

Nonconvex Matrix Completion with Linearly Parameterized Factors

Ji Chen

LJICHEN@UCDAVIS.EDU

*Department of Mathematics
University of California, Davis
Davis, CA 95616, USA*

Xiaodong Li

XDGLI@UCDAVIS.EDU

*Department of Statistics
University of California, Davis
Davis, CA 95616, USA*

Zongming Ma

ZONGMING@WHARTON.UPENN.EDU

*Department of Statistics, The Wharton School
University of Pennsylvania
Philadelphia, PA 19104, USA*

Editor: Animashree Anandkumar

Abstract

Techniques of matrix completion aim to impute a large portion of missing entries in a data matrix through a small portion of observed ones. In practice, prior information and special structures are usually employed in order to improve the accuracy of matrix completion. In this paper, we propose a unified nonconvex optimization framework for matrix completion with linearly parameterized factors. In particular, by introducing a condition referred to as Correlated Parametric Factorization, we conduct a unified geometric analysis for the nonconvex objective by establishing uniform upper bounds for low-rank estimation resulting from any local minimizer. Perhaps surprisingly, the condition of Correlated Parametric Factorization holds for important examples including subspace-constrained matrix completion and skew-symmetric matrix completion. The effectiveness of our unified nonconvex optimization method is also empirically illustrated by extensive numerical simulations.

Keywords: matrix completion, nonconvex geometry, nonconvex optimization, skew-symmetric, subspace constraints

1. Introduction

Matrix completion techniques predict missing entries in a data matrix from partially observed ones. Their most successful application includes collaborative filtering (Rennie and Srebro, 2005; Candès and Recht, 2009) in which unobserved user-item ratings are predicted with the available ones. To put the problem in mathematical terms: Let $\mathbf{M}^*$ be an $n_1 \times n_2$ matrix whose rank is $r \ll \min(n_1, n_2)$. We aim to estimate the whole matrix from a small proportion of noisy observed entries. To be specific, let $\Omega \subseteq [n_1] \times [n_2]$ be the index set that supports all observed entries. The observation is represented by

$$\mathcal{P}_\Omega(\mathbf{M}) = \mathcal{P}_\Omega(\mathbf{M}^* + \mathbf{N}), \quad (1)$$

where $\mathbf{N}$ is a matrix consisting of noise, and the operator $\mathcal{P}_\Omega(\cdot)$ preserves the entries on Ω while placing zeros in the entries on Ω^c .

Note that any rank- r matrix can be parameterized through the factorization $\mathbf{X}\mathbf{Y}^\top$, where both $\mathbf{X}$ and $\mathbf{Y}$ have r columns. With this parameterization, the regularized least squares objective proposed in Ge et al. (2016, 2017) is

$$f(\mathbf{X}, \mathbf{Y}) := \frac{1}{2p} \|\mathcal{P}_\Omega(\mathbf{X}\mathbf{Y}^\top - \mathbf{M})\|_F^2 + \frac{1}{8} \|\mathbf{X}^\top \mathbf{X} - \mathbf{Y}^\top \mathbf{Y}\|_F^2 + \lambda(G_\alpha(\mathbf{X}) + G_\alpha(\mathbf{Y})), \quad (2)$$

where

$$G_\alpha(\mathbf{X}) := \sum_{i=1}^n [(\|\mathbf{X}_{i,\cdot}\|_2 - \alpha)_+]^4 \quad (3)$$

and p denotes the sampling rate, while both α and λ can be viewed as tuning parameters.

Due to nonconvexity, standard methods such as gradient descent may converge to a local minimizer of the above objective function. A series of papers in the literature including Ge et al. (2016, 2017) and Chen and Li (2019) intend to understand the nonconvex geometry of (2). In particular, they focus on figuring out sufficient conditions for the ground-truth low rank matrix $\mathbf{M}^\star$ as well as the size and pattern of Ω , under which any local minimizer $(\widehat{\mathbf{X}}, \widehat{\mathbf{Y}})$ of (2) leads to an accurate estimate of $\mathbf{M}^\star$ by $\widehat{\mathbf{M}} = \widehat{\mathbf{X}}\widehat{\mathbf{Y}}^\top$. For example, in the noiseless case where the noise matrix $\mathbf{N} = \mathbf{0}$, if $\mathbf{M}^\star$ is a rank- r well-conditioned positive semidefinite matrix and its eigenspace incoherence parameter defined in Candès and Recht (2009) is well-bounded, then it has been shown in Chen and Li (2019) that with high probability, any local minimizer of the positive semidefinite analog of (2) yields exact low-rank recovery, i.e., there is no spurious local minimum, provided that the sampling pattern is i.i.d. and the sampling rate satisfies $p \gtrsim (r^2 \log n)/n$.

The goal of the present paper is to extend the nonconvex matrix completion objective (2) as well as its geometric analysis to cases in which prior information on the low-rank matrix is known and explored. Two examples are subspace-constrained and skew-symmetric matrix completions. We aim to understand how to adapt the nonconvex objective (2) to such cases with prior information, how to conduct corresponding geometrical analysis, and how to build a uniform framework to incorporate both cases as specific examples.

A critical observation is that both examples can be represented in the form $\mathbf{M}^\star = \mathbf{X}(\boldsymbol{\xi})\mathbf{Y}(\boldsymbol{\xi})^\top$, where the factors $\mathbf{X}(\boldsymbol{\xi})$ and $\mathbf{Y}(\boldsymbol{\xi})$ are linear mappings in $\boldsymbol{\xi} \in \mathbb{R}^d$. We elaborate on such parameterized factorizations as follows.

- Suppose $\mathbf{M}^\star$ is known to be constrained in some pre-specified column and row spaces, with dimensions s_1 and s_2 , respectively. Let $\tilde{\mathbf{U}}$ (and $\tilde{\mathbf{V}}$) be an $n_1 \times s_1$ (and $n_2 \times s_2$) matrix whose columns form an orthogonal basis for the given column (and row) space constraint for $\mathbf{M}^\star$. Given the rank of $\mathbf{M}^\star$, we know there must exist some (not unique) $\boldsymbol{\Xi}_A \in \mathbb{R}^{s_1 \times r}$ and $\boldsymbol{\Xi}_B \in \mathbb{R}^{s_2 \times r}$, such that

$$\mathbf{M}^\star = (\tilde{\mathbf{U}}\boldsymbol{\Xi}_A)(\tilde{\mathbf{V}}\boldsymbol{\Xi}_B)^\top.$$

Denote by $\boldsymbol{\theta} = \text{vec}(\boldsymbol{\Theta}_A, \boldsymbol{\Theta}_B)$ an $((s_1 + s_2)r)$ -dimensional vector that contains all entries in $\boldsymbol{\Theta}_A$ and $\boldsymbol{\Theta}_B$ (e.g., in the lexicographic order), and define the two linear mappings

$$\mathbf{X}(\boldsymbol{\theta}) = \tilde{\mathbf{U}}\boldsymbol{\Theta}_A \in \mathbb{R}^{n_1 \times r} \quad \text{and} \quad \mathbf{Y}(\boldsymbol{\theta}) = \tilde{\mathbf{V}}\boldsymbol{\Theta}_B \in \mathbb{R}^{n_2 \times r}. \quad (4)$$

Without loss of generality, assume that both $\tilde{\mathbf{U}}$ and $\tilde{\mathbf{V}}$ consist of orthonormal basis, i.e., $\tilde{\mathbf{U}}^\top \tilde{\mathbf{U}} = \mathbf{I}_{s_1}$ and $\tilde{\mathbf{V}}^\top \tilde{\mathbf{V}} = \mathbf{I}_{s_2}$. Then the above parameterized factorization becomes $\mathbf{M}^\star = \mathbf{X}(\boldsymbol{\xi})\mathbf{Y}(\boldsymbol{\xi})^\top$ with $\boldsymbol{\xi} = \text{vec}(\boldsymbol{\Xi}_A, \boldsymbol{\Xi}_B)$.

- If $\mathbf{M}^\star$ is an $n \times n$ rank- r skew-symmetric matrix (which implies that r is even), by the Youla decomposition (Youla, 1961), it can be represented (not uniquely) as

$$\mathbf{M}^\star = \boldsymbol{\Xi}_A \boldsymbol{\Xi}_B^\top - \boldsymbol{\Xi}_B \boldsymbol{\Xi}_A^\top,$$

where $\boldsymbol{\Xi}_A, \boldsymbol{\Xi}_B \in \mathbb{R}^{n \times \frac{r}{2}}$. Again, denote by $\boldsymbol{\theta} = \text{vec}(\boldsymbol{\Theta}_A, \boldsymbol{\Theta}_B)$ an (nr) -dimensional vector that contains all entries in $\boldsymbol{\Theta}_A$ and $\boldsymbol{\Theta}_B$, and define the linear mappings

$$\mathbf{X}(\boldsymbol{\theta}) = [\boldsymbol{\Theta}_A, -\boldsymbol{\Theta}_B] \in \mathbb{R}^{n \times r} \quad \text{and} \quad \mathbf{Y}(\boldsymbol{\theta}) = [\boldsymbol{\Theta}_B, \boldsymbol{\Theta}_A] \in \mathbb{R}^{n \times r}. \quad (5)$$

We also have the factorization $\mathbf{M}^\star = \mathbf{X}(\boldsymbol{\xi})\mathbf{Y}(\boldsymbol{\xi})^\top$ with $\boldsymbol{\xi} = \text{vec}(\boldsymbol{\Xi}_A, \boldsymbol{\Xi}_B)$.

Therefore, both subspace-constrained matrix completion and skew-symmetric matrix completion are special cases of the following *parameterized matrix completion* problem: If both factors for the low-rank matrix $\mathbf{M}^\star = \mathbf{X}\mathbf{Y}^\top$ can be parameterized by linear mappings $\mathbf{X}(\boldsymbol{\xi})$ and $\mathbf{Y}(\boldsymbol{\xi})$, how can we estimate $\mathbf{M}^\star$ from the partial and noisy observation $\mathcal{P}_\Omega(\mathbf{M}) = \mathcal{P}_\Omega(\mathbf{M}^\star + \mathbf{N})$?

Naturally, by substituting the parametric forms $\mathbf{X} = \mathbf{X}(\boldsymbol{\theta})$ and $\mathbf{Y} = \mathbf{Y}(\boldsymbol{\theta})$ into the nonconvex optimization (2), we obtain the following optimization with the argument $\boldsymbol{\theta}$

$$\begin{aligned} \tilde{f}(\boldsymbol{\theta}) &:= f(\mathbf{X}(\boldsymbol{\theta}), \mathbf{Y}(\boldsymbol{\theta})) \\ &= \frac{1}{2p} \|\mathcal{P}_\Omega(\mathbf{X}(\boldsymbol{\theta})\mathbf{Y}(\boldsymbol{\theta})^\top - \mathbf{M})\|_F^2 + \frac{1}{8} \|\mathbf{X}(\boldsymbol{\theta})^\top \mathbf{X}(\boldsymbol{\theta}) - \mathbf{Y}(\boldsymbol{\theta})^\top \mathbf{Y}(\boldsymbol{\theta})\|_F^2 \\ &\quad + \lambda(G_\alpha(\mathbf{X}(\boldsymbol{\theta})) + G_\alpha(\mathbf{Y}(\boldsymbol{\theta}))). \end{aligned} \quad (6)$$

The nonconvex objective (6) can be applied to various specific parameterized matrix completion problems. For example, we can substitute the parametric forms (4) into (6) to derive a nonconvex objective for subspace-constrained matrix completion, and substitute the parametric forms (5) into (6) to derive a nonconvex objective for skew-symmetric matrix completion. Details are elaborated on as follows:

- For subspace-constrained matrix completion, the parametric forms $\mathbf{X}(\boldsymbol{\theta})$ and $\mathbf{Y}(\boldsymbol{\theta})$ are defined in (4) with $\boldsymbol{\theta} = \text{vec}(\boldsymbol{\Theta}_A, \boldsymbol{\Theta}_B)$. Recall that we have assumed $\tilde{\mathbf{U}}^\top \tilde{\mathbf{U}} = \mathbf{I}_{s_1}$ and $\tilde{\mathbf{V}}^\top \tilde{\mathbf{V}} = \mathbf{I}_{s_2}$. Then the parameterization (4) implies the following

$$\begin{cases} \mathbf{X}(\boldsymbol{\theta})\mathbf{Y}(\boldsymbol{\theta})^\top = \tilde{\mathbf{U}}\boldsymbol{\Theta}_A\boldsymbol{\Theta}_B^\top\tilde{\mathbf{V}}^\top, \\ \mathbf{X}(\boldsymbol{\theta})^\top \mathbf{X}(\boldsymbol{\theta}) = \boldsymbol{\Theta}_A^\top \tilde{\mathbf{U}}^\top \tilde{\mathbf{U}} \boldsymbol{\Theta}_A = \boldsymbol{\Theta}_A^\top \boldsymbol{\Theta}_A, \\ \mathbf{Y}(\boldsymbol{\theta})^\top \mathbf{Y}(\boldsymbol{\theta}) = \boldsymbol{\Theta}_B^\top \tilde{\mathbf{V}}^\top \tilde{\mathbf{V}} \boldsymbol{\Theta}_B = \boldsymbol{\Theta}_B^\top \boldsymbol{\Theta}_B. \end{cases}$$

Substituting them into (6), we have the objective function

$$\begin{aligned} f_{\text{subspace}}(\boldsymbol{\Theta}_A, \boldsymbol{\Theta}_B) &:= \frac{1}{2p} \|\mathcal{P}_\Omega(\tilde{\mathbf{U}}\boldsymbol{\Theta}_A\boldsymbol{\Theta}_B^\top\tilde{\mathbf{V}}^\top - \mathbf{M})\|_F^2 + \frac{1}{8} \|\boldsymbol{\Theta}_A^\top \boldsymbol{\Theta}_A - \boldsymbol{\Theta}_B^\top \boldsymbol{\Theta}_B\|_F^2 \\ &\quad + \lambda(G_\alpha(\tilde{\mathbf{U}}\boldsymbol{\Theta}_A) + G_\alpha(\tilde{\mathbf{V}}\boldsymbol{\Theta}_B)). \end{aligned} \quad (7)$$

- For skew-symmetric matrix completion, the parametric forms $\mathbf{X}(\boldsymbol{\theta})$ and $\mathbf{Y}(\boldsymbol{\theta})$ are defined in (5) with $\boldsymbol{\theta} = \text{vec}(\boldsymbol{\Theta}_A, \boldsymbol{\Theta}_B)$. Straightforward calculation gives

$$\begin{cases} \mathbf{X}(\boldsymbol{\theta})\mathbf{Y}(\boldsymbol{\theta})^\top = \boldsymbol{\Theta}_A\boldsymbol{\Theta}_B^\top - \boldsymbol{\Theta}_B\boldsymbol{\Theta}_A^\top, \\ \mathbf{X}(\boldsymbol{\theta})^\top\mathbf{X}(\boldsymbol{\theta}) = \begin{bmatrix} \boldsymbol{\Theta}_A^\top\boldsymbol{\Theta}_A & -\boldsymbol{\Theta}_A^\top\boldsymbol{\Theta}_B \\ -\boldsymbol{\Theta}_B^\top\boldsymbol{\Theta}_A & \boldsymbol{\Theta}_B^\top\boldsymbol{\Theta}_B \end{bmatrix}, \\ \mathbf{Y}(\boldsymbol{\theta})^\top\mathbf{Y}(\boldsymbol{\theta}) = \begin{bmatrix} \boldsymbol{\Theta}_B^\top\boldsymbol{\Theta}_B & \boldsymbol{\Theta}_B^\top\boldsymbol{\Theta}_A \\ \boldsymbol{\Theta}_A^\top\boldsymbol{\Theta}_B & \boldsymbol{\Theta}_A^\top\boldsymbol{\Theta}_A \end{bmatrix}. \end{cases}$$

Substituting them into (6), we have the objective function

$$\begin{aligned} f_{\text{skew}}(\boldsymbol{\Theta}_A, \boldsymbol{\Theta}_B) = & \frac{1}{2p} \|\mathcal{P}_\Omega(\boldsymbol{\Theta}_A\boldsymbol{\Theta}_B^\top - \boldsymbol{\Theta}_B\boldsymbol{\Theta}_A^\top - \mathbf{M})\|_F^2 + \frac{1}{4} \|\boldsymbol{\Theta}_A^\top\boldsymbol{\Theta}_A - \boldsymbol{\Theta}_B^\top\boldsymbol{\Theta}_B\|_F^2 \quad (8) \\ & + \frac{1}{4} \|\boldsymbol{\Theta}_A^\top\boldsymbol{\Theta}_B + \boldsymbol{\Theta}_B^\top\boldsymbol{\Theta}_A\|_F^2 + 2\lambda G_\alpha([\boldsymbol{\Theta}_B, \boldsymbol{\Theta}_A]). \end{aligned}$$

Here we use the fact $G_\alpha([\boldsymbol{\Theta}_B, \boldsymbol{\Theta}_A]) = G_\alpha([\boldsymbol{\Theta}_A, -\boldsymbol{\Theta}_B])$.

Inspired by the previous studies in Ge et al. (2016, 2017) and Chen and Li (2019) regarding standard matrix completion, we are interested in a unified geometric analysis for the generic nonconvex objective (6) for parameterized matrix completion. The crux for such extension hinges on two key assumptions on the parametric forms $(\mathbf{X}(\boldsymbol{\theta}), \mathbf{Y}(\boldsymbol{\theta}))$ and the ground truth $\mathbf{M}^\star$. As suggested by the examples of subspace-constrained and skew-symmetric matrix completions, we first assume that both $\mathbf{X}(\boldsymbol{\theta})$ and $\mathbf{Y}(\boldsymbol{\theta})$ are linear mappings of $\boldsymbol{\theta}$. The second assumption regarding the relationship between $(\mathbf{X}(\boldsymbol{\theta}), \mathbf{Y}(\boldsymbol{\theta}))$ and the ground truth $\mathbf{M}^\star$, referred to as *correlated parametric factorization*, is novel in the literature of matrix completion, and plays the most essential role in our unified geometric analysis. We defer its formal definition to Section 2.1.

The most intriguing part in this work is that the non-intuitive assumption of *correlated parametric factorization* holds for various examples including the aforementioned subspace-constrained factorization (4) and skew-symmetric factorization (5). In fact, bringing both examples into a unified framework by this assumption is unexpected to us, and it is still unclear whether there is some geometrical explanation for such unification. Nevertheless, we verify the assumption of correlated parametric factorization for those two examples in Sections 4.3.1 and 4.4.1 respectively with very different algebraic arguments.

1.1 Related Work

As with most high-dimensional problems where low-complexity structures need to be explored and exploited for statistical learning and inferences, low-rank structures are a common assumption for deriving and justifying matrix completion algorithms. By imposing nuclear norm regularization to recover low-rank structures (Recht et al., 2010), convex optimization methods have been widely used in the literature of matrix completion, and their theoretical properties have been well studied; see, e.g., Candès and Recht (2009); Candès

and Tao (2010); Candès et al. (2011); Recht (2011); Gross (2011); Negahban and Wainwright (2012); Hsu et al. (2011); Sun and Zhang (2012); Koltchinskii et al. (2011); Li (2013).

Though convex optimization methods could have near-optimal theoretical guarantees for matrix completion under certain incoherence conditions, they could be unscalable to large data matrices whose dimensions are in hundreds of thousands. In contrast, nonconvex optimization methods for low-rank recovery have been proposed and analyzed in the literature due to computational convenience. Examples include optimization over Grassmann manifolds (Keshavan et al., 2010a,b), singular value projections (SVP) (Jain et al., 2010; Ding and Chen, 2020), alternating minimization (Jain et al., 2013), penalized, projected, or thresholded gradient descent (Sun and Luo, 2016; Chen and Wainwright, 2015; Cai et al., 2016; Zheng and Lafferty, 2016; Yi et al., 2016), vanilla gradient descent (Candès et al., 2015; Ma et al., 2020; Chen et al., 2020), etc. From a more geometric perspective, several works (see, e.g., Sun et al. (2018); Ge et al. (2016, 2017); Chen and Li (2019)) have established favorable properties of the landscape such as the absence of spurious local minima. The present work belongs to this geometric line of studies.

Prior information and special structures have also been explored and employed in the literature of matrix completion. Subspace-constrained matrix completion has been shown to be effective in improving the performances of collaborative filtering, and has various methodological and theoretical developments (Jain and Dhillon, 2013; Xu et al., 2013; Yi et al., 2013; Natarajan and Dhillon, 2014; Chen, 2015; Si et al., 2016; Eftekhari et al., 2018). In fact, without the penalization terms, the least squares term in (7) can be represented in form of the nonconvex objective studied in Jain and Dhillon (2013), though their distributional assumptions on rank-one sampling do not apply to our settings. Motivated by applications such as pairwise ranking, skew-symmetric structures have also been exploited for matrix completion in the literature. See, e.g., Jiang et al. (2011), Gleich and Lim (2011), and Chatterjee (2015). For example, it has been shown in Gleich and Lim (2011, Theorem 3) that if the initial input is skew-symmetric, singular value projection (SVP) (Jain et al., 2010) maintains the skew-symmetry in all iterations.

