

# Subspace Estimation from Unbalanced and Incomplete Data Matrices: $\ell_{2,\infty}$ Statistical Guarantees

Changxiao Cai\*    Gen Li†    Yuejie Chi‡    H. Vincent Poor\*    Yuxin Chen\*

October 2019;    Revised: June 2020

## Abstract

This paper is concerned with estimating the column space of an unknown low-rank matrix  $\mathbf{A}^* \in \mathbb{R}^{d_1 \times d_2}$ , given noisy and partial observations of its entries. There is no shortage of scenarios where the observations — while being too noisy to support faithful recovery of the entire matrix — still convey sufficient information to enable reliable estimation of the column space of interest. This is particularly evident and crucial for the highly unbalanced case where the column dimension  $d_2$  far exceeds the row dimension  $d_1$ , which is the focal point of the current paper.

We investigate an efficient spectral method, which operates upon the sample Gram matrix with diagonal deletion. While this algorithmic idea has been studied before, we establish new statistical guarantees for this method in terms of both  $\ell_2$  and  $\ell_{2,\infty}$  estimation accuracy, which improve upon prior results if  $d_2$  is substantially larger than  $d_1$ . To illustrate the effectiveness of our findings, we derive matching minimax lower bounds with respect to the noise levels, and develop consequences of our general theory for three applications of practical importance: (1) tensor completion from noisy data, (2) covariance estimation / principal component analysis with missing data, and (3) community recovery in bipartite graphs. Our theory leads to improved performance guarantees for all three cases.

**Keywords:** spectral method, principal component analysis with missing data, tensor completion, covariance estimation, spectral clustering, leave-one-out analysis

## Contents

<b>1</b>	<b>Introduction</b>	<b>3</b>
<b>2</b>	<b>Problem formulation</b>	<b>4</b>
2.1	Models	4
2.2	Notation	5
<b>3</b>	<b>Main results</b>	<b>6</b>
3.1	Algorithm: a spectral method with diagonal deletion	6
3.2	Theoretical guarantees	7
3.3	Minimax lower bounds	9

---

Corresponding author: Yuxin Chen.

\*Department of Electrical Engineering, Princeton University, Princeton, NJ 08544, USA; email: {ccai, poor, yuxin.chen}@princeton.edu.

†Department of Electronic Engineering, Tsinghua University, Beijing, China 100086, USA; email: g-li16@mails.tsinghua.edu.cn.

‡Department of Electrical and Computer Engineering, Carnegie Mellon University, Pittsburgh, PA 15213, USA; email: yuejiechi@cmu.edu.

<b>4</b>	<b>Consequences for concrete applications</b>	<b>10</b>
4.1	Noisy tensor completion	10
4.2	Covariance estimation / principal component analysis with missing data	13
4.3	Community recovery in bipartite stochastic block models	15
4.4	Numerical experiments	17
<b>5</b>	<b>Further related work</b>	<b>19</b>
<b>6</b>	<b>Analysis</b>	<b>21</b>
6.1	Leave-one-out and leave-two-out estimates	21
6.2	Key lemmas	22
6.3	Proof of Theorem 1	24
6.3.1	Statistical accuracy measured by $\ \cdot\ $	24
6.3.2	Statistical accuracy measured by $\ \cdot\ _{2,\infty}$	25
<b>7</b>	<b>Discussion</b>	<b>27</b>
<b>A</b>	<b>Proofs for corollaries</b>	<b>27</b>
A.1	Proof of Corollary 1	27
A.2	Proof of Corollary 2	29
A.2.1	Estimation error of the principle subspace	30
A.2.2	Estimation error of the covariance matrix	31
A.3	Proof of Corollary 3	33
<b>B</b>	<b>Proofs for key lemmas</b>	<b>34</b>
B.1	Auxiliary quantities, notation, and preliminary facts	34
B.2	Proof of Lemma 1	35
B.2.1	Step 1: bounding the term $\mathcal{P}_{\text{off-diag}}(\mathbf{E}\mathbf{E}^\top)$	36
B.2.2	Step 2: bounding the term $\mathcal{P}_{\text{off-diag}}(\mathbf{A}^*\mathbf{E}^\top + \mathbf{E}\mathbf{A}^{*\top})$	37
B.2.3	Step 3: combining Step 1 and Step 2	40
B.3	Proof of Lemma 2	40
B.4	Proof of Lemma 3	42
B.5	Proof of Lemma 4	43
B.6	Proof of Lemma 5	44
<b>C</b>	<b>Proofs for auxiliary lemmas</b>	<b>46</b>
C.1	Auxiliary technical lemmas	46
C.2	Proof of Lemma 6	47
C.3	Proof of Lemma 7	48
C.4	Proof of Lemma 8	48
C.5	Proof of Lemma 9	50
C.5.1	Step 1: controlling $\ \mathbf{G}^{(m)} - \mathbf{G}^{(m,l)}\ $	51
C.5.2	Step 2: controlling $\ (\mathbf{G}^{(m)} - \mathbf{G}^{(m,l)})\mathbf{U}^{(m,l)}\ $	51
C.5.3	Step 3: combining Step 1 and Step 2	53
C.6	Proof of Lemma 10	54
<b>D</b>	<b>Proofs for lower bounds</b>	<b>56</b>
D.1	Proof of Theorem 2	56
D.2	Proof of Theorem 3	59
<b>E</b>	<b>A few more auxiliary lemmas</b>	<b>60</b>
E.1	Proof of Lemma 11	61
E.2	Proof of Lemma 12	61
E.3	Proof of Lemma 13	62
E.4	Proof of Lemma 14	62

# 1 Introduction

Consider the problem of estimating the *column space* of a low-rank matrix  $\mathbf{A}^* = [A_{i,j}^*]_{1 \leq i \leq d_1, 1 \leq j \leq d_2}$ , based on noisy and highly incomplete observations of its entries. To set the stage, suppose that we observe

$$A_{i,j} = A_{i,j}^* + N_{i,j}, \quad \forall (i, j) \in \Omega, \quad (1)$$

where  $\Omega \subseteq \{1, \dots, d_1\} \times \{1, \dots, d_2\}$  is the sampling set, and  $A_{i,j}$  denotes the observed entry at location  $(i, j)$ , which is corrupted by noise  $N_{i,j}$ . In contrast to the classical matrix completion problem that aims to fill in all missing entries [CR09, KMO10a, CLC19, CC18b], the current paper focuses solely on estimating the column space of  $\mathbf{A}^*$ , which is oftentimes a less stringent requirement.

**Motivating applications.** A problem of this kind arises in numerous applications. We immediately point out several representative examples as follows, with precise descriptions postponed to Section 4.

- *Tensor completion and estimation.* Imagine we seek to estimate a low-rank symmetric tensor from partial observations of its entries [BM16, XYZ17, MS18], a task that spans various applications like visual data inpainting [LMWY13] and medical imaging [SHKM14]. Consider, for example, an order-3 tensor  $\mathbf{T}^* = \sum_{s=1}^r \mathbf{w}_s^* \otimes \mathbf{w}_s^* \otimes \mathbf{w}_s^* \in \mathbb{R}^{d \times d \times d}$ , where  $\{\mathbf{w}_s^*\}$  represents a collection of tensor factors.<sup>1</sup> An alternative representation of  $\mathbf{T}^*$  can be obtained by unfolding the tensor of interest into a low-rank structured matrix  $\mathbf{A}^* \in \mathbb{R}^{d \times d^2}$ . Consequently, estimation of the subspace spanned by  $\{\mathbf{w}_s^*\}$  from partial noisy entries of  $\mathbf{T}^*$  — which serves as a common and crucial step for tensor completion [MS18, CLPC19] — is equivalent to estimating the column space of  $\mathbf{A}^*$  from incomplete and noisy data; see Section 4.1 for details. Notably, the unfolded matrix becomes extremely fat as the dimension  $d$  grows.
- *Covariance estimation/principal component analysis with missing data.* Suppose we have available a sequence of  $n$  independent sample vectors  $\{\mathbf{x}_i \in \mathbb{R}^d\}_{i=1}^n$ , whose covariance matrix exhibits certain low-dimensional structure. Several statistical models fall in the same vein of this model, e.g. the generalized spiked model [BY12] and the factor model [FWZZ18]. An important task amounts to estimating the principal subspace of the covariance matrix of interest, possibly in the presence of missing data (where we only get to see highly incomplete entries of  $\{\mathbf{x}_i\}$ ). If a substantial amount of data are missing, then individual sample vectors cannot possibly be recovered, thus enabling privacy protection for individual data. Fortunately, one might still hope to estimate the principal subspace, provided that a large number of sample vectors are queried (which might yield a fat data matrix  $\mathbf{X} := [\mathbf{x}_1, \dots, \mathbf{x}_n]$ ). See Section 4.2.
- *Community recovery in bipartite graphs.* Community recovery is often concerned with clustering a collection of individuals or nodes into different communities, based on similarities between pairs of nodes. In many complex networks, such pairwise interactions might only occur when the two nodes involved belong to two disjoint groups (denoted by  $\mathcal{U}$  and  $\mathcal{V}$  respectively). This calls for community recovery in bipartite networks (sometimes referred to as biclustering) [Dhi01, LCJ14, AHB14, ZA18]. As we shall detail in Section 4.3, the biclustering problem is tightly connected to subspace estimation; for instance, the column subspace of some bi-adjacency matrix  $\mathbf{A} \in \mathbb{R}^{|\mathcal{U}| \times |\mathcal{V}|}$  (which is a noisy copy of a low-rank matrix) reveals the community memberships in  $\mathcal{U}$ . When the size of  $\mathcal{V}$  is substantially larger than that of  $\mathcal{U}$ , one might encounter a situation where only the nodes in  $\mathcal{U}$  (rather than those in  $\mathcal{V}$ ) can be reliably clustered. This calls for development of “one-sided” community recovery algorithms, that is, the type of algorithms that guarantee reliable clustering of  $\mathcal{U}$  without worrying about the clustering accuracy in  $\mathcal{V}$ .

**Contributions.** Since we concentrate primarily on estimating the column space of  $\mathbf{A}^*$ , it is natural to expect a reduced sample complexity as well as a weaker requirement on the signal-to-noise ratio, in comparison to the conditions required for reliable reconstruction of the whole matrix — particularly for those highly unbalanced problems with drastically different dimensions  $d_1$  and  $d_2$ . Focusing on a spectral method applied to the Gram matrix  $\mathbf{A}\mathbf{A}^\top$  with diagonal deletion (whose variants have been studied in multiple contexts [Lou14, FP16, LW12, MS18, EvdG19]), we establish new statistical guarantees in terms of the sample complexity and the estimation accuracy, both of which strengthen prior theory. Our results deliver optimal

<sup>1</sup>For any vectors  $\mathbf{a}, \mathbf{b}, \mathbf{c} \in \mathbb{R}^d$ , we use  $\mathbf{a} \otimes \mathbf{b} \otimes \mathbf{c}$  to denote a  $d \times d \times d$  array whose  $(i, j, k)$ -th entry is given by  $a_i b_j c_k$ .

$\ell_{2,\infty}$  estimation risk bounds with respect to the noise level, which are previously unavailable. All this is accomplished via a powerful leave-one-out / leave-two-out analysis framework. Further, we develop minimax lower bounds under Gaussian noise, revealing that the sample complexity and the signal-to-noise ratio (SNR) required for spectral methods to achieve consistent estimation are both minimax optimal (up to some logarithmic factor). Finally, we develop concrete consequences of our general theory for all three applications mentioned above, leading to optimal performance guarantees that improve upon prior literature.

It is worth noting that low-rank subspace estimation from noisy and incomplete data has been extensively studied in a large number of prior work (e.g. [Lou14, AFWZ17, CKR17, ZCW18, ZWS19]). While many of these prior results allow  $d_1$  and  $d_2$  to differ, they typically fall short of establishing optimal dependency on  $d_1$  and  $d_2$  in the highly unbalanced scenarios. The focal point of this paper is thus to characterize the effect of such unbalancedness (as reflected by the aspect ratio  $d_2/d_1$ ) upon consistent subspace estimation.

**Paper organization.** The rest of this paper is organized as follows. Section 2 formulates the problem and introduces basic definitions and notations. In Section 3, we present our theoretical guarantees for a spectral method, as well as minimax lower bounds. Section 4 applies our general theorem to the aforementioned applications, and corroborate our theory by numerical experiments. Section 5 provides an overview of related prior works. The proof of our main theory is outlined in Section 6, with the proofs of many auxiliary lemmas postponed to the appendix. We conclude the paper with a discussion of future directions in Section 7.

## 2 Problem formulation

### 2.1 Models

**Low-rank matrix.** Suppose that the unknown matrix  $\mathbf{A}^* \in \mathbb{R}^{d_1 \times d_2}$  is rank- $r$ , where the row dimension  $d_1$  and the column dimension  $d_2$  are allowed to be drastically different. Assume that the (compact) singular value decomposition (SVD) of  $\mathbf{A}^*$  is given by

$$\mathbf{A}^* = \mathbf{U}^* \mathbf{\Sigma}^* \mathbf{V}^{*\top} = \sum_{i=1}^r \sigma_i^* \mathbf{u}_i^* \mathbf{v}_i^{*\top}. \quad (2)$$

Here,  $\sigma_1^* \geq \sigma_2^* \geq \dots \geq \sigma_r^* > 0$  represent the  $r$  nonzero singular values of  $\mathbf{A}^*$ , and  $\mathbf{\Sigma}^* \in \mathbb{R}^{r \times r}$  is a diagonal matrix whose diagonal entries are given by  $\{\sigma_1^*, \dots, \sigma_r^*\}$ . The columns of  $\mathbf{U}^* = [\mathbf{u}_1^*, \dots, \mathbf{u}_r^*] \in \mathbb{R}^{d_1 \times r}$  (resp.  $\mathbf{V}^* = [\mathbf{v}_1^*, \dots, \mathbf{v}_r^*] \in \mathbb{R}^{d_2 \times r}$ ) are orthonormal, which are the top- $r$  left (resp. right) singular vectors of  $\mathbf{A}^*$ . We define and denote the condition number of  $\mathbf{A}^*$  as follows

$$\kappa := \sigma_1^* / \sigma_r^*, \quad (3)$$

and take

$$d := \max \{d_1, d_2\}. \quad (4)$$

**Incoherence.** Further, we impose certain incoherence conditions on the unknown matrix  $\mathbf{A}^*$ , which are commonly adopted in the matrix completion literature (e.g. [CR09, KMO10a, CLC19]).

