_L)\|_F + \|\mathcal{P}_{\Omega^*}\mathcal{P}_T(\Delta_L)\|_F + \|\mathcal{P}_{\Omega^*}\mathcal{P}_{T^\perp}(\Delta_L)\|_F \\ &\leq \frac{\sigma}{n^7} + \frac{1}{2} \sqrt{\frac{c_{\text{inj}}}{\kappa}} p \|\mathcal{P}_T \Delta_L\|_F + \frac{\sigma}{n^{14}}. \end{aligned}$$

Substituting the above bounds into (E.19) gives

$$\frac{\sigma}{n^7} \geq \sqrt{\frac{c_{\text{inj}}}{\kappa}} p \|\mathcal{P}_T(\Delta_L)\|_F - \frac{1}{2} \sqrt{\frac{c_{\text{inj}}}{\kappa}} p \|\mathcal{P}_T(\Delta_L)\|_F - \frac{\sigma}{n^7} - \frac{\sigma}{n^6} - \frac{2\sigma}{n^{14}}$$

$$\geq \frac{1}{2} \sqrt{\frac{c_{\text{inj}}}{\kappa}} p \|\mathcal{P}_T(\Delta_L)\|_F - \frac{2\sigma}{n^6},$$

which further yields

$$\|\mathcal{P}_T(\Delta_L)\|_F \lesssim \frac{\sigma}{n^6} \sqrt{\frac{\kappa}{p}} \leq \frac{\sigma}{n^5},$$

provided that $n^2 p \gg \kappa$. This combined with (E.16) allows one to control the size of Δ_L :

$$\|\Delta_L\|_F \leq \|\mathcal{P}_T(\Delta_L)\|_F + \|\mathcal{P}_{T^\perp}(\Delta_L)\|_F \lesssim \frac{\sigma}{n^5}.$$

In view of (E.15) and the fact that Δ_S is supported on Ω_{obs} , we have

$$\begin{aligned} \|\Delta_S\|_F &= \|\mathcal{P}_{\Omega_{\text{obs}}}(\Delta_S)\|_F \leq \|\mathcal{P}_{\Omega_{\text{obs}}}(\Delta_L + \Delta_S)\|_F + \|\mathcal{P}_{\Omega_{\text{obs}}}(\Delta_L)\|_F \\ &\leq \|\mathcal{P}_{\Omega_{\text{obs}}}(\Delta_L + \Delta_S)\|_F + \|\Delta_L\|_F \lesssim \frac{\sigma}{n^5}, \end{aligned}$$

thus concluding the proof.

E.2.4 Proof of Claim 1

We first recall the facts that

$$\mathbf{S} = \mathcal{P}_{\Omega_{\text{obs}}} [\mathcal{S}_\tau(\mathbf{M} - \mathbf{X}\mathbf{Y}^\top)] \quad \text{and} \quad \mathbf{S} + \Delta_S = \mathcal{P}_{\Omega_{\text{obs}}} [\mathcal{S}_\tau(\mathbf{M} - \mathbf{X}\mathbf{Y}^\top - \Delta_L)],$$

where the second identity follows since $(\mathbf{L}_{\text{cvx}}, \mathbf{S}_{\text{cvx}}) = (\mathbf{X}\mathbf{Y}^\top + \Delta_L, \mathbf{S} + \Delta_S)$ is the optimizer of the convex program (1.3). These allow us to write

$$\begin{aligned} \Delta_S &= \mathcal{P}_{\Omega_{\text{obs}}} [\mathcal{S}_\tau(\mathbf{M} - \mathbf{X}\mathbf{Y}^\top - \Delta_L) - \mathcal{S}_\tau(\mathbf{M} - \mathbf{X}\mathbf{Y}^\top)] \\ &= \mathcal{P}_{\Omega_{\text{obs}}} [\mathcal{S}_\tau[\mathbf{M} - \mathbf{X}\mathbf{Y}^\top + \Delta_S - (\Delta_L + \Delta_S)] - \mathcal{S}_\tau(\mathbf{M} - \mathbf{X}\mathbf{Y}^\top)]. \end{aligned} \quad (\text{E.20})$$

This characterization of Δ_S turns out to be crucial when establishing the inclusion $\Omega_{\text{obs}} \setminus \Omega^* \subseteq \Omega_1$. Towards this end, we need to introduce another index subset

$$\Omega_2 := \left\{ (i, j) \in \Omega_{\text{obs}} : \tau - \|\mathcal{P}_{\Omega_{\text{obs}}}(\Delta_L + \Delta_S)\|_\infty \leq |(\mathbf{M} - \mathbf{X}\mathbf{Y}^\top)_{ij}| \leq \tau \right\}.$$

As it turns out, the sets Ω, Ω_1 and Ω_2 obey the following three conditions

$$\Omega_2 \cap \Omega = \emptyset, \quad \Omega_{\text{obs}} \setminus \Omega \subseteq \Omega_1 \cup \Omega_2, \quad \text{and} \quad \Omega \cup \Omega_2 \subseteq \Omega^*,$$

which immediately lead to

$$\Omega_{\text{obs}} \setminus \Omega^* \stackrel{(i)}{\subseteq} \Omega_{\text{obs}} \setminus (\Omega \cup \Omega_2) \stackrel{(ii)}{=} (\Omega_{\text{obs}} \setminus \Omega) \setminus \Omega_2 \stackrel{(iii)}{\subseteq} (\Omega_1 \cup \Omega_2) \setminus \Omega_2 \subseteq \Omega_1.$$

Here, (i) follows since $\Omega \cup \Omega_2 \subseteq \Omega^*$, (ii) holds true since $\Omega_2 \cap \Omega = \emptyset$, and (iii) results from the condition $\Omega_{\text{obs}} \setminus \Omega \subseteq \Omega_1 \cup \Omega_2$. It then boils down to proving each of the above three conditions.

1. The first one $\Omega_2 \cap \Omega = \emptyset$ is straightforward to establish. Note that for any $(i, j) \in \Omega_2$, one must have $|(\mathbf{M} - \mathbf{X}\mathbf{Y}^\top)_{ij}| \leq \tau$ and hence $[\mathcal{S}_\tau(\mathbf{M} - \mathbf{X}\mathbf{Y}^\top)]_{ij} = 0$, which means that $(i, j) \notin \Omega$. This proves the relation $\Omega_2 \cap \Omega = \emptyset$.
2. Moving on to the second one $\Omega_{\text{obs}} \setminus \Omega \subseteq \Omega_1 \cup \Omega_2$, we prove this via contradiction. Suppose that this inclusion is false, i.e. there exists an index $(i, j) \in \Omega_{\text{obs}} \setminus \Omega$ such that

$$|(\Delta_S)_{ij}| > \|\mathcal{P}_{\Omega_{\text{obs}}}(\Delta_L + \Delta_S)\|_\infty \quad \text{and} \quad |(\mathbf{M} - \mathbf{X}\mathbf{Y}^\top)_{ij}| < \tau - \|\mathcal{P}_{\Omega_{\text{obs}}}(\Delta_L + \Delta_S)\|_\infty. \quad (\text{E.21})$$

Here, we have taken into account the fact that

$$|(\mathbf{M} - \mathbf{X}\mathbf{Y}^\top)_{ij}| \leq \tau, \quad \text{for any } (i, j) \in \Omega_{\text{obs}} \setminus \Omega.$$

To reach contradiction, we find it convenient to state the following simple fact.

Fact 2. Suppose that $|a| \leq \tau$ and that $\mathcal{S}_\tau(a+b) \neq 0$. Then

$$|\mathcal{S}_\tau(a+b)| \leq |b| + |a| - \tau.$$

Proof. Given that $\mathcal{S}_\tau(a+b) \neq 0$, one necessarily has $|a+b| > \tau$. Without loss of generality, assume that $a+b > 0$, which gives

$$\mathcal{S}_\tau(a+b) = a+b-\tau > 0.$$

This together with the fact $\tau \geq |a|$ yields $|\mathcal{S}_\tau(a+b)| = a+b-\tau \leq |b| + |a| - \tau$. \square

With this fact in mind, we can deduce that

$$\begin{aligned} |(\Delta_S)_{ij}| &= \left| \{ \mathcal{S}_\tau [M - \mathbf{X}\mathbf{Y}^\top + \Delta_S - (\Delta_L + \Delta_S)] - \mathcal{S}_\tau (M - \mathbf{X}\mathbf{Y}^\top) \}_{ij} \right| \\ &\stackrel{(i)}{=} \left| \{ \mathcal{S}_\tau [M - \mathbf{X}\mathbf{Y}^\top + \Delta_S - (\Delta_L + \Delta_S)] \}_{ij} \right| \\ &\stackrel{(ii)}{\leq} \left| [\Delta_S - (\Delta_L + \Delta_S)]_{ij} \right| + |(M - \mathbf{X}\mathbf{Y}^\top)_{ij}| - \tau \\ &\stackrel{(iii)}{<} |(\Delta_S)_{ij}| + \|\mathcal{P}_{\Omega_{\text{obs}}}(\Delta_L + \Delta_S)\|_\infty - \|\mathcal{P}_{\Omega_{\text{obs}}}(\Delta_L + \Delta_S)\|_\infty \\ &= |(\Delta_S)_{ij}|, \end{aligned} \tag{E.22}$$

where (i) holds true since $[\mathcal{S}_\tau(M - \mathbf{X}\mathbf{Y}^\top)]_{ij} = 0$ for any $(i, j) \in \Omega_{\text{obs}} \setminus \Omega$, (ii) follows from Fact 2 (by taking $a = (M - \mathbf{X}\mathbf{Y}^\top)_{ij}$ and $b = [\Delta_S - (\Delta_L + \Delta_S)]_{ij}$), and (iii) is a consequence of (E.21) as well as the triangle inequality. The inequality (E.22), however, is clearly impossible. This establishes that $\Omega_{\text{obs}} \setminus \Omega \subseteq \Omega_1 \cup \Omega_2$.

3. We are left with the last one $\Omega \cup \Omega_2 \subseteq \Omega^*$, which is equivalent to saying $\Omega \subseteq \Omega^*$ and $\Omega_2 \subseteq \Omega^*$. First, for any $(i, j) \in \Omega$, one has

$$\begin{aligned} |S_{ij}| > 0 &\implies (i, j) \in \Omega_{\text{obs}} \quad \text{and} \quad |(\mathbf{L}^* + \mathbf{S}^* + \mathbf{E} - \mathbf{X}\mathbf{Y}^\top)_{ij}| > \tau \\ &\implies (i, j) \in \Omega_{\text{obs}} \quad \text{and} \quad |S_{ij}^*| > \tau - \|\mathbf{L}^* - \mathbf{X}\mathbf{Y}^\top\|_\infty - \|\mathbf{E}\|_\infty > 0. \end{aligned}$$

Here, the last step comes from the triangle inequality and Condition 1. This reveals that $\Omega \subseteq \Omega^*$. Similarly, for any $(i, j) \in \Omega_2$ we have

$$\begin{aligned} \tau - \|\mathcal{P}_{\Omega_{\text{obs}}}(\Delta_L + \Delta_S)\|_\infty &\leq |(M - \mathbf{X}\mathbf{Y}^\top)_{ij}| \\ \iff \tau - \|\mathcal{P}_{\Omega_{\text{obs}}}(\Delta_L + \Delta_S)\|_\infty &\leq |(\mathbf{S}^* + \mathbf{L}^* - \mathbf{X}\mathbf{Y}^\top + \mathbf{E})_{ij}| \\ \implies |S_{ij}^*| &\geq \tau - \|\mathbf{L}^* - \mathbf{X}\mathbf{Y}^\top\|_\infty - \|\mathbf{E}\|_\infty - \|\mathcal{P}_{\Omega_{\text{obs}}}(\Delta_L + \Delta_S)\|_F \geq \frac{\tau}{2} - \frac{\sigma}{n^7} > 0, \end{aligned}$$

where we have used Condition 1, the bound (E.15), and the fact that $\tau \gg \sigma$. This demonstrates that $\Omega_2 \subseteq \Omega^*$. We have therefore justified that $\Omega \cup \Omega_2 \subseteq \Omega^*$.

F Analysis of the nonconvex procedure (Proof of Theorem 5)

This section is devoted to establishing Theorem 5. For notational convenience, we introduce

$$\mathbf{F}^t := [\mathbf{X}^{t\top}, \mathbf{Y}^{t\top}]^\top \in \mathbb{R}^{2n \times r} \quad \text{and} \quad \mathbf{F}^* := [\mathbf{X}^{*\top}, \mathbf{Y}^{*\top}]^\top \in \mathbb{R}^{2n \times r}. \tag{F.1}$$

These allow us to express succinctly the rotation matrix \mathbf{H}^t defined in (3.10) as

$$\mathbf{H}^t = \arg \min_{\mathbf{R} \in \mathcal{O}^{r \times r}} \|\mathbf{F}^t \mathbf{R} - \mathbf{F}^*\|_F. \tag{F.2}$$

With the definitions of \mathbf{F}^t and \mathbf{H}^t in mind, it suffices to justify that: for all $0 \leq t \leq t_0 = n^{47}$, the following hypotheses

$$\|\mathbf{F}^t \mathbf{H}^t - \mathbf{F}^*\|_{\mathbb{F}} \leq C_{\mathbb{F}} \left(\frac{\sigma}{\sigma_{\min}} \sqrt{\frac{n}{p}} + \frac{\lambda}{p\sigma_{\min}} \right) \|\mathbf{X}^*\|_{\mathbb{F}}, \quad (\text{F.3a})$$

$$\|\mathbf{F}^t \mathbf{H}^t - \mathbf{F}^*\| \leq C_{\text{op}} \left(\frac{\sigma}{\sigma_{\min}} \sqrt{\frac{n}{p}} + \frac{\lambda}{p\sigma_{\min}} \right) \|\mathbf{X}^*\|, \quad (\text{F.3b})$$

$$\|\mathbf{F}^t \mathbf{H}^t - \mathbf{F}^*\|_{2,\infty} \leq C_{\infty} \kappa \left(\frac{\sigma}{\sigma_{\min}} \sqrt{\frac{n \log n}{p}} + \frac{\lambda}{p\sigma_{\min}} \right) \|\mathbf{F}^*\|_{2,\infty}, \quad (\text{F.3c})$$

$$\|\mathbf{X}^{t\top} \mathbf{X}^t - \mathbf{Y}^{t\top} \mathbf{Y}^t\|_{\mathbb{F}} \leq C_{\text{B}} \kappa \eta \left(\frac{\sigma}{\sigma_{\min}} \sqrt{\frac{n}{p}} + \frac{\lambda}{p\sigma_{\min}} \right) \sqrt{r} \sigma_{\max}^2, \quad (\text{F.3d})$$

$$\|\mathbf{S}^t - \mathbf{S}^*\| \leq C_{\text{S}} \sigma \sqrt{np} \quad (\text{F.3e})$$

hold for some universal constants $C_{\mathbb{F}}, C_{\text{op}}, C_{\infty}, C_{\text{B}}, C_{\text{S}} > 0$, and, in addition,

$$F(\mathbf{X}^t, \mathbf{Y}^t; \mathbf{S}^t) \leq F(\mathbf{X}^{t-1}, \mathbf{Y}^{t-1}; \mathbf{S}^{t-1}) - \frac{\eta}{2} \|\nabla f(\mathbf{X}^{t-1}, \mathbf{Y}^{t-1}; \mathbf{S}^{t-1})\|_{\mathbb{F}}^2 \quad (\text{F.4})$$

holds for all $1 \leq t \leq t_0 = n^{47}$.

Clearly, the bounds (3.11a), (3.11b), (3.11c), and (3.11d) in Theorem 5 follow immediately from (F.3a), (F.3b), (F.3c), and (F.3e), respectively. It remains to justify the small gradient bound (3.12) on the basis of (F.3) and (F.4), which is exactly the content of the following lemma.