1.2 Notation

Throughout this paper, bold uppercase/lowercase characters denote matrices/vectors, respectively. For a given matrix $\mathbf{A}$, its (i, j) -th entry, i -th row, and j -th column are denoted by $A_{i,j}$, $\mathbf{A}_{i,\cdot}$, and $\mathbf{A}_{\cdot,j}$, respectively. Its spectral, Frobenius, and $\ell_{2,\infty}$ norms are denoted by $\|\mathbf{A}\|$, $\|\mathbf{A}\|_F$ and $\|\mathbf{A}\|_{2,\infty} := \max_i \|\mathbf{A}_{i,\cdot}\|_2^2$, respectively. Denote by $\text{colspan}(\mathbf{A})/\text{rowspan}(\mathbf{A})$ the column/row space of $\mathbf{A}$. Denote by $\mathbf{P}_{\mathbf{A}}$ the Euclidean projector onto $\text{colspan}(\mathbf{A})$. Denote $\mathbf{A} \succeq \mathbf{0}$ if $\mathbf{A}$ is a symmetric or Hermitian positive semidefinite matrix. For any two matrices $\mathbf{A}$ and $\mathbf{B}$ of the same dimensions, their matrix inner product is denoted by $\langle \mathbf{A}, \mathbf{B} \rangle = \text{trace}(\mathbf{A}^\top \mathbf{B}) = \sum_i \sum_j A_{i,j} B_{i,j}$, and their Hadamard/entrywise product is denoted by $\mathbf{A} \circ \mathbf{B}$ with entries $[\mathbf{A} \circ \mathbf{B}]_{i,j} = A_{i,j} B_{i,j}$. For any two matrices $\mathbf{A}$ and $\mathbf{B}$, $\text{vec}(\mathbf{A}, \mathbf{B})$ denotes a vector consisting of all entries in $\mathbf{A}$ and $\mathbf{B}$ in some fixed order. Denote by $\mathbf{J}_{n_1 \times n_2}$ (or $\mathbf{J}$ when the dimensions are clear in the context) the $n_1 \times n_2$ matrix with all entries equal to one. Denote by $\mathcal{O}(r)$ the set of $r \times r$ orthogonal matrices. Let $n_{\min} := \min\{n_1, n_2\}$ and $n_{\max} := \max\{n_1, n_2\}$. Finally, denote by $C_0, C_1, \dots$ and $C_v, C_c, \dots$ some fixed positive absolute constants.

1.3 Organization of the Paper

The rest of the paper is organized as follows: In Section 2, we first introduce the key assumptions regarding the parameterized factorization formally, and then introduce our main results on the low-rank estimation with any local minimizer, followed by corollaries for subspace-constrained and skew-symmetric matrix completions. Some results of numerical simulations are presented in Section 3 so as to illustrate our theoretical findings. We give a summary of our contributions and propose some open questions for future research in Section 5. All proofs can be found in Sections 4 with technical details deferred to the appendices.

2. Main Results

The major contribution of this paper is on deriving theoretical properties of parameterized matrix completion with nonconvex optimization (6). We first introduce key assumptions on the low-rank matrix $\mathbf{M}^*$, the parameterized factors $(\mathbf{X}(\boldsymbol{\theta}), \mathbf{Y}(\boldsymbol{\theta}))$, and the observation support Ω . Under these assumptions, we show in Section 2.2 that the geometric analysis in Chen and Li (2019) for PSD matrix completion can be extended to parameterized matrix completion. In particular, we give a uniform upper bound for the estimation error $\|\mathbf{X}(\hat{\boldsymbol{\xi}})\mathbf{Y}(\hat{\boldsymbol{\xi}})^\top - \mathbf{M}^*\|_F^2$ with any local minimizer $\hat{\boldsymbol{\xi}}$ of the nonconvex objective (6). Corollaries of our meta-theorem give conditions for consistent subspace-constrained and skew-symmetric matrix completions via nonconvex optimization.

2.1 Assumptions

Let us now introduce the key assumptions on the parameterized factors $(\mathbf{X}(\boldsymbol{\theta}), \mathbf{Y}(\boldsymbol{\theta}))$ mentioned earlier. We first assume that both factors are linear mappings in $\boldsymbol{\theta}$.

Assumption 1 (Linearity). Both $\mathbf{X}(\boldsymbol{\theta}) \in \mathbb{R}^{n_1 \times r}$ and $\mathbf{Y}(\boldsymbol{\theta}) \in \mathbb{R}^{n_2 \times r}$ are linear mappings in $\boldsymbol{\theta}$, i.e. $\mathbf{X}(t_1\boldsymbol{\theta}_1 + t_2\boldsymbol{\theta}_2) = t_1\mathbf{X}(\boldsymbol{\theta}_1) + t_2\mathbf{X}(\boldsymbol{\theta}_2)$ and $\mathbf{Y}(t_1\boldsymbol{\theta}_1 + t_2\boldsymbol{\theta}_2) = t_1\mathbf{Y}(\boldsymbol{\theta}_1) + t_2\mathbf{Y}(\boldsymbol{\theta}_2)$ for all $\boldsymbol{\theta}_1, \boldsymbol{\theta}_2 \in \mathbb{R}^d$ and $t_1, t_2 \in \mathbb{R}$.

The next assumption, referred to as the *Correlated Parametric Factorization*, plays an essential role in the geometric analysis of the nonconvex objective (6). We will verify this assumption for subspace-constrained factorization (4) and skew-symmetric factorization (5) in Section 4.3.1 and Section 4.4.1, correspondingly.

Assumption 2 (Correlated Parametric Factorization of $\mathbf{M}^*$). The rank- r matrix $\mathbf{M}^*$ and the parameterized factors $(\mathbf{X}(\boldsymbol{\theta}), \mathbf{Y}(\boldsymbol{\theta}))$ defined in Assumption 1 are said to satisfy the correlated parameterized factorization, if for any $\boldsymbol{\theta} \in \mathbb{R}^d$, there exists $\boldsymbol{\xi} \in \mathbb{R}^d$ (not necessarily unique), such that

$$\begin{cases} \mathbf{M}^* = \mathbf{X}(\boldsymbol{\xi})\mathbf{Y}(\boldsymbol{\xi})^\top, \\ \mathbf{X}(\boldsymbol{\xi})^\top \mathbf{X}(\boldsymbol{\xi}) = \mathbf{Y}(\boldsymbol{\xi})^\top \mathbf{Y}(\boldsymbol{\xi}), \\ \mathbf{X}(\boldsymbol{\theta})^\top \mathbf{X}(\boldsymbol{\xi}) + \mathbf{Y}(\boldsymbol{\theta})^\top \mathbf{Y}(\boldsymbol{\xi}) \succeq \mathbf{0}. \end{cases} \quad (9)$$

Let us now introduce models of observation patterns $\Omega \subseteq [n_1] \times [n_2]$. We assume that the entries are observed independently with certain probability p for the following two scenarios of rectangular and square matrices, respectively.

Model 1. For rectangular matrix completion, the index set Ω is assumed to follow the independent $\text{Ber}(p)$ model, i.e., each entry is sampled independently with probability p .

Model 2. For square matrix completion ($n_1 = n_2 := n$), the index set Ω is assumed to follow the off-diagonal symmetric independent $\text{Ber}(p)$ model, i.e., Ω is the support of symmetric off-diagonal entries that are sampled independently with probability p . No diagonal entries are included in Ω .

2.2 A Meta-Theorem

Our main theorem establishes an upper bound on the error of any local minimizer of the nonconvex objective (6), under the correlated parametric factorization assumption.

Theorem 1. Assume that $\mathbf{M}^*$ has rank r and its reduced singular value decomposition (SVD) is $\mathbf{M}^* = \mathbf{U}^* \mathbf{\Lambda} \mathbf{V}^{*\top}$, where $\mathbf{U}^* \in \mathbb{R}^{n_1 \times r}$, $\mathbf{V}^* \in \mathbb{R}^{n_2 \times r}$ and $\mathbf{\Lambda} = \text{diag}(\sigma_1, \sigma_2, \dots, \sigma_r)$ with $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_r > 0$.¹ The condition number of $\mathbf{\Lambda}$ is denoted by $\kappa = \sigma_1/\sigma_r$. Moreover, following Candès and Recht (2009), the incoherence parameter μ for $\mathbf{M}^*$ determined by its row and column spaces is defined as

$$\mu = \max \left\{ \frac{n_1}{r} \|\mathbf{U}^*\|_{2,\infty}^2, \frac{n_2}{r} \|\mathbf{V}^*\|_{2,\infty}^2 \right\}. \quad (10)$$

Suppose that $\mathbf{M}^*$, $\mathbf{X}(\boldsymbol{\theta})$ and $\mathbf{Y}(\boldsymbol{\theta})$ satisfy Assumptions 1 and 2, and that Ω , the support of observed entries, follows either Model 1 or 2. Moreover, let the sampling rate p and the tuning parameters α and λ in (2) satisfy the following inequalities

$$\begin{cases} p \geq C_1 \max \left\{ \frac{1}{n_{\min}} \mu r \log n_{\max}, \frac{n_{\max}}{n_{\min}^2} \mu^2 r^2 \kappa^2 \right\}, \\ C_2 \sqrt{n_{\max}/p} \leq \lambda \leq 10C_2 \sqrt{n_{\max}/p}, \\ C_2 \sqrt{\mu r \sigma_1 / n_{\min}} \leq \alpha \leq 10C_2 \sqrt{\mu r \sigma_1 / n_{\min}}, \end{cases} \quad (11)$$

where $n_{\max} := \max(n_1, n_2)$ and $n_{\min} := \min(n_1, n_2)$. Then, on an event E with probability $\mathbb{P}[E] \geq 1 - (n_1 + n_2)^{-3}$, any local minimizer $\hat{\boldsymbol{\xi}}$ of (6) satisfies

$$\|\mathbf{M}^* - \widehat{\mathbf{M}}\|_F^2 \leq \frac{C_3 r}{p^2} \psi^2,$$

where $\widehat{\mathbf{M}} := \mathbf{X}(\hat{\boldsymbol{\xi}}) \mathbf{Y}(\hat{\boldsymbol{\xi}})^\top$ and

$$\psi := \max_{\boldsymbol{\theta}_1, \boldsymbol{\theta}_2 \in \mathbb{R}^d} \|\mathbf{P}_{\mathbf{X}(\boldsymbol{\theta}_1)} \mathcal{P}_\Omega(\mathbf{N}) \mathbf{P}_{\mathbf{Y}(\boldsymbol{\theta}_2)}\|. \quad (12)$$

Here C_1, C_2, C_3 are fixed absolute constants. Recall that $\mathbf{P}_{\mathbf{A}}$ is the Euclidean projection matrix onto $\text{colspan}(\mathbf{A})$.

Remark 1. In particular, if there is no noise, i.e., $\mathbf{N} = \mathbf{0}$, then on the same event any local minimizer $\hat{\boldsymbol{\xi}}$ gives exact low-rank recovery $\widehat{\mathbf{M}} = \mathbf{M}^*$, that is, there is no spurious local minimum.

1. In this paper, (reduced) SVD of matrices is repeatedly used. Note that (reduced) SVD may not be unique due to possible multiplicity in singular values. In that case, we simply choose one (reduced) SVD and keep it fixed throughout the discussion.

2.3 Implication for PSD Matrix Completion

Besides applications in some specific parametric matrix completion problems, existing geometric analysis for nonconvex matrix completion for more standard cases like general rectangular or positive semidefinite (PSD) matrices, can be implied by our meta-theorem.

Consider, for example, PSD matrix completion where we have $n_1 = n_2 = n$, and the underlying rank- r matrix can be decomposed as $\mathbf{M}^\star = \mathbf{\Xi}_0 \mathbf{\Xi}_0^\top$ for some $\mathbf{\Xi}_0 \in \mathbb{R}^{n \times r}$. For any $\mathbf{\Theta} \in \mathbb{R}^{n \times r}$, denote $\boldsymbol{\theta} := \text{vec}(\mathbf{\Theta})$ and define the linear mappings $\mathbf{X}(\boldsymbol{\theta}) = \mathbf{Y}(\boldsymbol{\theta}) = \mathbf{\Theta}$. The parameterized factorization for rank- r PSD matrices becomes $\mathbf{X}(\boldsymbol{\theta})\mathbf{Y}(\boldsymbol{\theta})^\top = \mathbf{\Theta}\mathbf{\Theta}^\top$, and the corresponding parameterized factorization for the ground truth is $\mathbf{M}^\star = \mathbf{X}(\boldsymbol{\xi}_0)\mathbf{Y}(\boldsymbol{\xi}_0)^\top$ for $\boldsymbol{\xi}_0 := \text{vec}(\mathbf{\Xi}_0)$.

It is obvious that $\mathbf{X}(\boldsymbol{\theta}) = \mathbf{Y}(\boldsymbol{\theta}) = \mathbf{\Theta}$ are linear mappings in $\boldsymbol{\theta}$, and so Assumption 1 holds. Let us now see how to construct $\mathbf{\Xi}$ based on any given $\mathbf{\Theta}$, such that the requirements in Assumption 2 can hold simultaneously. The second equality in (9) holds since $\mathbf{X}(\boldsymbol{\xi}) = \mathbf{Y}(\boldsymbol{\xi}) = \mathbf{\Xi}$. To ensure the first equality, we let $\mathbf{\Xi} = \mathbf{\Xi}_0 \mathbf{R}$ for some underdetermined orthogonal matrix $\mathbf{R}$, which gives $\mathbf{\Xi}\mathbf{\Xi}^\top = \mathbf{\Xi}_0 \mathbf{\Xi}_0^\top = \mathbf{M}^\star$. The remaining question is how to choose $\mathbf{R}$ such that $\mathbf{\Theta}^\top \mathbf{\Xi}$ is PSD. This can be easily obtained by considering the SVD of $\mathbf{\Theta}^\top \mathbf{\Xi}_0$. It is interesting that the construction of such $\mathbf{R}$ plays an important role in the literature of PSD matrix completion. See, e.g., Chen and Wainwright (2015, Lemma 1).

The nonconvex objective (6) for PSD matrix completion takes the form

$$f_{\text{psd}}(\boldsymbol{\theta}) := \frac{1}{2p} \left\| \mathcal{P}_\Omega(\mathbf{\Theta}\mathbf{\Theta}^\top - \mathbf{M}) \right\|_F^2 + 2\lambda G_\alpha(\boldsymbol{\theta}), \quad (13)$$

which is exactly the objective studied in Ge et al. (2016, 2017) and Chen and Li (2019). Implied by Theorem 1, with high probability, there is no spurious local minimum of (13), provided that the tuning parameters are suitably chosen and the sampling rate satisfies $p \geq \frac{C_1}{n} \max \{ \mu r \log n, \mu^2 r^2 \kappa^2 \}$, which is exactly the same as the state-of-the-art result in Chen and Li (2019). Furthermore, consider the special noisy case in which the entries of noise matrix $\mathbf{N}$ are i.i.d. Gaussian random variables with mean 0 and variance σ^2 . Then $\|\mathcal{P}_\Omega(\mathbf{N})\|^2 = O((np + \log^2 n)\sigma^2)$ (see, e.g., Chen and Wainwright (2015, Lemma 11)). Theorem 1 implies the estimation error bound $\|\mathbf{M}^\star - \widehat{\mathbf{M}}\|_F^2 = O((\frac{nr}{p} + \frac{r \log^2 n}{p^2})\sigma^2)$, which matches the state-of-the-art results in the literature of noisy matrix completion. See, e.g., Keshavan et al. (2010b), Chen and Wainwright (2015) and Ma et al. (2020).

In the next two subsections, we apply Theorem 1 to subspace-constrained matrix completion (7) and skew-symmetric matrix completion (8), respectively.

2.4 Nonconvex Subspace-Constrained Matrix Completion

For subspace-constrained matrix completion, the parametric factors have the form (4), i.e. $\mathbf{X}(\boldsymbol{\theta}) = \tilde{\mathbf{U}}\boldsymbol{\Theta}_A$ and $\mathbf{Y}(\boldsymbol{\theta}) = \tilde{\mathbf{V}}\boldsymbol{\Theta}_B$, where $\boldsymbol{\theta} = \text{vec}(\boldsymbol{\Theta}_A, \boldsymbol{\Theta}_B)$. Recall that we always assume $\tilde{\mathbf{U}}^\top \tilde{\mathbf{U}} = \mathbf{I}_{s_1}$ and $\tilde{\mathbf{V}}^\top \tilde{\mathbf{V}} = \mathbf{I}_{s_2}$. Here $\mathbf{X}(\boldsymbol{\theta})$ and $\mathbf{Y}(\boldsymbol{\theta})$ are linear mappings, thereby satisfying Assumption 1. The verification of Assumption 2 is summarized as the following lemma, the proof of which is deferred to Section 4.3.1.

Lemma 1. *Let $\mathbf{M}^\star \in \mathbb{R}^{n_1 \times n_2}$ be a rank- r matrix whose column space and row space are constrained in $\text{colspan}(\tilde{\mathbf{U}})$ and $\text{colspan}(\tilde{\mathbf{V}})$, where $\tilde{\mathbf{U}}^\top \tilde{\mathbf{U}} = \mathbf{I}_{s_1}$ and $\tilde{\mathbf{V}}^\top \tilde{\mathbf{V}} = \mathbf{I}_{s_2}$. Then the parameterized factors $\mathbf{X}(\boldsymbol{\theta})$, $\mathbf{Y}(\boldsymbol{\theta})$ defined in (4) satisfy Assumption 2 with $\mathbf{M}^\star \in \mathbb{R}^{n_1 \times n_2}$.*

Theorem 1 implies the following corollary for nonconvex subspace-constrained matrix completion with objective (7).

Corollary 1. *Let $\mathbf{M}^* \in \mathbb{R}^{n_1 \times n_2}$ be a rank- r matrix. The incoherence parameter μ and condition number κ are defined in Theorem 1. Assume that the columns of $\tilde{\mathbf{U}} \in \mathbb{R}^{n_1 \times s_1}$ constitute an orthonormal basis of the column space constraint for $\mathbf{M}^*$, while the columns of $\tilde{\mathbf{V}} \in \mathbb{R}^{n_2 \times s_2}$ constitute an orthonormal basis of the row space constraint. The support of observation, Ω , is assumed to follow Model 1. The entries of the noise matrix $\mathbf{N}$ are i.i.d. centered sub-exponential random variables satisfying the Bernstein condition (Wainwright, 2019, (2.15)) with parameter b and variance ν^2 .*

If the sampling rate p and the tuning parameters α, λ satisfy (11). Then, uniformly on an event E_{subspace} with probability $\mathbb{P}[E_{\text{subspace}}] \geq 1 - 2(n_1 + n_2)^{-3}$, any local minimizer $(\hat{\mathbf{\Xi}}_A, \hat{\mathbf{\Xi}}_B)$ of $f_{\text{subspace}}(\mathbf{\Theta}_A, \mathbf{\Theta}_B)$ in (7) satisfies

$$\|\tilde{\mathbf{U}}\hat{\mathbf{\Xi}}_A\hat{\mathbf{\Xi}}_B^\top\tilde{\mathbf{V}}^\top - \mathbf{M}^*\|_F^2 \leq \frac{C_4 r}{p^2} \left(\nu^2 p(s_1 + s_2) \log(n_1 + n_2) + b^2 \frac{\mu_{\tilde{\mathbf{U}}} \mu_{\tilde{\mathbf{V}}} s_1 s_2}{n_1 n_2} \log^2(n_1 + n_2) \right). \quad (14)$$

Here $\mu_{\tilde{\mathbf{U}}} = \frac{n_1}{s_1} \|\tilde{\mathbf{U}}\|_{2,\infty}^2$, $\mu_{\tilde{\mathbf{V}}} = \frac{n_2}{s_2} \|\tilde{\mathbf{V}}\|_{2,\infty}^2$, and C_4 is some fixed positive absolute constant.

To the best of our knowledge, existing investigations in the literature on matrix completion with subspace constraints are majorly focused on the noiseless case (Yi et al., 2013; Xu et al., 2013; Chen, 2015; Jain and Dhillon, 2013; Eftekhari et al., 2018), while error rates for noisy recovery are little-studied. Consider the scenario in which the noise matrix $\mathbf{N}$ consists of i.i.d. $\mathcal{N}(0, \sigma^2)$ entries. This gives $b = \sigma$ and variance $\nu^2 = \sigma^2$. For simplicity of discussion, assume $s_1 = s_2 := s$, $n_1 = n_2 := n$, $\mu = O(1)$, $\mu_{\tilde{\mathbf{U}}} = O(1)$, $\mu_{\tilde{\mathbf{V}}} = O(1)$ and $\kappa = O(1)$. Then Corollary 1 implies that as long as $p \gtrsim \frac{1}{n} \max\{r \log n, r^2\}$, with high probability, $\|\tilde{\mathbf{U}}\hat{\mathbf{\Xi}}_A\hat{\mathbf{\Xi}}_B^\top\tilde{\mathbf{V}}^\top - \mathbf{M}^*\|_F^2 \lesssim \sigma^2 s r (\log n)/p$ holds uniformly for any local minimizer. As explained in the previous subsection, the error rates for standard matrix completion are $O(\sigma^2 n r / p)$. Therefore, Corollary 1 indicates that the estimation error can be significantly reduced by exploring and employing subspace constraints.

In the noiseless case, we should admit that the sufficient condition on the sampling rate for exact matrix completion, $p \gtrsim \frac{1}{n} \max\{r \log n, r^2\}$, may be suboptimal. In fact, fewer samples are sufficient for exact low-rank recovery for convex optimization methods (Yi et al., 2013; Xu et al., 2013; Chen, 2015) as well as alternating minimization (Jain and Dhillon, 2013). Given our aim is to establish a unified methodological and theoretical framework for parameterized matrix completion via nonconvex optimization, improving sample size conditions for special cases is beyond the scope of the paper, but we are interested in studying this issue in future work.