**Definition 1** (Incoherence parameters). *Define the incoherence parameters  $\mu_0$ ,  $\mu_1$  and  $\mu_2$  as follows*

$$\mu_0 := \frac{d_1 d_2 \max_{1 \leq i \leq d_1, 1 \leq j \leq d_2} |A_{i,j}^*|^2}{\|\mathbf{A}^*\|_F^2}, \quad (5)$$

$$\mu_1 := \frac{d_1}{r} \max_{1 \leq i \leq d_1} \|\mathbf{U}^{*\top} \mathbf{e}_i\|_2^2 \quad \text{and} \quad \mu_2 := \frac{d_2}{r} \max_{1 \leq i \leq d_2} \|\mathbf{V}^{*\top} \mathbf{e}_i\|_2^2, \quad (6)$$

where  $\mathbf{e}_i$  is the  $i$ -th standard basis vector of compatible dimensionality.

Intuitively, when  $\mu_0$ ,  $\mu_1$  and  $\mu_2$  are all small, the energies of the matrices  $\mathbf{A}^*$ ,  $\mathbf{U}^*$  and  $\mathbf{V}^*$  are (nearly) evenly spread out across all entries, rows, and columns. For notational simplicity, we shall set

$$\mu := \max \{\mu_0, \mu_1, \mu_2\}. \quad (7)$$

**Random sampling and random noise.** Suppose that we have only collected noisy observations of the entries of  $\mathbf{A}^*$  over a sampling set  $\Omega \subseteq \{1, \dots, d_1\} \times \{1, \dots, d_2\}$ . Specifically, we observe

$$A_{i,j} = \begin{cases} A_{i,j}^* + N_{i,j}, & (i,j) \in \Omega, \\ 0, & \text{else,} \end{cases} \quad (8)$$

where  $N_{i,j}$  denotes the noise at location  $(i,j)$ . For notational simplicity, we shall write

$$\mathbf{A} = \mathcal{P}_\Omega(\mathbf{A}) = \mathcal{P}_\Omega(\mathbf{A}^*) + \mathcal{P}_\Omega(\mathbf{N}), \quad (9)$$

where  $\mathcal{P}_\Omega$  represents the Euclidean projection onto the subspace of matrices supported on  $\Omega$ . In addition, this paper concentrates on random sampling and random noise as follows.

**Assumption 1** (Random sampling). *Each  $(i,j)$  is included in the sampling set  $\Omega$  independently with probability  $p$ .*

**Assumption 2** (Random noise). *The noise  $N_{i,j}$ 's are independent random variables and satisfy the following conditions: for each  $1 \leq i \leq d_1, 1 \leq j \leq d_2$ ,*

(1) (**Zero mean**)  $\mathbb{E}[N_{i,j}] = 0$ ;

(2) (**Variance**)  $\text{Var}(N_{i,j}) \leq \sigma^2$ ;

(3) (**Magnitude**) *Each  $N_{i,j}$  satisfies either of the following condition:*

(a)  $|N_{i,j}| \leq R$ ;

(b)  $N_{i,j}$  has a symmetric distribution satisfying  $\mathbb{P}\{|N_{i,j}| > R\} \leq c_r d^{-12}$  for some universal constant  $c_r > 0$ .

Here,  $R$  is some quantity obeying

$$\frac{R^2}{\sigma^2} \leq C_r \frac{\min\{p\sqrt{d_1 d_2}, pd_2\}}{\log d} \quad (10)$$

for some universal constant  $C_r > 0$ .

As a remark, Assumption 2 allows the largest possible size  $R$  of each noise component to be substantially larger than its typical size  $\sigma$ . For example, if  $p \asymp 1$ , then  $R$  can be  $\min\{(d_1 d_2)^{1/4}, \sqrt{d_2}\}$  times larger than  $\sigma$  (ignoring any logarithmic factor). In addition, the noise  $N_{i,j}$ 's do not necessarily have identical variance; in fact, our formulation allows us to accommodate the heteroscedasticity of noise (i.e. the scenario where the noise has location-varying variance).

**Goal.** Given incomplete and noisy observations about  $\mathbf{A}^* \in \mathbb{R}^{d_1 \times d_2}$  (cf. (8)), we seek to estimate  $\mathbf{U}^* \in \mathbb{R}^{d_1 \times r}$  modulo some global rotation. We emphasize once again that the aim here is not to estimate the entire matrix. In truth, there are many unbalanced cases with  $d_2 \gg d_1$  such that (1) reliable estimation of  $\mathbf{U}^*$  is feasible, but (2) faithful estimation of the whole matrix  $\mathbf{A}^*$  is information theoretically impossible.

## 2.2 Notation

We denote  $[n] := \{1, \dots, n\}$ . For any matrix  $\mathbf{A} \in \mathbb{R}^{d_1 \times d_2}$ , we use  $\sigma_i(\mathbf{A})$  and  $\lambda_i(\mathbf{A})$  to represent the  $i$ -th largest singular value and the  $i$ -th largest eigenvalue of  $\mathbf{A}$ , respectively. Let  $\mathbf{A}_{i,\cdot}$  and  $\mathbf{A}_{\cdot,j}$  denote respectively the  $i$ -th row and the  $j$ -th column of  $\mathbf{A}$ . Let  $\|\mathbf{A}\|$  (resp.  $\|\mathbf{A}\|_{\text{F}}$ ) represent the spectral norm (resp. the Frobenius norm) of  $\mathbf{A}$ . We also denote by  $\|\mathbf{A}\|_{2,\infty} := \max_{i \in [d_1]} \|\mathbf{A}_{i,\cdot}\|_2$  and  $\|\mathbf{A}\|_{\infty} := \max_{i \in [d_1], j \in [d_2]} |A_{i,j}|$  the  $\ell_{2,\infty}$  norm and the entrywise  $\ell_{\infty}$  norm of  $\mathbf{A}$ , respectively. Similarly, for any tensor  $\mathbf{T}$ , we use  $\|\mathbf{T}\|_{\infty}$  to represent the largest magnitude of the entries of  $\mathbf{T}$ . Moreover, we denote by  $\mathcal{P}_{\text{diag}}$  the projection onto the subspace that vanish outside the diagonal, and define  $\mathcal{P}_{\text{off-diag}}$  such that  $\mathcal{P}_{\text{off-diag}}(\mathbf{A}) := \mathbf{A} - \mathcal{P}_{\text{diag}}(\mathbf{A})$ . Let  $\mathcal{O}^{r \times r}$  stand for the set of  $r \times r$  orthonormal matrices. In addition, we use  $\text{diag}(\mathbf{a})$  to represent a diagonal

matrix whose  $(i, i)$ -th entry is equal to  $a_i$ . Throughout this paper, the notations  $C, C_1, \dots, c, c_1, \dots$  denote absolute positive constants whose values may change from line to line.

Furthermore, for any real-valued functions  $f(d_1, d_2)$  and  $g(d_1, d_2)$ ,  $f(d_1, d_2) \lesssim g(d_1, d_2)$  or  $f(d_1, d_2) = O(g(d_1, d_2))$  mean that  $|f(d_1, d_2)/g(d_1, d_2)| \leq C_1$  for some constant  $C_1 > 0$ ;  $f(d_1, d_2) \gtrsim g(d_1, d_2)$  means that  $|f(d_1, d_2)/g(d_1, d_2)| \geq C_2$  for some universal constant  $C_2 > 0$ ;  $f(d_1, d_2) \asymp g(d_1, d_2)$  means that  $C_1 \leq |f(d_1, d_2)/g(d_1, d_2)| \leq C_2$  for some universal constants  $C_1, C_2 > 0$ ;  $f(d_1, d_2) = o(g(d_1, d_2))$  means that  $f(d_1, d_2)/g(d_1, d_2) \rightarrow 0$  as  $\min\{d_1, d_2\} \rightarrow \infty$ . In addition,  $f(d_1, d_2) \ll g(d_1, d_2)$  (resp.  $f(d_1, d_2) \gg g(d_1, d_2)$ ) means that there exists some sufficiently small (resp. large) constant  $c_1 > 0$  (resp.  $c_2 > 0$ ) such that  $f(d_1, d_2) \leq c_1 g(d_1, d_2)$  (resp.  $f(d_1, d_2) \geq c_2 g(d_1, d_2)$ ) holds true for all sufficiently large  $d_1$  and  $d_2$ .

### 3 Main results

#### 3.1 Algorithm: a spectral method with diagonal deletion

Recall that  $\mathbf{A} = [A_{i,j}]_{1 \leq i \leq d_1, 1 \leq j \leq d_2}$  is the zero-padded data matrix (see (8)). It is easily seen that, under our random sampling model (i.e. Assumption 1),  $p^{-1}\mathbf{A}$  serves as an unbiased estimator of  $\mathbf{A}^*$ . One might thus expect the left singular subspace of  $\mathbf{A}$  to form a reasonably good estimator of the subspace spanned by  $\mathbf{U}^*$ . As it turns out, when  $\mathbf{A}^*$  is a very fat matrix (namely,  $d_2 \gg d_1$ ), this approach might fail to work when the sample complexity is not sufficiently large or when the noise size is not sufficiently small.

This paper adopts an alternative route by resorting to the sample Gram matrix  $\mathbf{A}\mathbf{A}^\top$  (properly rescaled). Straightforward calculation reveals that

$$\frac{1}{p^2} \mathbb{E} [\mathbf{A}\mathbf{A}^\top] = \underbrace{\mathbf{A}^* \mathbf{A}^{*\top} + \left(\frac{1}{p} - 1\right) \mathcal{P}_{\text{diag}} (\mathbf{A}^* \mathbf{A}^{*\top})}_{\text{a diagonal matrix}} + \frac{1}{p} \text{diag} \left( \left[ \sum_{j=1}^{d_2} \text{Var}(N_{i,j}) \right]_{1 \leq i \leq d_1} \right), \quad (11)$$

where  $\text{diag}(\mathbf{a})$  with  $\mathbf{a} \in \mathbb{R}^{d_1}$  represents a diagonal matrix whose  $(i, i)$ -th entry equals  $a_i$ . The identity (11) implies that the diagonal components of  $p^{-2} \mathbb{E} [\mathbf{A}\mathbf{A}^\top]$  are significantly inflated, which might need to be properly suppressed.

In order to remedy the above-mentioned diagonal inflation issue, we adopt a simple strategy that zeros out all diagonal entries; that is, performing the spectral method on the following matrix

$$\mathbf{G} = \frac{1}{p^2} \mathcal{P}_{\text{off-diag}} (\mathbf{A}\mathbf{A}^\top) \quad (12)$$

with  $\mathcal{P}_{\text{off-diag}}(\mathbf{M}) := \mathbf{M} - \mathcal{P}_{\text{diag}}(\mathbf{M})$  denoting projection onto the set of zero-diagonal matrices. This clearly satisfies

$$\mathbb{E} [\mathbf{G}] = \mathcal{P}_{\text{off-diag}} (\mathbf{A}^* \mathbf{A}^{*\top}) = \mathcal{P}_{\text{off-diag}} (\mathbf{U}^* \boldsymbol{\Sigma}^{*2} \mathbf{U}^{*\top}).$$

If the diagonal entries of  $\mathbf{A}^* \mathbf{A}^{*\top}$  are not too large, then one has  $\mathbf{A}^* \mathbf{A}^{*\top} \approx \mathcal{P}_{\text{off-diag}} (\mathbf{A}^* \mathbf{A}^{*\top})$  and, as a result, the rank- $r$  eigen-subspace of  $\mathbf{G}$  might form a reliable estimate of the subspace spanned by  $\mathbf{U}^*$ . The procedure is summarized in Algorithm 1.

---

**Algorithm 1** The spectral method on the diagonal-deleted Gram matrix

---

- 1: **Input:** sampling set  $\Omega$ , observed entries  $\{A_{i,j} \mid (i, j) \in \Omega\}$ , sampling rate  $p$ , rank  $r$ .
- 2: **Compute** the (truncated) rank- $r$  eigen-decomposition  $\mathbf{U}\boldsymbol{\Lambda}\mathbf{U}^\top$  of  $\mathbf{G}$ , where  $\mathbf{U} \in \mathbb{R}^{d_1 \times r}$ ,  $\boldsymbol{\Lambda} \in \mathbb{R}^{r \times r}$ , and

$$\mathbf{G} := \mathcal{P}_{\text{off-diag}} \left( \frac{1}{p^2} \mathbf{A}\mathbf{A}^\top \right). \quad (13)$$

Here,  $\mathbf{A}$  is defined in (9) and  $\mathcal{P}_{\text{off-diag}}(\mathbf{M})$  zeros out the diagonal entries of a matrix  $\mathbf{M}$ .

- 3: **Output**  $\mathbf{U}$  as the subspace estimate, and  $\boldsymbol{\Sigma} = \boldsymbol{\Lambda}^{1/2}$  as the spectrum estimate.
- 

We remark that this is clearly not a new algorithmic idea. In fact, proper handling of the diagonal entries (e.g. diagonal deletion, diagonal reweighting) has already been recommended in several different applications, including bipartite stochastic block models [FP16], covariance estimation [Lou13, Lou14, LW12, EvdG19], tensor completion [MS18], to name just a few.

### 3.2 Theoretical guarantees

In general, one can only hope to estimate  $\mathbf{U}^*$  up to global rotation. With this in mind, we introduce the following rotation matrix

$$\mathbf{R} := \arg \min_{\mathbf{Q} \in \mathcal{O}^{r \times r}} \|\mathbf{U}\mathbf{Q} - \mathbf{U}^*\|_{\text{F}}, \quad (14)$$

where  $\mathcal{O}^{r \times r}$  stands for the set of  $r \times r$  orthonormal matrices. In words,  $\mathbf{R}$  is the global rotation matrix that best aligns  $\mathbf{U}$  and  $\mathbf{U}^*$ . Equipped with this notation, the following theorem delivers upper bounds on the difference between the obtained estimate  $\mathbf{U}$  and the ground truth  $\mathbf{U}^*$ . The proof is postponed to Section 6.