Lemma 10 (Small gradient). *Set $\lambda = C_{\lambda} \sigma \sqrt{np \log n}$ for some large constant $C_{\lambda} > 0$. Suppose that $n^2 p \gg \kappa^3 \mu r n \log^2 n$ and that the noise satisfies $\frac{\sigma}{\sigma_{\min}} \sqrt{\frac{n}{p}} \ll \frac{1}{\sqrt{\kappa^4 \mu r \log n}}$. Take $\eta \asymp 1/(n \kappa^3 \sigma_{\max})$. If the iterates satisfy (F.3) for all $0 \leq t \leq t_0$ and (F.4) for all $1 \leq t \leq t_0$, then with probability at least $1 - O(n^{-50})$, one has*

$$\min_{0 \leq t < t_0} \|\nabla f(\mathbf{X}^t, \mathbf{Y}^t; \mathbf{S}^t)\|_{\mathbb{F}} \leq \frac{1}{n^{20}} \frac{\lambda}{p} \sqrt{\sigma_{\min}}.$$

Proof. See Appendix F.3. □

The remainder of this section is thus dedicated to showing that (F.3) and (F.4) hold for $\{(\mathbf{F}^t, \mathbf{S}^t)\}_{0 \leq t \leq t_0}$, which we accomplish via mathematical induction. Throughout this section, we let $\mathbf{X}_{l\cdot}$ denote the l th row of a matrix \mathbf{X} .

F.1 Leave-one-out analysis

The above hypotheses (F.3) require, among other things, sharp control of the $\ell_{2,\infty}$ estimation errors, which calls for fine-grained statistical analyses. In order to decouple complicated statistical dependency, we resort to the following leave-one-out analysis framework that has been successfully applied to analyze other nonconvex algorithms [ZB18, MWCC17, CLL19, CCFM19, CCF+19, CLPC19, DC18].

Leave-one-out loss functions. For each $1 \leq l \leq n$, we define the following auxiliary loss functions

$$\begin{aligned} & F^{(l)}(\mathbf{X}, \mathbf{Y}, \mathbf{S}) \\ & := \underbrace{\frac{1}{2p} \|\mathcal{P}_{(\Omega_{\text{obs}})_{-l\cdot}}(\mathbf{X}\mathbf{Y}^{\top} + \mathbf{S} - \mathbf{M})\|_{\mathbb{F}}^2 + \frac{1}{2} \|\mathcal{P}_{l\cdot}(\mathbf{X}\mathbf{Y}^{\top} - \mathbf{L}^*)\|_{\mathbb{F}}^2 + \frac{\lambda}{2p} \|\mathbf{X}\|_{\mathbb{F}}^2 + \frac{\lambda}{2p} \|\mathbf{Y}\|_{\mathbb{F}}^2 + \frac{\tau}{p} \|\mathbf{S}\|_1}_{=: f^{(l)}(\mathbf{X}, \mathbf{Y}; \mathbf{S})}. \end{aligned}$$

Here, $\mathcal{P}_{(\Omega_{\text{obs}})_{-l\cdot}}(\cdot)$ (resp. $\mathcal{P}_{l\cdot}(\cdot)$) denotes orthogonal projection onto the space of matrices supported on the index set $\{(i, j) \in \Omega_{\text{obs}} \mid i \neq l\}$ (resp. $\{(i, j) \mid i = l\}$), namely, for any matrix $\mathbf{B} \in \mathbb{R}^{n \times n}$ one has

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

The above auxiliary loss function is obtained by dropping the randomness coming from the l th row of \mathbf{M} , which, as we shall see shortly, facilitates analysis in establishing the incoherence properties (F.3c). Similarly, we define for each $n + 1 \leq l \leq 2n$ that

$$F^{(l)}(\mathbf{X}, \mathbf{Y}, \mathbf{S}) := \underbrace{\frac{1}{2p} \left\| \mathcal{P}_{(\Omega_{\text{obs}}), -(l-n)}(\mathbf{X}\mathbf{Y}^\top + \mathbf{S} - \mathbf{M}) \right\|_{\text{F}}^2 + \frac{1}{2} \left\| \mathcal{P}_{\cdot, (l-n)}(\mathbf{X}\mathbf{Y}^\top - \mathbf{L}^*) \right\|_{\text{F}}^2 + \frac{\lambda}{2p} \|\mathbf{X}\|_{\text{F}}^2 + \frac{\lambda}{2p} \|\mathbf{Y}\|_{\text{F}}^2 + \tau \|\mathbf{S}\|_1}_{=: f^{(l)}(\mathbf{X}, \mathbf{Y}; \mathbf{S})},$$

where the projection operators $\mathcal{P}_{(\Omega_{\text{obs}}), -(l-n)}(\cdot)$ and $\mathcal{P}_{\cdot, (l-n)}(\cdot)$ are defined such that for any matrix $\mathbf{B} \in \mathbb{R}^{n \times n}$,

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

Again, this auxiliary loss function is produced in a way that is independent from the $(l-n)$ -th column of \mathbf{M} . In the above notation, $f^{(l)}(\mathbf{X}, \mathbf{Y}; \mathbf{S})$ is a function of \mathbf{X} and \mathbf{Y} with \mathbf{S} frozen.

Leave-one-out auxiliary sequences. For each $1 \leq l \leq 2n$, we construct a sequence of leave-one-out iterates $\{\mathbf{F}^{t,(l)}, \mathbf{S}^{t,(l)}\}_{t \geq 0}$ via Algorithm 2.

Algorithm 2 Construction of the l th leave-one-out sequences.

Initialization: $\mathbf{X}^{0,(l)} = \mathbf{X}^*$, $\mathbf{Y}^{0,(l)} = \mathbf{Y}^*$, $\mathbf{S}^{0,(l)} = \mathbf{S}^*$, $\mathbf{F}^{0,(l)} := \begin{bmatrix} \mathbf{X}^{0,(l)} \\ \mathbf{Y}^{0,(l)} \end{bmatrix}$, and the step size $\eta > 0$.

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

$$\mathbf{F}^{t+1,(l)} := \begin{bmatrix} \mathbf{X}^{t+1,(l)} \\ \mathbf{Y}^{t+1,(l)} \end{bmatrix} = \begin{bmatrix} \mathbf{X}^{t,(l)} - \eta \nabla_{\mathbf{X}} f^{(l)}(\mathbf{X}^{t,(l)}, \mathbf{Y}^{t,(l)}; \mathbf{S}^{t,(l)}) \\ \mathbf{Y}^{t,(l)} - \eta \nabla_{\mathbf{Y}} f^{(l)}(\mathbf{X}^{t,(l)}, \mathbf{Y}^{t,(l)}; \mathbf{S}^{t,(l)}) \end{bmatrix}; \quad (\text{F.5a})$$

$$\mathbf{S}^{t+1,(l)} := \begin{cases} \mathcal{S}_\tau [\mathcal{P}_{-l, \cdot}(\mathbf{M} - \mathbf{X}^{t+1,(l)}\mathbf{Y}^{t+1,(l)\top})] + \mathcal{P}_{l, \cdot}(\mathbf{S}^*), & \text{if } 1 \leq l \leq n, \\ \mathcal{S}_\tau [\mathcal{P}_{\cdot, -(l-n)}(\mathbf{M} - \mathbf{X}^{t+1,(l)}\mathbf{Y}^{t+1,(l)\top})] + \mathcal{P}_{\cdot, (l-n)}(\mathbf{S}^*). & \text{if } n+1 \leq l \leq 2n. \end{cases} \quad (\text{F.5b})$$

Properties of leave-one-out sequences. There are several features of the leave-one-out sequences that prove useful for our statistical analysis: (1) for the l th leave-one-out sequence, one can exploit the statistical independence to control the estimation error of $\mathbf{F}^{t,(l)}$ in the l th row; (2) the leave-one-out sequences and the original sequence $(\mathbf{F}^t, \mathbf{S}^t)$ are exceedingly close (since we have only discarded a small amount of information). These properties taken collectively allow us to control the estimation error of \mathbf{F}^t in each row. To formalize these features, we make an additional set of induction hypotheses

$$\|\mathbf{F}^t \mathbf{H}^t - \mathbf{F}^{t,(l)} \mathbf{R}^{t,(l)}\|_{\text{F}} \leq C_1 \left(\frac{\sigma}{\sigma_{\min}} \sqrt{\frac{n \log n}{p}} + \frac{\lambda}{p \sigma_{\min}} \right) \|\mathbf{F}^*\|_{2, \infty}, \quad (\text{F.6a})$$

$$\max_{1 \leq l \leq 2n} \|(\mathbf{F}^{t,(l)} \mathbf{H}^{t,(l)} - \mathbf{F}^*)_{l, \cdot}\|_2 \leq C_2 \kappa \left(\frac{\sigma}{\sigma_{\min}} \sqrt{\frac{n \log n}{p}} + \frac{\lambda}{p \sigma_{\min}} \right) \|\mathbf{F}^*\|_{2, \infty}, \quad (\text{F.6b})$$

$$\max_{1 \leq l \leq n} \|\mathcal{P}_{-l, \cdot}(\mathbf{S}^t - \mathbf{S}^{t,(l)})\|_{\text{F}} \leq C_3 \frac{\sigma}{\sigma_{\min}} \sqrt{n \log n} \|\mathbf{F}^*\| \|\mathbf{F}^*\|_{2, \infty}, \quad (\text{F.6c})$$

$$\max_{n < l \leq 2n} \|\mathcal{P}_{\cdot, -(l-n)}(\mathbf{S}^t - \mathbf{S}^{t,(l)})\|_{\text{F}} \leq C_3 \frac{\sigma}{\sigma_{\min}} \sqrt{n \log n} \|\mathbf{F}^*\| \|\mathbf{F}^*\|_{2, \infty}. \quad (\text{F.6d})$$

Here, the rotation matrices $\mathbf{H}^{t,(l)}$ and $\mathbf{R}^{t,(l)}$ are defined respectively by

$$\mathbf{H}^{t,(l)} := \arg \min_{\mathbf{R} \in \mathcal{O}^{r \times r}} \|\mathbf{F}^{t,(l)} \mathbf{R} - \mathbf{F}^*\|_{\text{F}}, \quad \text{and} \quad \mathbf{R}^{t,(l)} := \arg \min_{\mathbf{R} \in \mathcal{O}^{r \times r}} \|\mathbf{F}^{t,(l)} \mathbf{R} - \mathbf{F}^t\|_{\text{F}}.$$

F.2 Key lemmas for establishing the induction hypotheses

This subsection establishes the induction hypotheses made in Appendix F.1, namely (F.3), (F.4) and (F.6). Before continuing, we find it convenient to introduce another function of \mathbf{X} and \mathbf{Y} (with \mathbf{S} frozen) as follows

$$f_{\text{aug}}(\mathbf{X}, \mathbf{Y}; \mathbf{S}) := \frac{1}{2p} \|\mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{X}\mathbf{Y}^\top + \mathbf{S} - \mathbf{M})\|_{\text{F}}^2 + \frac{\lambda}{2p} \|\mathbf{X}\|_{\text{F}}^2 + \frac{\lambda}{2p} \|\mathbf{Y}\|_{\text{F}}^2 + \frac{1}{8} \|\mathbf{X}^\top \mathbf{X} - \mathbf{Y}^\top \mathbf{Y}\|_{\text{F}}^2. \quad (\text{F.7})$$

The difference between f_{aug} and f lies in the following balancing term

$$f_{\text{diff}}(\mathbf{X}, \mathbf{Y}) := -\frac{1}{8} \|\mathbf{X}^\top \mathbf{X} - \mathbf{Y}^\top \mathbf{Y}\|_{\text{F}}^2,$$

that is, $f = f_{\text{aug}} + f_{\text{diff}}$.

The following four lemmas, which are inherited from [CCF⁺19] with little modification, are concerned with local strong convexity as well as the hypotheses (F.3a), (F.3b), (F.3d), (F.6b) and (F.3c).

Lemma 11 (Restricted strong convexity). *Set $\lambda = C_\lambda \sigma \sqrt{np}$ for some large enough constant $C_\lambda > 0$. Suppose that the sample size obeys $n^2 p \gg \kappa \mu r n \log n$ and that the noise satisfies $\frac{\sigma}{\sigma_{\min}} \sqrt{\frac{n \log n}{p}} \ll 1$. Let the function f_{aug} be defined in (F.7). Then with probability at least $1 - O(n^{-100})$,*

$$\begin{aligned} \text{vec}(\Delta)^\top \nabla^2 f_{\text{aug}}(\mathbf{X}, \mathbf{Y}; \mathbf{S}) \text{vec}(\Delta) &\geq \frac{1}{10} \sigma_{\min} \|\Delta\|_{\text{F}}^2 \\ \max \{ \|\nabla^2 f_{\text{aug}}(\mathbf{X}, \mathbf{Y}; \mathbf{S})\|, \|\nabla^2 f(\mathbf{X}, \mathbf{Y})\| \} &\leq 10 \sigma_{\max} \end{aligned}$$

hold uniformly over all $\mathbf{X}, \mathbf{Y} \in \mathbb{R}^{n \times r}$, $\mathbf{S} \in \mathbb{R}^{n \times n}$ obeying

$$\left\| \begin{bmatrix} \mathbf{X} - \mathbf{X}^* \\ \mathbf{Y} - \mathbf{Y}^* \end{bmatrix} \right\|_{2, \infty} \leq \frac{1}{1000 \kappa \sqrt{n}} \|\mathbf{X}^*\|, \quad \|\mathbf{S} - \mathbf{S}^*\| \leq C_S \sigma \sqrt{np}$$

and all $\Delta = \begin{bmatrix} \Delta_{\mathbf{X}} \\ \Delta_{\mathbf{Y}} \end{bmatrix} \in \mathbb{R}^{2n \times r}$ lying in the set

$$\left\{ \left[\begin{array}{c} \mathbf{X}_1 \\ \mathbf{Y}_1 \end{array} \right] \hat{\mathbf{H}} - \left[\begin{array}{c} \mathbf{X}_2 \\ \mathbf{Y}_2 \end{array} \right] \left\| \left\| \begin{array}{c} \mathbf{X}_2 - \mathbf{X}^* \\ \mathbf{Y}_2 - \mathbf{Y}^* \end{array} \right\| \right\| \leq \frac{1}{500 \kappa} \|\mathbf{X}^*\|, \hat{\mathbf{H}} := \arg \min_{\mathbf{R} \in \mathcal{O}^{r \times r}} \left\| \left\| \begin{array}{c} \mathbf{X}_1 \\ \mathbf{Y}_1 \end{array} \right\| \mathbf{R} - \left[\begin{array}{c} \mathbf{X}_2 \\ \mathbf{Y}_2 \end{array} \right] \right\|_{\text{F}} \right\}.$$

Lemma 12 (Frobenius norm error w.r.t. \mathbf{F}). *Set $\lambda = C_\lambda \sigma \sqrt{np}$ for some large enough constant $C_\lambda > 0$. Suppose that the sample size obeys $n^2 p \gg \kappa \mu r n \log^2 n$ and that the noise satisfies $\frac{\sigma}{\sigma_{\min}} \sqrt{\frac{n}{p}} \ll \frac{1}{\sqrt{\kappa^4 \mu r \log n}}$. If the iterates satisfy (F.3) in the t th iteration, then with probability at least $1 - O(n^{-100})$,*

$$\|\mathbf{F}^{t+1} \mathbf{H}^{t+1} - \mathbf{F}^*\|_{\text{F}} \leq C_{\text{F}} \left(\frac{\sigma}{\sigma_{\min}} \sqrt{\frac{n}{p}} + \frac{\lambda}{p \sigma_{\min}} \right) \|\mathbf{X}^*\|_{\text{F}}$$

holds as long as $0 < \eta \ll 1/(\kappa^{5/2} \sigma_{\max})$.

Lemma 13 (Spectral norm error w.r.t. \mathbf{F}). *Set $\lambda = C_\lambda \sigma \sqrt{np}$ for some large enough constant $C_\lambda > 0$. Suppose that the sample size obeys $n^2 p \gg \kappa^4 \mu^2 r^2 n \log^2 n$ and that the noise satisfies $\frac{\sigma}{\sigma_{\min}} \sqrt{\frac{n}{p}} \ll \frac{1}{\sqrt{\kappa^4 \log n}}$. If the iterates satisfy (F.3) in the t th iteration, then with probability at least $1 - O(n^{-100})$, one has*

$$\|\mathbf{F}^{t+1} \mathbf{H}^{t+1} - \mathbf{F}^*\| \leq C_{\text{op}} \left(\frac{\sigma}{\sigma_{\min}} \sqrt{\frac{n}{p}} + \frac{\lambda}{p \sigma_{\min}} \right) \|\mathbf{X}^*\|$$

holds as long as $0 < \eta \ll 1/(\kappa^3 \sigma_{\max} \sqrt{r})$ and $C_{\text{op}} \gg 1$.