2.5 Nonconvex Skew-Symmetric Matrix Completion

For skew-symmetric matrix completion, the parameterized factors are defined in (5), which are $\mathbf{X}(\boldsymbol{\theta}) = [\mathbf{\Theta}_A, -\mathbf{\Theta}_B]$ and $\mathbf{Y}(\boldsymbol{\theta}) = [\mathbf{\Theta}_B, \mathbf{\Theta}_A]$ with $\boldsymbol{\theta} = \text{vec}(\mathbf{\Theta}_A, \mathbf{\Theta}_B)$. The factors $\mathbf{X}(\boldsymbol{\theta})$ and $\mathbf{Y}(\boldsymbol{\theta})$ are linear mappings and thereby satisfying Assumption 1. Assumption 2 is verified through the following result with the proof deferred to Section 4.4.1.

Lemma 2. *Let $\mathbf{M}^*$ be a rank- r skew-symmetric matrix. Then, the parameterization $(\mathbf{X}(\boldsymbol{\theta}), \mathbf{Y}(\boldsymbol{\theta}))$ defined in (5) and $\mathbf{M}^*$ satisfy Assumption 2.*

Theorem 1 implies the following noisy low-rank recovery result for skew-symmetric matrix completion via nonconvex objective (8).

Corollary 2. *Let $\mathbf{M}^* \in \mathbb{R}^{n \times n}$ be a rank- r skew-symmetric matrix. The incoherence parameter μ and the condition number κ are defined in Theorem 1. The support of the observed entries Ω is assumed to follow from Model 2. The skew-symmetric noise matrix $\mathbf{N}$ consists of i.i.d. upper triangular entries, which are centered sub-exponential random variables satisfying the Bernstein condition with parameter b and variance ν^2 . Suppose that the sampling rate p and the tuning parameters α and λ satisfy (11). Then, uniformly on an event E_{skew} with probability $\mathbb{P}[E_{skew}] \geq 1 - 2n^{-3}$, any local minimizer $(\hat{\mathbf{E}}_A, \hat{\mathbf{E}}_B)$ of $f_{skew}(\boldsymbol{\Theta}_A, \boldsymbol{\Theta}_B)$ defined in (8) satisfies*

$$\|\hat{\mathbf{E}}_A \hat{\mathbf{E}}_B^\top - \hat{\mathbf{E}}_B \hat{\mathbf{E}}_A^\top - \mathbf{M}^*\|_F^2 \leq \frac{C_5 r}{p^2} (\nu^2 p n \log n + b^2 \log^2 n).$$

Where C_5 is a fixed positive absolute constant.

As with the discussion in Section 2.4, if the upper triangular part of noise matrix $\mathbf{N}$ consists of i.i.d. Gaussian random variables with mean 0 and variance σ^2 , and the sampling rate satisfies $p \gtrsim \frac{1}{n} \max\{r \log n, r^2\}$, then the estimation error satisfies $\|\hat{\mathbf{E}}_A \hat{\mathbf{E}}_B^\top - \hat{\mathbf{E}}_B \hat{\mathbf{E}}_A^\top - \mathbf{M}^*\|_F^2 = O(\sigma^2 n r (\log n)/p)$, which is comparable to the aforementioned error rate $O(\sigma^2 n r/p)$ for PSD matrix completion up to a logarithmic factor.

3. Experiments

In this section, some numerical results are shown to demonstrate the performance of the proposed nonconvex optimization for subspace-constrained matrix completion (7) and skew-symmetric matrix completion (8).

In all simulations, the sampling rate p is replaced with the empirical analog $\hat{p} := \frac{|\Omega|}{n_1 n_2}$. The hyperparameters are set as $\lambda = 100 \sqrt{(n_1 + n_2)/\hat{p}}$ and $\alpha = 100$. In both examples, the nonconvex objective (6) is minimized by gradient descent, with initialization for $\boldsymbol{\theta}$ consisting with i.i.d. standard normal entries. At each step of the gradient descent, the step size is selected through line search. To be specific, at each update of $\boldsymbol{\theta}$, the step size is set to be $\max\{2^{-k}, 10^{-10}\}$ for $k := \min\{t \mid t = 0, 1, 2, 3, \dots, \tilde{f}(\boldsymbol{\theta} - 2^{-t} \nabla \tilde{f}(\boldsymbol{\theta})) \leq \tilde{f}(\boldsymbol{\theta})\}$. The gradient descent iteration is terminated either after 500 iterations or as soon as the update on $\boldsymbol{\theta}$ satisfied $\|\nabla \tilde{f}(\boldsymbol{\theta})\|_2^2 \leq 10^{-10}$.

3.1 Subspace-Constrained Matrix Completion

For simulation of nonconvex optimization for subspace-constrained matrix completion (7), the dimensions of $\mathbf{M}^*$ are set to be $n_1 = n_2 = 500$ and its rank is set as $r = 2$. We also set the dimensions of the prior column and row subspaces as $s_1 = s_2 := s \leq 40$. In preparation for the construction of the column and row subspace constraints with various dimensions, we generate $[\mathbf{u}_1, \dots, \mathbf{u}_{40}]$ and $[\mathbf{v}_1, \dots, \mathbf{v}_{40}]$ according to the Haar measure on the manifold of 500×40 orthonormal basis matrices. The matrix to recover is fixed as $\mathbf{M}^* = \mathbf{u}_1 \mathbf{v}_1^\top + \mathbf{u}_2 \mathbf{v}_2^\top$, which gives $\|\mathbf{M}^*\|_F^2 = 2$. In the case of noisy matrix completion, the noise matrix $\mathbf{N}$ consists of i.i.d. $\mathcal{N}(0, \sigma^2)$ entries with $\sigma^2 = 1/500^2$, which gives $\mathbb{E}\|\mathbf{N}\|_F^2 = 1$, and hence the SNR is $\|\mathbf{M}^*\|_F^2 / \mathbb{E}\|\mathbf{N}\|_F^2 = 2$. In the noiseless case we set $\mathbf{N} = \mathbf{0}$.

In the noisy case, to implement gradient descent for solving the optimization (7), we consider different scenarios in terms of sampling rates and dimensions of subspace constraints: $p = 0.005, 0.010, 0.015, \dots, 0.1$ and $s = 10, 20, 30, 40$. For each fixed pair of (p, s) , the sampling support is generated from Model 1, and the subspace constraints are set by $\tilde{\mathbf{U}} = [\mathbf{u}_1, \dots, \mathbf{u}_s]$ and $\tilde{\mathbf{V}} = [\mathbf{v}_1, \dots, \mathbf{v}_s]$. Then, gradient descent is implemented to solve (7) with the input $\mathcal{P}_\Omega(\mathbf{M}^* + \mathbf{N})$. Logarithmic relative estimation errors $\log_{10} \frac{\|\widehat{\mathbf{M}} - \mathbf{M}^*\|_F^2}{\|\mathbf{M}^*\|_F^2}$ are averaged over 10 independent generations of the support of observations Ω and the noise $\mathbf{N}$, which are shown in Figure 1a. The results indicate that estimation errors are reduced when the the constraining subspaces have lower dimensions. This dependency illustrates our theoretical result Theorem (14).

Our experiments for the noiseless case are conducted in a similar manner, but here the sampling rate is chosen as $p = 0.0001, 0.0002, \dots, 0.0020$. Rather than recording the relative errors for low-rank recovery, each experiment is viewed to be “successful” if and only if $\frac{\|\widehat{\mathbf{M}} - \mathbf{M}^*\|_F}{\|\mathbf{M}^*\|_F} \leq 10^{-3}$, and average rates of success are plotted in Figure 1b. The results clearly show that with more restrictive subspace constraints, the required sample size for exact matrix completion is smaller. We admit that this phenomenon is not explained by Theorem 14, in which the sampling size requirement is implied by the meta-theorem and hence not sufficiently adaptive to subspace-constrained models. As noted earlier, we intend to improve the sample size condition results in future work.

3.2 Skew-Symmetric Matrix Completion


For skew-symmetric matrix completion, we can either seek to recover the low-rank matrix by exploiting the skew-symmetric structures and solving the nonconvex optimization (8), or ignoring this structure and directly solving the standard matrix completion nonconvex objective (2). We are interested in understanding whether there is any empirical advantage to exploit the skew-symmetric structures.

For simplicity, our experiments focused on the noiseless case. The matrix size is fixed as $n = 500$ and the rank is chosen to be $r = 4, 10, 20$. For each r , $[\mathbf{u}_1, \dots, \mathbf{u}_{r/2}, \mathbf{v}_1, \dots, \mathbf{v}_{r/2}]$ is generated according to the Haar measure on the manifold of $500 \times r$ orthonormal basis matrices. The low-rank matrix is then constructed as $\mathbf{M}^* = \mathbf{u}_1 \mathbf{v}_1^\top - \mathbf{v}_1 \mathbf{u}_1^\top + \dots + \mathbf{u}_{r/2} \mathbf{v}_{r/2}^\top - \mathbf{v}_{r/2} \mathbf{u}_{r/2}^\top$. The sampling rate was fixed at $p = 0.01, 0.02, \dots, 0.20$. For each fixed pair (r, p) , 10 independent copies of $\Omega \in [500] \times [500]$ are generated from Model 2. Figure 2 plots the relative estimation errors as well as the corresponding medians in logarithmic scale by implementing (8) and (2) with gradient descent, respectively. The comparison indicates that when the rank r is not too small, exploiting skew-symmetric structures in nonconvex optimization is helpful in reducing the required sample size for low-rank recovery.


4. Proofs

In this section, we give a proof of Theorem 1, followed by proofs of the two corollaries on subspace-constrained and skew-symmetric matrix completion by verifying the assumption of correlated parametric factorization in respective settings.

In the proof of Theorem 1, for any local minimizer $\hat{\boldsymbol{\xi}}$ of (6), we aim to control the estimation error $\|\mathbf{M}^* - \mathbf{X}(\hat{\boldsymbol{\xi}})\mathbf{Y}(\hat{\boldsymbol{\xi}})^\top\|_F^2$. Given $\tilde{f}(\boldsymbol{\theta})$ is smooth, it is natural to study the



(a) Logarithm of relative estimation error $\log_{10} \frac{\|\widehat{\mathbf{M}} - \mathbf{M}^*\|_F^2}{\|\mathbf{M}^*\|_F^2}$ of nonconvex subspace constrained matrix completion. Each dot in the plot represents one trial of the numerical experiment, and the curves represent the mean of 10 independent trials for given s .



(b) Rates of success for noiseless nonconvex subspace constrained matrix completion. Here we set the dimension of ground truth $\mathbf{M}^* \in \mathbb{R}^{n_1 \times n_2}$ as $n_1 = n_2 = 500$, rank of $\mathbf{M}^*$ as $r = 2$, dimension of the column/row subspace constraint as $s_1 = s_2 = s$.

Figure 1: Numerical results for nonconvex subspace-constrained matrix completion in both noisy and noiseless cases.

first and second order optimality conditions, i.e., $\nabla \tilde{f}(\boldsymbol{\theta}) = \mathbf{0}$ and $\nabla^2 \tilde{f}(\boldsymbol{\theta}) \succeq \mathbf{0}_{d \times d}$. The strategy on employing these two conditions in the control of estimation error is key. To this end, we start with introducing an auxiliary function that connects the optimality conditions and the estimation error.




Figure 2: Logarithm of relative estimation error $\log_{10}(\|\widehat{\mathbf{M}} - \mathbf{M}^*\|_F^2 / \|\mathbf{M}^*\|_F^2)$. In this plot, skew-symmetric matrix completion and regular rectangular matrix completion are used to complete noiseless skew-symmetric matrices with rank $r = 4, 10, 20$. Here we set the dimension of ground truth $\mathbf{M}^* \in \mathbb{R}^{n \times n}$ as $n = 500$. Each dot in the plot represents one trial of the numerical experiment, and the curves represent the median of 10 independent trials.

4.1 An Auxiliary Function

We first recall a key component in geometric analysis for nonconvex matrix completion: the auxiliary function associated to the gradient and Hessian of $f(\mathbf{X}, \mathbf{Y})$ defined in (2), which has been used in Ge et al. (2017); Jin et al. (2017); Chen and Li (2019). For any $\mathbf{X} \in \mathbb{R}^{n_1 \times r}$, $\mathbf{Y} \in \mathbb{R}^{n_2 \times r}$, $\mathbf{D}_\mathbf{X} \in \mathbb{R}^{n_1 \times r}$ and $\mathbf{D}_\mathbf{Y} \in \mathbb{R}^{n_2 \times r}$, the auxiliary function associated to $f(\mathbf{X}, \mathbf{Y})$ is defined as

$$K_f(\mathbf{X}, \mathbf{Y}; \mathbf{D}_\mathbf{X}, \mathbf{D}_\mathbf{Y}) := \text{vec}(\mathbf{D}_\mathbf{X}, \mathbf{D}_\mathbf{Y})^\top \nabla^2 f(\mathbf{X}, \mathbf{Y}) \text{vec}(\mathbf{D}_\mathbf{X}, \mathbf{D}_\mathbf{Y}) - 4 \langle \nabla f(\mathbf{X}, \mathbf{Y}), \text{vec}(\mathbf{D}_\mathbf{X}, \mathbf{D}_\mathbf{Y}) \rangle. \quad (15)$$

Because any local minimizer $(\widehat{\mathbf{X}}, \widehat{\mathbf{Y}})$ of $f(\mathbf{X}, \mathbf{Y})$ satisfies the first and second order conditions, i.e., $\nabla f(\widehat{\mathbf{X}}, \widehat{\mathbf{Y}}) = \mathbf{0}$ and $\nabla^2 f(\widehat{\mathbf{X}}, \widehat{\mathbf{Y}}) \succeq \mathbf{0}$, we have

$$K_f(\widehat{\mathbf{X}}, \widehat{\mathbf{Y}}; \mathbf{D}_\mathbf{X}, \mathbf{D}_\mathbf{Y}) \geq 0, \quad \forall \mathbf{D}_\mathbf{X} \in \mathbb{R}^{n_1 \times r}, \mathbf{D}_\mathbf{Y} \in \mathbb{R}^{n_2 \times r}.$$

The ratio of coefficients $1 : (-4)$ for the two terms on the right hand side of (15) turns out to be essential for the nonconvex geometric analysis; see Jin et al. (2017); Ge et al. (2017); Chen and Li (2019).

Following this idea, we define an analogous auxiliary function associated to $\tilde{f}(\boldsymbol{\theta})$ in general parametric matrix completion (6). For any $\boldsymbol{\theta}, \boldsymbol{\delta}_\boldsymbol{\theta} \in \mathbb{R}^d$, define the auxiliary function associated with $\tilde{f}$ as

$$K_{\tilde{f}}(\boldsymbol{\theta}; \boldsymbol{\delta}_\boldsymbol{\theta}) := \boldsymbol{\delta}_\boldsymbol{\theta}^\top \nabla^2 \tilde{f}(\boldsymbol{\theta}) \boldsymbol{\delta}_\boldsymbol{\theta} - 4 \boldsymbol{\delta}_\boldsymbol{\theta}^\top \nabla \tilde{f}(\boldsymbol{\theta}). \quad (16)$$

Again, any minimizer $\hat{\boldsymbol{\xi}}$ of $\tilde{f}$ satisfies $\nabla \tilde{f}(\hat{\boldsymbol{\xi}}) = \mathbf{0}$ and $\nabla^2 \tilde{f}(\hat{\boldsymbol{\xi}}) \succeq \mathbf{0}$, which gives

$$K_{\tilde{f}}(\hat{\boldsymbol{\xi}}; \boldsymbol{\delta}_\theta) \geq 0, \quad \forall \boldsymbol{\delta}_\theta \in \mathbb{R}^d.$$

Control of estimation error $\|\mathbf{M}^* - \mathbf{X}(\hat{\boldsymbol{\xi}})\mathbf{Y}(\hat{\boldsymbol{\xi}})^\top\|_F^2$ relies on controlling $K_{\tilde{f}}(\boldsymbol{\theta}; \boldsymbol{\delta}_\theta)$, which in turn depends crucially on a careful choice of $\boldsymbol{\delta}_\theta$. To this end, Assumption 2 plays the key role and serves as the guiding principle.

We first give a lemma that characterizes the relationship between $K_f(\mathbf{X}, \mathbf{Y}; \mathbf{D}_\mathbf{X}, \mathbf{D}_\mathbf{Y})$ and $K_{\tilde{f}}(\boldsymbol{\theta}; \boldsymbol{\delta}_\theta)$.

Lemma 3. *For any $\boldsymbol{\theta}, \boldsymbol{\delta}_\theta \in \mathbb{R}^p$, it holds that*

$$K_{\tilde{f}}(\boldsymbol{\theta}; \boldsymbol{\delta}_\theta) = K_f(\mathbf{X}(\boldsymbol{\theta}), \mathbf{Y}(\boldsymbol{\theta}); \mathbf{X}(\boldsymbol{\delta}_\theta), \mathbf{Y}(\boldsymbol{\delta}_\theta)). \quad (17)$$

Proof Assumption 1 implies

$$\tilde{f}(\boldsymbol{\theta} + \boldsymbol{\delta}_\theta) = f(\mathbf{X}(\boldsymbol{\theta} + \boldsymbol{\delta}_\theta), \mathbf{Y}(\boldsymbol{\theta} + \boldsymbol{\delta}_\theta)) = f(\mathbf{X}(\boldsymbol{\theta}) + \mathbf{X}(\boldsymbol{\delta}_\theta), \mathbf{Y}(\boldsymbol{\theta}) + \mathbf{Y}(\boldsymbol{\delta}_\theta)). \quad (18)$$

Due to the linearity and homogeneity of $\mathbf{X}(\boldsymbol{\theta})$ and $\mathbf{Y}(\boldsymbol{\theta})$, considering the Taylor expansions of both sides in (18) at $\boldsymbol{\theta}$, we get

$$\boldsymbol{\delta}_\theta^\top \nabla \tilde{f}(\boldsymbol{\theta}) = \text{vec}(\mathbf{X}(\boldsymbol{\delta}_\theta), \mathbf{Y}(\boldsymbol{\delta}_\theta))^\top \nabla f(\mathbf{X}(\boldsymbol{\theta}), \mathbf{Y}(\boldsymbol{\theta})) \quad (19)$$

and

$$\boldsymbol{\delta}_\theta^\top \nabla^2 \tilde{f}(\boldsymbol{\theta}) \boldsymbol{\delta}_\theta = \text{vec}(\mathbf{X}(\boldsymbol{\delta}_\theta), \mathbf{Y}(\boldsymbol{\delta}_\theta))^\top \nabla^2 f(\mathbf{X}(\boldsymbol{\theta}), \mathbf{Y}(\boldsymbol{\theta})) \text{vec}(\mathbf{X}(\boldsymbol{\delta}_\theta), \mathbf{Y}(\boldsymbol{\delta}_\theta)) \quad (20)$$

By combining (15), (16), (19) and (20), the equality (17) is obtained. $\blacksquare$

Our choice of $\boldsymbol{\delta}_\theta$ is inspired by the choices of $\mathbf{D}_\mathbf{X}$ and $\mathbf{D}_\mathbf{Y}$ for K_f in Jin et al. (2017); Ge et al. (2017); Chen and Li (2019). For any $\boldsymbol{\theta} \in \mathbb{R}^d$, let $\boldsymbol{\xi} \in \mathbb{R}^d$ be selected (not necessarily uniquely) such that Assumption 2 is satisfied. Then let

$$\boldsymbol{\delta}_\theta := \boldsymbol{\theta} - \boldsymbol{\xi}.$$

In the rest of this subsection, we will explain why our choice of $\boldsymbol{\delta}_\theta$ is helpful in controlling the auxiliary function. For simplicity, we introduce the following abbreviations and notations

$$\begin{aligned} \mathbf{X} &= \mathbf{X}(\boldsymbol{\theta}) \in \mathbb{R}^{n_1 \times r}, \quad \mathbf{U} = \mathbf{X}(\boldsymbol{\xi}) \in \mathbb{R}^{n_1 \times r}, \quad \boldsymbol{\Delta}_\mathbf{X} = \mathbf{X}(\boldsymbol{\delta}_\theta) = \mathbf{X} - \mathbf{U} \in \mathbb{R}^{n_1 \times r}, \\ \mathbf{Y} &= \mathbf{Y}(\boldsymbol{\theta}) \in \mathbb{R}^{n_2 \times r}, \quad \mathbf{V} = \mathbf{Y}(\boldsymbol{\xi}) \in \mathbb{R}^{n_2 \times r}, \quad \boldsymbol{\Delta}_\mathbf{Y} = \mathbf{Y}(\boldsymbol{\delta}_\theta) = \mathbf{Y} - \mathbf{V} \in \mathbb{R}^{n_2 \times r}. \end{aligned} \quad (21)$$

Throughout our proof, $\mathbf{X}, \mathbf{U}, \boldsymbol{\Delta}_\mathbf{X}, \mathbf{Y}, \mathbf{V}$, and $\boldsymbol{\Delta}_\mathbf{Y}$ will refer to the matrices defined in (21) if not specified otherwise. Now Lemma 3 gives

$$K_{\tilde{f}}(\boldsymbol{\theta}; \boldsymbol{\delta}_\theta) = K_f(\mathbf{X}, \mathbf{Y}; \boldsymbol{\Delta}_\mathbf{X}, \boldsymbol{\Delta}_\mathbf{Y}).$$

Note that Assumption 2 implies

$$\mathbf{M}^* = \mathbf{U}\mathbf{V}^\top, \quad \mathbf{U}^\top \mathbf{U} = \mathbf{V}^\top \mathbf{V}, \quad \text{and} \quad \mathbf{X}^\top \mathbf{U} + \mathbf{Y}^\top \mathbf{V} \succeq \mathbf{0}. \quad (22)$$

The following lemma gives an explicit upper bound for $K_{\tilde{f}}(\boldsymbol{\theta}; \boldsymbol{\delta}_\theta)$.