**Theorem 1.** *Assume that the following conditions hold*

$$p \geq c_0 \max \left\{ \frac{\mu \kappa^4 r \log^2 d}{\sqrt{d_1 d_2}}, \frac{\mu \kappa^8 r \log^2 d}{d_2} \right\}, \quad \frac{\sigma}{\sigma_r^*} \leq c_1 \min \left\{ \frac{\sqrt{p}}{\kappa^4 \sqrt{d_1 d_2} \sqrt{\log d}}, \frac{1}{\kappa^3} \sqrt{\frac{p}{d_1 \log d}} \right\} \quad \text{and} \quad r \leq \frac{c_2 d_1}{\mu_1 \kappa^4}, \quad (15)$$

where  $c_0 > 0$  is some sufficiently large constant and  $c_1, c_2 > 0$  are some sufficiently small constants. Then with probability at least  $1 - O(d^{-10})$ , the matrices  $\mathbf{U}$  and  $\mathbf{\Sigma}$  returned by Algorithm 1 satisfy

$$\|\mathbf{U}\mathbf{R} - \mathbf{U}^*\| \lesssim \mathcal{E}_{\text{general}}, \quad (16a)$$

$$\|\mathbf{U}\mathbf{R} - \mathbf{U}^*\|_{2, \infty} \lesssim \kappa^2 \sqrt{\frac{\mu r}{d_1}} \cdot \mathcal{E}_{\text{general}}, \quad (16b)$$

$$\|\mathbf{\Sigma} - \mathbf{\Sigma}^*\| \lesssim \sigma_r^* \cdot \mathcal{E}_{\text{general}}, \quad (16c)$$

where  $\mathbf{R}$  is defined in (14), and

$$\mathcal{E}_{\text{general}} := \underbrace{\frac{\mu \kappa^2 r \log d}{\sqrt{d_1 d_2} p} + \sqrt{\frac{\mu \kappa^4 r \log d}{d_2 p}}}_{\text{missing data effect}} + \underbrace{\frac{\sigma^2 \sqrt{d_1 d_2} \log d}{\sigma_r^{*2} p} + \frac{\sigma \kappa}{\sigma_r^*} \sqrt{\frac{d_1 \log d}{p}}}_{\text{noise effect}} + \underbrace{\frac{\mu_1 \kappa^2 r}{d_1}}_{\text{diagonal deletion effect}}. \quad (17)$$

**Remark 1.** If there is no missing data, namely,  $p = 1$ , then Theorem 1 continues to hold if the first two terms on the right-hand side of (17) are removed.

In a nutshell, Theorem 1 asserts that Algorithm 1 produces reliable estimates of the column subspace of  $\mathbf{A}^*$  — with respect to both the spectral norm and the  $\|\cdot\|_{2, \infty}$  norm — under certain conditions imposed on the sample size and the noise size. For instance, consider the settings where  $\mu, \kappa \asymp 1$  and  $r \ll d_1 \leq d_2$ . Then as long as the following condition holds:

$$p \gtrsim \frac{r \log^2 d}{\sqrt{d_1 d_2}} \quad \text{and} \quad \frac{\sigma^2}{\sigma_r^{*2}} = o\left(\frac{p}{\sqrt{d_1 d_2} \log d}\right), \quad (18)$$

the proposed spectral method achieves consistent estimation with high probability, namely,

$$\min_{\mathbf{Q} \in \mathcal{O}^{r \times r}} \frac{\|\mathbf{U}\mathbf{Q} - \mathbf{U}^*\|}{\|\mathbf{U}^*\|} = o(1), \quad \min_{\mathbf{Q} \in \mathcal{O}^{r \times r}} \frac{\|\mathbf{U}\mathbf{Q} - \mathbf{U}^*\|_{2, \infty}}{\|\mathbf{U}^*\|_{2, \infty}} = o(1) \quad \text{and} \quad \frac{\|\mathbf{\Sigma} - \mathbf{\Sigma}^*\|}{\|\mathbf{\Sigma}^*\|} = o(1). \quad (19)$$

Our upper bound (17) on the spectral norm error contains five terms. The first two terms of (17) are incurred by missing data; the third and the fourth terms of (17) represent the influence of observation noise; and the last term of (17) arises due to the bias caused by diagonal deletion. In particular, the last term is expected to be vanishingly small in the low-rank and incoherent case. Interestingly, both the missing data effect and the noise effect are captured by two different terms, which we shall interpret in what follows. Note that a primary focus of this paper is to demonstrate the feasibility of obtaining a tight control of the  $\ell_{2, \infty}$  statistical error. This is particularly evident for the low-rank, incoherent, and well-conditioned case with  $r, \mu, \kappa = O(1)$ , in which our theory (cf. (16a) and (16b)) reveals that the  $\ell_{2, \infty}$  error can be a factor of  $\sqrt{1/d_1}$  smaller than the spectral norm error. The discussion below focuses on this case (namely,  $r, \mu, \kappa = O(1)$ ), with all logarithmic factors omitted for simplicity of presentation.



	large-noise regime (i.e. $\sigma/\sigma_r^* \gtrsim \sqrt{p/d_2}$ )	small-noise regime (i.e. $\sigma/\sigma_r^* \lesssim \sqrt{p/d_2}$ )
dominant term	$\frac{\sigma^2}{\sigma_r^{*2}} \frac{\sqrt{d_1 d_2}}{p}$	$\frac{\sigma}{\sigma_r^*} \sqrt{\frac{d_1}{p}}$

Table 1: The dominant term of the noise effect in  $\frac{\sigma^2}{\sigma_r^{*2}} \frac{\sqrt{d_1 d_2}}{p} + \frac{\sigma}{\sigma_r^*} \sqrt{\frac{d_1}{p}}$  if  $d_2 \geq d_1$  (omitting logarithmic factors and assuming  $r, \kappa, \mu \asymp 1$ ).

	high-missingness regime (i.e. $p \lesssim 1/d_1$ )	low-missingness regime (i.e. $p \gtrsim 1/d_1$ )
dominant term	$\frac{1}{\sqrt{d_1 d_2 p}}$	$\frac{1}{\sqrt{d_2 p}}$

Table 2: The dominant term of the missing data effect in  $\frac{1}{\sqrt{d_1 d_2 p}} + \frac{1}{\sqrt{d_2 p}}$  if  $d_2 \geq d_1$  (omitting logarithmic factors and assuming  $r, \kappa, \mu \asymp 1$ ).

- Let us first examine the influence of observation noise, which reads

$$\frac{\sigma^2}{\sigma_r^{*2}} \frac{\sqrt{d_1 d_2}}{p} + \frac{\sigma}{\sigma_r^*} \sqrt{\frac{d_1}{p}}. \quad (20)$$

This contains a quadratic term as well as a linear term w.r.t.  $\sigma/\sigma_r^*$ . To interpret this, consider, for example, the case without missing data (i.e.  $p = 1$ ) and decompose

$$\mathbf{A}\mathbf{A}^\top = \mathbf{A}^* \mathbf{A}^{*\top} + \underbrace{\mathbf{A}^* \mathbf{N}^\top + \mathbf{N} \mathbf{A}^{*\top}}_{\text{linear perturbation}} + \underbrace{\mathbf{N} \mathbf{N}^\top}_{\text{quadratic perturbation}},$$

which clearly explains why eigenspace perturbation bounds depend both linearly and quadratically on the noise magnitudes. In general, the quadratic term  $\frac{\sigma^2}{\sigma_r^{*2}} \frac{\sqrt{d_1 d_2}}{p}$  is dominant when the signal-to-noise ratio (SNR) is not large enough; as the noise decreases to a sufficiently low level, the linear term starts to enter the picture. See Table 1 for a more precise summary. As we shall demonstrate momentarily, the terms (20) match the minimax limits (up to some logarithmic factor), meaning that it is generally impossible to get rid of either the linear term or the quadratic term.

- Next, we examine the influence of missing data and assume  $\sigma = 0$  to simplify the discussion. If we view  $\mathbf{N}_{\text{missing}} = \frac{1}{p} \mathbf{A} - \mathbf{A}^*$  as a zero-mean perturbation matrix, then one can write

$$\frac{1}{p^2} \mathbf{A}\mathbf{A}^\top = \mathbf{A}^* \mathbf{A}^{*\top} + \underbrace{\mathbf{A}^* \mathbf{N}_{\text{missing}}^\top + \mathbf{N}_{\text{missing}} \mathbf{A}^{*\top}}_{\text{linear perturbation}} + \underbrace{\mathbf{N}_{\text{missing}} \mathbf{N}_{\text{missing}}^\top}_{\text{quadratic perturbation}}.$$

Similar to the above noisy case with  $p = 1$ , this decomposition explains why the influence of missing data contains two terms as well (see Table 2)

$$\underbrace{\frac{1}{\sqrt{d_1 d_2 p}}}_{\text{quadratic term in } 1/\sqrt{p}} + \underbrace{\frac{1}{\sqrt{d_2 p}}}_{\text{linear term in } 1/\sqrt{p}}.$$

**Comparison with prior results.** To demonstrate the effectiveness of our theory, we take a moment to compare them with several prior results. Once again, the discussion below focuses on the case with  $\max\{\mu, \kappa, r\} \asymp 1$ . To be fair, it is worth noting that most papers discussed below either have different objectives (e.g. aiming at matrix estimation rather than subspace estimation [KMO10a, CT10, CCF+19]), or work with different (and possibly more general) model assumptions (e.g. square matrices [AFWZ17], or heteroskedastic noise [ZCW18]). Our purpose here is not to argue that our results are always stronger than the previous ones, but rather to point out the insufficiency of prior theory when directly applied to some basic settings.



- To begin with, we compare our spectral norm bound with that required for matrix completion [KMO10a, CT10, AFWZ17, CLC19, CCF+19] in the noise-free case (i.e.  $\sigma = 0$ ), in order to show how much saving can be harvested when we move from matrix estimation to subspace estimation. Suppose that  $d_2 \geq d_1$ . As is well known, for both spectral and optimization-based methods, the sample complexities required for faithful matrix completion need to satisfy  $pd_1d_2 \gtrsim d_2 \text{poly log } d$ . In comparison, faithful estimation of the column subspace becomes feasible under the sample size  $pd_1d_2 \gtrsim \sqrt{d_1d_2} \text{poly log } d$ , which can be much lower than that required for matrix completion (i.e. by a factor of  $\sqrt{d_2/d_1}$ ). Further, we compare our  $\|\cdot\|_{2,\infty}$  bound with the theory derived in [AFWZ17] when  $d_2 \gtrsim d_1 \log^2 d$ . The theory in [AFWZ17, Theorem 3.4] requires the sample size and the noise level to satisfy  $p \gtrsim d_1^{-1} \log d$  and  $\sigma/\sigma_r^* \lesssim \sqrt{\frac{p}{d_2 \log d}}$ , both of which are more stringent requirements than ours (namely,  $p \gtrsim \frac{\log^2 d}{\sqrt{d_1d_2}}$  and  $\sigma/\sigma_r^* \lesssim \frac{\sqrt{p}}{\sqrt[3]{d_1d_2} \sqrt{\log d}}$ ). Again, this arises primarily because [AFWZ17] seeks to estimate the whole matrix as opposed to its column subspace.
- We then compare our results with [MS18], which studies a diagonal-rescaling algorithm for the noise-free case (i.e.  $\sigma = 0$ ). Combining [MS18, Theorem 6.2] with the standard Davis-Kahan matrix perturbation theory, we can easily see that their spectral norm bound for subspace estimation reads

$$\frac{\text{poly log } d}{\sqrt{d_1d_2}p} + \frac{\text{poly log } d}{\sqrt{d_2p}}.$$

This coincides with our bound except for the last term of (17) (due to the bias incurred by diagonal deletion). In comparison, our theory offers additional  $\ell_{2,\infty}$  statistical guarantees and covers the noisy case, thus strengthening the theory presented in [MS18].

- Additionally, we compare our spectral norm bound with the results derived in [ZCW18]. Consider the noiseless case where  $\sigma = 0$ . It is proven in [ZCW18, Theorem 6] (see also the remark that follows) that: if the sample size satisfies  $pd_1d_2 \gtrsim \max\{d_1^{1/3}d_2^{2/3}, d_1\} \text{poly log } d$ , then the HeteroPCA estimator is consistent in estimating the column subspace (namely, achieving a relative  $\ell_2$  estimation error not exceeding  $o(1)$ ). In comparison, our theory claims that Algorithm 1 is guaranteed to yield consistent column subspace estimation as long as the sample size obeys  $pd_1d_2 \gtrsim \sqrt{d_1d_2} \text{poly log } d$ . Consequently, if we omit logarithmic terms, then our sample complexity improves upon the theoretical support of HeteroPCA by a factor of  $(d_2/d_1)^{1/6}$  if  $d_2 \geq d_1$ . Once again, the comparison here focuses on the effect of the aspect ratio  $d_2/d_1$ , without accounting for the influence of other parameters like  $\mu, \kappa, r$ .

**SVD applied directly to  $\mathbf{A}$ ?** Finally, another natural spectral method that comes immediately into mind is to compute the rank- $r$  SVD of  $\mathbf{A}$ , and return the matrix containing the  $r$  left singular vectors as the column subspace estimate. The  $\ell_2$  risk analysis of this approach is typically based on classical matrix perturbation theory like Wedin’s theorem [Wed73]. We caution, however, that this approach becomes highly sub-optimal when the aspect ratio  $d_2/d_1$  grows. Take the case with Gaussian noise and no missing data (i.e.  $p = 1$ ) for example: in order for Wedin’s theorem to be applicable, a basic requirement is  $\|\mathbf{N}\| < \sigma_r^*$ , which translates to the condition

$$\frac{\sigma}{\sigma_r^*} \lesssim \frac{1}{\sqrt{d_2}} \tag{21}$$

since  $\|\mathbf{N}\| \lesssim \sigma\sqrt{d_2}$ . In comparison, our theory covers the range  $\frac{\sigma}{\sigma_r^*} \lesssim \frac{1}{(d_1d_2)^{1/4}}$  (modulo some log factor), which allows the noise level to be  $(d_2/d_1)^{1/4}$  times larger than the upper bound (21) derived for the above SVD approach. The sub-optimality of this approach can also be easily seen from numerical experiments as well; see Section 4 for details.

### 3.3 Minimax lower bounds

It is natural to wonder whether our theoretical guarantees are tight, and whether there are other estimators that can potentially improve the performance of Algorithm 1. To answer these questions, we develop the following minimax lower bounds under Gaussian noise; the proof is deferred to Appendix D.1.