Lemma 14 (Approximate balancedness). *Set $\lambda = C_\lambda \sigma \sqrt{np}$ for some large enough constant $C_\lambda > 0$. Suppose that the sample size obeys $n^2 p \gg \kappa^2 \mu^2 r^2 n \log n$ and that the noise satisfies $\frac{\sigma}{\sigma_{\min}} \sqrt{\frac{n}{p}} \ll \frac{1}{\sqrt{\kappa^2 \log n}}$. If the iterates satisfy (F.3) in the t th iteration, then with probability at least $1 - O(n^{-100})$,*

$$\|\mathbf{X}^{t+1 \top} \mathbf{X}^{t+1} - \mathbf{Y}^{t+1 \top} \mathbf{Y}^{t+1}\|_{\text{F}} \leq C_{\text{B}} \kappa \eta \left(\frac{\sigma}{\sigma_{\min}} \sqrt{\frac{n}{p}} + \frac{\lambda}{p \sigma_{\min}} \right) \sqrt{r} \sigma_{\max}^2$$

$$\max_{1 \leq l \leq 2n} \left\| \mathbf{X}^{t+1,(l)\top} \mathbf{X}^{t+1,(l)} - \mathbf{Y}^{t+1,(l)\top} \mathbf{Y}^{t+1,(l)} \right\|_{\text{F}} \leq C_{\text{B}} \kappa \eta \left(\frac{\sigma}{\sigma_{\min}} \sqrt{\frac{n}{p}} + \frac{\lambda}{p \sigma_{\min}} \right) \sqrt{r} \sigma_{\max}^2$$

hold for some sufficiently large constant $C_{\text{B}} \gg C_{\text{op}}^2$, provided that $0 < \eta < 1/\sigma_{\min}$.

Lemma 15 ($\ell_{2,\infty}$ norm error of leave-one-out sequences). *Set $\lambda = C_{\lambda} \sigma \sqrt{np}$ for some large enough constant $C_{\lambda} > 0$. Suppose that the sample size obeys $n^2 p \gg \kappa^4 \mu^2 r^2 n \log^3 n$ and that the noise satisfies $\frac{\sigma}{\sigma_{\min}} \sqrt{\frac{n}{p}} \ll \frac{1}{\sqrt{\kappa^2 \log n}}$. If the iterates satisfy (F.3) in the t th iteration, then with probability at least $1 - O(n^{-100})$,*

$$\max_{1 \leq l \leq 2n} \left\| (\mathbf{F}^{t+1,(l)} \mathbf{H}^{t+1,(l)} - \mathbf{F}^*)_{l,\cdot} \right\|_2 \leq C_2 \kappa \left(\frac{\sigma}{\sigma_{\min}} \sqrt{\frac{n \log n}{p}} + \frac{\lambda}{p \sigma_{\min}} \right) \|\mathbf{F}^*\|_{2,\infty}$$

holds, provided that $0 < \eta \ll 1/(\kappa^2 \sqrt{r} \sigma_{\max})$, $C_{\text{op}} \gg 1$ and $C_2 \gg C_{\text{op}}$.

Lemma 16 ($\ell_{2,\infty}$ norm error of the true sequence). *Set $\lambda = C_{\lambda} \sigma \sqrt{np}$ for some large enough constant $C_{\lambda} > 0$. Suppose that $n \geq \mu r$ and that the noise satisfies $\frac{\sigma}{\sigma_{\min}} \sqrt{\frac{n}{p}} \ll \frac{1}{\sqrt{\kappa^2 \log n}}$. If the iterates satisfy (F.3) and (F.6) in the t th iteration, then with probability at least $1 - O(n^{-99})$, one has*

$$\|\mathbf{F}^{t+1} \mathbf{H}^{t+1} - \mathbf{F}^*\|_{2,\infty} \leq C_{\infty} \kappa \left(\frac{\sigma}{\sigma_{\min}} \sqrt{\frac{n \log n}{p}} + \frac{\lambda}{p \sigma_{\min}} \right) \|\mathbf{F}^*\|_{2,\infty},$$

provided that $C_{\infty} \geq 5C_1 + C_2$.

Proof of Lemmas 11, 12, 13, 14, 15 and 16. As it turns out, Lemmas 11, 12, 13 and 14 follow immediately from the proofs of [CCF⁺19, Lemmas 17, 10, 11, 15] respectively. More specifically, the proofs can be accomplished by replacing \mathbf{E} in the proofs therein with $\tilde{\mathbf{E}} := \mathbf{E} + \mathbf{S}^* - \mathbf{S}^t$. To see this, we remark that the only property of the perturbation matrix \mathbf{E} utilized in the proofs therein is that $\|\mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{E})\| \lesssim \sigma \sqrt{np}$ with probability at least $1 - O(n^{-10})$; under our hypotheses, the new matrix $\tilde{\mathbf{E}}$ clearly satisfies this property since

$$\|\mathcal{P}_{\Omega_{\text{obs}}}(\tilde{\mathbf{E}})\| \leq \|\mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{E})\| + \|\mathbf{S}^* - \mathbf{S}^t\| \lesssim \sigma \sqrt{np}.$$

Regarding Lemma 15, we note that $\mathbf{S}_{l,\cdot}^{t,(l)} \equiv \mathbf{S}_{l,\cdot}^*$ by construction. Therefore, the update rule regarding the l th row of $\{\mathbf{X}_{l,\cdot}^{t,(l)}\}_{t \geq 0}$ and $\{\mathbf{Y}_{l,\cdot}^{t,(l)}\}_{t \geq 0}$ is exactly the same as that in the leave-one-out sequence introduced in [CCF⁺19]. Thus, Lemma 15 follows immediately from the proof of [CCF⁺19, Lemma 13].

Finally, the proof of Lemma 16 is exactly the same as the proof of [CCF⁺19, Lemma 14]. \square

Next, we justify the hypotheses (F.3e), (F.6a) and (F.6c) in the following three lemmas, which require more careful analysis of the properties about $\{\mathbf{S}^t\}$.

Lemma 17 (Spectral norm error w.r.t. \mathbf{S}). *Set $\tau = C_{\tau} \sigma \sqrt{\log n}$ for some large enough constant $C_{\tau} > 0$. Suppose that the sample size obeys $n^2 p \gg \kappa^4 \mu r n \log n$, the noise satisfies $\frac{\sigma}{\sigma_{\min}} \sqrt{\frac{n}{p}} \ll 1/\sqrt{\kappa^2 \log n}$, the outlier fraction satisfies $\rho_{\text{s}} \leq \rho_{\text{aug}} \ll 1/\sqrt{\kappa^5 \mu r \log^2 n}$ and $n^2 p \rho_{\text{aug}} \gg \mu r n \log^2 n$. If the iterates satisfy (F.3b) and (F.3c) in the $(t+1)$ -th iteration, then with probability at least $1 - O(n^{-100})$,*

$$\|\mathbf{S}^{t+1} - \mathbf{S}^*\| \leq C_{\text{S}} \sigma \sqrt{np}$$

holds for some constant $C_{\text{S}} > 0$ that does not rely on the choice of other constants.

Proof. See Appendix F.4. \square

Lemma 18 (Leave-one-out perturbation w.r.t. \mathbf{F}). *Set $\lambda = C_{\lambda} \sigma \sqrt{np}$ for some large enough constant $C_{\lambda} > 0$. Suppose that the sample size obeys $n^2 p \gg \kappa^4 \mu^2 r^2 n \log^4 n$, the noise satisfies $\frac{\sigma}{\sigma_{\min}} \sqrt{\frac{n}{p}} \ll$*

$1/\sqrt{\kappa^4 \mu r \log n}$, the outlier fraction satisfies $\rho_s \leq \rho_{\text{aug}} \ll 1/(\kappa^3 \mu r \log n)$ and $n^2 p \rho_{\text{aug}} \gg \mu r n \log n$. If the iterates satisfy (F.3) and (F.6) in the t th iteration, then with probability at least $1 - O(n^{-100})$,

$$\max_{1 \leq l \leq 2n} \|\mathbf{F}^{t+1} \mathbf{H}^{t+1} - \mathbf{F}^{t+1, (l)} \mathbf{R}^{t+1, (l)}\|_{\text{F}} \leq C_1 \left(\frac{\sigma}{\sigma_{\min}} \sqrt{\frac{n \log n}{p}} + \frac{\lambda}{p \sigma_{\min}} \right) \|\mathbf{F}^*\|_{2, \infty}$$

holds for some constant $C_1 > 0$, provided that $\eta \ll 1/(n \kappa^2 \sigma_{\max})$ and $C_1 \gg C_3$.

Proof. See Appendix F.5. □

Lemma 19 (Leave-one-out perturbation w.r.t. \mathbf{S}). Set $\tau = C_\tau \sigma \sqrt{\log n}$ for some large enough constant $C_\tau > 0$. Suppose that the sample size satisfies $n^2 p \gg \kappa^4 \mu^2 r^2 n \log n$, the noise obeys $\frac{\sigma}{\sigma_{\min}} \sqrt{\frac{n}{p}} \ll 1/\sqrt{\kappa^2 \log n}$ and the outlier fraction satisfies $\rho_s \leq \rho_{\text{aug}} \ll 1/\kappa$. If the iterates satisfy (F.3b), (F.3c) and (F.6a) in the $(t+1)$ -th iteration, then with probability at least $1 - O(n^{-100})$,

$$\begin{aligned} \max_{1 \leq l \leq n} \|\mathcal{P}_{-l, \cdot}(\mathbf{S}^{t+1} - \mathbf{S}^{t+1, (l)})\|_{\text{F}} &\leq C_3 \frac{\sigma}{\sigma_{\min}} \sqrt{n \log n} \|\mathbf{F}^*\| \|\mathbf{F}^*\|_{2, \infty}, \\ \max_{n < l \leq 2n} \|\mathcal{P}_{\cdot, -(l-n)}(\mathbf{S}^{t+1} - \mathbf{S}^{t+1, (l)})\|_{\text{F}} &\leq C_3 \frac{\sigma}{\sigma_{\min}} \sqrt{n \log n} \|\mathbf{F}^*\| \|\mathbf{F}^*\|_{2, \infty} \end{aligned}$$

hold for some constant C_3 that does not rely on the choice of other constants.

Proof. See Appendix F.6. □

Finally, it remains to justify (F.4), which is a straightforward consequence from standard gradient descent theory and implies the existence of a point with nearly zero gradient.

Lemma 20 (Monotonicity of the function values). Set $\lambda = C_\lambda \sigma \sqrt{n p}$ for some large enough constant $C_\lambda > 0$. Suppose that the noise satisfies $\frac{\sigma}{\sigma_{\min}} \sqrt{\frac{n}{p}} \ll 1$. If the iterates satisfy (F.3) in the t th iteration, then with probability at least $1 - O(n^{-100})$,

$$F(\mathbf{X}^{t+1}, \mathbf{Y}^{t+1}; \mathbf{S}^{t+1}) \leq F(\mathbf{X}^t, \mathbf{Y}^t; \mathbf{S}^t) - \frac{\eta}{2} \|\nabla f(\mathbf{X}^t, \mathbf{Y}^t; \mathbf{S}^t)\|_{\text{F}}^2$$

holds as long as $\eta \ll 1/(\kappa n \sigma_{\max})$.

Proof. See Appendix F.7. □

F.3 Proof of Lemma 10

Summing (F.4) over $t = 1, \dots, t_0$ gives

$$F(\mathbf{X}^{t_0}, \mathbf{Y}^{t_0}, \mathbf{S}^{t_0}) \leq F(\mathbf{X}^0, \mathbf{Y}^0, \mathbf{S}^0) - \frac{\eta}{2} \sum_{t=0}^{t_0-1} \|\nabla f(\mathbf{X}^t, \mathbf{Y}^t; \mathbf{S}^t)\|_{\text{F}}^2,$$

which further implies

$$\begin{aligned} \min_{0 \leq t < t_0} \|\nabla f(\mathbf{X}^t, \mathbf{Y}^t; \mathbf{S}^t)\|_{\text{F}}^2 &\leq \frac{1}{t_0} \sum_{t=0}^{t_0-1} \|\nabla f(\mathbf{X}^t, \mathbf{Y}^t; \mathbf{S}^t)\|_{\text{F}}^2 \\ &\leq \frac{2}{\eta t_0} [F(\mathbf{X}^*, \mathbf{Y}^*, \mathbf{S}^*) - F(\mathbf{X}^{t_0}, \mathbf{Y}^{t_0}, \mathbf{S}^{t_0})]. \end{aligned} \quad (\text{F.8})$$

Here, the last inequality results from our choice $(\mathbf{X}^0, \mathbf{Y}^0, \mathbf{S}^0) = (\mathbf{X}^*, \mathbf{Y}^*, \mathbf{S}^*)$. Therefore, it suffices to control $F(\mathbf{X}^*, \mathbf{Y}^*, \mathbf{S}^*) - F(\mathbf{X}^{t_0}, \mathbf{Y}^{t_0}, \mathbf{S}^{t_0})$.

We first decompose the difference into

$$F(\mathbf{X}^*, \mathbf{Y}^*, \mathbf{S}^*) - F(\mathbf{X}^{t_0}, \mathbf{Y}^{t_0}, \mathbf{S}^{t_0}) = f(\mathbf{X}^*, \mathbf{Y}^*; \mathbf{S}^*) + \frac{\tau}{p} \|\mathbf{S}^*\|_1 - f(\mathbf{X}^{t_0}, \mathbf{Y}^{t_0}; \mathbf{S}^{t_0}) - \frac{\tau}{p} \|\mathbf{S}^{t_0}\|_1$$

$$= \underbrace{f(\mathbf{X}^*, \mathbf{Y}^*; \mathbf{S}^*) - f(\mathbf{X}^{t_0}, \mathbf{Y}^{t_0}; \mathbf{S}^*)}_{=:\Delta_1} + \underbrace{f(\mathbf{X}^{t_0}, \mathbf{Y}^{t_0}; \mathbf{S}^*) - f(\mathbf{X}^{t_0}, \mathbf{Y}^{t_0}; \mathbf{S}^{t_0})}_{=:\Delta_2} + \underbrace{\frac{\tau}{p} \|\mathbf{S}^*\|_1 - \frac{\tau}{p} \|\mathbf{S}^{t_0}\|_1}_{=:\Delta_3}.$$

In what follows, we shall bound Δ_1, Δ_2 and Δ_3 separately.

1. In view of the proof of [CCF⁺19, Lemma 9], we have

$$|\Delta_1| \lesssim r\kappa^2 \left(\frac{\lambda}{p}\right)^2,$$

provided that $\frac{\sigma}{\sigma_{\min}} \sqrt{\frac{n}{p}} \ll \frac{1}{\sqrt{\kappa^4 \mu r \log n}}$.