Lemma 4. For any $\boldsymbol{\theta} \in \mathbb{R}^d$, with $\boldsymbol{\delta}_{\boldsymbol{\theta}} = \boldsymbol{\theta} - \boldsymbol{\xi}$ where $\boldsymbol{\xi}$ satisfies the conditions in (9), and $\mathbf{X}, \mathbf{U}, \boldsymbol{\Delta}_{\mathbf{X}}, \mathbf{Y}, \mathbf{V}, \boldsymbol{\Delta}_{\mathbf{Y}}$ defined as in (21), denote

$$\mathbf{Z} = \begin{bmatrix} \mathbf{X} \\ \mathbf{Y} \end{bmatrix}, \quad \mathbf{W} = \begin{bmatrix} \mathbf{U} \\ \mathbf{V} \end{bmatrix}, \quad \text{and} \quad \boldsymbol{\Delta}_{\mathbf{Z}} = \begin{bmatrix} \boldsymbol{\Delta}_{\mathbf{X}} \\ \boldsymbol{\Delta}_{\mathbf{Y}} \end{bmatrix} = \mathbf{Z} - \mathbf{W}.$$

Then the auxiliary function $K_{\tilde{f}}(\boldsymbol{\theta}; \boldsymbol{\delta}_{\boldsymbol{\theta}})$ defined in (16) can be upper bounded by

$$K_{\tilde{f}}(\boldsymbol{\theta}; \boldsymbol{\delta}_{\boldsymbol{\theta}}) \leq K_1 + K_2 + K_3 + K_4, \quad (23)$$

where

$$\begin{aligned} K_1 &:= \frac{1}{4} \left(\left\| \boldsymbol{\Delta}_{\mathbf{Z}} \boldsymbol{\Delta}_{\mathbf{Z}}^{\top} \right\|_F^2 - 3 \left\| \mathbf{Z} \mathbf{Z}^{\top} - \mathbf{W} \mathbf{W}^{\top} \right\|_F^2 \right), \\ K_2 &:= \left(\frac{1}{p} \left\| \mathcal{P}_{\Omega} \left(\boldsymbol{\Delta}_{\mathbf{X}} \boldsymbol{\Delta}_{\mathbf{Y}}^{\top} \right) \right\|_F^2 - \left\| \boldsymbol{\Delta}_{\mathbf{X}} \boldsymbol{\Delta}_{\mathbf{Y}}^{\top} \right\|_F^2 \right) \\ &\quad - \left(\frac{3}{p} \left\| \mathcal{P}_{\Omega} \left(\mathbf{X} \mathbf{Y}^{\top} - \mathbf{U} \mathbf{V}^{\top} \right) \right\|_F^2 - 3 \left\| \mathbf{X} \mathbf{Y}^{\top} - \mathbf{U} \mathbf{V}^{\top} \right\|_F^2 \right), \\ K_3 &:= \lambda \left[\text{vec}(\boldsymbol{\Delta}_{\mathbf{X}})^{\top} \nabla^2 G_{\alpha}(\mathbf{X}) \text{vec}(\boldsymbol{\Delta}_{\mathbf{X}}) - 4 \langle \nabla G_{\alpha}(\mathbf{X}), \boldsymbol{\Delta}_{\mathbf{X}} \rangle \right] \\ &\quad + \lambda \left[\text{vec}(\boldsymbol{\Delta}_{\mathbf{Y}})^{\top} \nabla^2 G_{\alpha}(\mathbf{Y}) \text{vec}(\boldsymbol{\Delta}_{\mathbf{Y}}) - 4 \langle \nabla G_{\alpha}(\mathbf{Y}), \boldsymbol{\Delta}_{\mathbf{Y}} \rangle \right], \\ K_4 &:= \frac{6}{p} \left\langle \boldsymbol{\Delta}_{\mathbf{X}} \boldsymbol{\Delta}_{\mathbf{Y}}^{\top}, \mathcal{P}_{\Omega}(\mathbf{N}) \right\rangle + \frac{4}{p} \left\langle \mathbf{U} \boldsymbol{\Delta}_{\mathbf{Y}}^{\top} + \boldsymbol{\Delta}_{\mathbf{X}} \mathbf{V}^{\top}, \mathcal{P}_{\Omega}(\mathbf{N}) \right\rangle. \end{aligned} \quad (24)$$

The proof of Lemma 4 follows closely that of Ge et al. (2017, Lemma 16). We give a complete proof in Appendix A.2 for being self-contained.

4.2 Proof of Theorem 1

Given the decomposition in Lemma 4, the proof of Theorem 1 depends on the following two auxiliary lemmas. Note that all results are established on the event E defined in Lemma 12. Their proofs are provided in Section A.3 and Section A.4, respectively.

Lemma 5. Let the quantities K_2 and K_3 be defined as in Lemma 4. Then the sum $K_2 + K_3$ is bounded as

$$K_2 + K_3 \leq 0.001 \left(\left\| \boldsymbol{\Delta}_{\mathbf{X}} \mathbf{V}^{\top} \right\|_F^2 + \left\| \mathbf{U} \boldsymbol{\Delta}_{\mathbf{Y}}^{\top} \right\|_F^2 \right).$$

Lemma 6. Let the quantity K_4 be defined as in Lemma 4. If $K_1 + K_2 + K_3 + K_4 \geq 0$, the following upper bounds hold

1. $\left\| \boldsymbol{\Delta}_{\mathbf{Z}} \boldsymbol{\Delta}_{\mathbf{Z}}^{\top} \right\|_F = \left\| \boldsymbol{\Delta}_{\mathbf{Z}}^{\top} \boldsymbol{\Delta}_{\mathbf{Z}} \right\|_F \leq 3\sqrt{|K_4|}, \quad \left\| \boldsymbol{\Delta}_{\mathbf{Z}}^{\top} \mathbf{W} \right\|_F \leq 2\sqrt{|K_4|}.$
2. $\left\| \boldsymbol{\Delta}_{\mathbf{X}} \mathbf{V}^{\top} \right\|_F^2 + \left\| \mathbf{U} \boldsymbol{\Delta}_{\mathbf{Y}}^{\top} \right\|_F^2 \leq 13|K_4|.$

By the definition of $\mathbf{X}, \mathbf{Y}, \mathbf{U}$ and $\mathbf{V}$ in (21), we note that the difference between any local minimizer $\widehat{\mathbf{M}}$ and the matrix $\mathbf{M}^{\star}$ can be written as

$$\left\| \widehat{\mathbf{M}} - \mathbf{M}^{\star} \right\|_F^2 = \left\| \mathbf{X} \mathbf{Y}^{\top} - \mathbf{U} \mathbf{V}^{\top} \right\|_F^2 \quad (25)$$

Next, by the definition of $\mathbf{Z}$ and $\mathbf{W}$ in Lemma 4, we expand $\mathbf{Z}\mathbf{Z}^\top - \mathbf{W}\mathbf{W}^\top$ to obtain the inequality

$$\|\mathbf{X}\mathbf{Y}^\top - \mathbf{U}\mathbf{V}^\top\|_F^2 \leq \|\mathbf{Z}\mathbf{Z}^\top - \mathbf{W}\mathbf{W}^\top\|_F^2. \quad (26)$$

Since $\widehat{\mathbf{M}} = \mathbf{X}(\hat{\boldsymbol{\xi}})\mathbf{Y}(\hat{\boldsymbol{\xi}})^\top$ is a local minimizer, $\hat{\boldsymbol{\xi}}$ is a local minimizer of $\tilde{f}$ and so it holds that $K_{\tilde{f}}(\hat{\boldsymbol{\xi}}; \delta_{\hat{\boldsymbol{\xi}}}) \geq 0$. By Lemma 4, it holds that

$$K_1 + K_2 + K_3 + K_4 \geq 0. \quad (27)$$

By Lemma 5 and the second claim of Lemma 6, $K_2 + K_3 + K_4 \leq 2|K_4|$. Together with the definition of K_1 in (24), (27) and the first claim of Lemma 6, this leads to

$$\frac{3}{4}\|\mathbf{Z}\mathbf{Z}^\top - \mathbf{W}\mathbf{W}^\top\|_F^2 \leq \frac{1}{4}\|\Delta_{\mathbf{Z}}\Delta_{\mathbf{Z}}^\top\|_F^2 + K_2 + K_3 + K_4 \leq \frac{17}{4}|K_4|. \quad (28)$$

If we can show that

$$|K_4| \leq \frac{10^4 r}{p^2} \psi^2, \quad (29)$$

then the desired conclusion of Theorem 1 is a direct consequence of (25), (26), (28), and (29). It remains only to prove (29), to which we now turn.

Proof of (29) Recall the fact that $\Delta_{\mathbf{Z}} = [\Delta_{\mathbf{X}}^\top \ \Delta_{\mathbf{Y}}^\top]^\top$, then $\|\Delta_{\mathbf{X}}\Delta_{\mathbf{Y}}^\top\|_F \leq \|\Delta_{\mathbf{Z}}\Delta_{\mathbf{Z}}^\top\|_F$. Therefore, by Lemma 6, and the definition of K_4 in (24), we have

$$\begin{aligned} |K_4| &\leq \frac{6}{p}\|\Delta_{\mathbf{X}}\Delta_{\mathbf{Y}}^\top\|_F\|\mathbf{P}_{\Delta_{\mathbf{X}}}\mathcal{P}_\Omega(\mathbf{N})\mathbf{P}_{\Delta_{\mathbf{Y}}}\|_F + \frac{4}{p}\|\mathbf{U}\Delta_{\mathbf{Y}}^\top\|_F\|\mathbf{P}_{\mathbf{U}}\mathcal{P}_\Omega(\mathbf{N})\mathbf{P}_{\Delta_{\mathbf{Y}}}\|_F \\ &\quad + \frac{4}{p}\|\Delta_{\mathbf{X}}\mathbf{V}^\top\|_F\|\mathbf{P}_{\Delta_{\mathbf{X}}}\mathcal{P}_\Omega(\mathbf{N})\mathbf{P}_{\mathbf{V}}\|_F \\ &\leq \frac{100\sqrt{|K_4|}}{p} \max\{\|\mathbf{P}_{\Delta_{\mathbf{X}}}\mathcal{P}_\Omega(\mathbf{N})\mathbf{P}_{\Delta_{\mathbf{Y}}}\|_F, \|\mathbf{P}_{\mathbf{U}}\mathcal{P}_\Omega(\mathbf{N})\mathbf{P}_{\Delta_{\mathbf{Y}}}\|_F, \|\mathbf{P}_{\Delta_{\mathbf{X}}}\mathcal{P}_\Omega(\mathbf{N})\mathbf{P}_{\mathbf{V}}\|_F\}. \end{aligned} \quad (30)$$

Due to the fact that $\mathbf{U}, \Delta_{\mathbf{X}} \in \mathbb{R}^{n_1 \times r}$ and $\mathbf{V}, \Delta_{\mathbf{Y}} \in \mathbb{R}^{n_2 \times r}$, $\mathbf{P}_{\Delta_{\mathbf{X}}}\mathcal{P}_\Omega(\mathbf{N})\mathbf{P}_{\Delta_{\mathbf{Y}}}$, $\mathbf{P}_{\mathbf{U}}\mathcal{P}_\Omega(\mathbf{N})\mathbf{P}_{\Delta_{\mathbf{Y}}}$ and $\mathbf{P}_{\Delta_{\mathbf{X}}}\mathcal{P}_\Omega(\mathbf{N})\mathbf{P}_{\mathbf{V}}$ are matrices with rank at most r . Therefore,

$$\begin{aligned} &\max\{\|\mathbf{P}_{\Delta_{\mathbf{X}}}\mathcal{P}_\Omega(\mathbf{N})\mathbf{P}_{\Delta_{\mathbf{Y}}}\|_F, \|\mathbf{P}_{\mathbf{U}}\mathcal{P}_\Omega(\mathbf{N})\mathbf{P}_{\Delta_{\mathbf{Y}}}\|_F, \|\mathbf{P}_{\Delta_{\mathbf{X}}}\mathcal{P}_\Omega(\mathbf{N})\mathbf{P}_{\mathbf{V}}\|_F\} \\ &\leq \sqrt{r} \max\{\|\mathbf{P}_{\Delta_{\mathbf{X}}}\mathcal{P}_\Omega(\mathbf{N})\mathbf{P}_{\Delta_{\mathbf{Y}}}\|, \|\mathbf{P}_{\mathbf{U}}\mathcal{P}_\Omega(\mathbf{N})\mathbf{P}_{\Delta_{\mathbf{Y}}}\|, \|\mathbf{P}_{\Delta_{\mathbf{X}}}\mathcal{P}_\Omega(\mathbf{N})\mathbf{P}_{\mathbf{V}}\|\} \\ &= \sqrt{r}\psi. \end{aligned}$$

Where the last line follows from (21) and (12). Therefore, (30) gives

$$|K_4| \leq \frac{100\sqrt{|K_4|}r}{p}\psi.$$

Rearranging terms in the last display yields (29).

4.3 Analysis for Subspace-Constrained Matrix Completion

We give a proof of Lemma 1 followed by that of Corollary 1.

4.3.1 PROOF OF LEMMA 1

In order to show the parameterization (4) satisfies Assumption 2, we want to show that for any $\boldsymbol{\theta} = \text{vec}(\boldsymbol{\Theta}_A, \boldsymbol{\Theta}_B) \in \mathbb{R}^{r(s_1+s_2)}$, there exists a $\boldsymbol{\xi} = \text{vec}(\boldsymbol{\Xi}_A, \boldsymbol{\Xi}_B) \in \mathbb{R}^{r(s_1+s_2)}$ that satisfies (9). Let $\boldsymbol{S} := \tilde{\boldsymbol{U}}^\top \boldsymbol{M}^* \tilde{\boldsymbol{V}} \in \mathbb{R}^{s_1 \times s_2}$. Recall that $\tilde{\boldsymbol{U}}$ consists of an orthonormal basis of the column space constraint, and $\tilde{\boldsymbol{V}}$ consists of an orthonormal basis of the column row constraint of $\boldsymbol{M}^*$. Therefore, $\boldsymbol{P}_{\tilde{\boldsymbol{U}}} = \tilde{\boldsymbol{U}}\tilde{\boldsymbol{U}}^\top$, $\boldsymbol{P}_{\tilde{\boldsymbol{V}}} = \tilde{\boldsymbol{V}}\tilde{\boldsymbol{V}}^\top$ and $\boldsymbol{M}^*$ can be represented as $\boldsymbol{M}^* = \tilde{\boldsymbol{U}}\boldsymbol{S}\tilde{\boldsymbol{V}}^\top$. Since $\boldsymbol{M}^*$ is of rank r , by the orthogonality of $\tilde{\boldsymbol{U}}$ and $\tilde{\boldsymbol{V}}$, $\text{rank}(\boldsymbol{S}) = r$. Let the reduced SVD of $\boldsymbol{S}$ be

$$\boldsymbol{S} = \boldsymbol{S}_L \boldsymbol{\Lambda} \boldsymbol{S}_R^\top, \quad (31)$$

where $\boldsymbol{S}_L \in \mathbb{R}^{s_1 \times r}$, $\boldsymbol{S}_R \in \mathbb{R}^{s_2 \times r}$, $\boldsymbol{S}_L^\top \boldsymbol{S}_L = \boldsymbol{I}_{s_1}$, $\boldsymbol{S}_R^\top \boldsymbol{S}_R = \boldsymbol{I}_{s_2}$ and $\boldsymbol{\Lambda} = \text{diag}(\sigma_1, \dots, \sigma_r)$ is a $r \times r$ diagonal matrix with $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_r$. Moreover, by letting $\boldsymbol{U}^* := \tilde{\boldsymbol{U}}\boldsymbol{S}_L \in \mathbb{R}^{n_1 \times r}$, $\boldsymbol{V}^* := \tilde{\boldsymbol{V}}\boldsymbol{S}_R \in \mathbb{R}^{n_2 \times r}$, we can verify that $\boldsymbol{M}^* = \boldsymbol{U}^* \boldsymbol{\Lambda} \boldsymbol{V}^{*\top}$ is a reduced SVD of $\boldsymbol{M}^*$.

Define

$$\boldsymbol{\Xi}_A^* := \boldsymbol{S}_L \boldsymbol{\Lambda}^{1/2} \in \mathbb{R}^{s_1 \times r}, \quad \boldsymbol{\Xi}_B^* := \boldsymbol{S}_R \boldsymbol{\Lambda}^{1/2} \in \mathbb{R}^{s_2 \times r}.$$

For any $\boldsymbol{\Theta}_A \in \mathbb{R}^{s_1 \times r}$ and $\boldsymbol{\Theta}_B \in \mathbb{R}^{s_2 \times r}$, by considering the SVD of $(\boldsymbol{\Theta}_A^\top \boldsymbol{\Xi}_A^* + \boldsymbol{\Theta}_B^\top \boldsymbol{\Xi}_B^*)$, we know there exists an $r \times r$ orthogonal matrix $\boldsymbol{T} \in O(r)$ (Chen and Wainwright, 2015, Lemma 1), such that

$$(\boldsymbol{\Theta}_A^\top \boldsymbol{\Xi}_A^* + \boldsymbol{\Theta}_B^\top \boldsymbol{\Xi}_B^*) \boldsymbol{T} \succeq \mathbf{0}.$$

Let $\boldsymbol{\xi} = \text{vec}(\boldsymbol{\Xi}_A^* \boldsymbol{T}, \boldsymbol{\Xi}_B^* \boldsymbol{T})$, then

$$\boldsymbol{X}(\boldsymbol{\xi}) = \tilde{\boldsymbol{U}} \boldsymbol{\Xi}_A^* \boldsymbol{T} \in \mathbb{R}^{n_1 \times r} \quad \text{and} \quad \boldsymbol{Y}(\boldsymbol{\xi}) = \tilde{\boldsymbol{V}} \boldsymbol{\Xi}_B^* \boldsymbol{T} \in \mathbb{R}^{n_2 \times r}.$$

Keeping in mind that both $\tilde{\boldsymbol{U}}$ and $\tilde{\boldsymbol{V}}$ are orthonormal basis matrices, the conditions in (9) can be verified one by one. First,

$$\boldsymbol{X}(\boldsymbol{\xi}) \boldsymbol{Y}(\boldsymbol{\xi})^\top = \tilde{\boldsymbol{U}} \boldsymbol{\Xi}_A^* \boldsymbol{T} (\tilde{\boldsymbol{V}} \boldsymbol{\Xi}_B^* \boldsymbol{T})^\top = \tilde{\boldsymbol{U}} \boldsymbol{\Xi}_A^* \boldsymbol{\Xi}_B^{*\top} \tilde{\boldsymbol{V}}^\top = \tilde{\boldsymbol{U}} \boldsymbol{S}_L \boldsymbol{\Lambda} \boldsymbol{S}_R^\top \tilde{\boldsymbol{V}}^\top = \boldsymbol{M}^*.$$

The last equality is by (31). Second,

$$\begin{aligned} \boldsymbol{X}(\boldsymbol{\xi})^\top \boldsymbol{X}(\boldsymbol{\xi}) &= (\tilde{\boldsymbol{U}} \boldsymbol{\Xi}_A^* \boldsymbol{T})^\top \tilde{\boldsymbol{U}} \boldsymbol{\Xi}_A^* \boldsymbol{T} = \boldsymbol{T}^\top \boldsymbol{\Xi}_A^{*\top} \boldsymbol{\Xi}_A^* \boldsymbol{T} = \boldsymbol{T}^\top \boldsymbol{\Lambda}^{1/2} \boldsymbol{S}_L^\top \boldsymbol{S}_L \boldsymbol{\Lambda}^{1/2} \boldsymbol{T} = \boldsymbol{T}^\top \boldsymbol{\Lambda} \boldsymbol{T} \\ &= \boldsymbol{T}^\top \boldsymbol{\Xi}_B^{*\top} \boldsymbol{\Xi}_B^* \boldsymbol{T} = (\tilde{\boldsymbol{V}} \boldsymbol{\Xi}_B^* \boldsymbol{T})^\top \tilde{\boldsymbol{V}} \boldsymbol{\Xi}_B^* \boldsymbol{T} = \boldsymbol{Y}(\boldsymbol{\xi})^\top \boldsymbol{Y}(\boldsymbol{\xi}), \end{aligned}$$

where we use the fact $\boldsymbol{S}_L^\top \boldsymbol{S}_L = \boldsymbol{S}_R^\top \boldsymbol{S}_R = \boldsymbol{I}_r$. Finally,

$$\begin{aligned} \boldsymbol{X}(\boldsymbol{\theta})^\top \boldsymbol{X}(\boldsymbol{\xi}) + \boldsymbol{Y}(\boldsymbol{\theta})^\top \boldsymbol{Y}(\boldsymbol{\xi}) &= (\tilde{\boldsymbol{U}} \boldsymbol{\Theta}_A)^\top \tilde{\boldsymbol{U}} \boldsymbol{\Xi}_A^* \boldsymbol{T} + (\tilde{\boldsymbol{V}} \boldsymbol{\Theta}_B)^\top \tilde{\boldsymbol{V}} \boldsymbol{\Xi}_B^* \boldsymbol{T} \\ &= (\boldsymbol{\Theta}_A^\top \boldsymbol{\Xi}_A^* + \boldsymbol{\Theta}_B^\top \boldsymbol{\Xi}_B^*) \boldsymbol{T} \succeq \mathbf{0}. \end{aligned}$$

Therefore, the parameterization $(\boldsymbol{X}(\boldsymbol{\theta}), \boldsymbol{Y}(\boldsymbol{\theta}))$ satisfies Assumption 2.