**Theorem 2.** Suppose  $1 \leq r \leq d_1/2$ , and  $N_{i,j} \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, \sigma^2)$ . Define

$$\mathcal{M}^* := \{ \mathbf{B} \in \mathbb{R}^{d_1 \times d_2} \mid \text{rank}(\mathbf{B}) = r, \sigma_r(\mathbf{B}) \in [0.9\sigma_r^*, 1.1\sigma_r^*] \}.$$

Denote by  $\mathbf{U}(\mathbf{B}) \in \mathbb{R}^{d_1 \times r}$  the matrix containing the  $r$  left singular vectors of  $\mathbf{B}$ . Then there exists some universal constant  $c_{\text{lb}} > 0$  such that

$$\inf_{\widehat{\mathbf{U}}} \sup_{\mathbf{A}^* \in \mathcal{M}^*} \mathbb{E} \left[ \min_{\mathbf{R} \in \mathcal{O}^{r \times r}} \|\widehat{\mathbf{U}}\mathbf{R} - \mathbf{U}(\mathbf{A}^*)\| \right] \geq c_{\text{lb}} \min \left\{ \frac{\sigma^2 \sqrt{d_1 d_2}}{\sigma_r^{*2} p} + \frac{\sigma}{\sigma_r^*} \sqrt{\frac{d_1}{p}}, 1 \right\}, \quad (22a)$$

$$\inf_{\widehat{\mathbf{U}}} \sup_{\mathbf{A}^* \in \mathcal{M}^*} \mathbb{E} \left[ \min_{\mathbf{R} \in \mathcal{O}^{r \times r}} \|\widehat{\mathbf{U}}\mathbf{R} - \mathbf{U}(\mathbf{A}^*)\|_{2,\infty} \right] \geq c_{\text{lb}} \min \left\{ \frac{\sigma^2 \sqrt{d_1 d_2}}{\sigma_r^{*2} p} + \frac{\sigma}{\sigma_r^*} \sqrt{\frac{d_1}{p}}, 1 \right\} \frac{1}{\sqrt{d_1}}, \quad (22b)$$

where the infimum is taken over all estimators for  $\mathbf{U}(\mathbf{A}^*)$  based on the observation  $\mathcal{P}_\Omega(\mathbf{A}^* + \mathbf{N})$ .

If we again consider the case where  $r, \kappa, \mu \asymp 1$ , then the above lower bounds (22) match the noise effect terms in Theorem 1 (or equivalently, (20)) up to logarithmic factors. This unveils a fundamental reason why the linear and the quadratic terms in (20) are both essential in determining the estimation risk.

Another information-theoretic limit that concerns only the influence of subsampling is supplied as follows; the proof is postponed to Appendix D.2.

**Theorem 3.** Suppose  $d_1 \leq d_2$  and  $p < \frac{1-\epsilon}{\sqrt{d_1 d_2}}$  for any small constant  $0 < \epsilon < 1$ . With probability approaching one, there exist unit vectors  $\mathbf{u}^*, \tilde{\mathbf{u}}^* \in \mathbb{R}^{d_1}$  and  $\mathbf{v}^*, \tilde{\mathbf{v}}^* \in \mathbb{R}^{d_2}$  such that

- $\min \|\mathbf{u}^* \pm \tilde{\mathbf{u}}^*\|_2 \asymp 1$  and  $\|\mathbf{u}^* \mathbf{v}^{*\top} - \tilde{\mathbf{u}}^* \tilde{\mathbf{v}}^{*\top}\|_{\text{F}} \asymp 1$ ;
- one cannot distinguish  $\mathbf{u}^* \mathbf{v}^{*\top}$  and  $\tilde{\mathbf{u}}^* \tilde{\mathbf{v}}^{*\top}$  from the entries in  $\Omega$ , i.e.  $\mathcal{P}_\Omega(\mathbf{u}^* \mathbf{v}^{*\top}) = \mathcal{P}_\Omega(\tilde{\mathbf{u}}^* \tilde{\mathbf{v}}^{*\top})$ .

In words, Theorem 3 asserts that one cannot hope to achieve consistent subspace estimation (in the sense of (19)) at all, as soon as the sampling rate  $p$  falls below the threshold  $1/\sqrt{d_1 d_2}$ . Putting Theorem 2-3 together reveals that: consistent estimation can by no means be guaranteed unless

$$p \gtrsim \frac{1}{\sqrt{d_1 d_2}} \quad \text{and} \quad \frac{\sigma^2}{\sigma_r^{*2}} \lesssim \frac{p}{\sqrt{d_1 d_2}}, \quad (23)$$

which agrees with our theoretical guarantees (18) (up to some logarithmic term). As a result, our minimax lower bounds confirm the near optimality of Algorithm 1 in enabling consistent estimation.

On the other hand, it is widely recognized that spectral methods are typically unable to achieve exact recovery or optimal estimation accuracy in the presence of missing data, even in the balanced case with  $d_1 = d_2$ . For instance, if there is no noise, namely  $\sigma = 0$ , the spectral methods fail to achieve perfect recovery as long as  $p < 1$  (basically the first two terms of (17) do not vanish) [KMO10b], whereas exact recovery might sometimes be feasible with the aid of optimization-based approaches [CR09]. More often than not, spectral methods are employed to produce a rough initial estimate that outperforms the random guess, which can then be refined via other algorithms (e.g. nonconvex optimization algorithms like gradient descent and alternating minimization [KMO10b, SL16, MWCC17, CLPC19]).

## 4 Consequences for concrete applications

We showcase the consequence of Theorem 1 in three concrete applications previously introduced in Section 1 in relatively simple settings. Rather than striving for full generality, our purpose is to highlight the broad applicability of our main results.

### 4.1 Noisy tensor completion

**Problem settings.** We begin by considering the problem of symmetric tensor completion. Consider an unknown order-3 tensor

$$\mathbf{T}^* = \sum_{s=1}^r \mathbf{w}_s^* \otimes \mathbf{w}_s^* \otimes \mathbf{w}_s^* := \sum_{s=1}^r (\mathbf{w}_s^*)^{\otimes 3} \in \mathbb{R}^{d \times d \times d},$$

with canonical polyadic (CP) rank  $r$ . The goal is to estimate the subspace spanned by  $\{\mathbf{w}_s^*\}_{s=1}^r$ , on the basis of the noisy tensor  $\mathbf{T} = [T_{i,j,k}]_{1 \leq i,j,k \leq d}$  obeying

$$T_{i,j,k} = \begin{cases} T_{i,j,k}^* + N_{i,j,k}, & (i,j,k) \in \Omega, \\ 0, & (i,j,k) \notin \Omega. \end{cases} \quad (24)$$

Here,  $T_{i,j,k}$  is the observed entry in location  $(i,j,k)$ ,  $N_{i,j,k}$  is the associated independent random noise satisfying Assumption 2, and  $\Omega \subseteq [d]^3$  stands for a sampling set obtained via uniform random sampling with sampling rate  $p$  (namely, each entry is observed independently with probability  $p$ ).

**Algorithm.** Observe that the mode-1 matricization of  $\mathbf{T}^*$  is given by<sup>2</sup>

$$\mathbf{A}^* = \sum_{s=1}^r \mathbf{w}_s^* (\mathbf{w}_s^* \otimes \mathbf{w}_s^*)^\top \in \mathbb{R}^{d \times d^2}, \quad (25)$$

indicating that the column subspace of  $\mathbf{A}^*$  is essentially the subspace spanned by the tensor factors  $\{\mathbf{w}_s^*\}_{s=1}^r$ . Therefore, if we denote by  $\mathbf{A} \in \mathbb{R}^{d \times d^2}$  the mode-1 matricization of  $\mathbf{T}$ , then we can invoke our general spectral method to estimate the column subspace of  $\mathbf{A}^*$  given  $\mathbf{A}$ . This procedure is summarized in Algorithm 2.

---

**Algorithm 2** The spectral method for tensor completion

---

- 1: **Input:** sampling set  $\Omega$ , observed entries  $\{T_{i,j,k} \mid (i,j,k) \in \Omega\}$ , sampling rate  $p$ , CP-rank  $r$ .
  - 2: Let  $\mathbf{A} \in \mathbb{R}^{d \times d^2}$  be the mode-1 matricization of the observed tensor  $\mathbf{T}$  (see (24)), namely, set  $A_{i,(j-1)d+k} = T_{i,j,k}$  for each  $(i,j,k) \in [d]^3$ , and employ  $\mathbf{A}$  as the input of Algorithm 1.
  - 3: **Output**  $\mathbf{U} \in \mathbb{R}^{d \times r}$  returned by Algorithm 1 as the subspace estimate.
- 

**Theoretical guarantees.** In order to provide theoretical support for Algorithm 2, we introduce a few more notation. First, we introduce the following quantities

$$\kappa_{\text{tc}} := \lambda_{\max}^* / \lambda_{\min}^*, \quad \lambda_{\min}^* := \min_{1 \leq i \leq r} \|\mathbf{w}_i^*\|_2^3, \quad \lambda_{\max}^* := \max_{1 \leq i \leq r} \|\mathbf{w}_i^*\|_2^3. \quad (26)$$

Note that  $\|\mathbf{w}_i^*\|_2^3$  is precisely the Frobenius norm of the rank-1 tensor  $\mathbf{w}_i^{*\otimes 3}$  — the  $i$ -th tensor component. Informally,  $\kappa_{\text{tc}}$  captures the condition number of the unknown tensor. Additionally, similar to matrix completion, we introduce the following incoherence definitions that enable efficient tensor completion:

**Definition 2** (Incoherence). *Define the incoherence parameters  $\mu_3, \mu_4, \mu_5$  for the tensor  $\mathbf{T}^*$  and its tensor factors  $\{\mathbf{w}_s^*\}_{s=1}^r$  as follows:*

$$\mu_3 := \frac{d^3 \|\mathbf{T}^*\|_\infty^2}{\|\mathbf{T}^*\|_{\text{F}}^2}, \quad \mu_4 := \max_{1 \leq i \leq r} \frac{d \|\mathbf{w}_i^*\|_\infty^2}{\|\mathbf{w}_i^*\|_2^2}, \quad \mu_5 := \max_{1 \leq i \neq j \leq r} \frac{d \langle \mathbf{w}_i^*, \mathbf{w}_j^* \rangle^2}{\|\mathbf{w}_i^*\|_2 \|\mathbf{w}_j^*\|_2} \quad (27a)$$

For notational convenience, we also set

$$\mu_{\text{tc}} := \max \{\mu_3, \mu_4^2\}. \quad (28)$$

Given that the tensor factors  $\{\mathbf{w}_s^*\}_{1 \leq s \leq r}$  are in general not orthogonal to each other, we introduce the following orthonormal matrix  $\mathbf{U}^* \in \mathbb{R}^{d \times r}$  to represent the subspace spanned by  $\{\mathbf{w}_s^*\}_{1 \leq s \leq r}$ :

$$\mathbf{U}^* := \mathbf{W}^* (\mathbf{W}^{*\top} \mathbf{W}^*)^{-1/2}, \quad \mathbf{W}^* := [\mathbf{w}_1^*, \dots, \mathbf{w}_r^*] \in \mathbb{R}^{d \times r}. \quad (29)$$

Note that the particular choice of  $\mathbf{U}^*$  in (29) is not pivotal, and can be replaced by any  $d_1 \times r$  orthonormal matrix that spans the same column space as  $\mathbf{W}^*$ . With these in place, we are now ready to quantify the estimation error of this spectral algorithm. The proof is deferred to Appendix A.1.

---

<sup>2</sup>We let  $\mathbf{a} \otimes \mathbf{b} := \begin{bmatrix} a_1 b \\ \vdots \\ a_d b \end{bmatrix}$  represent a  $d^2$ -dimensional vector.

**Corollary 1** (Symmetric tensor completion). *Consider the above tensor completion model. There exist some universal constants  $c_0, c_1, c_2 > 0$  such that if*

$$p \geq c_0 \max \left\{ \frac{\mu_{\text{tc}} \kappa_{\text{tc}}^4 r \log^2 d}{d^{3/2}}, \frac{\mu_{\text{tc}} \kappa_{\text{tc}}^8 r \log^2 d}{d^2} \right\}, \quad \frac{\sigma}{\lambda_{\min}^*} \leq c_1 \min \left\{ \frac{\sqrt{p}}{\kappa_{\text{tc}} d^{3/4} \sqrt{\log d}}, \frac{1}{\kappa_{\text{tc}}^3} \sqrt{\frac{p}{d \log d}} \right\}$$

$$\text{and } r \leq c_2 \min \left\{ \frac{d}{\kappa_{\text{tc}}^4 \mu_4}, \frac{1}{\kappa_{\text{tc}}^2} \sqrt{\frac{d}{\mu_5}} \right\},$$

then with probability exceeding  $1 - O(d^{-10})$ , Algorithm 2 yields

$$\|\mathbf{UR} - \mathbf{U}^*\| \lesssim \mathcal{E}_{\text{tc}}, \quad (30a)$$

$$\|\mathbf{UR} - \mathbf{U}^*\|_{2,\infty} \lesssim \kappa_{\text{tc}}^2 \sqrt{\frac{\mu_{\text{tc}} r}{d}} \cdot \mathcal{E}_{\text{tc}}, \quad (30b)$$

where  $\mathbf{R} := \arg \min_{\mathbf{Q} \in \mathcal{O}^{r \times r}} \|\mathbf{UQ} - \mathbf{U}^*\|_{\text{F}}$  and

$$\mathcal{E}_{\text{tc}} := \frac{\mu_{\text{tc}} \kappa_{\text{tc}}^2 r \log d}{d^{3/2} p} + \sqrt{\frac{\mu_{\text{tc}} \kappa_{\text{tc}}^4 r \log d}{d^2 p}} + \frac{\sigma^2}{\lambda_{\min}^{*2}} \frac{d^{3/2} \log d}{p} + \frac{\sigma \kappa_{\text{tc}}}{\lambda_{\min}^*} \sqrt{\frac{d \log d}{p}} + \frac{\mu_4 \kappa_{\text{tc}}^2 r}{d}. \quad (31)$$

As discussed in several related work (e.g. [XYZ17, JO14, MS18, XY17, CLPC19]), once we obtain reliable estimates of the subspace spanned by the tensor factors, we can further exploit the tensor structure to estimate the unknown tensor. Indeed, in many tensor completion algorithms, subspace estimation serves as a crucial initial step for tensor completion. Moreover, while prior works only provide  $\ell_2$  estimation error bounds, Corollary 1 further delivers  $\ell_{2,\infty}$  statistical guarantees, which reflect a stronger sense of statistical accuracy. We note that [XZ19, Theorem 4] derived an appealing  $\ell_{2,\infty}$  statistical error bound for an algorithm called HOSVD, under the tensor de-noising setting. In comparison to the Gaussian noise considered therein, our results accommodate the case with missing data and possibly spiky noise.

**Implications.** In what follows, we discuss the sample size and the signal-to-noise (SNR) required for achieving consistent tensor estimation (namely, obtaining an  $o(1)$  relative estimation error). For convenience of presentation, we again focus on the low-rank, incoherent, and well-conditioned case with  $r, \mu, \kappa_{\text{tc}} \asymp 1$ . In this case, our results in Corollary 1 indicate that

$$\min_{\mathbf{Q} \in \mathcal{O}^{r \times r}} \|\mathbf{UQ} - \mathbf{U}^*\| = o(1) \quad \text{and} \quad \min_{\mathbf{Q} \in \mathcal{O}^{r \times r}} \|\mathbf{UQ} - \mathbf{U}^*\|_{2,\infty} = o(1/\sqrt{d}) \quad (32)$$

with high probability, provided that the sample size and the noise satisfy

$$p \gtrsim \frac{\log^2 d}{d^{3/2}} \quad \text{and} \quad \frac{\sigma}{\lambda_{\min}^*} = o\left(\sqrt{\frac{p}{d^{3/2} \log d}}\right). \quad (33)$$

Several remarks are in order.