2. When it comes to Δ_2 , we deduce that

$$\begin{aligned} |\Delta_2| &= \left| \frac{1}{2} \|\mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{X}^{t_0} \mathbf{Y}^{t_0 \top} + \mathbf{S}^* - \mathbf{M})\|_{\text{F}}^2 + \frac{\lambda}{2p} \|\mathbf{F}^{t_0}\|_{\text{F}}^2 - \frac{1}{2} \|\mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{X}^{t_0} \mathbf{Y}^{t_0 \top} + \mathbf{S}^{t_0} - \mathbf{M})\|_{\text{F}}^2 - \frac{\lambda}{2p} \|\mathbf{F}^{t_0}\|_{\text{F}}^2 \right| \\ &= \left| \langle \mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{X}^{t_0} \mathbf{Y}^{t_0 \top} + \mathbf{S}^{t_0} - \mathbf{M}), \mathbf{S}^* - \mathbf{S}^{t_0} \rangle + \frac{1}{2} \|\mathbf{S}^* - \mathbf{S}^{t_0}\|_{\text{F}}^2 \right| \\ &\leq \|\mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{X}^{t_0} \mathbf{Y}^{t_0 \top} + \mathbf{S}^{t_0} - \mathbf{M})\|_{\text{F}} \|\mathbf{S}^* - \mathbf{S}^{t_0}\|_{\text{F}} + \frac{1}{2} \|\mathbf{S}^* - \mathbf{S}^{t_0}\|_{\text{F}}^2, \end{aligned} \quad (\text{F.9})$$

where the last step arises from the elementary inequality $\langle \mathbf{A}, \mathbf{B} \rangle \leq \|\mathbf{A}\|_{\text{F}} \|\mathbf{B}\|_{\text{F}}$ and the triangle inequality. It is straightforward to derive from (F.3e) that

$$\|\mathbf{S}^* - \mathbf{S}^{t_0}\|_{\text{F}} \leq \sqrt{n} \|\mathbf{S}^* - \mathbf{S}^{t_0}\| \lesssim \sigma n \sqrt{p}.$$

Moving on to the first term in (F.9), one has by the triangle inequality

$$\begin{aligned} \|\mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{X}^{t_0} \mathbf{Y}^{t_0 \top} + \mathbf{S}^{t_0} - \mathbf{M})\|_{\text{F}} &\leq \|\mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{X}^{t_0} \mathbf{Y}^{t_0 \top} - \mathbf{X}^* \mathbf{Y}^{* \top})\|_{\text{F}} + \|\mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{S}^{t_0} - \mathbf{S}^*)\|_{\text{F}} + \|\mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{E})\|_{\text{F}} \\ &\lesssim \|\mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{X}^{t_0} \mathbf{Y}^{t_0 \top} - \mathbf{X}^* \mathbf{Y}^{* \top})\|_{\text{F}} + \sigma n \sqrt{p}, \end{aligned}$$

where the last bound follows from Lemma 1 and the bound above $\|\mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{S}^{t_0} - \mathbf{S}^*)\|_{\text{F}} \leq \|\mathbf{S}^{t_0} - \mathbf{S}^*\|_{\text{F}} \lesssim \sigma n \sqrt{p}$. We can further decompose $\|\mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{X}^{t_0} \mathbf{Y}^{t_0 \top} - \mathbf{X}^* \mathbf{Y}^{* \top})\|_{\text{F}}$ into

$$\begin{aligned} \|\mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{X}^{t_0} \mathbf{Y}^{t_0 \top} - \mathbf{X}^* \mathbf{Y}^{* \top})\|_{\text{F}} &\leq \|\mathcal{P}_{\Omega_{\text{obs}}}[(\mathbf{X}^{t_0} \mathbf{H}^{t_0} - \mathbf{X}^*) \mathbf{Y}^{* \top}]\|_{\text{F}} + \left\| \mathcal{P}_{\Omega_{\text{obs}}} \left[\mathbf{X}^{t_0} \mathbf{H}^{t_0} (\mathbf{Y}^{t_0} \mathbf{H}^{t_0} - \mathbf{Y}^*)^{\top} \right] \right\|_{\text{F}} \\ &\stackrel{(i)}{\lesssim} \sqrt{\kappa p} \|(\mathbf{X}^{t_0} \mathbf{H}^{t_0} - \mathbf{X}^*) \mathbf{Y}^{* \top}\|_{\text{F}} + \sqrt{\kappa p} \left\| \mathbf{X}^{t_0} \mathbf{H}^{t_0} (\mathbf{Y}^{t_0} \mathbf{H}^{t_0} - \mathbf{Y}^*)^{\top} \right\|_{\text{F}} \\ &\lesssim \sqrt{\kappa p} \|\mathbf{X}^{t_0} \mathbf{H}^{t_0} - \mathbf{X}^*\|_{\text{F}} \|\mathbf{Y}^*\| + \sqrt{\kappa p} \|\mathbf{X}^{t_0} \mathbf{H}^{t_0}\| \|\mathbf{Y}^{t_0} \mathbf{H}^{t_0} - \mathbf{Y}^*\|_{\text{F}} \\ &\stackrel{(ii)}{\lesssim} \sqrt{\kappa p} \left(\frac{\sigma}{\sigma_{\min}} \sqrt{\frac{n}{p}} + \frac{\lambda}{p \sigma_{\min}} \right) \|\mathbf{X}^*\|_{\text{F}} \|\mathbf{X}^*\| \lesssim \kappa^{3/2} \frac{\lambda}{\sqrt{p}} \sqrt{r}. \end{aligned} \quad (\text{F.10})$$

Here, the relation (i) utilizes Lemma 4, and the facts that $(\mathbf{X}^{t_0} \mathbf{H}^{t_0} - \mathbf{X}^*) \mathbf{Y}^{* \top} \in T^*$ and that $\mathbf{X}^{t_0} \mathbf{H}^{t_0} (\mathbf{Y}^{t_0} \mathbf{H}^{t_0} - \mathbf{Y}^*)^{\top} \in T^{t_0}$, where T^{t_0} denotes the tangent space at $\mathbf{X}^{t_0} \mathbf{Y}^{t_0 \top}$. In addition, the last line (ii) holds because of the hypothesis (F.3a) and the simple fact $\|\mathbf{X}^{t_0} \mathbf{H}^{t_0}\| \leq 2\|\mathbf{X}^*\|$, which is an immediate consequence of the hypothesis (F.3b) provided that $\frac{\sigma}{\sigma_{\min}} \sqrt{\frac{n}{p}} \ll 1$. Collecting the bounds together, we arrive at

$$|\Delta_2| \lesssim \left(\kappa^{3/2} \frac{\lambda}{\sqrt{p}} \sqrt{r} + \sigma n \sqrt{p} \right) \cdot \sigma n \sqrt{p} + \sigma^2 n^2 p + \frac{\lambda}{p} \cdot \kappa \frac{\lambda}{p} r \lesssim \sigma^2 n^2 p,$$

with the proviso that $np \gg \kappa^3 r$.

3. In the end, we have the following upper bound on Δ_3 :

$$|\Delta_3| \leq \frac{\tau}{p} \|\mathbf{S}^{t_0} - \mathbf{S}^*\|_1 \leq \frac{\tau}{p} n \|\mathbf{S}^{t_0} - \mathbf{S}^*\|_{\text{F}} \lesssim \frac{1}{p} \frac{\lambda}{\sqrt{np/\log n}} \sigma n^2 \sqrt{p} \asymp \frac{\lambda}{p} \sigma n^{3/2} \sqrt{\log n},$$

where we have made use of the elementary fact that $\|\mathbf{A}\|_1 \leq n \|\mathbf{A}\|_{\text{F}}$ for all $\mathbf{A} \in \mathbb{R}^{n \times n}$.

Putting the above bounds together, one can reach

$$|F(\mathbf{X}^*, \mathbf{Y}^*, \mathbf{S}^*) - F(\mathbf{X}^{t_0}, \mathbf{Y}^{t_0}, \mathbf{S}^{t_0})| \lesssim r\kappa^2 \left(\frac{\lambda}{p}\right)^2 + \sigma^2 n^2 p + \frac{\lambda}{p} \sigma n^{3/2} \sqrt{\log n} \lesssim n \left(\frac{\lambda}{p}\right)^2 \sqrt{\log n}$$

as long as $n \gg \kappa^2 r$ and $\lambda \asymp \sigma \sqrt{np}$. Substitution into (F.8) allows us to conclude that

$$\min_{0 \leq t \leq t_0} \|\nabla f(\mathbf{X}^t, \mathbf{Y}^t; \mathbf{S}^t)\|_{\mathbb{F}} \lesssim \sqrt{\frac{1}{\eta t_0} n \left(\frac{\lambda}{p}\right)^2 \sqrt{\log n}} \leq \frac{1}{n^{20}} \frac{\lambda}{p} \sqrt{\sigma_{\min}},$$

provided that $\eta \asymp 1/(n\kappa^3 \sigma_{\max})$, $t_0 \geq n^{47}$ and $n \geq \kappa$.

F.4 Proof of Lemma 17

In view of the definitions $\Omega^* = \{(i, j) : S_{ij}^* \neq 0\} \subseteq \Omega_{\text{aug}} \subseteq \Omega_{\text{obs}}$ and $\mathbf{S}^{t+1} = \mathcal{S}_\tau[\mathcal{P}_{\Omega_{\text{obs}}}(\mathbf{L}^* + \mathbf{S}^* + \mathbf{E} - \mathbf{X}^{t+1} \mathbf{Y}^{t+1\top})]$, we have the decomposition

$$\begin{aligned} \mathbf{S}^{t+1} - \mathbf{S}^* &= \mathcal{P}_{\Omega_{\text{aug}}}(\mathbf{S}^{t+1}) - \mathcal{P}_{\Omega_{\text{aug}}}(\mathbf{S}^*) + \mathcal{P}_{\Omega_{\text{c}}}(\mathbf{S}^{t+1}) \\ &= \underbrace{\mathcal{S}_\tau[\mathcal{P}_{\Omega_{\text{aug}}}(\mathbf{X}^* \mathbf{Y}^{*\top} + \mathbf{S}^* + \mathbf{E} - \mathbf{X}^{t+1} \mathbf{Y}^{t+1\top})]}_{=: \mathbf{A}^{t+1}} - \mathcal{P}_{\Omega_{\text{aug}}}(\mathbf{S}^*) + \underbrace{\mathcal{S}_\tau[\mathcal{P}_{\Omega_{\text{obs}} \setminus \Omega_{\text{aug}}}(\mathbf{X}^* \mathbf{Y}^{*\top} + \mathbf{E} - \mathbf{X}^{t+1} \mathbf{Y}^{t+1\top})]}_{=: \mathbf{B}^{t+1}}. \end{aligned} \quad (\text{F.11})$$

We shall control $\|\mathbf{A}^{t+1}\|$ and $\|\mathbf{B}^{t+1}\|$ separately.

1. We begin by controlling the size of \mathbf{A}^{t+1} , which can be further decomposed into

$$\begin{aligned} \mathbf{A}^{t+1} &= \mathcal{P}_{\Omega_{\text{aug}}}(\mathbf{E}) + \underbrace{\mathcal{S}_\tau[\mathcal{P}_{\Omega_{\text{aug}}}(\mathbf{S}^* + \mathbf{E})]}_{=: \mathbf{A}_1} - \mathcal{P}_{\Omega_{\text{aug}}}(\mathbf{S}^* + \mathbf{E}) \\ &\quad + \underbrace{\mathcal{S}_\tau[\mathcal{P}_{\Omega_{\text{aug}}}(\mathbf{X}^* \mathbf{Y}^{*\top} - \mathbf{X}^{t+1} \mathbf{Y}^{t+1\top} + \mathbf{S}^* + \mathbf{E})]}_{=: \mathbf{A}_2^{t+1}} - \mathcal{S}_\tau[\mathcal{P}_{\Omega_{\text{aug}}}(\mathbf{S}^* + \mathbf{E})]. \end{aligned} \quad (\text{F.12})$$

First of all, we know that $\|\mathcal{P}_{\Omega_{\text{aug}}}(\mathbf{E})\| \lesssim \sigma \sqrt{np\rho_{\text{aug}}} \leq \sigma \sqrt{np}$, as long as $n^2 p \rho_{\text{aug}} \gg n \log^2 n$. This arises from standard concentration results for the spectral norm of sub-Gaussian random matrices (cf. Lemma 1). Regarding \mathbf{A}_1 , we know from the definition of $\mathcal{S}_\tau(\cdot)$ that $\|\mathbf{A}_1\|_\infty \leq \tau$. More precisely, we have

$$(\mathbf{A}_1)_{ij} = \begin{cases} -\tau & \text{if } S_{ij}^* + E_{ij} \geq \tau, \\ -S_{ij}^* - E_{ij} & \text{if } -\tau < S_{ij}^* + E_{ij} < \tau, \\ \tau & \text{if } S_{ij}^* + E_{ij} \leq -\tau. \end{cases}$$

Recall from Assumption 4 that \mathbf{S}^* has random signs on its support $\Omega^* \subseteq \Omega_{\text{aug}}$ and E_{ij} is symmetric around zero. It then follows from standard concentration results for the spectral norm of matrices with i.i.d. entries that

$$\|\mathbf{A}_1\| \lesssim \tau \sqrt{np\rho_{\text{aug}}} = C_\tau \sigma \sqrt{np\rho_{\text{aug}}} \log n,$$

provided that $n^2 p \rho_{\text{aug}} \gg n \log^2 n$. Moving on to \mathbf{A}_2^{t+1} , since it is supported on Ω_{aug} , we can further decompose $\|\mathbf{A}_2^{t+1}\|$ into

$$\|\mathbf{A}_2^{t+1}\| = \|\mathcal{P}_{\Omega_{\text{aug}}}(\mathbf{A}_2^{t+1})\| \leq p\rho_{\text{aug}} \|\mathbf{A}_2^{t+1}\| + \|\mathcal{P}_{\Omega_{\text{aug}}}(\mathbf{A}_2^{t+1}) - p\rho_{\text{aug}} \mathbf{A}_2^{t+1}\|. \quad (\text{F.13})$$

Invoking Lemma 5 with $\mathbf{A} = \mathbf{A}_2^{t+1}$, $\mathbf{B} = \mathbf{I}_n$ and $\rho_0 = p\rho_{\text{aug}}$, we have

$$\|\mathcal{P}_{\Omega_{\text{aug}}}(\mathbf{A}_2^{t+1}) - p\rho_{\text{aug}} \mathbf{A}_2^{t+1}\| \leq C \sqrt{np\rho_{\text{aug}}} \|\mathbf{A}_2^{t+1}\|_{2,\infty}, \quad (\text{F.14})$$

with the proviso that $n^2 p \rho_{\text{aug}} \gg n \log n$. Combine the above bounds to reach

$$\|\mathbf{A}_2^{t+1}\| \leq p\rho_{\text{aug}} \|\mathbf{A}_2^{t+1}\| + C \sqrt{np\rho_{\text{aug}}} \|\mathbf{A}_2^{t+1}\|_{2,\infty} \leq \frac{1}{2} \|\mathbf{A}_2^{t+1}\| + C \sqrt{np\rho_{\text{aug}}} \|\mathbf{A}_2^{t+1}\|_{2,\infty},$$

as soon as $\rho_s \leq \rho_{\text{aug}} \leq 1/2$. We are then in need of an upper bound on $\|\mathbf{A}_2^{t+1}\|_{2,\infty}$, which is supplied in the following fact.

Fact 3. Suppose that $n^2 p \rho_{\text{aug}} \gg \mu r n \log n$ and $\frac{\sigma}{\sigma_{\min}} \sqrt{\frac{n \log n}{p}} \ll 1/\kappa$. Then with probability exceeding $1 - O(n^{-100})$, one has

$$\|\mathbf{A}_2^{t+1}\|_{2,\infty} \leq \sqrt{40\kappa p \rho_{\text{aug}}} (C_\infty \kappa + 2C_{\text{op}}) \left(\frac{\sigma}{\sigma_{\min}} \sqrt{\frac{n \log n}{p}} + \frac{\lambda}{p \sigma_{\min}} \right) \|\mathbf{F}^*\|_{2,\infty} \|\mathbf{X}^*\|.$$

With the help of Fact 3, we can continue the upper bound as follows

$$\|\mathbf{A}_2^{t+1}\| \leq 2C \sqrt{n p \rho_{\text{aug}}} \|\mathbf{A}_2^{t+1}\|_{2,\infty} \lesssim (C_\infty + C_{\text{op}}) \sqrt{n p} \kappa^{3/2} \rho_{\text{aug}} \frac{\sigma}{\sigma_{\min}} \sqrt{n \log n} \|\mathbf{F}^*\|_{2,\infty} \|\mathbf{X}^*\|.$$

All in all, we obtain the following bound on \mathbf{A}^{t+1} :

$$\begin{aligned} \|\mathbf{A}^{t+1}\| &\leq \|\mathcal{P}_{\Omega_{\text{aug}}}(\mathbf{E})\| + \|\mathbf{A}_1\| + \|\mathbf{A}_2^{t+1}\| \\ &\lesssim \sigma \sqrt{np} + C_\tau \sigma \sqrt{n p \rho_{\text{aug}} \log n} + (C_\infty + C_{\text{op}}) \sqrt{n p} \kappa^{3/2} \rho_{\text{aug}} \frac{\sigma}{\sigma_{\min}} \sqrt{n \log n} \|\mathbf{F}^*\|_{2,\infty} \|\mathbf{X}^*\| \\ &\leq C_S \sigma \sqrt{np}, \end{aligned}$$

with the proviso that $\rho_s \leq \rho_{\text{aug}} \ll 1/\sqrt{\kappa^5 \mu r \log^2 n}$. Here, the last line uses the incoherence assumption $\|\mathbf{F}^*\|_{2,\infty} \leq \sqrt{\mu r/n} \|\mathbf{X}^*\|$ (cf. (B.1)).