4.3.2 PROOF OF COROLLARY 1

Since the assumptions of Theorem 1 are satisfied, on the event E defined in Theorem 1, there holds

$$\|\boldsymbol{M}^* - \widehat{\boldsymbol{M}}\|_F^2 \leq \frac{C_3 r}{p^2} \psi^2.$$

Therefore, it suffices to show that

$$\psi^2 \leq \frac{C_4}{C_3} \left(p s_{\max} \log n_{\max} \nu^2 + b^2 \frac{\mu_{\tilde{U}} \mu_{\tilde{V}} s_1 s_2}{n_1 n_2} \log^2 n_{\max} \right). \quad (32)$$

By (21) and (4), for any $\boldsymbol{\theta} \in \mathbb{R}^d$, we have

$$\text{colspan}(\mathbf{X}(\boldsymbol{\theta})) \subseteq \text{colspan}(\tilde{\mathbf{U}}) \quad \text{and} \quad \text{colspan}(\mathbf{Y}(\boldsymbol{\theta})) \subseteq \text{colspan}(\tilde{\mathbf{V}}).$$

Therefore, for any $\boldsymbol{\theta}_1, \boldsymbol{\theta}_2 \in \mathbb{R}^d$, there holds

$$\|\mathbf{P}_{\mathbf{X}(\boldsymbol{\theta}_1)} \mathcal{P}_{\Omega}(\mathbf{N}) \mathbf{P}_{\mathbf{Y}(\boldsymbol{\theta}_2)}\| = \|\mathbf{P}_{\mathbf{X}(\boldsymbol{\theta}_1)} \mathbf{P}_{\tilde{\mathbf{U}}} \mathcal{P}_{\Omega}(\mathbf{N}) \mathbf{P}_{\tilde{\mathbf{V}}} \mathbf{P}_{\mathbf{Y}(\boldsymbol{\theta}_2)}\| \leq \|\mathbf{P}_{\tilde{\mathbf{U}}} \mathcal{P}_{\Omega}(\mathbf{N}) \mathbf{P}_{\tilde{\mathbf{V}}}\|.$$

So we have

$$\psi \leq \|\mathbf{P}_{\tilde{\mathbf{U}}} \mathcal{P}_{\Omega}(\mathbf{N}) \mathbf{P}_{\tilde{\mathbf{V}}}\|.$$

Therefore, (32) can be proved by the following lemma.

Lemma 7. *Assume that the support of observation Ω follows from Model 1. We assume that the entries of the noise matrix $\mathbf{N}$ are i.i.d. centered sub-exponential random variables satisfying the Bernstein condition with parameter b and variance ν^2 . $\tilde{\mathbf{U}}$ and $\tilde{\mathbf{V}}$ are defined in Corollary 1. Then in an event $E_{\text{subspace_noise}}$ with probability $\mathbb{P}[E_{\text{subspace_noise}}] \geq 1 - (n_1 + n_2)^{-3}$, we have*

$$\|\mathbf{P}_{\tilde{\mathbf{U}}} \mathcal{P}_{\Omega}(\mathbf{N}) \mathbf{P}_{\tilde{\mathbf{V}}}\| \leq C_w \left(\sqrt{p(s_1 + s_2) \log(n_1 + n_2)} \nu + b \sqrt{\frac{\mu_{\tilde{\mathbf{U}}} \mu_{\tilde{\mathbf{V}}} s_1 s_2}{n_1 n_2} \log(n_1 + n_2)} \right)$$

for some absolute constant C_w defined in the proof.

The proof of Lemma 7 is mainly following the discussion in Wainwright (2019, Example 6.18) as well as Wainwright (2019, Example 6.14) and is deferred to Appendix B.

Letting $E_{\text{subspace}} = E \cap E_{\text{subspace_noise}}$, and $C_4 = 2C_3 C_w^2$ finishes the proof.

4.4 Analysis for Skew-Symmetric Matrix Completion

In this section, we first give a proof of Lemma 2. Then we give a proof of Corollary 2.

4.4.1 PROOF OF LEMMA 2

Recall that $\mathbf{M}^*$ is a rank- r skew-symmetric matrix, where r is even. Then its Youla decomposition (Youla, 1961) can be written as

$$\mathbf{M}^* := \lambda_1 \boldsymbol{\phi}_1 \boldsymbol{\psi}_1^\top - \lambda_1 \boldsymbol{\psi}_1 \boldsymbol{\phi}_1^\top + \lambda_2 \boldsymbol{\phi}_2 \boldsymbol{\psi}_2^\top - \lambda_2 \boldsymbol{\psi}_2 \boldsymbol{\phi}_2^\top + \dots + \lambda_{r/2} \boldsymbol{\phi}_{r/2} \boldsymbol{\psi}_{r/2}^\top - \lambda_{r/2} \boldsymbol{\psi}_{r/2} \boldsymbol{\phi}_{r/2}^\top,$$

where $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_{r/2} > 0$ and $\boldsymbol{\phi}_1, \dots, \boldsymbol{\phi}_{r/2}, \boldsymbol{\psi}_1, \dots, \boldsymbol{\psi}_{r/2}$ are unit vectors in $\mathbb{R}^n$. Moreover, $\boldsymbol{\phi}_i$'s and $\boldsymbol{\psi}_i$'s are pairwise perpendicular to each other, i.e., for any $i, j \in [r/2]$, $\boldsymbol{\phi}_i^\top \boldsymbol{\psi}_j = 0$, $\boldsymbol{\phi}_i^\top \boldsymbol{\phi}_j = 0$ if $i \neq j$, and $\boldsymbol{\psi}_i^\top \boldsymbol{\psi}_j = 0$ if $i \neq j$.

Let

$$\boldsymbol{\Xi}_A^* = [\sqrt{\lambda_1} \boldsymbol{\phi}_1, \dots, \sqrt{\lambda_{r/2}} \boldsymbol{\phi}_{r/2}] \in \mathbb{R}^{n \times \frac{r}{2}} \quad \text{and} \quad \boldsymbol{\Xi}_B^* = [\sqrt{\lambda_1} \boldsymbol{\psi}_1, \dots, \sqrt{\lambda_{r/2}} \boldsymbol{\psi}_{r/2}] \in \mathbb{R}^{n \times \frac{r}{2}}.$$

It is straightforward to verify

$$M^* = \Xi_A^* \Xi_B^{*\top} - \Xi_B^* \Xi_A^{*\top}.$$

Recall the fact that for any $i, j \in [r/2]$, $\phi_i^\top \psi_j = 0$; $\phi_i^\top \phi_j = 0$ if $i \neq j$ and $\phi_i^\top \phi_j = 1$ if $i = j$; $\psi_i^\top \psi_j = 0$ if $i \neq j$ and $\psi_i^\top \psi_j = 1$ if $i = j$. Therefore,

$$\Xi_A^{*\top} \Xi_B^* = \mathbf{0} \quad \text{and} \quad \Xi_A^{*\top} \Xi_A^* = \Xi_B^{*\top} \Xi_B^* = \text{diag}(\lambda_1, \dots, \lambda_{r/2}). \quad (33)$$

For any $\theta = \text{vec}(\Theta_A, \Theta_B)$ with $\Theta_A, \Theta_B \in \mathbb{R}^{n \times \frac{r}{2}}$, consider the singular value decomposition of the complex matrix $(\Theta_A + \sqrt{-1}\Theta_B)^H (\Xi_A^* + \sqrt{-1}\Xi_B^*)$ (A^H is conjugate transpose of complex matrix A)

$$(\Theta_A + \sqrt{-1}\Theta_B)^H (\Xi_A^* + \sqrt{-1}\Xi_B^*) = ADB^H,$$

where $A, B \in \mathbb{C}^{\frac{r}{2} \times \frac{r}{2}}$ are complex unitary matrices and $D \in \mathbb{R}^{\frac{r}{2} \times \frac{r}{2}}$ is a real diagonal matrix. Therefore, BA^H is also a complex unitary matrix, decompose it as

$$BA^H = R_1 + \sqrt{-1}R_2$$

with $R_1, R_2 \in \mathbb{R}^{\frac{r}{2} \times \frac{r}{2}}$. Therefore,

$$(\Theta_A + \sqrt{-1}\Theta_B)^H (\Xi_A^* + \sqrt{-1}\Xi_B^*) (R_1 + \sqrt{-1}R_2) = ADB^H BA^H = ADA^H \succeq \mathbf{0},$$

that is, it is a Hermitian positive semidefinite matrix. Let

$$\Xi_A = \Xi_A^* R_1 - \Xi_B^* R_2 \quad \text{and} \quad \Xi_B = \Xi_A^* R_2 + \Xi_B^* R_1.$$

Then it holds that

$$(\Xi_A^* + \sqrt{-1}\Xi_B^*) (R_1 + \sqrt{-1}R_2) = \Xi_A + \sqrt{-1}\Xi_B$$

and

$$(\Theta_A + \sqrt{-1}\Theta_B)^H (\Xi_A + \sqrt{-1}\Xi_B) \succeq \mathbf{0},$$

which is equivalent to

$$\begin{bmatrix} \Theta_A^\top \Xi_A + \Theta_B^\top \Xi_B & \Theta_B^\top \Xi_A - \Theta_A^\top \Xi_B \\ \Theta_A^\top \Xi_B - \Theta_B^\top \Xi_A & \Theta_A^\top \Xi_A + \Theta_B^\top \Xi_B \end{bmatrix} \succeq \mathbf{0}.$$

Also, since $R_1 + \sqrt{-1}R_2$ is unitary, we have

$$R := \begin{bmatrix} R_1 & -R_2 \\ R_2 & R_1 \end{bmatrix} \in O(r).$$

Let $\xi = \text{vec}(\Xi_A, \Xi_B)$. Then we have

$$\begin{aligned} X(\xi) &= [\Xi_A, -\Xi_B] = [\Xi_A^* R_1 - \Xi_B^* R_2, -\Xi_A^* R_2 - \Xi_B^* R_1] \\ &= [\Xi_A^*, -\Xi_B^*] \begin{bmatrix} R_1 & -R_2 \\ R_2 & R_1 \end{bmatrix} = [\Xi_A^*, -\Xi_B^*] R, \end{aligned}$$

and similarly

$$\mathbf{Y}(\boldsymbol{\xi}) = [\boldsymbol{\Xi}_B, \boldsymbol{\Xi}_A] = [\boldsymbol{\Xi}_A^* \mathbf{R}_2 + \boldsymbol{\Xi}_B^* \mathbf{R}_1, \boldsymbol{\Xi}_A^* \mathbf{R}_1 - \boldsymbol{\Xi}_B^* \mathbf{R}_2] = [\boldsymbol{\Xi}_B^*, \boldsymbol{\Xi}_A^*] \begin{bmatrix} \mathbf{R}_1 & -\mathbf{R}_2 \\ \mathbf{R}_2 & \mathbf{R}_1 \end{bmatrix} = [\boldsymbol{\Xi}_B^*, \boldsymbol{\Xi}_A^*] \mathbf{R}.$$

It is then straightforward to verify that

$$\mathbf{X}(\boldsymbol{\xi})\mathbf{Y}(\boldsymbol{\xi})^\top = [\boldsymbol{\Xi}_A, -\boldsymbol{\Xi}_B] [\boldsymbol{\Xi}_B, \boldsymbol{\Xi}_A]^\top = [\boldsymbol{\Xi}_A^*, -\boldsymbol{\Xi}_B^*] [\boldsymbol{\Xi}_B^*, \boldsymbol{\Xi}_A^*]^\top = \mathbf{M}^*.$$

In order to further verify $\mathbf{X}(\boldsymbol{\xi})^\top \mathbf{X}(\boldsymbol{\xi}) = \mathbf{Y}(\boldsymbol{\xi})^\top \mathbf{Y}(\boldsymbol{\xi})$, it suffices to prove

$$[\boldsymbol{\Xi}_A^*, -\boldsymbol{\Xi}_B^*]^\top [\boldsymbol{\Xi}_A^*, -\boldsymbol{\Xi}_B^*] = [\boldsymbol{\Xi}_B^*, \boldsymbol{\Xi}_A^*]^\top [\boldsymbol{\Xi}_B^*, \boldsymbol{\Xi}_A^*],$$

which is guaranteed by $\boldsymbol{\Xi}_A^{*\top} \boldsymbol{\Xi}_B^* = \mathbf{0}$ and $\boldsymbol{\Xi}_A^{*\top} \boldsymbol{\Xi}_A^* = \boldsymbol{\Xi}_B^{*\top} \boldsymbol{\Xi}_B^*$ as was shown in (33).

Finally, straightforward calculation gives

$$\begin{aligned} \mathbf{X}(\boldsymbol{\theta})^\top \mathbf{X}(\boldsymbol{\xi}) + \mathbf{Y}(\boldsymbol{\theta})^\top \mathbf{Y}(\boldsymbol{\xi}) &= [\boldsymbol{\Theta}_A, -\boldsymbol{\Theta}_B]^\top [\boldsymbol{\Xi}_A, -\boldsymbol{\Xi}_B] + [\boldsymbol{\Theta}_B, \boldsymbol{\Theta}_A]^\top [\boldsymbol{\Xi}_B, \boldsymbol{\Xi}_A] \\ &= \begin{bmatrix} \boldsymbol{\Theta}_A^\top \boldsymbol{\Xi}_A + \boldsymbol{\Theta}_B^\top \boldsymbol{\Xi}_B & \boldsymbol{\Theta}_B^\top \boldsymbol{\Xi}_A - \boldsymbol{\Theta}_A^\top \boldsymbol{\Xi}_B \\ \boldsymbol{\Theta}_A^\top \boldsymbol{\Xi}_B - \boldsymbol{\Theta}_B^\top \boldsymbol{\Xi}_A & \boldsymbol{\Theta}_A^\top \boldsymbol{\Xi}_A + \boldsymbol{\Theta}_B^\top \boldsymbol{\Xi}_B \end{bmatrix} \succeq \mathbf{0}. \end{aligned}$$

4.4.2 PROOF OF COROLLARY 2

Following the lines in Section 4.3.2, it suffices to show that

$$\psi^2 \leq \frac{C_5}{C_3} (pn \log n \nu^2 + b^2 \log^2 n).$$

Recall the fact that $\psi \leq \|\mathcal{P}_\Omega(\mathbf{N})\|$, then the proof can be done by employing the following Lemma.

Lemma 8. *Let the support of the observed entries Ω satisfy Model 2. We assume that the noise matrix $\mathbf{N}$ is a skew-symmetric matrix, and upper triangular part of $\mathbf{N}$ consists of i.i.d. centered sub-exponential random variables satisfying the Bernstein condition with parameter b and variance ν^2 . Then in an event $E_{\text{skew-noise}}$ with probability $\mathbb{P}[E_{\text{skew-noise}}] \geq 1 - n^{-3}$, we have*

$$\|\mathcal{P}_\Omega(\mathbf{N})\| \leq C_{w'} \left(\sqrt{pn \log n \nu} + b \log n \right).$$

for some absolute constant $C_{w'}$.

The proof is almost exactly the same with proof of Lemma 7, thus omitted here. We can finish the proof of Corollary 2 by letting $E_{\text{skew}} = E \cap E_{\text{skew-noise}}$ and $C_5 = 2C_3 C_{w'}^2$.

5. Discussion

This paper proposes a unified nonconvex optimization framework for matrix completion with linearly parameterized factors. Examples include subspace-constrained and skew-symmetric matrix completion, with applications in collaborative filtering with side information and pairwise comparisons. We also conduct a unified geometric analysis for this nonconvex optimization framework, where a novel assumption referred to as *correlated parametric factorization* plays a key role. In particular, for the noiseless case, we give a sufficient condition

on the sampling rate to guarantee that with high probability no spurious local minimum exists. Our meta-theorem is applicable to examples including subspace-constrained and skew-symmetric matrix completion, given in either case the assumption of correlated parametric factorization is verified.

For future work, we are particularly interested in extending the current framework to more general parameterized factorization, such as extension from linear mappings to affine mappings or even nonlinear mappings. The geometrical interpretation of the assumption of correlated parametric factorization is still intriguing. Its verification for subspace-constrained matrix completion and skew-symmetric matrix completion relies on very different algebraic arguments. Are there any underlying geometric reasons behind such coincidence? Due to its applications in practice, we are particularly interested in the geometric analysis for nonconvex subspace-constrained matrix completion. As explained in Section 2.4 and illustrated in our simulations, our current sampling rate result is suboptimal and we are interested in improving upon it in future.

Acknowledgements

X. Li and J. Chen acknowledge support from the NSF CAREER Award DMS-1848575. J. Chen also acknowledges support from the NSF HDR TRIPODS grant CCF-1934568. We would like to thank an anonymous reviewer who provides ideas for reformulating the proof of Theorem 1.

Appendix A. Proofs of Auxiliary Results in the Proof of Theorem 1

A.1 Preliminaries

First, we list two useful lemmas controlling the difference between the random sampled matrix inner product (or Frobenius norm) and its expectation.

Lemma 9 (Chen and Li 2019). *It holds that uniformly for all $\mathbf{A}, \mathbf{B} \in \mathbb{R}^{n_1 \times r}$, $\mathbf{C}, \mathbf{D} \in \mathbb{R}^{n_2 \times r}$*

$$\begin{aligned} & |\langle \mathcal{P}_\Omega(\mathbf{A}\mathbf{C}^\top), \mathcal{P}_\Omega(\mathbf{B}\mathbf{D}^\top) \rangle - p \langle \mathbf{A}\mathbf{C}^\top, \mathbf{B}\mathbf{D}^\top \rangle| \\ & \leq \|\mathbf{\Omega} - p\mathbf{J}\| \sqrt{\sum_{k=1}^{n_1} \|\mathbf{A}_{k,\cdot}\|_2^2 \|\mathbf{B}_{k,\cdot}\|_2^2} \sqrt{\sum_{k=1}^{n_2} \|\mathbf{C}_{k,\cdot}\|_2^2 \|\mathbf{D}_{k,\cdot}\|_2^2} \\ & \leq \frac{1}{2} \|\mathbf{\Omega} - p\mathbf{J}\| \left(\sum_{k=1}^{n_1} \|\mathbf{A}_{k,\cdot}\|_2^2 \|\mathbf{B}_{k,\cdot}\|_2^2 + \sum_{k=1}^{n_2} \|\mathbf{C}_{k,\cdot}\|_2^2 \|\mathbf{D}_{k,\cdot}\|_2^2 \right). \end{aligned} \quad (34)$$

Here $\mathbf{J}$ denotes the matrix with all entries equal to one.

Lemma 10 (Candès and Recht 2009). *Denote*

$$\mathcal{T} := \left\{ \mathbf{M} \in \mathbb{R}^{n_1 \times n_2} \mid \mathbf{M} = \mathbf{U}^* \mathbf{A}^\top + \mathbf{B} \mathbf{V}^{*\top}, \mathbf{A} \in \mathbb{R}^{n_2 \times r}, \mathbf{B} \in \mathbb{R}^{n_1 \times r} \right\}.$$

If $p \geq C_c \frac{\mu r \log n_{\max}}{n_{\min}}$ for some absolute constant C_c , then in an event E_c with probability $\mathbb{P}[E_c] \geq 1 - (n_1 + n_2)^{-5}$,

$$\left\| \frac{1}{p} \mathcal{P}_\mathcal{T} \mathcal{P}_\Omega \mathcal{P}_\mathcal{T} - \mathcal{P}_\mathcal{T} \right\| \leq 10^{-4}.$$

Here $\mathcal{P}_\Omega$ is determined as in Models 1 and 2, where $\mathcal{P}_\mathcal{T}$ is the orthogonal projector on $\mathcal{T}$ with the Euclidean spaced defined by the standard matrix inner product. In other words, uniformly for all $\mathbf{M} \in \mathcal{T}$,

$$\left| \left\| \frac{1}{p} \mathcal{P}_\Omega(\mathbf{M}) \right\|_F^2 - \|\mathbf{M}\|_F^2 \right| \leq 10^{-4} \|\mathbf{M}\|_F^2. \quad (35)$$

The following lemma on $\|\Omega - p\mathbf{J}\|$ is well known in the literature, see, e.g., Vu (2018) and Bandeira and Van Handel (2016):

Lemma 11. *There is an absolute constant $C_v > 0$, such that if $p \geq C_v \frac{\log n_{\max}}{n_{\min}}$, on an event with probability at least $1 - (n_1 + n_2)^{-5}$, there holds*

$$\|\Omega - p\mathbf{J}\| < C_v \sqrt{n_{\max} p}. \quad (36)$$

We use the following lemma to specify the event E in Theorem 1.

Lemma 12. *Under the assumptions in Theorem 1, in an event E with probability $\mathbb{P}[E] \geq 1 - (n_1 + n_2)^{-3}$,*

$$\begin{cases} p \geq C_0 \frac{\mu r \log n_{\max}}{n_{\min}}, \\ \alpha \geq 100 \sqrt{\frac{\mu r \sigma_1}{n_{\min}}}, \\ \sigma_r \geq C_0 \left(\frac{\mu r \sigma_1 \|\Omega - p\mathbf{J}\|}{p n_{\min}} + \lambda \alpha^2 \right), \\ p\lambda \geq C_0 \|\Omega - p\mathbf{J}\| \end{cases} \quad (37)$$

as well as (35) in Lemma 10 and (36) in Lemma 11 hold. Here $C_0 = 5 \times 10^5$.