- *Sample complexity.* It is widely conjectured that the sample complexity  $pd^3$  required to reconstruct a order-3 tensor in polynomial time — even in the noiseless case — is at least  $d^{3/2}$  (or equivalently,  $p \gtrsim 1/d^{3/2}$ ) [BM16, XYZ17, MS18]. Therefore, our theory reveals that spectral methods achieve consistent estimation (w.r.t. both  $\|\cdot\|$  and  $\|\cdot\|_{2,\infty}$ ), as long as the sample size is slightly above the (conjectured) computational limit. Moreover, it is easily seen that the bias incurred by deleting the diagonal is much smaller than the error due to missing data, which justifies the rationale that diagonal deletion does not harm the performance by much.
- *Noise size.* We now comment on the noise size requirement. It is easily seen that the maximum magnitude of the entries of  $\mathbf{T}^*$  in this case is  $\|\mathbf{T}^*\|_{\infty} \asymp \lambda_{\max}^*/d^{3/2}$ . As a result, the noise size condition in (33) is equivalent to

$$\frac{\sigma}{\|\mathbf{T}^*\|_{\infty}} \lesssim \sqrt{pd^{3/2}}.$$

Taken together with our sample size requirement  $p \gtrsim \frac{\log^2 d}{d^{3/2}}$ , this condition allows the noise magnitude in each observed entry to significantly exceed the size of the corresponding entry, which covers a broad range of scenarios of practical interest. In addition, in the fully-observed case (i.e.  $p = 1$ ) with i.i.d. Gaussian noise, the authors in [ZX18] showed that the noise size condition (33) — up to some log factor — is necessary for any polynomial-time algorithm to achieve consistent estimation, provided that a certain hypergraphic planted clique conjecture holds.

Finally, we remark that in the fully-observed case (i.e.  $p = 1$ ) with i.i.d. Gaussian noise, it can be seen from [ZX18, Theorem 1] that (30a) is suboptimal; in fact, the minimax risk consists only of the linear term in  $\sigma$  (namely,  $\frac{\sigma}{\lambda_{\min}^*} \sqrt{d}$ , if we omit log factors and assume  $r, \mu_{\text{tc}}, \mu_{\text{S}}, \kappa \asymp 1$ ). This is a typical drawback of the spectral method for tensor estimation, since it falls short of exploiting the low-complexity structure in the row subspace. However, the spectral estimate offers a reasonably good initial estimate for this problem, and one can often employ optimization-based iterative refinement paradigms (like gradient descent [CLPC19]) to obtain minimax optimal estimates.

## 4.2 Covariance estimation / principal component analysis with missing data

**Model and algorithm.** Next, we study covariance estimation with missing data, as previously introduced in Section 1. For concreteness, imagine a set of independent sample vectors obeying

$$\mathbf{x}_i = \mathbf{B}^* \mathbf{f}_i^* + \boldsymbol{\eta}_i \in \mathbb{R}^d, \quad \mathbf{f}_i^* \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(\mathbf{0}, \mathbf{I}_r), \quad 1 \leq i \leq n.$$

Here  $\mathbf{B}^* \in \mathbb{R}^{d \times r}$  encodes the  $r$ -dimensional principal subspace underlying the data (sometimes referred to as the factor loading matrix in factor models [LM62, FWZZ18]),  $\mathbf{f}_i^* \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_r)$  represents some random coefficients, and the noise vector  $\boldsymbol{\eta}_i = [\eta_{i,j}]_{1 \leq j \leq d}$  consists of independent Gaussian components<sup>3</sup> obeying

$$\mathbb{E}[\eta_{i,j}] = 0 \quad \text{and} \quad \text{Var}[\eta_{i,j}] \leq \sigma^2.$$

What we observe is a partial set of entries of  $\mathbf{x}_i = [x_{i,j}]_{1 \leq j \leq d}$ , namely, we only observe  $x_{i,j}$  for any  $(i, j) \in \Omega$ , where  $\Omega$  is obtained by random sampling with rate  $p$ . The goal is to estimate the subspace spanned by  $\mathbf{B}^*$ , or even  $\mathbf{B}^* \mathbf{B}^{*\top}$ .

If we write  $\mathbf{F}^* = [\mathbf{f}_1^*, \dots, \mathbf{f}_n^*] \in \mathbb{R}^{r \times n}$  and  $\mathbf{N} = [\boldsymbol{\eta}_1, \dots, \boldsymbol{\eta}_n] \in \mathbb{R}^{d \times n}$ , then it boils down to estimating the column space of  $\mathbf{A}^* := \mathbf{B}^* \mathbf{F}^*$  from the data  $\mathcal{P}_\Omega(\mathbf{X}) = \mathcal{P}_\Omega(\mathbf{B}^* \mathbf{F}^* + \mathbf{N})$ . Our spectral method for covariance estimation is summarized in Algorithm 3.

---

### Algorithm 3 The spectral method for covariance estimation

---

- 1: **Input:** sampling set  $\Omega$ , observed entries  $\{X_{i,j} \mid (i, j) \in \Omega\}$ , sampling rate  $p$ , rank  $r$ .
  - 2: Let  $\mathbf{A} = \mathcal{P}_\Omega(\mathbf{X}) \in \mathbb{R}^{d \times n}$  with  $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_n]$ , and use  $\mathbf{A}$  as the input of Algorithm 1. Let  $\mathbf{U} \in \mathbb{R}^{d \times r}$  and  $\boldsymbol{\Sigma} \in \mathbb{R}^{r \times r}$  be the estimates returned by Algorithm 1, and set  $\mathbf{B} := \frac{1}{\sqrt{n}} \mathbf{U} \boldsymbol{\Sigma}$ .
  - 3: **Output**  $\mathbf{U}$  as the subspace estimate and  $\mathbf{S} := \mathbf{B} \mathbf{B}^\top$  as the covariance estimate.
- 

**Theoretical guarantees.** In order to present our theory, we make a few more definitions. Without loss of generality, we shall define

$$\mathbf{S}^* := \mathbf{B}^* \mathbf{B}^{*\top} = \mathbf{U}^* \boldsymbol{\Lambda}^* \mathbf{U}^{*\top} \quad \text{and} \quad \mathbf{B}^* = \mathbf{U}^* \boldsymbol{\Lambda}^{*1/2}, \quad (34)$$

where  $\mathbf{U}^* \in \mathbb{R}^{d \times r}$  consists of orthonormal columns, and  $\boldsymbol{\Lambda}^* = \text{diag}(\lambda_1^*, \dots, \lambda_r^*) \in \mathbb{R}^{r \times r}$  is a diagonal matrix with  $\lambda_1^* \geq \dots \geq \lambda_r^* \geq 0$ . We also define the condition number and the incoherence parameter as

$$\kappa_{\text{ce}} := \lambda_1^* / \lambda_r^* \quad \text{and} \quad \mu_{\text{ce}} := \frac{d}{r} \|\mathbf{U}^*\|_{2, \infty}^2. \quad (35)$$

We are now positioned to derive statistical estimation guarantees using our general theorem. The following result is a consequence of Theorem 1; the proof is postponed to Appendix A.2.

---

<sup>3</sup>Here, we assume  $\mathbf{f}_i^*$  and  $\boldsymbol{\eta}_i$  to be Gaussian for simplicity of presentation. The results in this subsection continue to hold if they are sub-Gaussian random vectors.

**Corollary 2** (Covariance estimation). *Consider the above covariance estimation model with missing data. There exist universal constants  $c_0, c_1 > 0$  such that if  $r \leq c_1 \frac{d}{\mu_{ce} \kappa_{ce}^2}$  and*

$$n \geq c_0 \max \left\{ \frac{\mu_{ce}^2 \kappa_{ce}^6 r^2 \log^6(n+d)}{dp^2}, \frac{\mu_{ce} \kappa_{ce}^5 r \log^3(n+d)}{p}, \frac{\sigma^4 \kappa_{ce}^2 d \log^2(n+d)}{\lambda_r^{*2} p^2}, \frac{\sigma^2 \kappa_{ce}^3 d \log(n+d)}{\lambda_r^* p} \right\}, \quad (36)$$

then with probability exceeding  $1 - O((n+d)^{-10})$ , Algorithm 3 yields

$$\|\mathbf{UR} - \mathbf{U}^*\| \lesssim \mathcal{E}_{ce}, \quad (37a)$$

$$\|\mathbf{UR} - \mathbf{U}^*\|_{2,\infty} \lesssim \kappa_{ce}^{3/2} \sqrt{\frac{\mu_{ce} r \log(n+d)}{d}} \cdot \mathcal{E}_{ce}, \quad (37b)$$

$$\|\mathbf{S} - \mathbf{S}^*\| \lesssim \kappa_{ce} \lambda_1^* \cdot \mathcal{E}_{ce}, \quad (37c)$$

$$\|\mathbf{S} - \mathbf{S}^*\|_\infty \lesssim \frac{\kappa_{ce} \mu_{ce} r \log(n+d)}{d} \lambda_1^* \cdot \mathcal{E}_{ce}. \quad (37d)$$

Here,  $\mathbf{R} := \arg \min_{\mathbf{Q} \in \mathcal{O}^{r \times r}} \|\mathbf{UQ} - \mathbf{U}^*\|_F$  and

$$\mathcal{E}_{ce} := \frac{\mu_{ce} \kappa_{ce}^2 r \log^2(n+d)}{\sqrt{dn} p} + \sqrt{\frac{\mu_{ce} \kappa_{ce}^3 r \log^2(n+d)}{np}} + \frac{\sigma^2}{\lambda_r^*} \sqrt{\frac{d}{n}} \frac{\log(n+d)}{p} + \frac{\sigma}{\sqrt{\lambda_r^*}} \sqrt{\frac{d}{n}} \sqrt{\frac{\kappa_{ce} \log(n+d)}{p}} + \frac{\mu_{ce} \kappa_{ce} r}{d}. \quad (38)$$

**Remark 2.** In the full observation case (i.e.  $p = 1$ ), the first two terms on the right-hand side of (38) can be removed.

**Remark 3.** We make note of a scaling issue that one shall bear in mind when comparing this result with our main theorem. In the settings of Theorem 1, the singular values  $\{\sigma_i^*\}_{i=1}^r$  of the truth  $\mathbf{A}^*$  do not change as the column dimension  $d_2$  grows. In contrast, in the settings of Corollary 2, the singular values of the sample covariance matrix keep growing as we collect more sample vectors, which is equivalent to saying that these singular values scale with the column dimension.

**Discussion.** To facilitate interpretation, let us again focus on the case where  $\mu_{ce}, \kappa_{ce} \asymp 1$ . Corollary 2 demonstrates that for any given sampling rate  $p$ , we can achieve consistent estimation<sup>4</sup> as long as the number  $n$  of samples satisfies

$$n \gtrsim \max \left\{ \frac{r^2}{dp^2}, \frac{r}{p}, \frac{\sigma^4 d}{\lambda_r^{*2} p^2}, \frac{\sigma^2 d}{\lambda_r^* p} \right\} \text{poly log } d. \quad (39)$$

Throughout this subsection, the sample size refers to  $n$  — the number of sample vectors  $\{\mathbf{x}_i\}_{1 \leq i \leq n}$  available.

Next, we compare our  $\ell_{2,\infty}$  bounds with several prior work for the case with  $r \asymp 1$ . We emphasize again that the foci and model assumptions of these prior papers might be quite different from ours (e.g. [ZWS19] is able to accommodate inhomogeneous sampling patterns), and the advantages of our results discussed below are restricted to the settings considered in this paper. For simplicity, we ignore all log factors.

- Suppose that  $\sigma = 1$ . In this setting, [ZWS19, Theorem 4] demonstrates that if

$$n \gtrsim \max \left\{ \frac{1}{p^2}, \frac{d^2}{\lambda_r^{*2} p^2}, \frac{d}{\lambda_r^* p^2} \right\} \text{poly log } d,$$

then with high probability one has

$$\min_{\mathbf{Q} \in \mathcal{O}^{r \times r}} \|\mathbf{UQ} - \mathbf{U}^*\|_{2,\infty} \lesssim \frac{1}{p\sqrt{n}} \left( \frac{1}{\sqrt{\lambda_r^*}} + \frac{1}{\lambda_r^*} \right) \left( 1 + \frac{\sqrt{d}}{\lambda_r^*} \right) \text{poly log } d$$

In comparison, our sample size requirement for consistent estimation improves upon [ZWS19, Theorem 4] by a factor of  $\min\{d, p^{-1}\}$ . Moreover, our estimation error bound improves upon [ZWS19, Theorem 4] by a factor of  $\min\{\sqrt{\lambda_r^*}, \frac{1}{\sqrt{p}}\}$  if  $\sqrt{d} \ll \lambda_r^* \lesssim d$ , by a factor of  $\frac{\sqrt{d}}{\lambda_r^*}$  when  $\lambda_r^* \lesssim 1$ , and by a factor of  $\min\left\{\frac{\sqrt{d}}{\lambda_r^* \sqrt{p}}, \sqrt{\frac{d}{\lambda_r^*}}\right\}$  if  $1 \ll \lambda_r^* \lesssim \sqrt{d}$ .

<sup>4</sup>Here, consistent estimation is declared if  $\min_{\mathbf{Q} \in \mathcal{O}^{r \times r}} \|\mathbf{UQ} - \mathbf{U}^*\| = o(1)$  and  $\|\mathbf{S} - \mathbf{S}^*\| = o(\lambda_r^*)$ .

- In the absence of missing data, the  $\ell_{2,\infty}$  error bound presented in [CTP19b, Theorem 1.1] reads (ignoring logarithmic terms)

$$\min_{\mathbf{Q} \in \mathcal{O}^{r \times r}} \|\mathbf{U}\mathbf{Q} - \mathbf{U}^*\|_{2,\infty} \lesssim \begin{cases} \sqrt{\frac{1}{nd}}, & \text{for } \frac{\sigma}{\sqrt{\lambda_r^*}} \lesssim \frac{1}{\sqrt{d}} \\ \frac{\sigma^2}{\lambda_r^*} \sqrt{\frac{d}{n}}, & \text{for } \frac{1}{\sqrt{d}} \ll \frac{\sigma}{\sqrt{\lambda_r^*}} \lesssim 1. \end{cases}$$

Consequently, our result improves upon the above error bound by a factor of  $\frac{\sigma\sqrt{d}}{\sqrt{\lambda_r^*}}$  if  $\frac{1}{\sqrt{d}} \ll \frac{\sigma}{\sqrt{\lambda_r^*}} \lesssim 1$ , while being able to handle the case with larger noise (namely,  $\frac{\sigma}{\sqrt{\lambda_r^*}} \gg 1$ ).