2. When it comes to \mathbf{B}^{t+1} , we first note that

$$\begin{aligned} \|\mathbf{X}^* \mathbf{Y}^{*\top} - \mathbf{X}^{t+1} \mathbf{Y}^{t+1\top}\|_\infty &= \left\| (\mathbf{X}^* - \mathbf{X}^{t+1} \mathbf{H}^{t+1}) \mathbf{Y}^{*\top} + \mathbf{X}^{t+1} \mathbf{H}^{t+1} (\mathbf{Y}^* - \mathbf{Y}^{t+1} \mathbf{H}^{t+1})^\top \right\|_\infty \\ &\leq \|\mathbf{X}^* - \mathbf{X}^{t+1} \mathbf{H}^{t+1}\|_{2,\infty} \|\mathbf{Y}^*\|_{2,\infty} + \|\mathbf{X}^{t+1} \mathbf{H}^{t+1}\|_{2,\infty} \|\mathbf{Y}^* - \mathbf{Y}^{t+1} \mathbf{H}^{t+1}\|_{2,\infty} \\ &\leq 3C_\infty \kappa \left(\frac{\sigma}{\sigma_{\min}} \sqrt{\frac{n \log n}{p}} + \frac{\lambda}{p \sigma_{\min}} \right) \|\mathbf{F}^*\|_{2,\infty}^2 \leq 3C_\infty \kappa \left(\frac{\sigma}{\sigma_{\min}} \sqrt{\frac{n \log n}{p}} + \frac{\lambda}{p \sigma_{\min}} \right) \frac{\mu r}{n} \sigma_{\max}. \quad (\text{F.15}) \end{aligned}$$

Here, we have plugged in (F.3c) for the $(t+1)$ -th iteration and its immediate consequence $\|\mathbf{X}^{t+1} \mathbf{H}^{t+1}\|_{2,\infty} \leq \|\mathbf{F}^{t+1}\|_{2,\infty} \leq 2\|\mathbf{F}^*\|_{2,\infty}$, as long as $\frac{\sigma}{\sigma_{\min}} \sqrt{\frac{n \log n}{p}} \ll 1/\kappa$. As a result, for all (i, j) we have

$$\begin{aligned} \left| (\mathbf{M} - \mathbf{X}^{t+1} \mathbf{Y}^{t+1\top})_{ij} \right| &= \left| (\mathbf{X}^* \mathbf{Y}^{*\top} + \mathbf{E} - \mathbf{X}^{t+1} \mathbf{Y}^{t+1\top})_{ij} \right| \leq |E_{ij}| + \|\mathbf{X}^* \mathbf{Y}^{*\top} - \mathbf{X}^{t+1} \mathbf{Y}^{t+1\top}\|_\infty \\ &\stackrel{(i)}{\leq} |E_{ij}| + 3C_\infty \kappa \left(\frac{\sigma}{\sigma_{\min}} \sqrt{\frac{n \log n}{p}} + \frac{\lambda}{p \sigma_{\min}} \right) \frac{\mu r}{n} \sigma_{\max} \\ &\stackrel{(ii)}{\leq} C_\lambda \sigma \sqrt{\log n} = \tau. \end{aligned}$$

Here, the inequality (i) comes from (F.15), and the last line (ii) relies on the property of sub-Gaussian random variables (namely, $|E_{ij}| \leq \tau/2$ with probability exceeding $1 - O(n^{-102})$) and the sample size condition $n^2 p \gg \kappa^4 \mu^2 r^2 n \log n$. An immediate consequence is that with probability at least $1 - O(n^{-100})$,

$$\mathbf{B}^{t+1} = \mathcal{S}_\tau [\mathcal{P}_{\Omega_{\text{obs}} \setminus \Omega_{\text{aug}}}(\mathbf{X}^* \mathbf{Y}^{*\top} + \mathbf{E} - \mathbf{X}^{t+1} \mathbf{Y}^{t+1\top})] = \mathbf{0}. \quad (\text{F.16})$$

Substituting the above two bounds into (F.11), we conclude that $\|\mathbf{S}^{t+1} - \mathbf{S}^*\| \leq C_S \sigma \sqrt{np}$ as claimed.

Proof of Fact 3. In view of the definition of \mathbf{A}^{t+1} in (F.12), we have

$$\|\mathbf{A}_2^{t+1}\|_{2,\infty} \leq \|\mathcal{P}_{\Omega_{\text{aug}}}(\mathbf{X}^{t+1} \mathbf{Y}^{t+1\top} - \mathbf{X}^* \mathbf{Y}^{*\top})\|_{2,\infty},$$

where we use the non-expansiveness of the proximal operator $\mathcal{S}_\tau(\cdot)$. Apply a similar argument as in bounding (F.10) to obtain

$$\|\mathcal{P}_{\Omega_{\text{aug}}}(\mathbf{X}^{t+1} \mathbf{Y}^{t+1\top} - \mathbf{X}^* \mathbf{Y}^{*\top})\|_{2,\infty}$$

$$\begin{aligned}
&\leq \|\mathcal{P}_{\Omega_{\text{aug}}}[(\mathbf{X}^{t+1}\mathbf{H}^{t+1} - \mathbf{X}^*)\mathbf{Y}^{\star\top}]\|_{2,\infty} + \|\mathcal{P}_{\Omega_{\text{aug}}}[\mathbf{X}^{t+1}\mathbf{H}^{t+1}(\mathbf{Y}^{t+1}\mathbf{H}^{t+1} - \mathbf{Y}^*)^\top]\|_{2,\infty} \\
&\leq \sqrt{40\kappa p\rho_{\text{aug}}}\|(\mathbf{X}^{t+1}\mathbf{H}^{t+1} - \mathbf{X}^*)\mathbf{Y}^{\star\top}\|_{2,\infty} + \sqrt{40\kappa p\rho_{\text{aug}}}\|\mathbf{X}^{t+1}\mathbf{H}^{t+1}(\mathbf{Y}^{t+1}\mathbf{H}^{t+1} - \mathbf{Y}^*)^\top\|_{2,\infty} \\
&\leq \sqrt{40\kappa p\rho_{\text{aug}}}(C_\infty\kappa + 2C_{\text{op}})\left(\frac{\sigma}{\sigma_{\min}}\sqrt{\frac{n\log n}{p}} + \frac{\lambda}{p\sigma_{\min}}\right)\|\mathbf{F}^*\|_{2,\infty}\|\mathbf{X}^*\|,
\end{aligned}$$

as long as $n^2 p \rho_{\text{aug}} \gg \mu r n \log n$. Here, the last line uses the induction hypotheses (F.3b) and (F.3c) for the $(t+1)$ -th iteration and their immediate consequence $\|\mathbf{X}^{t+1}\mathbf{H}^{t+1}\|_{2,\infty} \leq 2\|\mathbf{F}^*\|_{2,\infty}$, as long as $\frac{\sigma}{\sigma_{\min}}\sqrt{\frac{n\log n}{p}} \ll 1/\kappa$. Taking the preceding two bounds together concludes the proof. \square

F.5 Proof of Lemma 18

Without loss of generality, we only consider the case when $1 \leq l \leq n$. The case with $n+1 \leq l \leq 2n$ can be derived similarly with very minor modification, and hence we omit it for the sake of brevity.

To begin with, since $(\mathbf{H}^{t+1}, \mathbf{R}^{t+1,(l)})$ is the choice of the rotation matrix that best aligns \mathbf{F}^{t+1} and $\mathbf{F}^{t+1,(l)}$, we have

$$\|\mathbf{F}^{t+1}\mathbf{H}^{t+1} - \mathbf{F}^{t+1,(l)}\mathbf{R}^{t+1,(l)}\|_{\text{F}} \leq \|\mathbf{F}^{t+1}\mathbf{H}^t - \mathbf{F}^{t+1,(l)}\mathbf{R}^{t,(l)}\|_{\text{F}}.$$

In view of the gradient update rule, one has

$$\begin{aligned}
&\mathbf{F}^{t+1}\mathbf{H}^t - \mathbf{F}^{t+1,(l)}\mathbf{R}^{t,(l)} \\
&= [\mathbf{F}^t - \eta\nabla f(\mathbf{F}^t; \mathbf{S}^t)]\mathbf{H}^t - [\mathbf{F}^{t,(l)} - \eta\nabla f^{(l)}(\mathbf{F}^{t,(l)}; \mathbf{S}^{t,(l)})]\mathbf{R}^{t,(l)} \\
&= \mathbf{F}^t\mathbf{H}^t - \eta\nabla f(\mathbf{F}^t\mathbf{H}^t; \mathbf{S}^t) - [\mathbf{F}^{t,(l)}\mathbf{R}^{t,(l)} - \eta\nabla f^{(l)}(\mathbf{F}^{t,(l)}\mathbf{R}^{t,(l)}; \mathbf{S}^{t,(l)})] \\
&= \underbrace{\mathbf{F}^t\mathbf{H}^t - \mathbf{F}^{t,(l)}\mathbf{R}^{t,(l)} - \eta[\nabla f_{\text{aug}}(\mathbf{F}^t\mathbf{H}^t; \mathbf{S}^t) - \nabla f_{\text{aug}}(\mathbf{F}^{t,(l)}\mathbf{R}^{t,(l)}; \mathbf{S}^t)]}_{=: \mathbf{C}_1} - \underbrace{\eta[\nabla f_{\text{diff}}(\mathbf{F}^t\mathbf{H}^t) - \nabla f_{\text{diff}}(\mathbf{F}^{t,(l)}\mathbf{R}^{t,(l)})]}_{=: \mathbf{C}_2} \\
&\quad + \underbrace{\eta[\nabla f^{(l)}(\mathbf{F}^{t,(l)}\mathbf{R}^{t,(l)}; \mathbf{S}^{t,(l)}) - \nabla f(\mathbf{F}^{t,(l)}\mathbf{R}^{t,(l)}; \mathbf{S}^{t,(l)})]}_{=: \mathbf{C}_3} + \underbrace{\eta[\nabla f(\mathbf{F}^{t,(l)}\mathbf{R}^{t,(l)}; \mathbf{S}^{t,(l)}) - \nabla f(\mathbf{F}^{t,(l)}\mathbf{R}^{t,(l)}; \mathbf{S}^t)]}_{=: \mathbf{C}_4}.
\end{aligned}$$

Here, the second identity relies on the facts that $\nabla f(\mathbf{F}; \mathbf{S})\mathbf{R} = \nabla f(\mathbf{F}\mathbf{R}; \mathbf{S})$ and $\nabla f^{(l)}(\mathbf{F}; \mathbf{S})\mathbf{R} = \nabla f^{(l)}(\mathbf{F}\mathbf{R}; \mathbf{S})$ for any orthonormal matrix $\mathbf{R} \in \mathcal{O}^{r \times r}$. We shall then control \mathbf{C}_1 , \mathbf{C}_2 , \mathbf{C}_3 and \mathbf{C}_4 separately.

Employing the same strategy used to bound \mathbf{A}_1 and \mathbf{A}_2 in the proof of [CCF+19, Lemma 12], we can demonstrate that

$$\|\mathbf{C}_1\|_{\text{F}} \leq \left(1 - \frac{\sigma_{\min}}{20}\eta\right)\|\mathbf{F}^t\mathbf{H}^t - \mathbf{F}^{t,(l)}\mathbf{R}^{t,(l)}\|_{\text{F}} \quad \text{and} \quad \|\mathbf{C}_2\|_{\text{F}} \leq \eta\left(\sigma\sqrt{\frac{n}{p}} + \frac{\lambda}{p}\right)\|\mathbf{F}^*\|_{2,\infty},$$

provided that $\frac{\sigma}{\sigma_{\min}}\sqrt{\frac{n}{p}} \ll 1/\sqrt{\kappa^4\mu r \log n}$ and $\eta \ll 1/(n\kappa^2\sigma_{\max})$. With regards to \mathbf{C}_3 , it is seen from the definitions of ∇f and $\nabla f^{(l)}$ that

$$\mathbf{C}_3 = \eta \left[\begin{aligned} &[\mathcal{P}_{l,\cdot}(\mathbf{X}^{t,(l)}\mathbf{Y}^{t,(l)} - \mathbf{L}^*) - p^{-1}\mathcal{P}_{(\Omega_{\text{obs}})l,\cdot}(\mathbf{X}^{t,(l)}\mathbf{Y}^{t,(l)} - \mathbf{L}^*)]^\top \mathbf{Y}^{t,(l)}\mathbf{R}^{t,(l)} + p^{-1}\mathcal{P}_{(\Omega_{\text{obs}})l,\cdot}(\mathbf{E})\mathbf{Y}^{t,(l)}\mathbf{R}^{t,(l)} \\ &[\mathcal{P}_{l,\cdot}(\mathbf{X}^{t,(l)}\mathbf{Y}^{t,(l)} - \mathbf{L}^*) - p^{-1}\mathcal{P}_{(\Omega_{\text{obs}})l,\cdot}(\mathbf{X}^{t,(l)}\mathbf{Y}^{t,(l)} - \mathbf{L}^*)]^\top \mathbf{X}^{t,(l)}\mathbf{R}^{t,(l)} + p^{-1}\mathcal{P}_{(\Omega_{\text{obs}})l,\cdot}(\mathbf{E})^\top \mathbf{X}^{t,(l)}\mathbf{R}^{t,(l)} \end{aligned} \right],$$

which has the same form as \mathbf{A}_3 in the proof of [CCF+19, Lemma 12]. It thus follows from [CCF+19, Claim 5, 6 and 7] that

$$\|\mathbf{C}_3\|_{\text{F}} \lesssim \eta\sigma\sqrt{\frac{n\log n}{p}}\|\mathbf{F}^*\|_{2,\infty} + \eta\sqrt{\frac{\mu^2 r^2 \log n}{np}}\|\mathbf{F}^{t,(l)}\mathbf{R}^{t,(l)} - \mathbf{F}^*\|_{2,\infty}\sigma_{\max},$$

provided that $\frac{\sigma}{\sigma_{\min}}\sqrt{\frac{n}{p}} \ll \frac{1}{\sqrt{\kappa^2 \log n}}$ and that $n^2 p \gg n \log^3 n$.

We are then left with controlling the term C_4 . Towards this, we invoke the definition of f to decompose

$$\begin{aligned} C_4 &= \eta \begin{bmatrix} p^{-1} \mathcal{P}_{\Omega_{\text{obs}}} (\mathbf{S}^{t,(l)} - \mathbf{S}^t) \mathbf{Y}^{t,(l)} \mathbf{R}^{t,(l)} \\ p^{-1} \mathcal{P}_{\Omega_{\text{obs}}} (\mathbf{S}^{t,(l)} - \mathbf{S}^t)^\top \mathbf{X}^{t,(l)} \mathbf{R}^{t,(l)} \end{bmatrix} \\ &= \frac{\eta}{p} \underbrace{\begin{bmatrix} \mathcal{P}_{-l,\cdot} (\mathbf{S}^{t,(l)} - \mathbf{S}^t) \mathbf{Y}^{t,(l)} \mathbf{R}^{t,(l)} \\ [\mathcal{P}_{-l,\cdot} (\mathbf{S}^{t,(l)} - \mathbf{S}^t)]^\top \mathbf{X}^{t,(l)} \mathbf{R}^{t,(l)} \end{bmatrix}}_{=: D_1} + \frac{\eta}{p} \underbrace{\begin{bmatrix} \mathcal{P}_{l,\cdot} (\mathbf{S}^{t,(l)} - \mathbf{S}^t) \mathbf{Y}^{t,(l)} \mathbf{R}^{t,(l)} \\ [\mathcal{P}_{l,\cdot} (\mathbf{S}^{t,(l)} - \mathbf{S}^t)]^\top \mathbf{X}^{t,(l)} \mathbf{R}^{t,(l)} \end{bmatrix}}_{=: D_2}. \end{aligned}$$

Here, we have used the fact that both $\mathbf{S}^{t,(l)}$ and \mathbf{S}^t are supported on $\Omega^* \subseteq \Omega_{\text{obs}}$. Regarding the first matrix D_1 , we have the following fact.