Finally, we collect some useful properties of $\mathbf{U} = \mathbf{X}(\boldsymbol{\xi})$ and $\mathbf{V} = \mathbf{Y}(\boldsymbol{\xi})$. The proof is left to Section A.6.

Proposition 1. *For any $\boldsymbol{\theta}$, the matrices $\mathbf{U} = \mathbf{X}(\boldsymbol{\xi})$ and $\mathbf{V} = \mathbf{Y}(\boldsymbol{\xi})$ defined in (21) satisfy the following basic properties:*

- $\text{colspan}(\mathbf{U}) = \text{colspan}(\mathbf{U}^*)$ and $\text{colspan}(\mathbf{V}) = \text{colspan}(\mathbf{V}^*)$;
- The largest singular values of both $\mathbf{U}$ and $\mathbf{V}$ are $\sqrt{\sigma_1}$;
- The r -th singular values of both $\mathbf{U}$ and $\mathbf{V}$ are $\sqrt{\sigma_r}$.
- $\|\mathbf{U}\|_{2,\infty}^2 \leq \frac{\mu r}{n_1} \sigma_1$ and $\|\mathbf{V}\|_{2,\infty}^2 \leq \frac{\mu r}{n_2} \sigma_1$.

A.2 Proof of Lemma 4

Lemma 4 is essentially Ge et al. (2017, Lemma 16) (with noise). Here we give a sketch of the proof for the purpose of being self-contained.

First, denote $f_{\text{clean}}(\mathbf{X}(\boldsymbol{\theta}), \mathbf{Y}(\boldsymbol{\theta}))$ as

$$\begin{aligned} f_{\text{clean}}(\mathbf{X}(\boldsymbol{\theta}), \mathbf{Y}(\boldsymbol{\theta})) &= \frac{1}{2p} \|\mathcal{P}_\Omega(\mathbf{X}(\boldsymbol{\theta})\mathbf{Y}(\boldsymbol{\theta})^\top - \mathbf{M}^*)\|_F^2 + \frac{1}{8} \|\mathbf{X}(\boldsymbol{\theta})^\top \mathbf{X}(\boldsymbol{\theta}) - \mathbf{Y}(\boldsymbol{\theta})^\top \mathbf{Y}(\boldsymbol{\theta})\|_F^2 \\ &\quad + \lambda(G_\alpha(\mathbf{X}(\boldsymbol{\theta})) + G_\alpha(\mathbf{Y}(\boldsymbol{\theta}))). \end{aligned}$$

Compare with (6), and use the simplified notations introduced in (21). We can see

$$f(\mathbf{X}, \mathbf{Y}) = f_{\text{clean}}(\mathbf{X}, \mathbf{Y}) - \frac{1}{p} \langle \mathcal{P}_\Omega(\mathbf{X}\mathbf{Y}^\top - \mathbf{M}^*), \mathcal{P}_\Omega(\mathbf{N}) \rangle + \frac{1}{2p} \|\mathcal{P}_\Omega(\mathbf{N})\|_F^2.$$

Therefore,

$$\begin{aligned} \langle \nabla f(\mathbf{X}, \mathbf{Y}), \text{vec}(\Delta_{\mathbf{X}}, \Delta_{\mathbf{Y}}) \rangle &= \langle \nabla f_{\text{clean}}(\mathbf{X}, \mathbf{Y}), \text{vec}(\Delta_{\mathbf{X}}, \Delta_{\mathbf{Y}}) \rangle \\ &\quad - \frac{1}{p} \langle \mathcal{P}_\Omega(\Delta_{\mathbf{X}}\mathbf{Y}^\top + \mathbf{X}\Delta_{\mathbf{Y}}^\top), \mathcal{P}_\Omega(\mathbf{N}) \rangle \end{aligned}$$

and

$$\begin{aligned} &\text{vec}(\Delta_{\mathbf{X}}, \Delta_{\mathbf{Y}})^\top \nabla^2 f(\mathbf{X}, \mathbf{Y}) \text{vec}(\Delta_{\mathbf{X}}, \Delta_{\mathbf{Y}}) \\ &= \text{vec}(\Delta_{\mathbf{X}}, \Delta_{\mathbf{Y}})^\top \nabla^2 f_{\text{clean}}(\mathbf{X}, \mathbf{Y}) \text{vec}(\Delta_{\mathbf{X}}, \Delta_{\mathbf{Y}}) - \frac{2}{p} \langle \mathcal{P}_\Omega(\Delta_{\mathbf{X}}\Delta_{\mathbf{Y}}^\top), \mathcal{P}_\Omega(\mathbf{N}) \rangle. \end{aligned}$$

Therefore, we only need to concern about $f_{\text{clean}}(\mathbf{X}, \mathbf{Y})$ now, which has already been discussed in Ge et al. (2017). Interested readers can refer to Ge et al. (2017) for the detail.

By Ge et al. (2017, Lemma 16), we have

$$\begin{aligned} &\text{vec}(\Delta_{\mathbf{X}}, \Delta_{\mathbf{Y}})^\top \nabla^2 f_{\text{clean}}(\mathbf{X}, \mathbf{Y}) \text{vec}(\Delta_{\mathbf{X}}, \Delta_{\mathbf{Y}}) - 4 \langle \nabla f_{\text{clean}}(\mathbf{X}, \mathbf{Y}), \text{vec}(\Delta_{\mathbf{X}}, \Delta_{\mathbf{Y}}) \rangle \\ &\leq \frac{1}{4} \left\{ \left\| \Delta_{\mathbf{Z}} \Delta_{\mathbf{Z}}^\top \right\|_F^2 - 3 \left\| \mathbf{Z} \mathbf{Z}^\top - \mathbf{W} \mathbf{W}^\top \right\|_F^2 \right\} + \left(\frac{1}{p} \left\| \mathcal{P}_\Omega(\Delta_{\mathbf{X}} \Delta_{\mathbf{Y}}^\top) \right\|_F^2 - \left\| \Delta_{\mathbf{X}} \Delta_{\mathbf{Y}}^\top \right\|_F^2 \right) \\ &\quad - \left(\frac{3}{p} \left\| \mathcal{P}_\Omega(\mathbf{X} \mathbf{Y}^\top - \mathbf{U} \mathbf{V}^\top) \right\|_F^2 - 3 \left\| \mathbf{X} \mathbf{Y}^\top - \mathbf{U} \mathbf{V}^\top \right\|_F^2 \right) \\ &\quad + \lambda \left[\text{vec}(\Delta_{\mathbf{X}})^\top \nabla^2 G_\alpha(\mathbf{X}) \text{vec}(\Delta_{\mathbf{X}}) - 4 \langle \nabla G_\alpha(\mathbf{X}), \Delta_{\mathbf{X}} \rangle \right] \\ &\quad + \lambda \left[\text{vec}(\Delta_{\mathbf{Y}})^\top \nabla^2 G_\alpha(\mathbf{Y}) \text{vec}(\Delta_{\mathbf{Y}}) - 4 \langle \nabla G_\alpha(\mathbf{Y}), \Delta_{\mathbf{Y}} \rangle \right]. \end{aligned}$$

Therefore,

$$\begin{aligned} &\text{vec}(\Delta_{\mathbf{X}}, \Delta_{\mathbf{Y}})^\top \nabla^2 f(\mathbf{X}, \mathbf{Y}) \text{vec}(\Delta_{\mathbf{X}}, \Delta_{\mathbf{Y}}) - 4 \langle \nabla f(\mathbf{X}, \mathbf{Y}), \text{vec}(\Delta_{\mathbf{X}}, \Delta_{\mathbf{Y}}) \rangle \\ &= \text{vec}(\Delta_{\mathbf{X}}, \Delta_{\mathbf{Y}})^\top \nabla^2 f_{\text{clean}}(\mathbf{X}, \mathbf{Y}) \text{vec}(\Delta_{\mathbf{X}}, \Delta_{\mathbf{Y}}) - \frac{2}{p} \langle \mathcal{P}_\Omega(\Delta_{\mathbf{X}} \Delta_{\mathbf{Y}}^\top), \mathcal{P}_\Omega(\mathbf{N}) \rangle \\ &\quad - 4 \langle \nabla f_{\text{clean}}(\mathbf{X}, \mathbf{Y}), \text{vec}(\Delta_{\mathbf{X}}, \Delta_{\mathbf{Y}}) \rangle + \frac{4}{p} \langle \mathcal{P}_\Omega(\Delta_{\mathbf{X}} \mathbf{Y}^\top + \mathbf{X} \Delta_{\mathbf{Y}}^\top), \mathcal{P}_\Omega(\mathbf{N}) \rangle \\ &\leq \frac{4}{p} \langle \mathcal{P}_\Omega(\Delta_{\mathbf{X}} \mathbf{V}^\top + \mathbf{U} \Delta_{\mathbf{Y}}^\top), \mathcal{P}_\Omega(\mathbf{N}) \rangle + \frac{6}{p} \langle \mathcal{P}_\Omega(\Delta_{\mathbf{X}} \Delta_{\mathbf{Y}}^\top), \mathcal{P}_\Omega(\mathbf{N}) \rangle \\ &\quad + \frac{1}{4} \left\{ \left\| \Delta_{\mathbf{Z}} \Delta_{\mathbf{Z}}^\top \right\|_F^2 - 3 \left\| \mathbf{Z} \mathbf{Z}^\top - \mathbf{W} \mathbf{W}^\top \right\|_F^2 \right\} + \left(\frac{1}{p} \left\| \mathcal{P}_\Omega(\Delta_{\mathbf{X}} \Delta_{\mathbf{Y}}^\top) \right\|_F^2 - \left\| \Delta_{\mathbf{X}} \Delta_{\mathbf{Y}}^\top \right\|_F^2 \right) \\ &\quad - \left(\frac{3}{p} \left\| \mathcal{P}_\Omega(\mathbf{X} \mathbf{Y}^\top - \mathbf{U} \mathbf{V}^\top) \right\|_F^2 - 3 \left\| \mathbf{X} \mathbf{Y}^\top - \mathbf{U} \mathbf{V}^\top \right\|_F^2 \right) \\ &\quad + \lambda \left[\text{vec}(\Delta_{\mathbf{X}})^\top \nabla^2 G_\alpha(\mathbf{X}) \text{vec}(\Delta_{\mathbf{X}}) - 4 \langle \nabla G_\alpha(\mathbf{X}), \Delta_{\mathbf{X}} \rangle \right] \\ &\quad + \lambda \left[\text{vec}(\Delta_{\mathbf{Y}})^\top \nabla^2 G_\alpha(\mathbf{Y}) \text{vec}(\Delta_{\mathbf{Y}}) - 4 \langle \nabla G_\alpha(\mathbf{Y}), \Delta_{\mathbf{Y}} \rangle \right]. \end{aligned}$$

Combining with (15) and Lemma 3 finishes the proof.

A.3 Proof of Lemma 5

This section is meant to control $K_2 + K_3$. The arguments follow closely those used in Chen and Li (2019) with modifications pertinent to the correlated parametric factorization structure in Assumption 2. Recall that we assume event E defined in Lemma 12 holds, so that (34), (35), (36) and (37) hold simultaneously. In what follows, we shall control K_2 and K_3 respectively first before piecing the bounds together.

A.3.1 CONTROL OF K_2

By the way we define Δ_X, Δ_Y in (21),

$$\begin{aligned} \mathbf{X}\mathbf{Y}^\top - \mathbf{U}\mathbf{V}^\top &= (\mathbf{U} + \Delta_X)(\mathbf{V} + \Delta_Y)^\top - \mathbf{U}\mathbf{V}^\top \\ &= \Delta_X\mathbf{V}^\top + \mathbf{U}\Delta_Y^\top + \Delta_X\Delta_Y^\top. \end{aligned}$$

Therefore,

$$\begin{aligned} &\left| \frac{1}{p} \left\| \mathcal{P}_\Omega \left(\mathbf{X}\mathbf{Y}^\top - \mathbf{U}\mathbf{V}^\top \right) \right\|_F^2 - \left\| \mathbf{X}\mathbf{Y}^\top - \mathbf{U}\mathbf{V}^\top \right\|_F^2 \right| \\ &= \left| \frac{1}{p} \left\| \mathcal{P}_\Omega \left(\Delta_X\mathbf{V}^\top + \mathbf{U}\Delta_Y^\top + \Delta_X\Delta_Y^\top \right) \right\|_F^2 - \left\| \Delta_X\mathbf{V}^\top + \mathbf{U}\Delta_Y^\top + \Delta_X\Delta_Y^\top \right\|_F^2 \right| \\ &\leq \textcircled{1} + \textcircled{2} + \textcircled{3} + \textcircled{4}. \end{aligned}$$

Where

$$\begin{aligned} \textcircled{1} &:= \left| \frac{1}{p} \left\| \mathcal{P}_\Omega \left(\Delta_X\mathbf{V}^\top + \mathbf{U}\Delta_Y^\top \right) \right\|_F^2 - \left\| \Delta_X\mathbf{V}^\top + \mathbf{U}\Delta_Y^\top \right\|_F^2 \right|, \\ \textcircled{2} &:= \left| \frac{1}{p} \left\| \mathcal{P}_\Omega \left(\Delta_X\Delta_Y^\top \right) \right\|_F^2 - \left\| \Delta_X\Delta_Y^\top \right\|_F^2 \right|, \\ \textcircled{3} &:= \left| \frac{2}{p} \left\langle \mathcal{P}_\Omega \left(\Delta_X\mathbf{V}^\top \right), \mathcal{P}_\Omega \left(\Delta_X\Delta_Y^\top \right) \right\rangle - 2 \left\langle \Delta_X\mathbf{V}^\top, \Delta_X\Delta_Y^\top \right\rangle \right|, \quad \text{and} \\ \textcircled{4} &:= \left| \frac{2}{p} \left\langle \mathcal{P}_\Omega \left(\mathbf{U}\Delta_Y^\top \right), \mathcal{P}_\Omega \left(\Delta_X\Delta_Y^\top \right) \right\rangle - 2 \left\langle \mathbf{U}\Delta_Y^\top, \Delta_X\Delta_Y^\top \right\rangle \right|. \end{aligned}$$

By Proposition 1, the matrix $\Delta_X\mathbf{V}^\top + \mathbf{U}\Delta_Y^\top$ belongs to the subspace $\mathcal{T}$ defined in Lemma 10. Therefore, Lemma 10 implies

$$\textcircled{1} \leq 0.0001 \left\| \Delta_X\mathbf{V}^\top + \mathbf{U}\Delta_Y^\top \right\|_F^2 \leq 0.0002 \left(\left\| \Delta_X\mathbf{V}^\top \right\|_F^2 + \left\| \mathbf{U}\Delta_Y^\top \right\|_F^2 \right).$$

Applying Lemma 9, we further have

$$\begin{aligned} \textcircled{2} &\leq \frac{\|\Omega - p\mathbf{J}\|}{2p} \left(\sum_{k=1}^{n_1} \|(\Delta_X)_{k,\cdot}\|_2^4 + \sum_{k=1}^{n_2} \|(\Delta_Y)_{k,\cdot}\|_2^4 \right), \\ \textcircled{3} &\leq \frac{\|\Omega - p\mathbf{J}\|}{p} \left(\sum_{k=1}^{n_1} \|(\Delta_X)_{k,\cdot}\|_2^4 + \sum_{k=1}^{n_2} \|\mathbf{V}_{k,\cdot}\|_2^2 \|(\Delta_Y)_{k,\cdot}\|_2^2 \right), \\ \textcircled{4} &\leq \frac{\|\Omega - p\mathbf{J}\|}{p} \left(\sum_{k=1}^{n_2} \|(\Delta_Y)_{k,\cdot}\|_2^4 + \sum_{k=1}^{n_1} \|\mathbf{U}_{k,\cdot}\|_2^2 \|(\Delta_X)_{k,\cdot}\|_2^2 \right). \end{aligned}$$

By using Proposition 1, $\|\mathbf{U}\|_{2,\infty}^2 \leq \frac{\mu r}{n_1} \sigma_1$ and $\|\mathbf{V}\|_{2,\infty}^2 \leq \frac{\mu r}{n_2} \sigma_1$. Hence, we further bound (3) and (4) as

$$\begin{aligned} \textcircled{3} &\leq \frac{\|\mathbf{\Omega} - p\mathbf{J}\|}{p} \left(\sum_{k=1}^{n_1} \|(\mathbf{\Delta}_X)_{k,\cdot}\|_2^4 + \frac{\mu r}{n_2} \sigma_1 \|\mathbf{\Delta}_Y\|_F^2 \right), \\ \textcircled{4} &\leq \frac{\|\mathbf{\Omega} - p\mathbf{J}\|}{p} \left(\sum_{k=1}^{n_2} \|(\mathbf{\Delta}_Y)_{k,\cdot}\|_2^4 + \frac{\mu r}{n_1} \sigma_1 \|\mathbf{\Delta}_X\|_F^2 \right). \end{aligned}$$

We combine the foregoing inequalities to obtain

$$\begin{aligned} K_2 &\leq \left| \frac{1}{p} \left\| \mathcal{P}_\Omega \left(\mathbf{\Delta}_X \mathbf{\Delta}_Y^\top \right) \right\|_F^2 - \|\mathbf{\Delta}_X \mathbf{\Delta}_Y^\top\|_F^2 \right| \\ &\quad + \left| \frac{3}{p} \left\| \mathcal{P}_\Omega \left(\mathbf{X} \mathbf{Y}^\top - \mathbf{U} \mathbf{V}^\top \right) \right\|_F^2 - 3 \|\mathbf{X} \mathbf{Y}^\top - \mathbf{U} \mathbf{V}^\top\|_F^2 \right| \\ &\leq \textcircled{2} + 3(\textcircled{1} + \textcircled{2} + \textcircled{3} + \textcircled{4}) \\ &\leq 0.0006 \left(\left\| \mathbf{\Delta}_X \mathbf{V}^\top \right\|_F^2 + \left\| \mathbf{U} \mathbf{\Delta}_Y^\top \right\|_F^2 \right) \\ &\quad + \frac{\|\mathbf{\Omega} - p\mathbf{J}\|}{p} \left(5 \sum_{k=1}^{n_1} \|(\mathbf{\Delta}_X)_{k,\cdot}\|_2^4 + 5 \sum_{k=1}^{n_2} \|(\mathbf{\Delta}_Y)_{k,\cdot}\|_2^4 + 3 \frac{\mu r}{n_1} \sigma_1 \|\mathbf{\Delta}_X\|_F^2 + 3 \frac{\mu r}{n_2} \sigma_1 \|\mathbf{\Delta}_Y\|_F^2 \right). \end{aligned}$$

A.3.2 CONTROL OF K_3

For K_3 , when $\alpha \geq 100 \sqrt{\frac{\mu r \sigma_1}{n_{\min}}}$, we have

$$K_3 \leq 200 \lambda \alpha^2 (\|\mathbf{\Delta}_X\|_F^2 + \|\mathbf{\Delta}_Y\|_F^2) - 0.3 \lambda \left(\sum_{k=1}^{n_1} \|(\mathbf{\Delta}_X)_{k,\cdot}\|_2^4 + \sum_{k=1}^{n_2} \|(\mathbf{\Delta}_Y)_{k,\cdot}\|_2^4 \right).$$

The remaining argument is the same as that in Ge et al. (2017, Lemma 11), which has also been employed in Chen and Li (2019, Lemma 14).