### 4.3 Community recovery in bipartite stochastic block models

As it turns out, if we denote by  $\mathbf{A} \in \mathbb{R}^{|\mathcal{U}| \times |\mathcal{V}|}$  the bi-adjacency matrix of the observed random bipartite graph or its centered version, then  $\mathbf{A}^* := \mathbb{E}[\mathbf{A}]$  exhibits a low-rank structure (as we shall elaborate momentarily). Perhaps more importantly, the column subspace of  $\mathbf{A}^*$  reveals the community memberships of all nodes in  $\mathcal{U}$ . As a result, this biclustering problem is tightly connected to subspace estimation given noisy observations of a low-rank matrix. In particular, when the size of  $\mathcal{V}$  is substantially larger than that of  $\mathcal{U}$ , one might encounter a situation where only the nodes in  $\mathcal{U}$  (rather than those in  $\mathcal{V}$ ) can be reliably clustered. This calls for development of “one-sided” community recovery algorithms, that is, the type of algorithms that guarantee reliable clustering of  $\mathcal{U}$  without worrying about the clustering accuracy in  $\mathcal{V}$ .

**Model.** This subsection investigates the problem of biclustering, by considering a bipartite stochastic block model (BSBM) with two disjoint groups of nodes  $\mathcal{U}$  and  $\mathcal{V}$ . Suppose that the nodes in  $\mathcal{U}$  (resp.  $\mathcal{V}$ ) form two clusters. For each pair of nodes  $(i, j) \in (\mathcal{U}, \mathcal{V})$ , there is an edge connecting them with probability depending only on the community memberships of  $i$  and  $j$ . To be more specific:

- *Biclustering structure.* Consider two disjoint collections of nodes  $\mathcal{U}$  and  $\mathcal{V}$ , which are of size  $n_u$  and  $n_v$ , respectively. Suppose that each collection of nodes can be clustered into two communities. To be more precise, let  $\mathcal{I}_1 \subseteq \mathcal{U}$  and  $\mathcal{I}_2 = \mathcal{U} \setminus \mathcal{I}_1$  (resp.  $\mathcal{J}_1 \subseteq \mathcal{V}$  and  $\mathcal{J}_2 = \mathcal{V} \setminus \mathcal{J}_1$ ) be two non-overlapping communities in  $\mathcal{U}$  (resp.  $\mathcal{V}$ ) that contain  $n_u/2$  (resp.  $n_v/2$ ) nodes each. Without loss of generality, we assume that  $\mathcal{I}_1$  contains the first  $n_u/2$  nodes of  $\mathcal{U}$ , and  $\mathcal{J}_1$  contains the first  $n_v/2$  nodes of  $\mathcal{V}$ ; these are of course *a priori* unknown.
- *Measurement model.* What we observe is a random bipartite graph generated based on the community memberships of the nodes. In the simplest version of BSBMs, a pair of nodes  $(i, j) \in (\mathcal{U}, \mathcal{V})$  is connected by an edge independently with probability  $q_{\text{in}}$  if either  $(i, j) \in (\mathcal{I}_1, \mathcal{J}_1)$  or  $(i, j) \in (\mathcal{I}_2, \mathcal{J}_2)$  holds, and with probability  $q_{\text{out}}$  otherwise. Here,  $0 \leq q_{\text{out}} \leq q_{\text{in}} \leq 1$  represent the edge densities. If we denote by  $\mathbf{C} \in \{0, 1\}^{n_u \times n_v}$  the bi-adjacency matrix of this random bipartite graph, then one has

$$\mathbb{P}\{C_{i,j} = 1\} \stackrel{\text{ind.}}{=} \begin{cases} q_{\text{in}}, & \text{if } (i, j) \in (\mathcal{I}_1, \mathcal{J}_1) \text{ or } (i, j) \in (\mathcal{I}_2, \mathcal{J}_2), \\ q_{\text{out}}, & \text{otherwise.} \end{cases}$$

Our goal is to recover the community memberships of the nodes in  $\mathcal{U}$ , based on the above random bipartite graph. In what follows, we define

$$n := n_u + n_v, \tag{40}$$

and declare exact community recovery of  $\mathcal{U}$  if the partition of the nodes returned by our algorithm coincides precisely with the true partition  $(\mathcal{I}_1, \mathcal{I}_2)$ .

While our theory covers a broad range of  $n_u$  and  $n_v$ , we emphasize the case where  $n_v \gg n_u$  (namely,  $\mathcal{V}$  contains far more nodes than  $\mathcal{U}$ ). In such a case, it is not uncommon to encounter a situation where one can only hope to recover the community memberships of the nodes in  $\mathcal{U}$  but not those in  $\mathcal{V}$ .



**Algorithm.** To attempt community recovery, we look at a centered version of the bi-adjacency matrix<sup>5</sup>

$$\mathbf{A} := \mathbf{C} - \frac{q_{\text{in}} + q_{\text{out}}}{2} \mathbf{1}_{n_u} \mathbf{1}_{n_v}^\top. \quad (41)$$

Recognizing that

$$\mathbf{A}^* := \mathbb{E}[\mathbf{A}] = \frac{q_{\text{in}} - q_{\text{out}}}{2} \begin{bmatrix} \mathbf{1}_{n_u/2} \mathbf{1}_{n_v/2}^\top & -\mathbf{1}_{n_u/2} \mathbf{1}_{n_v/2}^\top \\ -\mathbf{1}_{n_u/2} \mathbf{1}_{n_v/2}^\top & \mathbf{1}_{n_u/2} \mathbf{1}_{n_v/2}^\top \end{bmatrix} = \frac{q_{\text{in}} - q_{\text{out}}}{2} \begin{bmatrix} \mathbf{1}_{n_u/2} \\ -\mathbf{1}_{n_u/2} \end{bmatrix} \begin{bmatrix} \mathbf{1}_{n_v/2}^\top & -\mathbf{1}_{n_v/2}^\top \end{bmatrix}, \quad (42)$$

we see that the leading singular vectors of  $\mathbf{A}^*$  reveals the community memberships of all nodes. Motivated by this observation, our algorithm for recovering the community memberships in  $\mathcal{U}$  proceeds as follows:

---

**Algorithm 4** The spectral method for BSBM

---

- 1: **Input:** observed bi-adjacency matrix  $\mathbf{C}$ , edge probabilities  $q_{\text{in}}, q_{\text{out}}$ .
  - 2: Employ  $\mathbf{A}$  (cf. (41)) as the input of Algorithm 1, and let  $\mathbf{u} = [u_i] \in \mathbb{R}^{n_u}$  be the output returned by Algorithm 1 (which serves as the estimate of the leading left singular subspace of  $\mathbf{A}^*$ ).
  - 3: **Output:** for any  $i \in \mathcal{U}$ , we claim that  $i$  belongs to the first community if  $u_i > 0$ , and the second community otherwise.
- 

**Theoretical guarantees and implications.** We are now ready to invoke our general theory to demonstrate the effectiveness of the above algorithm, as asserted by the following result.

**Corollary 3** (Bipartite stochastic block model). *Consider the above bipartite stochastic block model. There exists some universal constant  $c_0 > 0$  such that if*

$$\frac{(q_{\text{in}} - q_{\text{out}})^2}{q_{\text{in}}} \geq c_0 \max \left\{ \frac{\log n}{\sqrt{n_u n_v}}, \frac{\log n}{n_v} \right\}, \quad (43)$$

then Algorithm 4 achieves exact community recovery of  $\mathcal{U}$  with probability exceeding  $1 - O(n^{-10})$ .

We then take a moment to discuss the implications of Corollary 3. For simplicity of presentation, we shall focus on the scenario with  $q_{\text{in}} \asymp q_{\text{out}} = o(1)$  and  $n_u \leq n_v$ .

- *Exact recovery via the spectral method alone.* Consider the following sparse regime, where

$$q_{\text{in}} = \frac{a \log n}{\sqrt{n_u n_v}} \quad \text{and} \quad q_{\text{out}} = \frac{b \log n}{\sqrt{n_u n_v}}$$

for some absolute positive constants  $a \geq b$ . Corollary 3 demonstrates that we can achieve exact recovery when  $\frac{(a-b)^2}{a} \gtrsim 1$ . This improves upon prior results presented in [FP16]. More specifically, the results in [FP16] only guaranteed *almost* exact recovery of community memberships (namely, obtaining correct community memberships for a fraction  $1 - o(1)$  of the nodes). In comparison, our results assert that the spectral estimates alone are sufficient to reveal exact community memberships for all nodes in  $\mathcal{U}$ ; there is no need to invoke further refinement procedures to clean up the remaining errors.

- *Near optimality.* In the balanced case where  $n_u \asymp n_v$ , the condition  $\frac{(a-b)^2}{a} \gtrsim 1$  above is known to be information-theoretically optimal up to a constant factor. In the unbalanced case with  $n_v \geq n_u$ , prior work has identified a sharp threshold for *detection* — the problem of recovering a fraction  $1/2 + \epsilon$  of the community memberships for an arbitrarily small fixed constant  $\epsilon > 0$ . Specifically, such results reveal a fundamental lower limit that requires  $\frac{(q_{\text{in}} - q_{\text{out}})^2}{q_{\text{in}}} \gtrsim \frac{1}{\sqrt{n_u n_v}}$  [FPV15, FP16], thus implying the information-theoretic optimality of the spectral method (up to a logarithmic factor).

---

<sup>5</sup>Here, we assume prior knowledge about  $q_{\text{in}}$  and  $q_{\text{out}}$ . Otherwise, the quantity  $\frac{q_{\text{in}} + q_{\text{out}}}{2}$  can also be easily estimated.




Figure 1: Relative estimation errors of the subspace estimate  $\mathbf{U}$  for both Algorithm 1 and the vanilla spectral method. The results are reported for (a) relative error vs. sampling rate  $p$  (where  $d_1 = 100$ ,  $d_2 = 1000$ ,  $r = 4$ ,  $\sigma = 1$ ), (b) relative error vs. column dimension  $d_2$  (where  $d_1 = 100$ ,  $r = 4$ ,  $\sigma = 1$ ,  $p = \frac{2r \log(d_1 + d_2)}{\sqrt{d_1 d_2}}$ ), and (c) relative error vs. noise standard deviation  $\sigma$  (where  $d_1 = 100$ ,  $d_2 = 1000$ ,  $r = 4$ ,  $p = 0.1$ ).




Figure 2: Absolute estimation errors of the subspace estimate  $\mathbf{U}$  for both Algorithm 1 and the vanilla spectral method. The results are reported for (a) absolute error vs. sampling rate  $p$  (where  $d_1 = 100$ ,  $d_2 = 1000$ ,  $r = 4$ ,  $\sigma = 1$ ), (b) absolute error vs. column dimension  $d_2$  (where  $d_1 = 100$ ,  $r = 4$ ,  $\sigma = 1$ ,  $p = \frac{2r \log(d_1 + d_2)}{\sqrt{d_1 d_2}}$ ), and (c) absolute error vs. noise standard deviation  $\sigma$  (where  $d_1 = 100$ ,  $d_2 = 1000$ ,  $r = 4$ ,  $p = 0.1$ ).

#### 4.4 Numerical experiments

To confirm the applicability of our algorithm and the theoretical findings, we conduct a series of numerical experiments. All results reported in this subsection are averaged over 100 independent Monte Carlo trials. For the sake of comparisons, we also report the numerical performance of the vanilla spectral method (namely, returning the  $r$ -dimensional principal column subspace of  $\mathbf{A}$  directly without proper diagonal deletion).

**Subspace estimation for random low-rank data matrices.** We start with subspace estimation for a randomly generated matrix  $\mathbf{A}^*$ . Specifically, generate  $\mathbf{A}^* = \mathbf{Z}_1 \mathbf{Z}_2^\top$ , where  $\mathbf{Z}_1 \in \mathbb{R}^{d_1 \times r}$ ,  $\mathbf{Z}_2 \in \mathbb{R}^{d_2 \times r}$  consist of i.i.d. standard Gaussian entries. The noise matrix contains i.i.d. Gaussian entries, namely,  $N_{i,j} \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, \sigma^2)$  for each  $(i, j) \in [d_1] \times [d_2]$ . Figures 1 and 2 plot respectively the numerical estimation errors of the estimate  $\mathbf{U}$  vs. the sampling rate  $p$ , the column dimension  $d_2$ , and the standard deviation  $\sigma$  of noise. Two types of estimation errors are reported: (1) the absolute spectral norm error  $\|\mathbf{U}\mathbf{R} - \mathbf{U}^*\|$  and the relative spectral norm error  $\|\mathbf{U}\mathbf{R} - \mathbf{U}^*\| / \|\mathbf{U}^*\|$ ; (2) the absolute  $\ell_{2,\infty}$  norm error  $\|\mathbf{U}\mathbf{R} - \mathbf{U}^*\|_{2,\infty}$  and the relative  $\ell_{2,\infty}$  norm error  $\|\mathbf{U}\mathbf{R} - \mathbf{U}^*\|_{2,\infty} / \|\mathbf{U}^*\|_{2,\infty}$ , where  $\mathbf{R} := \arg \min_{\mathbf{Q} \in \mathcal{O}_{r \times r}} \|\mathbf{U}\mathbf{Q} - \mathbf{U}^*\|_{\text{F}}$ . As can be seen from the plots, Algorithm 1 yields reasonably good estimates in terms of both the spectral norm and the  $\ell_{2,\infty}$  norm, outperforming the vanilla spectral method in all experiments.

**Tensor completion from noise data.** Next, we consider numerically the problem of tensor completion from noisy observations of its entries. Recall the notations in Section 4.1. We generate  $\mathbf{W}^* \in \mathbb{R}^{d \times r}$  with




Figure 3: Relative estimation errors of the subspace  $\mathbf{U}$  spanned by tensor components in tensor completion. The results are plotted for (a) relative error vs. sampling rate  $p$  (where  $d = 100$ ,  $r = 4$ ,  $\sigma = 2$ ), and (b) relative error vs. noise standard deviation  $\sigma$  (where  $d = 100$ ,  $r = 4$ ,  $p = 0.1$ ).