Fact 4. *Suppose that the sample size obeys $n^2 p \gg \kappa^4 \mu^2 r^2 n \log n$, the noise satisfies $\frac{\sigma}{\sigma_{\min}} \sqrt{\frac{n \log n}{p}} \ll 1/\kappa$, the outlier fraction satisfies $\rho_s \leq \rho_{\text{aug}} \ll 1/\kappa^3$ and $n^2 p \rho_{\text{aug}} \gg \mu r n \log n$ hold. Then with probability at least $1 - O(n^{-100})$, we have*

$$\|D_1\|_{\text{F}} \lesssim \eta \sigma \sqrt{\frac{n \log n}{p}} \|\mathbf{F}^*\|_{2,\infty}.$$

With regards to D_2 , recall that $\mathbf{S}_{l,\cdot}^{t,(l)} = \mathbf{S}_{l,\cdot}^*$. Using the decomposition (F.11) in the proof of Lemma 17, and recalling that $\mathbf{B}^{t+1} = \mathbf{0}$ from the proof of Lemma 17, we obtain

$$\mathcal{P}_{l,\cdot} (\mathbf{S}^{t,(l)} - \mathbf{S}^t) \mathbf{Y}^{t,(l)} \mathbf{R}^{t,(l)} = \mathcal{P}_{l,\cdot} (\mathbf{A}_1 + \mathbf{E}) \mathbf{Y}^{t,(l)} \mathbf{R}^{t,(l)} + \mathcal{P}_{l,\cdot} (\mathbf{A}_2^t) \mathbf{Y}^{t,(l)} \mathbf{R}^{t,(l)}, \quad (\text{F.17})$$

where

$$\begin{aligned} \mathbf{A}_1 &= \mathcal{S}_\tau [\mathcal{P}_{\Omega_{\text{aug}}} (\mathbf{S}^* + \mathbf{E})] - \mathcal{P}_{\Omega_{\text{aug}}} (\mathbf{S}^* + \mathbf{E}); \\ \mathbf{A}_2^t &:= \mathcal{S}_\tau [\mathcal{P}_{\Omega_{\text{aug}}} (\mathbf{X}^* \mathbf{Y}^{*\top} - \mathbf{X}^t \mathbf{Y}^{t\top} + \mathbf{S}^* + \mathbf{E})] - \mathcal{S}_\tau [\mathcal{P}_{\Omega_{\text{aug}}} (\mathbf{S}^* + \mathbf{E})]. \end{aligned}$$

For the first term $\mathcal{P}_{l,\cdot} (\mathbf{A}_1 + \mathbf{E}) \mathbf{Y}^{t,(l)} \mathbf{R}^{t,(l)}$, the independence between $\mathbf{Y}^{t,(l)} \mathbf{R}^{t,(l)}$ and the l -th row of $\mathbf{A}_1 + \mathbf{E}$ allows us to obtain the following bound.

Fact 5. *Suppose that $\rho_s \leq \rho_{\text{aug}} \ll 1/\log n$ and that $n^2 p \gg n \log^4 n$. Then with probability at least $1 - O(n^{-100})$, we have*

$$\left\| \mathcal{P}_{l,\cdot} (\mathbf{A}_1) \mathbf{Y}^{t,(l)} \mathbf{R}^{t,(l)} \right\|_{\text{F}} \lesssim \sigma \sqrt{np \log n} \|\mathbf{Y}^*\|_{2,\infty}.$$

The term involving \mathbf{A}_2^t is controlled in the following claim, which relies heavily on the small scale of the entries in \mathbf{A}_2^t .

Fact 6. *Suppose that $n \gg \kappa \mu r$, $\frac{\sigma}{\sigma_{\min}} \sqrt{\frac{n \log n}{p}} \ll 1/\kappa$, $\rho_s \leq \rho_{\text{aug}} \ll 1/(\kappa \mu r)$ and that $n^2 p \rho_{\text{aug}} \gg n \log n$. Then with probability at least $1 - O(n^{-100})$, we have*

$$\left\| \mathcal{P}_{l,\cdot} (\mathbf{A}_2^t) \mathbf{Y}^{t,(l)} \mathbf{R}^{t,(l)} \right\|_{\text{F}} \lesssim \sigma \sqrt{np \log n} \|\mathbf{Y}^*\|_{2,\infty}.$$

Combining the two bounds in Facts 5 and 6 gives

$$\left\| \mathcal{P}_{l,\cdot} (\mathbf{S}^{t,(l)} - \mathbf{S}^t) \mathbf{Y}^{t,(l)} \mathbf{R}^{t,(l)} \right\|_{\text{F}} \lesssim \sigma \sqrt{np \log n} \|\mathbf{Y}^*\|_{2,\infty}.$$

The same bound applies to $\left\| \mathcal{P}_{l,\cdot} (\mathbf{S}^{t,(l)} - \mathbf{S}^t)^\top \mathbf{X}^{t,(l)} \mathbf{R}^{t,(l)} \right\|_{\text{F}}$ via the same technique. As a result, we have

$$\|D_2\|_{\text{F}} \lesssim \eta \sigma \sqrt{\frac{n \log n}{p}} \|\mathbf{F}^*\|_{2,\infty}.$$

Putting the above bounds together yields

$$\left\| \mathbf{F}^{t+1} \mathbf{H}^{t+1} - \mathbf{F}^{t+1,(l)} \mathbf{R}^{t+1,(l)} \right\|_{\text{F}} \leq \|C_1\|_{\text{F}} + \|C_2\|_{\text{F}} + \|C_3\|_{\text{F}} + \|D_1\|_{\text{F}} + \|D_2\|_{\text{F}}$$

$$\begin{aligned}
&\leq \left(1 - \frac{\sigma_{\min}}{20}\eta\right) \left\| \mathbf{F}^t \mathbf{H}^t - \mathbf{F}^{t,(l)} \mathbf{R}^{t,(l)} \right\|_{\mathbb{F}} + \eta \left(\sigma \sqrt{\frac{n}{p}} + \frac{\lambda}{p} \right) \|\mathbf{F}^*\|_{2,\infty} \\
&\quad + \tilde{C} \left(\eta \sigma \sqrt{\frac{n \log n}{p}} \|\mathbf{F}^*\|_{2,\infty} + \eta \sqrt{\frac{\mu^2 r^2 \log n}{np}} \left\| \mathbf{F}^{t,(l)} \mathbf{R}^{t,(l)} - \mathbf{F}^* \right\|_{2,\infty} \sigma_{\max} \right) + \tilde{C} \eta \sigma \sqrt{\frac{n \log n}{p}} \|\mathbf{F}^*\|_{2,\infty} \\
&\stackrel{(i)}{\leq} \left(1 - \frac{\sigma_{\min}}{20}\eta\right) C_1 \left(\frac{\sigma}{\sigma_{\min}} \sqrt{\frac{n \log n}{p}} + \frac{\lambda}{p \sigma_{\min}} \right) \|\mathbf{F}^*\|_{2,\infty} + \eta \left(\sigma \sqrt{\frac{n}{p}} + \frac{\lambda}{p} \right) \|\mathbf{F}^*\|_{2,\infty} + \tilde{C} \frac{\lambda}{p} \|\mathbf{F}^*\|_{2,\infty} \\
&\quad + \tilde{C} \eta \sqrt{\frac{\mu^2 r^2 \log n}{np}} (C_{\infty} \kappa + C_1) \left(\frac{\sigma}{\sigma_{\min}} \sqrt{\frac{n \log n}{p}} + \frac{\lambda}{p \sigma_{\min}} \right) \|\mathbf{F}^*\|_{2,\infty} \sigma_{\max} + \tilde{C} \eta \sigma \sqrt{\frac{n \log n}{p}} \|\mathbf{F}^*\|_{2,\infty} \\
&\stackrel{(ii)}{\leq} C_1 \left(\frac{\sigma}{\sigma_{\min}} \sqrt{\frac{n \log n}{p}} + \frac{\lambda}{p \sigma_{\min}} \right) \|\mathbf{F}^*\|_{2,\infty},
\end{aligned}$$

where (i) invokes (F.6a) and its immediate consequence that

$$\left\| \mathbf{F}^{t,(l)} \mathbf{R}^{t,(l)} - \mathbf{F}^* \right\|_{2,\infty} \leq \left\| \mathbf{F}^t \mathbf{H}^t - \mathbf{F}^{t,(l)} \mathbf{R}^{t,(l)} \right\|_{\mathbb{F}} + \left\| \mathbf{F}^t \mathbf{H}^t - \mathbf{F}^* \right\|_{2,\infty} \quad (\text{F.18})$$

$$\leq (C_{\infty} \kappa + C_1) \left(\frac{\sigma}{\sigma_{\min}} \sqrt{\frac{n \log n}{p}} + \frac{\lambda}{p \sigma_{\min}} \right) \|\mathbf{F}^*\|_{2,\infty}. \quad (\text{F.19})$$

The last line (ii) holds as long as $n^2 p \gg \kappa^4 \mu^2 r^2 n \log n$ and C_1 is large enough.

Proof of Fact 4. First notice that \mathbf{S}^t is supported on Ω_{aug} , which is a consequence of (F.11) and (F.16) as long as $\frac{\sigma}{\sigma_{\min}} \sqrt{\frac{n \log n}{p}} \ll 1/\kappa$ and $n^2 p \gg \kappa^4 \mu^2 r^2 n \log n$. By replacing \mathbf{X}^{t+1} (resp. \mathbf{Y}^{t+1}) with $\mathbf{X}^{t+1,(l)}$ (resp. $\mathbf{Y}^{t+1,(l)}$) and invoking (F.19) instead of (F.3c), the same arguments yield the fact that $\mathbf{S}^{t,(l)}$ is also supported on Ω_{aug} . Define $\omega_{ij} := \mathbf{1}_{(i,j) \in \Omega_{\text{aug}}}$. The Frobenius norm of the upper block of \mathbf{D}_1 can be bounded by

$$\begin{aligned}
\left\| \mathcal{P}_{-l, \cdot} (\mathbf{S}^{t,(l)} - \mathbf{S}^t) \mathbf{Y}^{t,(l)} \mathbf{R}^{t,(l)} \right\|_{\mathbb{F}}^2 &= \sum_{i:i \neq l} \sum_{j=1}^r \left[\sum_{k=1}^n (\mathbf{S}^{t,(l)} - \mathbf{S}^t)_{ik} Y_{kj}^{t,(l)} \right]^2 \\
&= \sum_{i:i \neq l} \sum_{j=1}^r \left[\sum_{k=1}^n \omega_{ik} (\mathbf{S}^{t,(l)} - \mathbf{S}^t)_{ik} Y_{kj}^{t,(l)} \right]^2 \leq \sum_{i:i \neq l} \sum_{j=1}^r \left[\sum_{k=1}^n \omega_{ik} (\mathbf{S}^{t,(l)} - \mathbf{S}^t)_{ik}^2 \right] \left[\sum_{k=1}^n \omega_{ik} (Y_{kj}^{t,(l)})^2 \right],
\end{aligned}$$

where we use the Cauchy-Schwarz inequality in the last step. Converting to the matrix notation, we obtain

$$\sum_{k=1}^n \omega_{ik} (Y_{kj}^{t,(l)})^2 = \left\| \mathcal{P}_{\Omega_{\text{aug}}} \left(\mathbf{e}_i \mathbf{e}_j^{\top} \mathbf{Y}^{t,(l)\top} \right) \right\|_{\mathbb{F}}^2.$$

Applying a similar argument as in bounding (F.10), one can obtain from Lemma 4 that

$$\sum_{j=1}^r \left\| \mathcal{P}_{\Omega_{\text{aug}}} \left(\mathbf{e}_i \mathbf{e}_j^{\top} \mathbf{Y}^{t,(l)\top} \right) \right\|_{\mathbb{F}}^2 \lesssim \kappa p \rho_{\text{aug}} \sum_{j=1}^r \left\| \mathbf{Y}_{\cdot,j}^{t,(l)} \right\|_{\mathbb{F}}^2 = \kappa p \rho_{\text{aug}} \|\mathbf{Y}^{t,(l)}\|_{\mathbb{F}}^2,$$

provided that $n^2 p \rho_{\text{aug}} \gg \mu r n \log n$. This allows us to reach

$$\begin{aligned}
\left\| \mathcal{P}_{-l, \cdot} (\mathbf{S}^{t,(l)} - \mathbf{S}^t) \mathbf{Y}^{t,(l)} \mathbf{R}^{t,(l)} \right\|_{\mathbb{F}} &\lesssim \sqrt{\sum_{i:i \neq l} \left[\sum_{k=1}^n \omega_{ik} (\mathbf{S}^{t,(l)} - \mathbf{S}^t)_{ik}^2 \right]} \cdot \sqrt{\kappa p \rho_{\text{aug}}} \|\mathbf{Y}^{t,(l)}\|_{\mathbb{F}} \\
&= \sqrt{\kappa p \rho_{\text{aug}}} \left\| \mathcal{P}_{-l, \cdot} (\mathbf{S}^{t,(l)} - \mathbf{S}^t) \right\|_{\mathbb{F}} \|\mathbf{Y}^{t,(l)}\|_{\mathbb{F}}
\end{aligned}$$

$$\begin{aligned}
&\leq C_3 \sqrt{\kappa p \rho_{\text{aug}}} \frac{\sigma}{\sigma_{\min}} \sqrt{n \log n} \|\mathbf{F}^*\|^2 \|\mathbf{F}^*\|_{2,\infty} \\
&\lesssim \sigma \sqrt{np \log n} \|\mathbf{F}^*\|_{2,\infty}.
\end{aligned}$$

Here, the penultimate step comes from the hypothesis (F.6c), whereas the last step holds as long as $\rho_s \leq \rho_{\text{aug}} \ll 1/\kappa^3$. The Frobenius norm of the lower block of \mathbf{D}_1 admits the same bound. As a result, we obtain $\|\mathbf{D}_1\|_{\text{F}} \lesssim \eta \sigma \sqrt{\frac{n \log n}{p}} \|\mathbf{F}^*\|_{2,\infty}$ as claimed. \square