A.3.3 PUTTING K_2 AND K_3 TOGETHER

Combining the upper bounds on K_2 and K_3 , we obtain that

$$\begin{aligned} K_2 + K_3 &\leq 0.0006 \left(\left\| \mathbf{\Delta}_X \mathbf{V}^\top \right\|_F^2 + \left\| \mathbf{U} \mathbf{\Delta}_Y^\top \right\|_F^2 \right) \\ &\quad + \left(\frac{3 \mu r \sigma_1 \|\mathbf{\Omega} - p\mathbf{J}\|}{p n_{\min}} + 200 \lambda \alpha^2 \right) (\|\mathbf{\Delta}_X\|_F^2 + \|\mathbf{\Delta}_Y\|_F^2) \\ &\quad + \left(\frac{5 \|\mathbf{\Omega} - p\mathbf{J}\|}{p} - 0.3 \lambda \right) \left(\sum_{k=1}^{n_1} \|(\mathbf{\Delta}_X)_{k,\cdot}\|_2^4 + \sum_{k=1}^{n_2} \|(\mathbf{\Delta}_Y)_{k,\cdot}\|_2^4 \right). \end{aligned}$$

By the third inequality in (37),

$$\left(\frac{3 \mu r \sigma_1 \|\mathbf{\Omega} - p\mathbf{J}\|}{p n_{\min}} + 200 \lambda \alpha^2 \right) \leq \frac{200 \sigma_r}{C_0}.$$

In addition, by the fourth inequality in (37),

$$\left(\frac{5\|\boldsymbol{\Omega} - p\mathbf{J}\|}{p} - 0.3\lambda \right) \leq \left(\frac{5}{C_0} - 0.3 \right) \lambda.$$

Therefore, with sufficiently large C_0 (e.g., $C_0 = 200/0.0004 = 5 \times 10^5$),

$$\begin{aligned} K_2 + K_3 &\leq 0.0006 \left(\|\boldsymbol{\Delta}_X \mathbf{V}^\top\|_F^2 + \|\mathbf{U} \boldsymbol{\Delta}_Y^\top\|_F^2 \right) + 0.0004\sigma_r (\|\boldsymbol{\Delta}_X\|_F^2 + \|\boldsymbol{\Delta}_Y\|_F^2) \\ &\quad + (10^{-5} - 0.3)\lambda \left(\sum_{k=1}^{n_1} \|(\boldsymbol{\Delta}_X)_{k,\cdot}\|_2^4 + \sum_{k=1}^{n_2} \|(\boldsymbol{\Delta}_Y)_{k,\cdot}\|_2^4 \right) \\ &\leq 0.0006 \left(\|\boldsymbol{\Delta}_X \mathbf{V}^\top\|_F^2 + \|\mathbf{U} \boldsymbol{\Delta}_Y^\top\|_F^2 \right) + 0.0004\sigma_r (\|\boldsymbol{\Delta}_X\|_F^2 + \|\boldsymbol{\Delta}_Y\|_F^2), \end{aligned}$$

where the last inequality uses the facts that $\lambda \geq 0$ and that

$$\sum_{k=1}^{n_1} \|(\boldsymbol{\Delta}_X)_{k,\cdot}\|_2^4 + \sum_{k=1}^{n_2} \|(\boldsymbol{\Delta}_Y)_{k,\cdot}\|_2^4 \geq 0.$$

By Proposition 1, we have

$$\|\mathbf{U} \boldsymbol{\Delta}_Y^\top\|_F^2 \geq \sigma_r^2(\mathbf{U}) \|\boldsymbol{\Delta}_Y\|_F^2 = \sigma_r \|\boldsymbol{\Delta}_Y\|_F^2$$

and

$$\|\mathbf{V} \boldsymbol{\Delta}_X^\top\|_F^2 \geq \sigma_r^2(\mathbf{V}) \|\boldsymbol{\Delta}_X\|_F^2 = \sigma_r \|\boldsymbol{\Delta}_X\|_F^2.$$

Together with the last display, they lead to

$$K_2 + K_3 \leq 0.001 \left(\|\boldsymbol{\Delta}_X \mathbf{V}^\top\|_F^2 + \|\mathbf{U} \boldsymbol{\Delta}_Y^\top\|_F^2 \right). \quad (38)$$

This completes the proof of Lemma 5.

A.4 Proof of Lemma 6

By the way we define $\boldsymbol{\Delta}_Z$,

$$\begin{aligned} &\|\mathbf{Z}\mathbf{Z}^\top - \mathbf{W}\mathbf{W}^\top\|_F^2 \\ &= \|\boldsymbol{\Delta}_Z \mathbf{W}^\top + \mathbf{W} \boldsymbol{\Delta}_Z^\top + \boldsymbol{\Delta}_Z \boldsymbol{\Delta}_Z^\top\|_F^2 \\ &= \|\boldsymbol{\Delta}_Z \boldsymbol{\Delta}_Z^\top\|_F^2 + 2\|\boldsymbol{\Delta}_Z \mathbf{W}^\top\|_F^2 + 2\langle \boldsymbol{\Delta}_Z \mathbf{W}^\top, \mathbf{W} \boldsymbol{\Delta}_Z^\top \rangle + 4\langle \boldsymbol{\Delta}_Z \mathbf{W}^\top, \boldsymbol{\Delta}_Z \boldsymbol{\Delta}_Z^\top \rangle \\ &= \|\boldsymbol{\Delta}_Z^\top \boldsymbol{\Delta}_Z\|_F^2 + 2\langle \boldsymbol{\Delta}_Z^\top \boldsymbol{\Delta}_Z, \mathbf{W}^\top \mathbf{W} \rangle + 2\langle \boldsymbol{\Delta}_Z^\top \mathbf{W}, \mathbf{W}^\top \boldsymbol{\Delta}_Z \rangle + 4\langle \boldsymbol{\Delta}_Z^\top \boldsymbol{\Delta}_Z, \boldsymbol{\Delta}_Z^\top \mathbf{W} \rangle. \end{aligned} \quad (39)$$

Here we have used the facts that $\langle \mathbf{A}, \mathbf{B} \rangle = \text{trace}(\mathbf{A}^\top \mathbf{B})$ and $\langle \mathbf{A}, \mathbf{B} \rangle = \langle \mathbf{B}, \mathbf{A} \rangle$. By the definition of K_1 in (24), (39) further implies

$$\begin{aligned} K_1 &= -\frac{1}{2} \|\boldsymbol{\Delta}_Z^\top \boldsymbol{\Delta}_Z\|_F^2 - \frac{3}{2} \langle \boldsymbol{\Delta}_Z^\top \boldsymbol{\Delta}_Z, \mathbf{W}^\top \mathbf{W} \rangle - \frac{3}{2} \langle \boldsymbol{\Delta}_Z^\top \mathbf{W}, \mathbf{W}^\top \boldsymbol{\Delta}_Z \rangle \\ &\quad - 3 \langle \boldsymbol{\Delta}_Z^\top \boldsymbol{\Delta}_Z, \boldsymbol{\Delta}_Z^\top \mathbf{W} \rangle. \end{aligned} \quad (40)$$

Moreover, condition (22) leads to

$$\mathbf{Z}^\top \mathbf{W} = \mathbf{X}^\top \mathbf{U} + \mathbf{Y}^\top \mathbf{V} \succeq \mathbf{0}, \quad (41)$$

which implies the symmetricity of $\mathbf{W}^\top \Delta_{\mathbf{Z}} = \mathbf{W}^\top \mathbf{Z} - \mathbf{Z}^\top \mathbf{Z}$ (this is a crucial step for the analysis in Jin et al. (2017) and Ge et al. (2017)). Hence,

$$\langle \Delta_{\mathbf{Z}}^\top \mathbf{W}, \mathbf{W}^\top \Delta_{\mathbf{Z}} \rangle = \|\Delta_{\mathbf{Z}}^\top \mathbf{W}\|_F^2. \quad (42)$$

Combining (42) with (40) we have

$$K_1 = -0.5\|\Delta_{\mathbf{Z}}^\top \Delta_{\mathbf{Z}}\|_F^2 - 1.5\langle \Delta_{\mathbf{Z}}^\top \Delta_{\mathbf{Z}}, \mathbf{W}^\top \mathbf{W} \rangle - 1.5\|\Delta_{\mathbf{Z}}^\top \mathbf{W}\|_F^2 - 3\langle \Delta_{\mathbf{Z}}^\top \Delta_{\mathbf{Z}}, \Delta_{\mathbf{Z}}^\top \mathbf{W} \rangle.$$

Therefore, based on (38), we are able to upper bound the righthand side of (23) as

$$\begin{aligned} & K_1 + K_2 + K_3 + K_4 \\ & \leq -0.5\|\Delta_{\mathbf{Z}}^\top \Delta_{\mathbf{Z}}\|_F^2 - 1.5\langle \Delta_{\mathbf{Z}}^\top \Delta_{\mathbf{Z}}, \mathbf{W}^\top \mathbf{W} \rangle - 1.5\|\Delta_{\mathbf{Z}}^\top \mathbf{W}\|_F^2 - 3\langle \Delta_{\mathbf{Z}}^\top \Delta_{\mathbf{Z}}, \Delta_{\mathbf{Z}}^\top \mathbf{W} \rangle \\ & \quad + 0.001 \left(\|\Delta_{\mathbf{X}} \mathbf{V}^\top\|_F^2 + \|\mathbf{U} \Delta_{\mathbf{Y}}^\top\|_F^2 \right) + |K_4|. \end{aligned}$$

Furthermore, due to cyclic property of trace,

$$\begin{aligned} \langle \Delta_{\mathbf{Z}}^\top \Delta_{\mathbf{Z}}, \mathbf{W}^\top \mathbf{W} \rangle &= \text{trace}(\Delta_{\mathbf{Z}}^\top \Delta_{\mathbf{Z}} \mathbf{W}^\top \mathbf{W}) = \|\mathbf{W} \Delta_{\mathbf{Z}}^\top\|_F^2 \\ &\geq \|\Delta_{\mathbf{X}} \mathbf{V}^\top\|_F^2 + \|\mathbf{U} \Delta_{\mathbf{Y}}^\top\|_F^2. \end{aligned} \quad (43)$$

Together with the second last display, (43), we further have

$$\begin{aligned} & K_1 + K_2 + K_3 + K_4 \\ & \leq -0.5\|\Delta_{\mathbf{Z}}^\top \Delta_{\mathbf{Z}}\|_F^2 - 1.499\langle \Delta_{\mathbf{Z}}^\top \Delta_{\mathbf{Z}}, \mathbf{W}^\top \mathbf{W} \rangle - 1.5\|\Delta_{\mathbf{Z}}^\top \mathbf{W}\|_F^2 \\ & \quad - 3\langle \Delta_{\mathbf{Z}}^\top \Delta_{\mathbf{Z}}, \Delta_{\mathbf{Z}}^\top \mathbf{W} \rangle + |K_4|. \end{aligned} \quad (44)$$

Now (41) further implies that

$$\langle \Delta_{\mathbf{Z}}^\top \Delta_{\mathbf{Z}}, \mathbf{W}^\top \mathbf{W} \rangle + \langle \Delta_{\mathbf{Z}}^\top \Delta_{\mathbf{Z}}, \Delta_{\mathbf{Z}}^\top \mathbf{W} \rangle = \langle \Delta_{\mathbf{Z}}^\top \Delta_{\mathbf{Z}}, \mathbf{Z}^\top \mathbf{W} \rangle \geq 0,$$

in which we use the fact that the inner product of two PSD matrices is nonnegative. Thus,

$$\begin{aligned} & K_1 + K_2 + K_3 + K_4 \\ & \leq -0.5\|\Delta_{\mathbf{Z}}^\top \Delta_{\mathbf{Z}}\|_F^2 - 1.499 \left(\langle \Delta_{\mathbf{Z}}^\top \Delta_{\mathbf{Z}}, \mathbf{W}^\top \mathbf{W} \rangle + \langle \Delta_{\mathbf{Z}}^\top \Delta_{\mathbf{Z}}, \Delta_{\mathbf{Z}}^\top \mathbf{W} \rangle \right) - 1.5\|\Delta_{\mathbf{Z}}^\top \mathbf{W}\|_F^2 \\ & \quad - 1.501\langle \Delta_{\mathbf{Z}}^\top \Delta_{\mathbf{Z}}, \Delta_{\mathbf{Z}}^\top \mathbf{W} \rangle + |K_4| \\ & \leq -0.5\|\Delta_{\mathbf{Z}}^\top \Delta_{\mathbf{Z}}\|_F^2 - 1.5\|\Delta_{\mathbf{Z}}^\top \mathbf{W}\|_F^2 - 1.501\langle \Delta_{\mathbf{Z}}^\top \Delta_{\mathbf{Z}}, \Delta_{\mathbf{Z}}^\top \mathbf{W} \rangle + |K_4| \\ & \leq -0.5\|\Delta_{\mathbf{Z}}^\top \Delta_{\mathbf{Z}}\|_F^2 - 1.5\|\Delta_{\mathbf{Z}}^\top \mathbf{W}\|_F^2 + 1.501\|\Delta_{\mathbf{Z}}^\top \Delta_{\mathbf{Z}}\|_F \|\Delta_{\mathbf{Z}}^\top \mathbf{W}\|_F + |K_4|. \end{aligned} \quad (45)$$

Note that (45) implies

$$0.5\|\Delta_{\mathbf{Z}}^\top \Delta_{\mathbf{Z}}\|_F^2 + 1.5\|\Delta_{\mathbf{Z}}^\top \mathbf{W}\|_F^2 - 1.501\|\Delta_{\mathbf{Z}}^\top \Delta_{\mathbf{Z}}\|_F \|\Delta_{\mathbf{Z}}^\top \mathbf{W}\|_F \leq |K_4|.$$

We apply Young's inequality to the LHS of the inequality and subsequently deduce the upper bounds

$$\|\Delta_Z \Delta_Z^\top\|_F = \|\Delta_Z^\top \Delta_Z\|_F \leq 3\sqrt{|K_4|}, \quad \|\Delta_Z^\top \mathbf{W}\|_F \leq 2\sqrt{|K_4|}. \quad (46)$$

By (44) as well as (27), we have

$$\begin{aligned} & 1.499 \langle \Delta_Z^\top \Delta_Z, \mathbf{W}^\top \mathbf{W} \rangle \\ & \leq -0.5 \|\Delta_Z^\top \Delta_Z\|_F^2 - 1.5 \|\Delta_Z^\top \mathbf{W}\|_F^2 - 3 \langle \Delta_Z^\top \Delta_Z, \Delta_Z^\top \mathbf{W} \rangle + |K_4| \\ & \leq 3 \|\Delta_Z^\top \Delta_Z\|_F \|\Delta_Z^\top \mathbf{W}\|_F + |K_4| \\ & \leq 19 |K_4|. \end{aligned}$$

Combining with (43) we have claim 2:

$$\|\Delta_X \mathbf{V}^\top\|_F^2 + \|\mathbf{U} \Delta_Y^\top\|_F^2 \leq 13 |K_4|. \quad (47)$$

This concludes the proof of Lemma 6.

A.5 Proof of Lemma 12

First of all, for the first line and second line of (37), by first line and third line of assumption (11), we have

$$p \geq C_1 \frac{\mu r \log n_{\max}}{n_{\min}}$$

and

$$\alpha \geq C_2 \sqrt{\frac{\mu r \sigma_1}{n_{\min}}}.$$

Furthermore, for the fourth line of (37), by second line of (11), $p\lambda \geq C_2 \sqrt{n_{\max} p}$. From Lemma 11, if $p \geq C_v \frac{\log n_{\max}}{n_{\min}}$, in an event E_v with probability $\mathbb{P}[E_v] \geq 1 - (n_1 + n_2)^{-5}$, $\|\mathbf{\Omega} - p\mathbf{J}\| \leq C_v \sqrt{n_{\max} p}$, where C_v is defined in Lemma 11. Therefore, in the event E_v , $p\lambda \geq \frac{C_2}{C_v} \|\mathbf{\Omega} - p\mathbf{J}\|$.

Finally, by the fact that $\lambda \leq 10C_2 \sqrt{\frac{n_{\max}}{p}}$, $\alpha \leq 10C_2 \sqrt{\frac{\mu r \sigma_1}{n_{\min}}}$ and $p\lambda \geq \frac{C_2}{C_v} \|\mathbf{\Omega} - p\mathbf{J}\|$, we have

$$\frac{\mu r \sigma_1 \|\mathbf{\Omega} - p\mathbf{J}\|}{pn_{\min}} \leq \frac{C_v \mu r \sigma_1 p \lambda}{C_2 p n_{\min}} \leq \frac{C_v}{C_2} \frac{\mu r \sigma_1}{n_{\min}} 10C_2 \sqrt{\frac{n_{\max}}{p}} = 10C_v \sigma_r \sqrt{\frac{\mu^2 r^2 \kappa^2 n_{\max}}{n_{\min}^2 p}}$$

and

$$\lambda \alpha^2 \leq 10^3 C_2^3 \sigma_r \sqrt{\frac{\mu^2 r^2 \kappa^2 n_{\max}}{n_{\min}^2 p}}.$$

By the first line of (11),

$$p \geq \frac{C_1}{n_{\min}^2} n_{\max} \mu^2 r^2 \kappa^2.$$

Therefore,

$$\frac{\mu r \sigma_1 \|\mathbf{\Omega} - p\mathbf{J}\|}{pn_{\min}} + \lambda \alpha^2 \leq \frac{10C_v + 10^3 C_2^3}{\sqrt{C_1}} \sigma_r.$$

In other words,

$$\sigma_r \geq \frac{\sqrt{C_1}}{10C_v + 10^3 C_2^3} \left(\frac{\mu r \sigma_1 \|\mathbf{\Omega} - p\mathbf{J}\|}{pn_{\min}} + \lambda \alpha^2 \right).$$

Therefore, by choosing

$$C_1 = \max\{C_0, C_c, C_0^2(10C_v + 10^3 C_2^3)^2\}$$

and

$$C_2 = \max\{100, C_0 C_v\}$$

finishes the proof of the first part of the Lemma.

Recall by the way we define C_1 , if (11) is satisfied, by Lemma 10, in an event E_c with probability $\mathbb{P}[E_c] \geq 1 - (n_1 + n_2)^{-5}$, (35) holds. Therefore, let $E = E_v \cap E_c$, then by union bound, $\mathbb{P}[E] \geq 1 - (n_1 + n_2)^{-3}$.

A.6 Proof of Proposition 1

First, since $\mathbf{M}^\star$ has SVD $\mathbf{M}^\star = \mathbf{U}^\star \mathbf{\Lambda} \mathbf{V}^{\star\top}$, we have

$$\text{colspan}(\mathbf{U}^\star) = \text{colspan}(\mathbf{M}^\star) \quad \text{and} \quad \text{colspan}(\mathbf{V}^\star) = \text{rowspan}(\mathbf{M}^\star)$$

as well as

$$\dim(\text{colspan}(\mathbf{M}^\star)) = \dim(\text{rowspan}(\mathbf{M}^\star)) = r.$$

From (9), we also have

$$\text{colspan}(\mathbf{M}^\star) \subseteq \text{colspan}(\mathbf{U}) \quad \text{and} \quad \text{rowspan}(\mathbf{M}^\star) \subseteq \text{colspan}(\mathbf{V}).$$

By the way we define $\mathbf{U}$ and $\mathbf{V}$, we have $\dim(\text{colspan}(\mathbf{U})) \leq r$ and $\dim(\text{colspan}(\mathbf{V})) \leq r$. Therefore, $\text{colspan}(\mathbf{U}) = \text{colspan}(\mathbf{U}^\star)$ and $\text{colspan}(\mathbf{V}) = \text{colspan}(\mathbf{V}^\star)$.

From second equation in (9), $\mathbf{U}^\top \mathbf{U} = \mathbf{V}^\top \mathbf{V}$, therefore,

$$\sigma_i(\mathbf{U}) = \sqrt{\lambda_i(\mathbf{U}^\top \mathbf{U})} = \sqrt{\lambda_i(\mathbf{V}^\top \mathbf{V})} = \sigma_i(\mathbf{V}), \quad i = 1, 2, \dots, r.$$

Moreover, suppose $\mathbf{U}^\top \mathbf{U} = \mathbf{V}^\top \mathbf{V} = \mathbf{B} \mathbf{D}^2 \mathbf{B}^\top$ be a fixed eigenvalue decomposition of $\mathbf{U}^\top \mathbf{U}$, with $\mathbf{B} \in \text{O}(r)$ and $\mathbf{D} \in \mathbb{R}^{r \times r}$ diagonal matrix. Then the reduced SVD of $\mathbf{U}$ and $\mathbf{V}$ can be written as

$$\mathbf{U} = \mathbf{A}_U \mathbf{D} \mathbf{B}^\top, \quad \mathbf{V} = \mathbf{A}_V \mathbf{D} \mathbf{B}^\top$$

with $\mathbf{A}_U \in \mathbb{R}^{n_1 \times r}$, $\mathbf{A}_V \in \mathbb{R}^{n_2 \times r}$ satisfying $\mathbf{A}_U^\top \mathbf{A}_U = \mathbf{I}$ and $\mathbf{A}_V^\top \mathbf{A}_V = \mathbf{I}$. Therefore, $\mathbf{M}^\star = \mathbf{U} \mathbf{V}^\top = \mathbf{A}_U \mathbf{D}^2 \mathbf{A}_V^\top$. It is a reduced SVD of $\mathbf{M}^\star$ by the way we define $\mathbf{A}_U$, $\mathbf{A}_V$ and $\mathbf{D}$. Therefore, $\sigma_1(\mathbf{U}) = \sigma_1(\mathbf{V}) = \sqrt{\sigma_1}$ and $\sigma_r(\mathbf{U}) = \sigma_r(\mathbf{V}) = \sqrt{\sigma_r}$.

Moreover, there is $\mathbf{R}_U, \mathbf{R}_V \in \text{O}(r)$ such that $\mathbf{A}_U = \mathbf{U}^\star \mathbf{R}_U$, $\mathbf{A}_V = \mathbf{V}^\star \mathbf{R}_V$. Therefore,

$$\begin{aligned} \|\mathbf{U}\|_{2,\infty}^2 &= \|\mathbf{A}_U \mathbf{D} \mathbf{B}^\top\|_{2,\infty}^2 = \|\mathbf{A}_U \mathbf{D}\|_{2,\infty}^2 \\ &\leq \|\mathbf{A}_U\|_{2,\infty}^2 \|\mathbf{D}\|_{\ell_\infty}^2 = \sigma_1 \|\mathbf{U}^\star \mathbf{R}_U\|_{2,\infty}^2 = \sigma_1 \|\mathbf{U}^\star\|_{2,\infty}^2 \leq \frac{\mu r}{n_1} \sigma_1. \end{aligned}$$

Similarly, we also have $\|\mathbf{V}\|_{2,\infty}^2 \leq \frac{\mu r}{n_2} \sigma_1$.