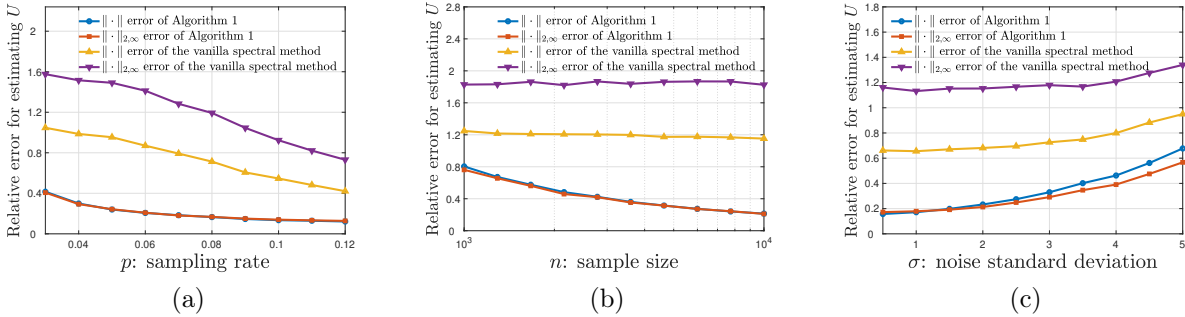


Figure 4: Relative error of the estimate  $\mathbf{U}$  for covariance estimation with missing data. The results are shown for (a) relative error vs. sampling rate  $p$  (where  $d = 100$ ,  $n = 5000$ ,  $r = 4$ ,  $\sigma = 1$ ), (b) relative error vs. sample size  $n$  (where  $d = 100$ ,  $r = 4$ ,  $\sigma = 1$ ,  $p = 0.05$ ), and (c) relative error vs. noise standard deviation  $\sigma$  (where  $d = 100$ ,  $n = 5000$ ,  $r = 4$ ,  $p = 0.1$ ).

i.i.d. standard Gaussian entries, and generate  $N_{i,j,k} \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, \sigma^2)$  independently for each  $(i, j, k) \in [d]^3$ . Figure 3(a) and Figure 3(b) illustrate the relative estimation errors of the subspace estimate  $\mathbf{U}$  vs. the sampling rate  $p$  and noise standard deviation  $\sigma$ , respectively. Encouragingly, Figure 3 shows that Algorithm 2 accurately recovers the subspace spanned by the tensor factors of interest (with respect to both the spectral norm and the  $\ell_{2,\infty}$  norm); in particular, it is capable of producing faithful subspace estimates even when the vanilla spectral method fails.

**Covariance estimation with missing data.** The next series of experiments is concerned with covariance estimation with missing data. Recall the notations in Section 4.2. We draw  $\mathbf{x}_i \sim \mathcal{N}(\mathbf{0}, \Sigma^*)$  independently with  $\Sigma^* = \mathbf{U}^* \mathbf{U}^{*\top}$ , where  $\mathbf{U}^* \in \mathbb{R}^{d \times r}$  is a i.i.d. standard Gaussian random matrix in  $\mathbb{R}^{d \times r}$ , and  $\varepsilon_i \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I}_d)$  for each  $1 \leq i \leq n$ . We first consider the estimation error of the subspace. The numerical estimation errors of the estimate  $\mathbf{U}$  vs. the sampling rate  $p$ , the sample size  $n$  and the noise standard deviation  $\sigma$  are plotted in Figure 4(a) – Figure 4(c), respectively. We then turn to the estimation accuracy of the covariance matrix. The numerical estimation errors of the estimate  $\mathbf{S}$  of Algorithm 3 and the vanilla spectral method vs. the sampling rate  $p$ , the sample size  $n$  and the noise standard deviation  $\sigma$  are plotted in Figure 5, respectively. Similar to previous experiments, Algorithm 3 produces reliable estimates both in terms of the spectral norm, the  $\ell_{2,\infty}$  norm and the  $\ell_\infty$  norm accuracy.

**Community recovery in bipartite stochastic block model.** Finally, we conduct numerical experiments for community recovery in bipartite stochastic block models. The parameters are chosen to be  $q_{\text{in}} = \frac{a \log(n_u + n_v)}{\sqrt{n_u n_v}}$  and  $q_{\text{out}} = \frac{b \log(n_u + n_v)}{\sqrt{n_u n_v}}$  for some constants  $a > b > 0$ . Figure 6(a) reveals a phase transition

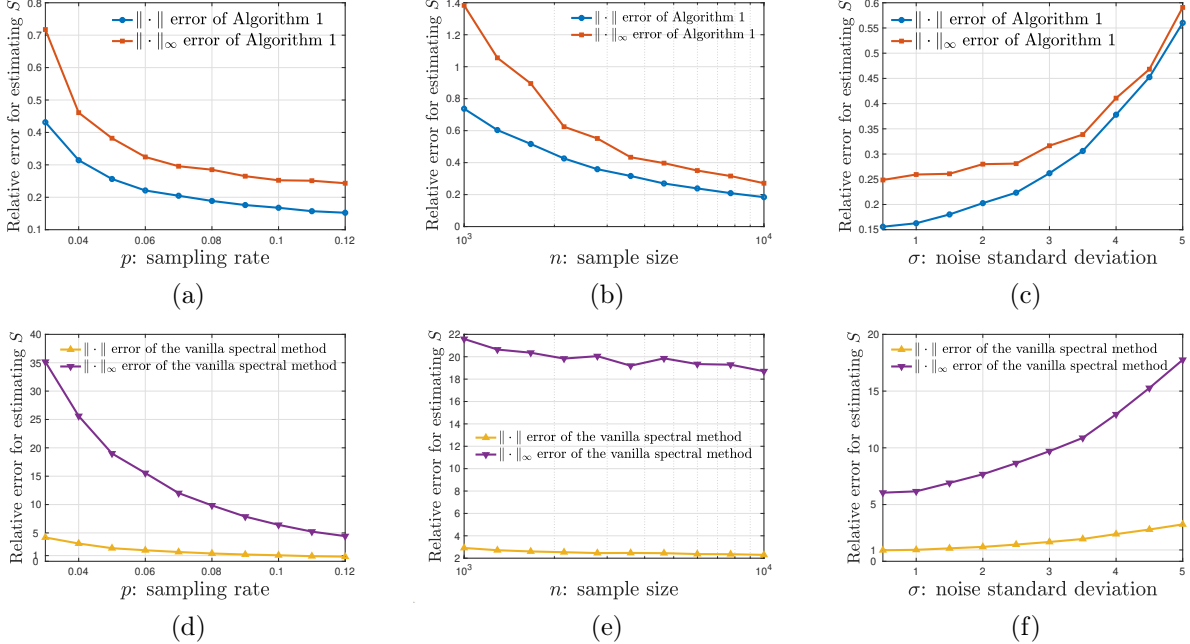


Figure 5: Relative error of the estimate  $S$  of Algorithm 3 (top row) and the vanilla spectral method (bottom row) for covariance estimation with missing data. The results are shown for (a,d) relative error vs. sampling rate  $p$  (where  $d = 100$ ,  $n = 5000$ ,  $r = 4$ ,  $\sigma = 1$ ), (b,e) relative error vs. sample size  $n$  (where  $d = 100$ ,  $r = 4$ ,  $\sigma = 1$ ,  $p = 0.05$ ), and (c,f) relative error vs. noise standard deviation  $\sigma$  (where  $d = 100$ ,  $n = 5000$ ,  $r = 4$ ,  $p = 0.1$ ). It is worth noting the different scales of the  $y$ -axis when plotting the errors of the two algorithms.

phenomenon concerned with exact community recovery. As can be seen, Algorithm 4 always succeeds in achieving exact recovery once  $a$  — or equivalently  $q_{\text{in}}$  — exceeds a certain threshold, which outperforms the vanilla spectral method. In Figure 6(b), we vary the number  $n_v$  nodes in  $\mathcal{V}$  and plot the empirical success rates for exact recovery. The advantage of Algorithm 4 compared to the vanilla spectral method can be clearly seen from the plot.

## 5 Further related work

A natural class of spectral algorithms to estimate the leading singular subspace of a matrix — when given a noisy and sub-sampled copy of the true matrix — is to compute the leading left singular subspace of the observed data matrix. Despite the simplicity of this idea, this type of spectral methods provably achieves appealing performances for multiple statistical problems when the true matrix is (nearly) square. A partial list of examples includes low-rank matrix estimation and completion [KMO10a, JNS13, Cha15, CW15, MWCC17], community detection [RCY11, YP14, AFWZ17, Lei19], and synchronization and alignment [Sin11, AFWZ17, CC18a, SHSS16]. The  $\ell_2$  statistical analysis of such algorithms relies heavily on the matrix perturbation theory such as the Davis-Kahan  $\sin \Theta$  theorem [DK70], the Wedin theorem [Wed72], and their extensions [Vu11, YWS14, OVW18, CZ18, CLC19, ZCW18].

However, the above-mentioned approach might lead to sub-optimal performance when the row dimension and the column dimension of the matrix differ dramatically. This issue has already been recognized in multiple contexts, including but not limited to unfolding-based spectral methods for tensor estimation [HSS16, XY17, MS18, XZ19, ZX18] and spectral methods for biclustering [FP16]. Motivated by this sub-optimality issue, an alternative is to look at the “sample Gram matrix” which, as one expects, shares the same leading left singular space as the original observed data matrix. However, in the highly noisy or highly subsampled regime, the diagonal entries of the sample Gram matrix are highly biased, thus requiring special care. Several different treatments of diagonal components have been adopted for different contexts, includ-

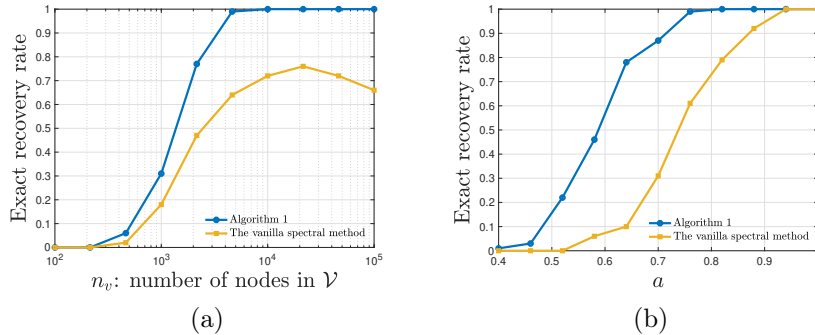


Figure 6: Empirical success rates for exact community recovery in bipartite stochastic block models, where  $q_{\text{in}} = \frac{a \log(n_u + n_v)}{\sqrt{n_u n_v}}$  and  $q_{\text{out}} = \frac{b \log(n_u + n_v)}{\sqrt{n_u n_v}}$ . The results are shown for (a) empirical success rate vs. the number  $n_v$  of nodes in  $\mathcal{V}$  (where  $n_u = 100$ ,  $a = 0.8$ ,  $b = 0.01$ ), and (b) empirical exact recovery rate vs.  $a$  (where  $n_u = 100$ ,  $n_v = 10000$ ,  $b = 0.01$ ).

ing proper rescaling [MS18, Lou14, GRESS16], deletion [FP16], and iterative updates [ZCW18]. The deletion strategy is perhaps the simplest of this kind, as it does not require estimation of noise parameters. We note, however, that performing more careful iterative updates might be beneficial for certain heteroskedastic noise scenarios; see [ZCW18] for detailed discussions.

An important application of our work is the problem of tensor completion and estimation [GRY11, LMWY13, RPP13, KOKC13, MHWG14, RM14, GSSB17, SRK09, YZ16, YZ17, HZC18]. Despite its similarity to matrix completion, tensor completion is considerably more challenging; for concreteness and simplicity, we shall only discuss order-3 symmetric rank- $r$  tensors in  $\mathbb{R}^{d^3}$ . Motivated by the success of matrix completion, a simple strategy for tensor completion / denoising is to unfold the observed tensor into a  $d \times d^2$  matrix and to apply standard matrix completion methods for completion. However, existing statistical guarantees derived in the matrix completion literature [CR09, KMO10a, Gro11] do not lead to useful bounds unless the sample size exceeds the order of  $rd^2$ , which far exceeds the requirement for other methods such as the sum-of-squares (SOS) hierarchy [BM16, PS17]. The work by [MS18] demonstrates that unfolding-based spectral algorithms can also lead to useful estimates under minimal sample complexity, as long as we look at the ‘‘Gram matrix’’ instead. In addition, such spectral algorithms also play an important role in initializing other nonconvex optimization methods [XY17, XYZ17, CLPC19, CPC20].

In addition, there is an enormous literature on covariance estimation and principal component analysis (PCA). For instance, the classical spiked covariance model [Joh01] has been extensively studied; in particular, the high-dimensional setting has inspired much investigation from both algorithmic and analysis perspectives [JL09, Pau07, BL08, Nad08, CY12, Ma13, CMW13, CMW15]. More recently, a computationally efficient algorithm called *HeteroPCA* has been proposed by [ZCW18] to achieve rate-optimal statistical guarantees for PCA in the presence of heteroskedastic noise. When it comes to incomplete data, a variety of methods have been introduced [Kie97, EvdG19, JH12]. For instance, Lounici considered estimating the top eigenvector in the setting of sparse PCA in [Lou13], and further proposed an estimator for the covariance matrix in [Lou14]. In [CZ16], bandable and sparse covariance matrices are considered. In addition, most of the prior work considered uniform random subsampling, and the recent work [ZWS19, PO19] began to account for heterogeneous missingness patterns.

Turning to the problem of community recovery or graph clustering, we note that extensive research has been carried out on stochastic block models or censored block models, which can be viewed as special cases of uni-partite networks [MNS14, Mas14, ABH15, MNS15, HWX16, CJSX14, CKST16, CRV15, GV16, CSG16, JMRT16, CLX18, CL15, GMZZ17]. The algorithms that enable exact community recovery in these block models include two-stage approaches [ABH15, MNS14] and semidefinite programming [HWX16, AL18, ABKK17, Ban18, GV16]. In addition, spectral clustering algorithms have been extensively studied as well [CO06, CO10, RCY11, YP14, LR15, YP16, AFWZ17, GMZZ17, Vu18, OVW18, STFP12, LMX15]. While this class of algorithms was originally developed to yield almost exact recovery (e.g. [ABH15]), the recent work by [AFWZ17, Lei19] uncovered that spectral methods alone are sufficient to achieve optimal exact community

recovery (a.k.a. achieving strong consistency) for stochastic block models. The interested reader is referred to [Abb18] for an in-depth overview. Various extensions of the SBMs have been introduced and studied in the last few years. Our work contributes to this growing literature by justifying the optimality of spectral methods in bipartite stochastic block models [FPV15, FP16, GLMZ16].