Proof of Fact 5. Regarding the first term on the right-hand side of (F.17), we can write

$$\left\| \mathcal{P}_{l,\cdot} (\mathbf{A}_1 + \mathbf{E}) \mathbf{Y}^{t,(l)} \mathbf{R}^{t,(l)} \right\|_{\text{F}} = \left\| \sum_{j=1}^n (\mathbf{A}_1 + \mathbf{E})_{lj} \mathbf{Y}_{j,\cdot}^{t,(l)} \right\|_2 = \left\| \sum_{j=1}^n \underbrace{\omega_{lj} [\mathcal{S}_\tau(S_{lj}^* + E_{ij}) - S_{lj}^*]}_{=: \mathbf{u}_j} \mathbf{Y}_{j,\cdot}^{t,(l)} \right\|_2,$$

where $\omega_{lj} := \mathbb{1}\{(l, j) \in \Omega_{\text{aug}}\}$ is a Bernoulli random variable with mean $p\rho_{\text{aug}}$. Since $\mathbf{Y}^{t,(l)}$ is independent of $\{\omega_{lj}\}_{1 \leq j \leq n}$ and $\mathbf{S}_{l,\cdot}^*$, the vectors $\{\mathbf{u}_j\}_{j=1}^n$ are statistically independent conditional on $\mathbf{Y}^{t,(l)}$. We can thus apply the matrix Bernstein inequality to control this term. Specifically, conditional on $\mathbf{Y}^{t,(l)}$, we have

$$\begin{aligned}
\|\mathbf{u}_j\|_2 \|\cdot\|_{\psi_1} &\leq \|\mathbf{Y}^{t,(l)}\|_{2,\infty} \|\omega_{lj} [\mathcal{S}_\tau(S_{lj}^* + E_{ij}) - S_{lj}^*]\|_{\psi_1} \stackrel{(i)}{\lesssim} \tau \|\mathbf{Y}^{t,(l)}\|_{2,\infty}, \\
V &:= \left\| \mathbb{E} \left[\sum_{j=1}^n \omega_{lj}^2 (\mathcal{S}_\tau(S_{lj}^* + E_{ij}) - S_{lj}^*)^2 \mathbf{Y}_{j,\cdot}^{t,(l)} \mathbf{Y}_{j,\cdot}^{t,(l)\top} \right] \right\| \stackrel{(ii)}{\lesssim} p\rho_{\text{aug}} \tau^2 \|\mathbf{Y}^{t,(l)}\|_{\text{F}}^2,
\end{aligned}$$

where $\|\cdot\|_{\psi_1}$ denotes the sub-exponential norm [Ver17]. Here, the relation (i) holds since

$$\begin{aligned}
\|\omega_{lj} [\mathcal{S}_\tau(S_{lj}^* + E_{ij}) - S_{lj}^*]\|_{\psi_1} &\leq \|\mathcal{S}_\tau(S_{lj}^* + E_{ij}) - S_{lj}^*\|_{\psi_1} \leq \|\mathcal{S}_\tau(S_{lj}^* + E_{ij}) - (S_{lj}^* + E_{ij})\|_{\psi_1} + \|E_{ij}\|_{\psi_1} \\
&\leq |\mathcal{S}_\tau(S_{lj}^* + E_{ij}) - (S_{lj}^* + E_{ij})| + \|E_{ij}\|_{\psi_2} \leq 2\tau,
\end{aligned}$$

where we have used the fact that $|\mathcal{S}_\tau(x) - x| \leq \tau$ and $\|E_{ij}\|_{\psi_1} \leq \|E_{ij}\|_{\psi_2} \leq \sigma \leq \tau$. In addition, the second inequality (ii) comes from the identity $\mathbb{E}[\omega_{lj}^2] = p\rho_{\text{aug}}$ and the fact that

$$\mathbb{E} \left[(\mathcal{S}_\tau(S_{lj}^* + E_{ij}) - S_{lj}^*)^2 \right] \leq 2\mathbb{E} \left[(\mathcal{S}_\tau(S_{lj}^* + E_{ij}) - S_{lj}^* - E_{ij})^2 \right] + 2\mathbb{E} [E_{ij}^2] \lesssim \tau^2.$$

With the aid of the above bounds, we can invoke the matrix Bernstein inequality [KLT11, Proposition 2] to reach

$$\begin{aligned}
\left\| \sum_{j=1}^n \mathbf{u}_j \right\|_2 &\lesssim \sqrt{V \log n} + \|\mathbf{u}_j\|_2 \|\cdot\|_{\psi_1} \log^2 n \\
&\lesssim \sqrt{p\rho_{\text{aug}} \tau^2 \|\mathbf{Y}^{t,(l)}\|_{\text{F}}^2 \log n} + \tau \|\mathbf{Y}^{t,(l)}\|_{2,\infty} \log^2 n \\
&\lesssim \left(\tau \sqrt{np\rho_{\text{aug}} \log n} + \tau \log^2 n \right) \|\mathbf{Y}^{t,(l)}\|_{2,\infty}
\end{aligned}$$

with probability at least $1 - O(n^{-10})$. Here, the last inequality arises from $\|\mathbf{Y}^{t,(l)}\|_{\text{F}}^2 \leq n \|\mathbf{Y}^{t,(l)}\|_{2,\infty}^2$. Consequently, we conclude that, with high probability,

$$\left\| \mathcal{P}_{l,\cdot} (\mathbf{A}_1 + \mathbf{E}) \mathbf{Y}^{t,(l)} \mathbf{R}^{t,(l)} \right\|_{\text{F}} \lesssim \left(\tau \sqrt{np\rho_{\text{aug}} \log n} + \tau \log^2 n \right) \|\mathbf{Y}^{t,(l)}\|_{2,\infty} \lesssim \sigma \sqrt{np \log n} \|\mathbf{Y}^*\|_{2,\infty},$$

with the proviso that $\rho_s \leq \rho_{\text{aug}} \ll 1/\log n$ and $n^2 p \gg n \log^4 n$. \square

Proof of Fact 6. Regarding the second term on the right-hand side of (F.17), we have

$$\begin{aligned}
\left\| \mathcal{P}_{l,\cdot} (\mathbf{A}_2^t) \mathbf{Y}^{t,(l)} \mathbf{R}^{t,(l)} \right\|_{\text{F}} &= \left\| \sum_{j=1}^n (\mathbf{A}_2^t)_{lj} \mathbf{Y}_{j,\cdot}^{t,(l)} \right\|_2 \\
&\stackrel{(i)}{\leq} 2np\rho_{\text{aug}} \|\mathbf{A}_2^t\|_{\infty} \|\mathbf{Y}^{t,(l)}\|_{2,\infty}
\end{aligned}$$

$$\begin{aligned}
&\stackrel{\text{(ii)}}{\leq} 12np\rho_{\text{aug}}C_\infty \left(\frac{\sigma}{\sigma_{\min}} \sqrt{\frac{n \log n}{p}} + \frac{\lambda}{p\sigma_{\min}} \right) \frac{\mu r}{n} \sigma_{\max} \|\mathbf{Y}^*\|_{2,\infty} \\
&\stackrel{\text{(iii)}}{\lesssim} \sigma \sqrt{np \log n} \|\mathbf{Y}^*\|_{2,\infty}.
\end{aligned}$$

Here, the first upper bound (i) arises from the fact that $\{j \mid (\mathbf{A}_2^t)_{lj} \neq 0\} \subseteq \{j \mid (l, j) \in \Omega_{\text{aug}}\}$, whose cardinality is upper bounded by $2np\rho_{\text{aug}}$ with high probability as long as $np\rho_{\text{aug}} \gg \log n$. The second inequality (ii) comes from the simple fact that $\|\mathbf{Y}^{t,(l)}\|_{2,\infty} \leq 2\|\mathbf{Y}^*\|_{2,\infty}$ as well as the bound

$$\begin{aligned}
\|\mathbf{A}_2^t\|_\infty &= \|\mathcal{S}_\tau [\mathcal{P}_{\Omega_{\text{aug}}} (\mathbf{X}^* \mathbf{Y}^{*\top} - \mathbf{X}^t \mathbf{Y}^{t\top} + \mathbf{S}^* + \mathbf{E})] - \mathcal{S}_\tau [\mathcal{P}_{\Omega_{\text{aug}}} (\mathbf{S}^* + \mathbf{E})]\|_\infty \\
&\leq \|\mathcal{P}_{\Omega_{\text{aug}}} (\mathbf{X}^* \mathbf{Y}^{*\top} - \mathbf{X}^t \mathbf{Y}^{t\top} + \mathbf{S}^* + \mathbf{E}) - \mathcal{P}_{\Omega_{\text{aug}}} (\mathbf{S}^* + \mathbf{E})\|_\infty \\
&\leq \|\mathbf{X}^* \mathbf{Y}^{*\top} - \mathbf{X}^t \mathbf{Y}^{t\top}\|_\infty \\
&\leq 3C_\infty \left(\frac{\sigma}{\sigma_{\min}} \sqrt{\frac{n \log n}{p}} + \frac{\lambda}{p\sigma_{\min}} \right) \frac{\mu r}{n} \sigma_{\max},
\end{aligned}$$

where we use the non-expansiveness of $\mathcal{S}_\tau(\cdot)$ and the established bound (F.15), which holds as long as $\frac{\sigma}{\sigma_{\min}} \sqrt{\frac{n \log n}{p}} \ll 1/\kappa$. Last but not least, the relation (iii) holds as long as $\rho_s \leq \rho_{\text{aug}} \ll 1/(\kappa\mu r)$ and $n \gg \kappa\mu r$. \square

F.6 Proof of Lemma 19

Without loss of generality, we assume $1 \leq l \leq n$. Following the definitions of $\mathbf{S}^{t+1,(l)}$ and \mathbf{S}^{t+1} , we have

$$\begin{aligned}
\|\mathcal{P}_{-l,\cdot}(\mathbf{S}^{t+1,(l)} - \mathbf{S}^{t+1})\|_{\text{F}} &= \left\| \mathcal{P}_{-l,\cdot} \left[\mathcal{S}_\tau(\mathbf{M} - \mathbf{X}^{t+1,(l)} \mathbf{Y}^{t+1,(l)\top}) - \mathcal{S}_\tau(\mathbf{M} - \mathbf{X}^{t+1} \mathbf{Y}^{t+1\top}) \right] \right\|_{\text{F}} \\
&\leq \|\mathcal{P}_{\Omega_{\text{aug}}}(\Delta)\|_{\text{F}} + \|\mathcal{P}_{\Omega_{\text{aug}}^c}(\Delta)\|_{\text{F}}, \tag{F.20}
\end{aligned}$$

where we denote $\Delta := \mathcal{S}_\tau(\mathbf{M} - \mathbf{X}^{t+1,(l)} \mathbf{Y}^{t+1,(l)\top}) - \mathcal{S}_\tau(\mathbf{M} - \mathbf{X}^{t+1} \mathbf{Y}^{t+1\top})$. Recall from Appendix A that each (i, j) is included in Ω_{aug} independently with probability $p\rho_{\text{aug}}$, where $1 \geq \rho_{\text{aug}} \geq \rho_s$.

1. For the first term $\|\mathcal{P}_{\Omega_{\text{aug}}}(\Delta)\|_{\text{F}}$, the non-expansiveness of the proximal operator $\mathcal{S}_\tau(\cdot)$ yields

$$\|\mathcal{P}_{\Omega_{\text{aug}}}(\Delta)\|_{\text{F}} \leq \left\| \mathcal{P}_{\Omega_{\text{aug}}}(\mathbf{X}^{t+1,(l)} \mathbf{Y}^{t+1,(l)\top} - \mathbf{X}^{t+1} \mathbf{Y}^{t+1\top}) \right\|_{\text{F}}.$$

Apply Lemma 4 and a similar argument in bounding (F.10) to obtain

$$\begin{aligned}
\|\mathcal{P}_{\Omega_{\text{aug}}}(\Delta)\|_{\text{F}} &\leq \left\| \mathcal{P}_{\Omega_{\text{aug}}} \left[\mathbf{X}^{t+1} \mathbf{H}^{t+1} \left(\mathbf{Y}^{t+1,(l)} \mathbf{R}^{t+1,(l)} - \mathbf{Y}^{t+1} \mathbf{H}^{t+1} \right)^\top \right] \right\|_{\text{F}} \\
&\quad + \left\| \mathcal{P}_{\Omega_{\text{aug}}} \left[\left(\mathbf{X}^{t+1,(l)} \mathbf{R}^{t+1,(l)} - \mathbf{X}^{t+1} \mathbf{H}^{t+1} \right) \mathbf{R}^{t+1,(l)\top} \mathbf{Y}^{t+1,(l)\top} \right] \right\|_{\text{F}} \\
&\lesssim \sqrt{\kappa p \rho_{\text{aug}}} \|\mathbf{X}^{t+1} \mathbf{H}^{t+1}\| \left\| \mathbf{Y}^{t+1,(l)} \mathbf{R}^{t+1,(l)} - \mathbf{Y}^{t+1} \mathbf{H}^{t+1} \right\|_{\text{F}} \\
&\quad + \sqrt{\kappa p \rho_{\text{aug}}} \left\| \mathbf{Y}^{t+1,(l)} \mathbf{H}^{t+1,(l)} \right\| \left\| \mathbf{X}^{t+1,(l)} \mathbf{R}^{t+1,(l)} - \mathbf{X}^{t+1} \mathbf{H}^{t+1} \right\|_{\text{F}},
\end{aligned}$$

with the proviso that $n^2 p \rho_{\text{aug}} \gg \mu r n \log n$. In view of (F.6a) and the simple facts $\|\mathbf{X}^{t+1} \mathbf{H}^{t+1}\| \leq 2\|\mathbf{X}^*\|$, $\|\mathbf{Y}^{t+1,(l)} \mathbf{H}^{t+1,(l)}\| \leq 2\|\mathbf{X}^*\|$, one has

$$\begin{aligned}
\|\mathcal{P}_{\Omega_{\text{aug}}}(\Delta)\|_{\text{F}} &\lesssim \sqrt{\kappa p \rho_{\text{aug}}} \|\mathbf{X}^*\| \left(\frac{\sigma}{\sigma_{\min}} \sqrt{\frac{n \log n}{p}} + \frac{\lambda}{p\sigma_{\min}} \right) \|\mathbf{F}^*\|_{2,\infty} \\
&\leq C_3 \frac{\sigma}{\sigma_{\min}} \sqrt{n \log n} \|\mathbf{F}^*\|_{2,\infty} \|\mathbf{F}^*\|,
\end{aligned}$$

provided that $\rho_{\text{aug}} \ll 1/\kappa$.

2. Regarding the second term $\|\mathcal{P}_{\Omega_{\text{aug}}^c}(\Delta)\|_{\text{F}}$, we first recall from (F.16) that

$$\mathcal{S}_\tau \left[\mathcal{P}_{\Omega_{\text{aug}}^c}(\mathbf{M} - \mathbf{X}^{t+1} \mathbf{Y}^{t+1\top}) \right] = \mathbf{0}.$$

By replacing \mathbf{X}^{t+1} (resp. \mathbf{Y}^{t+1}) with $\mathbf{X}^{t+1,(l)}$ (resp. $\mathbf{Y}^{t+1,(l)}$) and invoking (F.19) instead of (F.3c), the same arguments that we used to prove (F.16) also allow us to demonstrate

$$\mathcal{S}_\tau \left[\mathcal{P}_{\Omega_{\text{aug}}^c}(\mathbf{M} - \mathbf{X}^{t+1,(l)} \mathbf{Y}^{t+1,(l)\top}) \right] = \mathbf{0}$$

provided that $n^2 p \gg \kappa^4 \mu^2 r^2 n \log n$ and $\frac{\sigma}{\sigma_{\min}} \sqrt{\frac{n \log n}{p}} \ll 1/\kappa$. Consequently, we have $\mathcal{P}_{\Omega_{\text{aug}}^c}(\Delta) = \mathbf{0}$.

Substituting the above two bounds into (F.20), we conclude that

$$\left\| \mathcal{P}_{-l}(\mathbf{S}^{t+1,(l)} - \mathbf{S}^{t+1}) \right\|_{\text{F}} \leq \|\mathcal{P}_{\Omega_{\text{aug}}}(\Delta)\|_{\text{F}} \leq C_3 \frac{\sigma}{\sigma_{\min}} \sqrt{n \log n} \|\mathbf{F}^*\|_{2,\infty} \|\mathbf{F}^*\|.$$

F.7 Proof of Lemma 20

Following [CCF⁺19, Lemma 18], we already know that

$$f(\mathbf{X}^{t+1}, \mathbf{Y}^{t+1}; \mathbf{S}^t) \leq f(\mathbf{X}^t, \mathbf{Y}^t; \mathbf{S}^t) - \frac{\eta}{2} \|\nabla f(\mathbf{X}^t, \mathbf{Y}^t; \mathbf{S}^t)\|_{\text{F}}^2. \quad (\text{F.21})$$

As a result, one has

$$\begin{aligned} F(\mathbf{X}^{t+1}, \mathbf{Y}^{t+1}, \mathbf{S}^{t+1}) &\stackrel{(i)}{\leq} F(\mathbf{X}^{t+1}, \mathbf{Y}^{t+1}, \mathbf{S}^t) = f(\mathbf{X}^{t+1}, \mathbf{Y}^{t+1}; \mathbf{S}^t) + \tau \|\mathbf{S}^t\|_1 \\ &\stackrel{(ii)}{\leq} f(\mathbf{X}^t, \mathbf{Y}^t; \mathbf{S}^t) - \frac{\eta}{2} \|\nabla f(\mathbf{X}^t, \mathbf{Y}^t; \mathbf{S}^t)\|_{\text{F}}^2 + \tau \|\mathbf{S}^t\|_1 \\ &= F(\mathbf{X}^t, \mathbf{Y}^t, \mathbf{S}^t) - \frac{\eta}{2} \|\nabla f(\mathbf{X}^t, \mathbf{Y}^t; \mathbf{S}^t)\|_{\text{F}}^2, \end{aligned}$$

where (i) follows since, by construction, \mathbf{S}^{t+1} is the minimizer of $F(\mathbf{X}^{t+1}, \mathbf{Y}^{t+1}, \mathbf{S})$ for any given $(\mathbf{X}^{t+1}, \mathbf{Y}^{t+1})$, and (ii) arises from (F.21).