Appendix B. Proof of Lemma 7

Recall $\tilde{\mathbf{U}}$ and $\tilde{\mathbf{V}}$ are orthonormal basis matrices, $\mathbf{P}_{\tilde{\mathbf{U}}} = \tilde{\mathbf{U}}\tilde{\mathbf{U}}^\top$, $\mathbf{P}_{\tilde{\mathbf{V}}} = \tilde{\mathbf{V}}\tilde{\mathbf{V}}^\top$. Therefore,

$$\|\mathbf{P}_{\tilde{\mathbf{U}}}\mathcal{P}_\Omega(\mathbf{N})\mathbf{P}_{\tilde{\mathbf{V}}}\| = \|\tilde{\mathbf{U}}\tilde{\mathbf{U}}^\top\mathcal{P}_\Omega(\mathbf{N})\tilde{\mathbf{V}}\tilde{\mathbf{V}}^\top\| = \|\tilde{\mathbf{U}}^\top\mathcal{P}_\Omega(\mathbf{N})\tilde{\mathbf{V}}\|.$$

The last equality uses the fact that $\tilde{\mathbf{U}}$ and $\tilde{\mathbf{V}}$ are orthonormal basis matrices, therefore

$$\|\tilde{\mathbf{U}}\mathbf{A}\| = \|\mathbf{A}\|, \quad \|\mathbf{B}\tilde{\mathbf{V}}^\top\| = \|\mathbf{B}\|$$

for any $\mathbf{A}, \mathbf{B}$ with suitable size.

Due to the fact that Ω follows from Model 1, entries of $\mathcal{P}_\Omega(\mathbf{N})$ can be written as $[\mathcal{P}_\Omega(\mathbf{N})]_{i,j} = \delta_{i,j}N_{i,j}$, where $\delta_{i,j}$'s are i.i.d. Bernoulli random variables such that

$$\delta_{i,j} = \begin{cases} 1 & \text{with probability } p \\ 0 & \text{with probability } 1 - p. \end{cases}$$

And $N_{i,j}$'s are i.i.d. centered sub-exponential random variables. Moreover, $\delta_{i,j}$'s and $N_{i,j}$'s are mutually independent. Therefore,

$$\|\mathbf{P}_{\tilde{\mathbf{U}}}\mathcal{P}_\Omega(\mathbf{N})\mathbf{P}_{\tilde{\mathbf{V}}}\| = \|\tilde{\mathbf{U}}^\top\mathcal{P}_\Omega(\mathbf{N})\tilde{\mathbf{V}}\| = \left\| \tilde{\mathbf{U}}^\top \left(\sum_{i,j} \delta_{i,j} N_{i,j} \mathbf{e}_i \mathbf{e}_j^\top \right) \tilde{\mathbf{V}} \right\| = \left\| \sum_{i,j} \delta_{i,j} N_{i,j} \tilde{\mathbf{U}}_{i,\cdot} \tilde{\mathbf{V}}_{j,\cdot}^\top \right\|.$$

Now let

$$\mathbf{Q}_{i,j} := \delta_{i,j} N_{i,j} \begin{bmatrix} \mathbf{0} & \tilde{\mathbf{U}}_{i,\cdot} \tilde{\mathbf{V}}_{j,\cdot}^\top \\ \tilde{\mathbf{V}}_{j,\cdot} \tilde{\mathbf{U}}_{i,\cdot}^\top & \mathbf{0} \end{bmatrix}.$$

Therefore,

$$\|\mathbf{P}_{\tilde{\mathbf{U}}}\mathcal{P}_\Omega(\mathbf{N})\mathbf{P}_{\tilde{\mathbf{V}}}\| = \|\tilde{\mathbf{U}}^\top\mathcal{P}_\Omega(\mathbf{N})\tilde{\mathbf{V}}\| = \left\| \sum_{i,j} \mathbf{Q}_{i,j} \right\|$$

and $\mathbb{E}[\mathbf{Q}_{i,j}] = \mathbf{0}$. By following the symmetrization argument in Wainwright (2019, Example 6.14), without loss of generality, we can assume that $N_{i,j}$'s are symmetric random variable, i.e., $N_{i,j} \stackrel{d}{=} -N_{i,j}$. Now we want to verify the Bernstein's condition (Wainwright, 2019, Definition 6.10) for $\mathbf{Q}_{i,j}$'s. For $k \geq 3$,

$$\mathbb{E}[\mathbf{Q}_{i,j}^k] = \mathbb{E} \left[\delta_{i,j}^k N_{i,j}^k \begin{bmatrix} \mathbf{0} & \tilde{\mathbf{U}}_{i,\cdot} \tilde{\mathbf{V}}_{j,\cdot}^\top \\ \tilde{\mathbf{V}}_{j,\cdot} \tilde{\mathbf{U}}_{i,\cdot}^\top & \mathbf{0} \end{bmatrix}^k \right] = p \mathbb{E}[N_{i,j}^k] \begin{bmatrix} \mathbf{0} & \tilde{\mathbf{U}}_{i,\cdot} \tilde{\mathbf{V}}_{j,\cdot}^\top \\ \tilde{\mathbf{V}}_{j,\cdot} \tilde{\mathbf{U}}_{i,\cdot}^\top & \mathbf{0} \end{bmatrix}^k.$$

Due to the symmetry of $N_{i,j}$, $\mathbb{E}[N_{i,j}^k] = 0$ when $k \geq 3$ is odd, therefore, $\mathbb{E}[\mathbf{Q}_{i,j}^k] = 0$. For $k \geq 2$ even, we have

$$\begin{aligned} \begin{bmatrix} \mathbf{0} & \tilde{\mathbf{U}}_{i,\cdot} \tilde{\mathbf{V}}_{j,\cdot}^\top \\ \tilde{\mathbf{V}}_{j,\cdot} \tilde{\mathbf{U}}_{i,\cdot}^\top & \mathbf{0} \end{bmatrix}^k &= \begin{bmatrix} (\tilde{\mathbf{U}}_{i,\cdot} \tilde{\mathbf{V}}_{j,\cdot}^\top \tilde{\mathbf{V}}_{j,\cdot} \tilde{\mathbf{U}}_{i,\cdot}^\top)^{k/2} & \mathbf{0} \\ \mathbf{0} & (\tilde{\mathbf{V}}_{j,\cdot} \tilde{\mathbf{U}}_{i,\cdot}^\top \tilde{\mathbf{U}}_{i,\cdot} \tilde{\mathbf{V}}_{j,\cdot}^\top)^{k/2} \end{bmatrix} \\ &= \|\tilde{\mathbf{U}}_{i,\cdot}\|_2^k \|\tilde{\mathbf{V}}_{j,\cdot}\|_2^k \begin{bmatrix} \frac{1}{\|\tilde{\mathbf{U}}_{i,\cdot}\|_2^2} \tilde{\mathbf{U}}_{i,\cdot} \tilde{\mathbf{U}}_{i,\cdot}^\top & \mathbf{0} \\ \mathbf{0} & \frac{1}{\|\tilde{\mathbf{V}}_{j,\cdot}\|_2^2} \tilde{\mathbf{V}}_{j,\cdot} \tilde{\mathbf{V}}_{j,\cdot}^\top \end{bmatrix}, \end{aligned}$$

which is a positive semidefinite matrix. And due to the fact that $N_{i,j}$'s satisfy the Bernstein condition, for $k \geq 2$,

$$\mathbb{E}[N_{i,j}^k] \leq \frac{1}{2} k! \nu^2 b^{k-2}.$$

Therefore, for $k \geq 3$ even,

$$\mathbb{E}[\mathbf{Q}_{i,j}^k] \preceq \frac{1}{2} k! \nu^2 b^{k-2} p \|\tilde{\mathbf{U}}_{i,\cdot}\|_2^k \|\tilde{\mathbf{V}}_{j,\cdot}\|_2^k \begin{bmatrix} \frac{1}{\|\tilde{\mathbf{U}}_{i,\cdot}\|_2^2} \tilde{\mathbf{U}}_{i,\cdot} \tilde{\mathbf{U}}_{i,\cdot}^\top & \mathbf{0} \\ \mathbf{0} & \frac{1}{\|\tilde{\mathbf{V}}_{j,\cdot}\|_2^2} \tilde{\mathbf{V}}_{j,\cdot} \tilde{\mathbf{V}}_{j,\cdot}^\top \end{bmatrix}.$$

And we also have

$$\begin{aligned} \mathbb{V}[\mathbf{Q}_{i,j}] &= \mathbb{E}[\mathbf{Q}_{i,j}^2] = p \mathbb{E}[N_{i,j}^2] \begin{bmatrix} \mathbf{0} & \tilde{\mathbf{U}}_{i,\cdot} \tilde{\mathbf{V}}_{j,\cdot}^\top \\ \tilde{\mathbf{V}}_{j,\cdot} \tilde{\mathbf{U}}_{i,\cdot}^\top & \mathbf{0} \end{bmatrix}^2 \\ &= p \nu^2 \|\tilde{\mathbf{U}}_{i,\cdot}\|_2^2 \|\tilde{\mathbf{V}}_{j,\cdot}\|_2^2 \begin{bmatrix} \frac{1}{\|\tilde{\mathbf{U}}_{i,\cdot}\|_2^2} \tilde{\mathbf{U}}_{i,\cdot} \tilde{\mathbf{U}}_{i,\cdot}^\top & \mathbf{0} \\ \mathbf{0} & \frac{1}{\|\tilde{\mathbf{V}}_{j,\cdot}\|_2^2} \tilde{\mathbf{V}}_{j,\cdot} \tilde{\mathbf{V}}_{j,\cdot}^\top \end{bmatrix}. \end{aligned}$$

Therefore, for $k \geq 3$,

$$\mathbb{E}[\mathbf{Q}_{i,j}^k] \preceq \frac{1}{2} k! b^{k-2} \|\tilde{\mathbf{U}}_{i,\cdot}\|_2^{k-2} \|\tilde{\mathbf{V}}_{j,\cdot}\|_2^{k-2} \mathbb{V}[\mathbf{Q}_{i,j}].$$

Therefore, $\mathbf{Q}_{i,j}$ satisfies Bernstein condition with parameter $b \|\tilde{\mathbf{U}}_{i,\cdot}\|_2 \|\tilde{\mathbf{V}}_{j,\cdot}\|_2 \leq b \sqrt{\frac{\mu_{\tilde{\mathbf{U}}} \mu_{\tilde{\mathbf{V}}} s_1 s_2}{n_1 n_2}}$. Furthermore,

$$\begin{aligned} \frac{1}{n_1 n_2} \left\| \sum_{(i,j) \in [n_1] \times [n_2]} \mathbb{V}[\mathbf{Q}_{i,j}] \right\| &= \frac{1}{n_1 n_2} p \nu^2 \left\| \sum_{(i,j) \in [n_1] \times [n_2]} \begin{bmatrix} \|\tilde{\mathbf{V}}_{j,\cdot}\|_2^2 \tilde{\mathbf{U}}_{i,\cdot} \tilde{\mathbf{U}}_{i,\cdot}^\top & \mathbf{0} \\ \mathbf{0} & \|\tilde{\mathbf{U}}_{i,\cdot}\|_2^2 \tilde{\mathbf{V}}_{j,\cdot} \tilde{\mathbf{V}}_{j,\cdot}^\top \end{bmatrix} \right\| \\ &= \frac{1}{n_1 n_2} p \nu^2 \left\| \begin{bmatrix} \|\tilde{\mathbf{V}}\|_F^2 \tilde{\mathbf{U}}^\top \tilde{\mathbf{U}} & \mathbf{0} \\ \mathbf{0} & \|\tilde{\mathbf{U}}\|_F^2 \tilde{\mathbf{V}}^\top \tilde{\mathbf{V}} \end{bmatrix} \right\| \\ &\leq \frac{1}{n_1 n_2} p \nu^2 (s_1 + s_2). \end{aligned}$$

Where the last equality uses the fact that $\tilde{\mathbf{U}}^\top \tilde{\mathbf{U}} = \mathbf{I}$, $\tilde{\mathbf{V}}^\top \tilde{\mathbf{V}} = \mathbf{I}$. Then by Wainwright (2019, Theorem 6.17), for all $t > 0$,

$$\mathbb{P} \left[\frac{1}{n_1 n_2} \left\| \sum_{i,j} \mathbf{Q}_{i,j} \right\| \geq t \right] \leq 2(n_1 + n_2) \exp \left(- \frac{n_1 n_2 t^2}{2 \left(\frac{1}{n_1 n_2} p \nu^2 (s_1 + s_2) + b \sqrt{\frac{\mu_{\tilde{\mathbf{U}}} \mu_{\tilde{\mathbf{V}}} s_1 s_2}{n_1 n_2}} t \right)} \right).$$

Therefore, by choosing t as

$$t = C_w \frac{1}{n_1 n_2} \left(\sqrt{p \nu^2 (s_1 + s_2) \log(n_1 + n_2)} + b \sqrt{\frac{\mu_{\tilde{\mathbf{U}}} \mu_{\tilde{\mathbf{V}}} s_1 s_2}{n_1 n_2} \log(n_1 + n_2)} \right)$$

with absolute constant C_w sufficiently large, say $C_w = 10$, then

$$\begin{aligned} & \mathbb{P} \left[\left\| \mathbf{P}_{\tilde{\mathbf{U}}} \mathcal{P}_{\Omega}(\mathbf{N}) \mathbf{P}_{\tilde{\mathbf{V}}} \right\| \geq C_w \left(\sqrt{p\nu^2(s_1 + s_2) \log(n_1 + n_2)} + b \sqrt{\frac{\mu_{\tilde{\mathbf{U}}} \mu_{\tilde{\mathbf{V}}} s_1 s_2}{n_1 n_2} \log(n_1 + n_2)} \right) \right] \\ &= \mathbb{P} \left[\left\| \sum_{i,j} \mathbf{Q}_{i,j} \right\| \geq C_w \left(\sqrt{p\nu^2(s_1 + s_2) \log(n_1 + n_2)} + b \sqrt{\frac{\mu_{\tilde{\mathbf{U}}} \mu_{\tilde{\mathbf{V}}} s_1 s_2}{n_1 n_2} \log(n_1 + n_2)} \right) \right] \\ &\leq (n_1 + n_2)^{-3}. \end{aligned}$$

References

- Afonso S Bandeira and Ramon Van Handel. Sharp nonasymptotic bounds on the norm of random matrices with independent entries. *The Annals of Probability*, 44(4):2479–2506, 2016.
- T Tony Cai, Xiaodong Li, and Zongming Ma. Optimal rates of convergence for noisy sparse phase retrieval via thresholded wirtinger flow. *The Annals of Statistics*, 44(5):2221–2251, 2016.
- Emmanuel J Candès and Benjamin Recht. Exact matrix completion via convex optimization. *Foundations of Computational mathematics*, 9(6):717, 2009.
- Emmanuel J Candès and Terence Tao. The power of convex relaxation: Near-optimal matrix completion. *IEEE Transactions on Information Theory*, 56(5):2053–2080, 2010.
- Emmanuel J Candès, Xiaodong Li, Yi Ma, and John Wright. Robust principal component analysis? *Journal of the ACM (JACM)*, 58(3):11, 2011.
- Emmanuel J Candes, Xiaodong Li, and Mahdi Soltanolkotabi. Phase retrieval via wirtinger flow: Theory and algorithms. *IEEE Transactions on Information Theory*, 61(4):1985–2007, 2015.
- Sourav Chatterjee. Matrix estimation by universal singular value thresholding. *The Annals of Statistics*, 43(1):177–214, 2015.
- Ji Chen and Xiaodong Li. Model-free nonconvex matrix completion: Local minima analysis and applications in memory-efficient kernel PCA. *Journal of Machine Learning Research*, 20(142):1–39, 2019.
- Ji Chen, Dekai Liu, and Xiaodong Li. Nonconvex rectangular matrix completion via gradient descent without $\ell_{2,\infty}$ regularization. *IEEE Transactions on Information Theory*, 66(9):5806–5841, 2020.
- Yudong Chen. Incoherence-optimal matrix completion. *IEEE Transactions on Information Theory*, 61(5):2909–2923, 2015.
- Yudong Chen and Martin J Wainwright. Fast low-rank estimation by projected gradient descent: General statistical and algorithmic guarantees. *arXiv preprint arXiv:1509.03025*, 2015.

- Lijun Ding and Yudong Chen. Leave-one-out approach for matrix completion: Primal and dual analysis. *IEEE Transactions on Information Theory*, 66(11):7274–7301, 2020.
- Armin Eftekhari, Dehui Yang, and Michael B Wakin. Weighted matrix completion and recovery with prior subspace information. *IEEE Transactions on Information Theory*, 64(6):4044–4071, 2018.
- Rong Ge, Jason D Lee, and Tengyu Ma. Matrix completion has no spurious local minimum. In *Advances in Neural Information Processing Systems*, pages 2973–2981, 2016.
- Rong Ge, Chi Jin, and Yi Zheng. No spurious local minima in nonconvex low rank problems: A unified geometric analysis. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 1233–1242. JMLR. org, 2017.
- David F Gleich and Lek-Heng Lim. Rank aggregation via nuclear norm minimization. In *Proceedings of the 17th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 60–68. ACM, 2011.
- David Gross. Recovering low-rank matrices from few coefficients in any basis. *IEEE Transactions on Information Theory*, 57(3):1548–1566, 2011.
- Daniel Hsu, Sham M Kakade, and Tong Zhang. Robust matrix decomposition with sparse corruptions. *IEEE Transactions on Information Theory*, 57(11):7221–7234, 2011.
- Prateek Jain and Inderjit Dhillon. Provable inductive matrix completion. *arXiv preprint arXiv:1306.0626*, 06 2013.
- Prateek Jain, Raghu Meka, and Inderjit S Dhillon. Guaranteed rank minimization via singular value projection. In *Advances in Neural Information Processing Systems*, pages 937–945, 2010.
- Prateek Jain, Praneeth Netrapalli, and Sujay Sanghavi. Low-rank matrix completion using alternating minimization. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*, pages 665–674. ACM, 2013.
- Xiaoye Jiang, Lek-Heng Lim, Yuan Yao, and Yinyu Ye. Statistical ranking and combinatorial hodge theory. *Mathematical Programming*, 127(1):203–244, 2011.
- Chi Jin, Rong Ge, Praneeth Netrapalli, Sham M Kakade, and Michael I Jordan. How to escape saddle points efficiently. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 1724–1732. JMLR. org, 2017.
- Raghunandan H Keshavan, Andrea Montanari, and Sewoong Oh. Matrix completion from a few entries. *IEEE transactions on information theory*, 56(6):2980–2998, 2010a.
- Raghunandan H Keshavan, Andrea Montanari, and Sewoong Oh. Matrix completion from noisy entries. *Journal of Machine Learning Research*, 11(Jul):2057–2078, 2010b.
- Vladimir Koltchinskii, Karim Lounici, and Alexandre B Tsybakov. Nuclear-norm penalization and optimal rates for noisy low-rank matrix completion. *The Annals of Statistics*, 39(5):2302–2329, 2011.

- Xiaodong Li. Compressed sensing and matrix completion with constant proportion of corruptions. *Constructive Approximation*, 37(1):73–99, 2013.
- 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. *Foundations of Computational Mathematics*, 20(3), 2020.
- Nagarajan Natarajan and Inderjit S Dhillon. Inductive matrix completion for predicting gene–disease associations. *Bioinformatics*, 30(12):i60–i68, 2014.
- Sahand Negahban and Martin J Wainwright. Restricted strong convexity and weighted matrix completion: Optimal bounds with noise. *Journal of Machine Learning Research*, 13(May):1665–1697, 2012.
- Benjamin Recht. A simpler approach to matrix completion. *Journal of Machine Learning Research*, 12(Dec):3413–3430, 2011.
- Benjamin Recht, Maryam Fazel, and Pablo A Parrilo. Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization. *SIAM review*, 52(3):471–501, 2010.
- Jasson DM Rennie and Nathan Srebro. Fast maximum margin matrix factorization for collaborative prediction. In *Proceedings of the 22nd International Conference on Machine Learning*, pages 713–719. ACM, 2005.
- Si Si, Kai-Yang Chiang, Cho-Jui Hsieh, Nikhil Rao, and Inderjit S Dhillon. Goal-directed inductive matrix completion. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 1165–1174, 2016.
- Ju Sun, Qing Qu, and John Wright. A geometric analysis of phase retrieval. *Foundations of Computational Mathematics*, 18(5):1131–1198, 2018.
- Ruoyu Sun and Zhi-Quan Luo. Guaranteed matrix completion via non-convex factorization. *IEEE Transactions on Information Theory*, 62(11):6535–6579, 2016.
- Tingni Sun and Cun-Hui Zhang. Calibrated elastic regularization in matrix completion. In *Advances in Neural Information Processing Systems*, pages 863–871, 2012.
- Van Vu. A simple svd algorithm for finding hidden partitions. *Combinatorics, Probability and Computing*, 27(1):124–140, 2018.
- Martin J Wainwright. *High-Dimensional Statistics: A Non-Asymptotic Viewpoint*, volume 48. Cambridge University Press, 2019.
- Miao Xu, Rong Jin, and Zhi-Hua Zhou. Speedup matrix completion with side information: Application to multi-label learning. In *Advances in Neural Information Processing Systems*, pages 2301–2309, 2013.

- Jinfeng Yi, Lijun Zhang, Rong Jin, Qi Qian, and Anil Jain. Semi-supervised clustering by input pattern assisted pairwise similarity matrix completion. In *International Conference on Machine Learning*, pages 1400–1408, 2013.
- Xinyang Yi, Dohyung Park, Yudong Chen, and Constantine Caramanis. Fast algorithms for robust PCA via gradient descent. In *Advances in Neural Information Processing Systems*, pages 4152–4160, 2016.
- DC Youla. A normal form for a matrix under the unitary congruence group. *Canadian Journal of Mathematics*, 13:694–704, 1961.
- Qinqing Zheng and John Lafferty. Convergence analysis for rectangular matrix completion using burer-monteiro factorization and gradient descent. *arXiv preprint arXiv:1605.07051*, 2016.