References

- [ANW12] Alekh Agarwal, Sahand Negahban, and Martin J Wainwright. Noisy matrix decomposition via convex relaxation: Optimal rates in high dimensions. *The Annals of Statistics*, 40(2):1171–1197, 2012.
- [ARR14] Ali Ahmed, Benjamin Recht, and Justin Romberg. Blind deconvolution using convex programming. *IEEE Transactions on Information Theory*, 60(3):1711–1732, 2014.
- [CC14] Yuxin Chen and Yuejie Chi. Robust spectral compressed sensing via structured matrix completion. *IEEE Transactions on Information Theory*, 60(10):6576 – 6601, 2014.
- [CC17] Yuxin Chen and Emmanuel J. Candès. Solving random quadratic systems of equations is nearly as easy as solving linear systems. *Comm. Pure Appl. Math.*, 70(5):822–883, 2017.
- [CC18] Yuxin Chen and Emmanuel Candès. The projected power method: An efficient algorithm for joint alignment from pairwise differences. *Communications on Pure and Applied Mathematics*, 71(8):1648–1714, 2018.
- [CCD⁺19] Vasileios Charisopoulos, Yudong Chen, Damek Davis, Mateo Díaz, Lijun Ding, and Dmitriy Drusvyatskiy. Low-rank matrix recovery with composite optimization: good conditioning and rapid convergence. *arXiv preprint arXiv:1904.10020*, 2019.

- [CCF⁺19] Yuxin Chen, Yuejie Chi, Jianqing Fan, Cong Ma, and Yuling Yan. Noisy matrix completion: Understanding statistical guarantees for convex relaxation via nonconvex optimization. *arXiv:1902.07698*, 2019.
- [CCFM19] Yuxin Chen, Yuejie Chi, Jianqing Fan, and Cong Ma. Gradient descent with random initialization: Fast global convergence for nonconvex phase retrieval. *Mathematical Programming*, 176(1-2):5–37, July 2019.
- [CCW19] HanQin Cai, Jian-Feng Cai, and Ke Wei. Accelerated alternating projections for robust principal component analysis. *The Journal of Machine Learning Research*, 20(1):685–717, 2019.
- [CDDD19] Vasileios Charisopoulos, Damek Davis, Mateo Díaz, and Dmitriy Drusvyatskiy. Composite optimization for robust blind deconvolution. *arXiv preprint arXiv:1901.01624*, 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.
- [CGH14] Y. Chen, L. J. Guibas, and Q. Huang. Near-optimal joint optimal matching via convex relaxation. *International Conference on Machine Learning (ICML)*, pages 100 – 108, June 2014.
- [CGJ17] Yeshwanth Cherapanamjeri, Kartik Gupta, and Prateek Jain. Nearly optimal robust matrix completion. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 797–805. JMLR. org, 2017.
- [Che15] Yudong Chen. Incoherence-optimal matrix completion. *IEEE Transactions on Information Theory*, 61(5):2909–2923, 2015.
- [CJSC13] Yudong Chen, Ali Jalali, Sujay Sanghavi, and Constantine Caramanis. Low-rank matrix recovery from errors and erasures. *IEEE Transactions on Information Theory*, 59(7):4324–4337, 2013.
- [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, October 2019.
- [CLL19] Ji Chen, Dekai Liu, and Xiaodong Li. Nonconvex rectangular matrix completion via gradient descent without $\ell_{2,\infty}$ regularization. *arXiv:1901.06116v1*, 2019.
- [CLMW11] Emmanuel Candès, Xiaodong Li, Yi Ma, and John Wright. Robust principal component analysis? *Journal of ACM*, 58(3):11:1–11:37, Jun 2011.
- [CLPC19] Changxiao Cai, Gen Li, H Vincent Poor, and Yuxin Chen. Nonconvex low-rank symmetric tensor completion from noisy data. In *Advances in Neural Information Processing Systems*, pages 1861–1872, 2019.
- [CLS15] E. Candès, X. Li, and M. Soltanolkotabi. Phase retrieval via Wirtinger flow: Theory and algorithms. *IEEE Transactions on Information Theory*, 61(4):1985–2007, April 2015.
- [CMW13] T Tony Cai, Zongming Ma, and Yihong Wu. Sparse PCA: Optimal rates and adaptive estimation. *The Annals of Statistics*, 41(6):3074–3110, 2013.
- [CP10] Emmanuel Candès and Yaniv Plan. Matrix completion with noise. *Proceedings of the IEEE*, 98(6):925 –936, June 2010.
- [CPW12] Venkat Chandrasekaran, Pablo A Parrilo, and Alan S Willsky. Latent variable graphical model selection via convex optimization. *Annals of Statistics*, 40(4):1935–1967, 2012.
- [CR09] Emmanuel Candès and Benjamin Recht. Exact matrix completion via convex optimization. *Foundations of Computational Mathematics*, 9(6):717–772, April 2009.

- [CSPW11] Venkat Chandrasekaran, Sujay Sanghavi, Pablo A Parrilo, and Alan S Willsky. Rank-sparsity incoherence for matrix decomposition. *SIAM Journal on Optimization*, 21(2):572–596, 2011.
- [CT10] Emmanuel Candès and Terence Tao. The power of convex relaxation: Near-optimal matrix completion. *IEEE Transactions on Information Theory*, 56(5):2053–2080, May 2010.
- [CW15] Yudong Chen and Martin J Wainwright. Fast low-rank estimation by projected gradient descent: General statistical and algorithmic guarantees. *arXiv:1509.03025*, 2015.
- [CW18] Jian-Feng Cai and Ke Wei. Solving systems of phaseless equations via riemannian optimization with optimal sampling complexity. *arXiv preprint arXiv:1809.02773*, 2018.
- [DC18] Lijun Ding and Yudong Chen. The leave-one-out approach for matrix completion: Primal and dual analysis. *arXiv preprint arXiv:1803.07554*, 2018.
- [DG14] David Donoho and Matan Gavish. Minimax risk of matrix denoising by singular value thresholding. *The Annals of Statistics*, 42(6):2413–2440, 2014.
- [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.
- [FFL08] Jianqing Fan, Yingying Fan, and Jinchu Lv. High dimensional covariance matrix estimation using a factor model. *Journal of Econometrics*, 147(1):186–197, 2008.
- [FLM13] J. Fan, Y. Liao, and M. Mincheva. Large covariance estimation by thresholding principal orthogonal complements. *Journal of the Royal Stat. Society: Series B*, 75(4):603–680, 2013.
- [FSZZ18] Jianqing Fan, Qiang Sun, Wen-Xin Zhou, and Ziwei Zhu. Principal component analysis for big data. *arXiv preprint arXiv:1801.01602*, 2018.
- [FWZ18] Jianqing Fan, Weichen Wang, and Yiqiao Zhong. An ℓ_∞ eigenvector perturbation bound and its application. *Journal of Machine Learning Research*, 18(207):1–42, 2018.
- [FWZ19] Jianqing Fan, Weichen Wang, and Yiqiao Zhong. Robust covariance estimation for approximate factor models. *Journal of econometrics*, 208(1):5–22, 2019.
- [FXY13] Jiashi Feng, Huan Xu, and Shuicheng Yan. Online robust PCA via stochastic optimization. In *Advances in Neural Information Processing Systems*, pages 404–412, 2013.
- [GMS13] Donald Goldfarb, Shiqian Ma, and Katya Scheinberg. Fast alternating linearization methods for minimizing the sum of two convex functions. *Mathematical Programming*, 141(1-2):349–382, 2013.
- [GQV14] Han Guo, Chenlu Qiu, and Namrata Vaswani. An online algorithm for separating sparse and low-dimensional signal sequences from their sum. *IEEE Transactions on Signal Processing*, 62(16):4284–4297, 2014.
- [Gro11] David Gross. Recovering low-rank matrices from few coefficients in any basis. *IEEE Transactions on Information Theory*, 57(3):1548–1566, March 2011.
- [GWL⁺10] Arvind Ganesh, John Wright, Xiaodong Li, Emmanuel J Candes, and Yi Ma. Dense error correction for low-rank matrices via principal component pursuit. In *2010 IEEE international symposium on information theory*, pages 1513–1517. IEEE, 2010.
- [HG13] Q. Huang and L. Guibas. Consistent shape maps via semidefinite programming. *Computer Graphics Forum*, 32(5):177–186, 2013.
- [HKZ11] Daniel Hsu, Sham M Kakade, and Tong Zhang. Robust matrix decomposition with sparse corruptions. *IEEE Transactions on Information Theory*, 57(11):7221–7234, 2011.

- [JCSX11] Ali Jalali, Yudong Chen, Sujay Sanghavi, and Huan Xu. Clustering partially observed graphs via convex optimization. In *International Conference on Machine Learning*, volume 11, pages 1001–1008, 2011.
- [JNS13] P. Jain, P. Netrapalli, and S. Sanghavi. Low-rank matrix completion using alternating minimization. In *ACM symposium on Theory of computing*, pages 665–674, 2013.
- [Jol11] Ian Jolliffe. *Principal component analysis*. Springer, 2011.
- [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. *Ann. Statist.*, 39(5):2302–2329, 2011.
- [KLT17] Olga Klopp, Karim Lounici, and Alexandre B Tsybakov. Robust matrix completion. *Probability Theory and Related Fields*, 169(1-2):523–564, 2017.
- [KMO10] R. H. Keshavan, A. Montanari, and S. Oh. Matrix completion from a few entries. *IEEE Transactions on Information Theory*, 56(6):2980–2998, June 2010.
- [KS19] Felix Kraher and Dominik Stöger. On the convex geometry of blind deconvolution and matrix completion. *arXiv preprint arXiv:1902.11156*, 2019.
- [Li13] Xiaodong Li. Compressed sensing and matrix completion with constant proportion of corruptions. *Constructive Approximation*, 37:73–99, 2013.
- [LMCC18] Yuanxin Li, Cong Ma, Yuxin Chen, and Yuejie Chi. Nonconvex matrix factorization from rank-one measurements. *arXiv:1802.06286, accepted to AISTATS*, 2018.
- [MA18] Shiqian Ma and Necdet Serhat Aybat. Efficient optimization algorithms for robust principal component analysis and its variants. *Proceedings of the IEEE*, 106(8):1411–1426, 2018.
- [MHT10] R. Mazumder, T. Hastie, and R. Tibshirani. Spectral regularization algorithms for learning large incomplete matrices. *Journal of machine learning research*, 11(Aug):2287–2322, 2010.
- [MWCC17] Cong Ma, Kaizheng Wang, Yuejie Chi, and Yuxin Chen. Implicit regularization in nonconvex statistical estimation: Gradient descent converges linearly for phase retrieval, matrix completion and blind deconvolution. *arXiv preprint arXiv:1711.10467, accepted to Foundations of Computational Mathematics*, 2017.
- [NJS13] P. Netrapalli, P. Jain, and S. Sanghavi. Phase retrieval using alternating minimization. *Advances in Neural Information Processing Systems (NIPS)*, 2013.
- [NNS⁺14] P. Netrapalli, U. Niranjan, S. Sanghavi, A. Anandkumar, and P. Jain. Non-convex robust PCA. In *Advances in Neural Information Processing Systems*, pages 1107–1115, 2014.
- [NW12] S. Negahban and M.J. Wainwright. Restricted strong convexity and weighted matrix completion: Optimal bounds with noise. *Journal of Machine Learning Research*, pages 1665–1697, May 2012.
- [Pea01] Karl Pearson. Liii. on lines and planes of closest fit to systems of points in space. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 2(11):559–572, 1901.
- [QV10] Chenlu Qiu and Namrata Vaswani. Real-time robust principal components’ pursuit. In *2010 48th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pages 591–598. IEEE, 2010.
- [QVLH14] Chenlu Qiu, Namrata Vaswani, Brian Lois, and Leslie Hogben. Recursive robust pca or recursive sparse recovery in large but structured noise. *IEEE Transactions on Information Theory*, 60(8):5007–5039, 2014.

- [Sin11] Amit Singer. Angular synchronization by eigenvectors and semidefinite programming. *Applied and computational harmonic analysis*, 30(1):20–36, 2011.
- [SL16] Ruoyu Sun and Zhi-Quan Luo. Guaranteed matrix completion via non-convex factorization. *IEEE Transactions on Information Theory*, 62(11):6535–6579, 2016.
- [SS05] Nathan Srebro and Adi Shraibman. Rank, trace-norm and max-norm. In *International Conference on Computational Learning Theory*, pages 545–560. Springer, 2005.
- [SWZ14] Yuan Shen, Zaiwen Wen, and Yin Zhang. Augmented lagrangian alternating direction method for matrix separation based on low-rank factorization. *Optimization Methods and Software*, 29(2):239–263, 2014.
- [TY11] Min Tao and Xiaoming Yuan. Recovering low-rank and sparse components of matrices from incomplete and noisy observations. *SIAM Journal on Optimization*, 21(1):57–81, 2011.
- [Ver12] R. Vershynin. Introduction to the non-asymptotic analysis of random matrices. *Compressed Sensing, Theory and Applications*, pages 210 – 268, 2012.
- [Ver17] Roman Vershynin. High dimensional probability, 2017.
- [VN18] Namrata Vaswani and Praneeth Narayanamurthy. Static and dynamic robust pca and matrix completion: A review. *Proceedings of the IEEE*, 106(8):1359–1379, 2018.
- [WCCL16] K. Wei, J.F. Cai, T. Chan, and S. Leung. Guarantees of Riemannian optimization for low rank matrix recovery. *SIAM Journal on Matrix Analysis and Applications*, 37(3):1198–1222, 2016.
- [WGE17] Gang Wang, Georgios B Giannakis, and Yonina C Eldar. Solving systems of random quadratic equations via truncated amplitude flow. *IEEE Transactions on Information Theory*, 2017.
- [WL17] Raymond KW Wong and Thomas Lee. Matrix completion with noisy entries and outliers. *The Journal of Machine Learning Research*, 18(1):5404–5428, 2017.
- [YPCC16] Xinyang Yi, Dohyung Park, Yudong Chen, and Constantine Caramanis. Fast algorithms for robust PCA via gradient descent. In *NIPS*, pages 4152–4160, 2016.
- [ZB18] Yiqiao Zhong and Nicolas Boumal. Near-optimal bound for phase synchronization. *SIAM Journal on Optimization*, 2018.
- [ZCL16] Huishuai Zhang, Yuejie Chi, and Yingbin Liang. Provable non-convex phase retrieval with outliers: Median truncated Wirtinger flow. In *International conference on machine learning*, pages 1022–1031, 2016.
- [ZL16] Qinqing Zheng and John Lafferty. Convergence analysis for rectangular matrix completion using Burer-Monteiro factorization and gradient descent. *arXiv:1605.07051*, 2016.
- [ZLGV16] Jinchun Zhan, Brian Lois, Han Guo, and Namrata Vaswani. Online (and offline) robust PCA: Novel algorithms and performance guarantees. In *Artificial intelligence and statistics*, pages 1488–1496, 2016.
- [ZLW⁺10] Z. Zhou, X. Li, J. Wright, E. Candès, and Y. Ma. Stable principal component pursuit. In *International Symposium on Information Theory*, pages 1518–1522, 2010.