Further, entrywise statistical analysis has recently received significant attention for various statistical problems [FWZ16, AFWZ17, MSC17, CFMW19, CCF18, ZB18, ASSS19, Lei19, EBW17, CTP19b, CTP19a, XZ19, RV15, PW19]. For instance, entrywise guarantees for spectral methods are obtained in [CCF18, EBW17] based on an algebraic Neumann trick, while the results in [ZB18, AFWZ17, CFMW19] were established based on a leave-one-out analysis. The work by [KL16, KX16, CCF18] went one step further by controlling an arbitrary linear form of the eigenvectors or singular vectors of interest. These results, however, typically lead to suboptimal performance guarantees when the row dimension and the column dimension of the matrix are substantially different.

Finally, we recently became aware of an unpublished work by Abbe, Fan and Wang [AFW19], which also considers statistical guarantees of PCA beyond the usual  $\ell_2$  analysis; in particular, they develop an analysis framework that delivers tight  $\ell_p$  perturbation bounds. Note, however, that their results are very different from the ones presented here. For instance, the results presented herein emphasize the scenarios with drastically different  $d_1$  and  $d_2$ , which are not the main focus of [AFW19].

## 6 Analysis

In this section, we discuss in detail the analysis techniques employed to establish Theorem 1. This is built upon a leave-one-out (as well as a leave-two-out) analysis strategy that is particularly effective in controlling entrywise and  $\ell_{2,\infty}$  estimation errors [EKBB<sup>+</sup>13, EK15, ZB18, CFMW19, AFWZ17, SCC17, CCFM19, CFMY19, CLL19, LBEK18, PW19].

### 6.1 Leave-one-out and leave-two-out estimates

In order to facilitate the analysis when bounding  $\|\mathbf{UR} - \mathbf{U}^*\|_{2,\infty}$ , we introduce a set of auxiliary leave-one-out matrices — a powerful analysis technique that has been employed to decouple complicated statistical dependency. It is worth emphasizing that these procedures are never executed in practice. Specifically, for each  $1 \leq m \leq d_1$ , we introduce an auxiliary matrix

$$\mathbf{A}^{(m)} = \mathcal{P}_{-m,:}(\mathbf{A}) + p\mathcal{P}_{m,:}(\mathbf{A}^*), \quad (44)$$

where  $\mathcal{P}_{-m,:}$  (resp.  $\mathcal{P}_{m,:}$ ) represents the projection onto the subspace of matrices supported on the index subset  $\{[d_1] \setminus \{m\}\} \times [d_2]$  (resp.  $\{m\} \times [d_2]$ ). In other words,  $\mathbf{A}^{(m)}$  is obtained by replacing all entries in the  $m$ -th row by their expected values (taking into account the sampling rate). By construction, (1)  $\mathbf{A}^{(m)}$  is statistically independent of the data in the  $m$ -th row of  $\mathbf{A}$ , and (2)  $\mathbf{A}^{(m)}$  is expected to be quite close to  $\mathbf{A}$ , as we only discard a small fraction of data when constructing  $\mathbf{A}^{(m)}$ . These two observations taken together allow for optimal control of the estimation error in the  $m$ -th row of  $\mathbf{U}$ .

Armed with the leave-one-out matrices, we are ready to introduce auxiliary leave-one-out procedures for subspace estimation. Similar to the matrix  $\mathbf{G}$  in Algorithm 1 (whose eigenspace serves as an estimate of the column space of  $\mathbf{U}^*$ ), we define an auxiliary matrix  $\mathbf{G}^{(m)} \in \mathbb{R}^{d_1 \times d_1}$  as follows:

$$\mathbf{G}^{(m)} = \mathcal{P}_{\text{off-diag}}\left(\frac{1}{p^2}\mathbf{A}^{(m)}\mathbf{A}^{(m)\top}\right), \quad (45)$$

where  $\mathcal{P}_{\text{off-diag}}(\cdot)$  (as already defined in Section 2.2) extracts out all off-diagonal entries from a matrix. The auxiliary procedure, which is summarized in Algorithm 5, is very similar to Algorithm 1 except that it operates upon  $\mathbf{G}^{(m)}$ .

Given that  $\mathbf{A}^{(m)}$  (resp.  $\mathbf{G}^{(m)}$ ) is very close to  $\mathbf{A}$  (resp.  $\mathbf{G}$ ), one would naturally expect  $\mathbf{U}^{(m)}$  — the  $r$ -dimensional principal eigenspace of  $\mathbf{G}^{(m)}$  — to stay extremely close to the original estimate  $\mathbf{U}$ . This fact will be formalized shortly.

As it turns out, given that the spectral method is applied to the Gram matrix (which is a quadratic form of the original data matrix), introducing the leave-one-out sequences alone is not yet sufficient for our



---

**Algorithm 5** The  $m$ -th leave-one-out sequence

---

- 1: **Input:** sampling set  $\Omega$ , observed entries  $\{A_{i,j} \mid (i,j) \in \Omega\}$ , true entries  $\{A_{m,j}^* \mid j \in [d_2]\}$ , sampling rate  $p$ , rank  $r$ .
  - 2: Let  $\mathbf{U}^{(m)} \mathbf{\Lambda}^{(m)} \mathbf{U}^{(m)\top}$  be the (truncated) rank- $r$  eigen-decomposition of  $\mathbf{G}^{(m)}$ . Here,  $\mathbf{G}^{(m)}$  and  $\mathbf{A}^{(m)}$  are defined respectively in (45) and (44).
  - 3: **Output**  $\mathbf{U}^{(m)}$  as the subspace estimate and  $\mathbf{\Sigma}^{(m)} = (\mathbf{\Lambda}^{(m)})^{1/2}$  as the spectrum estimate.
- 

purpose; we still need to introduce an additional set of “leave-two-out” matrices, in the hope of simultaneously handling the row-wise and the column-wise statistical dependency. Specifically, for each  $1 \leq m \leq d_1$  and each  $1 \leq l \leq d_2$ , define the following auxiliary matrices:

$$\mathbf{A}^{(m,l)} := \mathcal{P}_{-m,-l}(\mathbf{A}) + p\mathcal{P}_{m,l}(\mathbf{A}^*), \quad (46a)$$

$$\mathbf{G}^{(m,l)} := \mathcal{P}_{\text{off-diag}}\left(\frac{1}{p^2}\mathbf{A}^{(m,l)}\mathbf{A}^{(m,l)\top}\right), \quad (46b)$$

where  $\mathcal{P}_{-m,-l}$  (resp.  $\mathcal{P}_{m,l}$ ) denotes the projection onto the subspace of matrices supported on  $\{[d_1] \setminus \{m\}\} \times \{[d_2] \setminus \{l\}\}$  (resp.  $\{m\} \times \{l\}$ ). Similar to  $\mathbf{A}^{(m)}$ ,  $\mathbf{A}^{(m,l)}$  is generated by replacing all data lying on the  $m$ -th row and the  $l$ -th column of  $\mathbf{A}$  by their expected values (taking into account the sampling rate). The precise procedure is summarized in Algorithm 6. Similar to the leave-one-out estimates, one expects the new leave-two-out estimates  $\mathbf{U}^{(m,l)}$  to be extremely close to  $\mathbf{U}^{(m)}$  (and hence  $\mathbf{U}$ ).

---

**Algorithm 6** The  $(m,l)$ -th leave-two-out sequence

---

- 1: **Input:** sampling set  $\Omega$ , observed entries  $\{A_{i,j} \mid (i,j) \in \Omega\}$ , true entries  $\{A_{m,j}^* \mid j \in [d_2]\} \cup \{A_{i,l}^* \mid i \in [d_1]\}$ , sampling rate  $p$ , rank  $r$ .
  - 2: Let  $\mathbf{U}^{(m,l)} \mathbf{\Lambda}^{(m,l)} \mathbf{U}^{(m,l)\top}$  be the (truncated) rank- $r$  eigen-decomposition of  $\mathbf{G}^{(m,l)}$ . Here,  $\mathbf{G}^{(m,l)}$  and  $\mathbf{A}^{(m,l)}$  are defined respectively in (46b) and (46a).
  - 3: **Output**  $\mathbf{U}^{(m,l)}$  as the subspace estimate and  $\mathbf{\Sigma}^{(m,l)} = (\mathbf{\Lambda}^{(m,l)})^{1/2}$  as the spectrum estimate.
- 

## 6.2 Key lemmas

In this subsection, we provide several lemmas that play a crucial role in establishing our main theorem. These lemmas are primarily concerned with the proximity between the original estimate, the leave-one-out estimates, and the ground truth. Throughout this section, we let

$$\mathbf{G}^* := \mathbf{A}^* \mathbf{A}^{*\top} = \mathbf{U}^* \mathbf{\Sigma}^{*2} \mathbf{U}^{*\top}. \quad (47)$$

To begin with, we demonstrate that  $\mathbf{G}$  is sufficiently close to  $\mathbf{G}^*$  when the difference is measured by the spectral norm. In view of standard matrix perturbation theory (which we shall make precise later), the proximity of  $\mathbf{G}$  and  $\mathbf{G}^*$  is crucial in bounding the difference between  $\mathbf{U}$  and  $\mathbf{U}^*$ . The proof is deferred to Appendix B.2.

**Lemma 1.** *Instate the assumptions of Theorem 1. With probability at least  $1 - O(d^{-10})$ , one has*

$$\|\mathbf{G} - \mathbf{G}^*\| \lesssim \underbrace{\frac{\mu r \sigma_1^{*2} \log d}{\sqrt{d_1 d_2} p} + \sqrt{\frac{\mu r \sigma_1^{*4} \log d}{d_2 p}} + \frac{\sigma^2 \sqrt{d_1 d_2} \log d}{p} + \sigma \sigma_1^* \sqrt{\frac{d_1 \log d}{p}}}_{=:\zeta_{\text{op}}} + \|\mathbf{A}^*\|_{2,\infty}^2. \quad (48)$$

In order to get a better sense of the term  $\zeta_{\text{op}}$  appearing above, we make note of a straightforward yet useful fact, which reveals that  $\zeta_{\text{op}}$  is much smaller than any nonzero eigenvalue of  $\mathbf{G}^*$ .

**Fact 1.** *Instate the assumptions of Theorem 1. Then the quantity  $\zeta_{\text{op}}$  as defined in (48) obeys*

$$\zeta_{\text{op}} + \|\mathbf{A}^*\|_{2,\infty}^2 \leq \frac{\mu r \sigma_1^{*2} \log d}{\sqrt{d_1 d_2} p} + \sqrt{\frac{\mu r \sigma_1^{*4} \log d}{d_2 p}} + \frac{\sigma^2}{\sigma_r^{*2}} \frac{\sqrt{d_1 d_2} \log d}{p} + \sigma \sigma_1^* \sqrt{\frac{d_1 \log d}{p}} + \frac{\mu_1 r \sigma_1^{*2}}{d_1}$$



$$\ll \frac{\sigma_r^{*2}}{\kappa^2},$$

where  $\|\mathbf{A}^*\|_{2,\infty}^2 \leq \frac{\mu_1 r \sigma_1^{*2}}{d_1}$  (cf. Lemma 11).

Further, the following lemma upper bounds the difference between  $\mathbf{G}$  and  $\mathbf{G}^*$  in the  $m$ -th row, when projected onto the subspace represented by  $\mathbf{U}^*$ ; the proof is postponed to Appendix B.3. This result gives a more refined control of the difference between  $\mathbf{G}$  and  $\mathbf{G}^*$ .

**Lemma 2.** *Instate the assumptions of Theorem 1. With probability at least  $1 - O(d^{-10})$ , the following holds simultaneously for all  $1 \leq m \leq d_1$ :*

$$\|(\mathbf{G} - \mathbf{G}^*)_{m,:} \mathbf{U}^*\|_2 \lesssim (\zeta_{\text{op}} + \|\mathbf{A}^*\|_{2,\infty}^2) \sqrt{\frac{\mu r}{d_1}},$$

where  $\zeta_{\text{op}}$  is defined in (48).

The next step, which is also the most challenging and crucial step, lies in showing that: every row of  $\mathbf{U}$ , under certain global linear transformation, serves as a good approximation of the corresponding row of  $\mathbf{U}^*$ . Towards this end, we begin with the following preparations:

- We first introduce the following matrix  $\mathbf{H}$  to represent the linear transformation we have in mind:

$$\mathbf{H} := \mathbf{U}^\top \mathbf{U}^*. \quad (49)$$

While this is not a rotation matrix, it is quite close to the rotation matrix  $\mathbf{R}$  defined in (14).

- In addition, we find it convenient to express

$$\mathbf{U}^* = \mathbf{G}^* \mathbf{U}^* (\boldsymbol{\Sigma}^*)^{-2}.$$

Combining this with Lemma 2, one would expect  $\mathbf{U}^*$  and  $\mathbf{G} \mathbf{U}^* (\boldsymbol{\Sigma}^*)^{-2}$  to be reasonably close, namely,

$$\mathbf{U}^* \approx \mathbf{G} \mathbf{U}^* (\boldsymbol{\Sigma}^*)^{-2}. \quad (50)$$

With these in hand, the following lemma (together with Lemma 2) asserts that

$$\mathbf{U} \mathbf{H} \approx \mathbf{G} \mathbf{U}^* (\boldsymbol{\Sigma}^*)^{-2} \approx \mathbf{U}^*$$

in an  $\ell_{2,\infty}$  sense.

**Lemma 3.** *Instate the assumptions of Theorem 1, and recall the definition of  $\zeta_{\text{op}}$  in (48). With probability at least  $1 - O(d^{-10})$ , one has*

$$\|\mathbf{U} \mathbf{H} - \mathbf{G} \mathbf{U}^* (\boldsymbol{\Sigma}^*)^{-2}\|_{2,\infty} \lesssim \frac{\kappa^2 (\zeta_{\text{op}} + \|\mathbf{A}^*\|_{2,\infty}^2)}{\sigma_r^{*2}} \left( \|\mathbf{U} \mathbf{H}\|_{2,\infty} + \sqrt{\frac{\mu r}{d_1}} \right).$$

The proof of this lemma, however, goes far beyond conventional matrix perturbation theory, and requires delicate decoupling of statistical dependencies. This is accomplished via leave-one-out and leave-two-out analysis arguments. In what follows, we take a moment to explain the high-level idea.

To establish Lemma 3, we first learn from standard matrix perturbation theory [AFWZ17, Lemma 1] that: for each  $1 \leq m \leq d_1$ ,

$$\left\| (\mathbf{U} \mathbf{H} - \mathbf{G} \mathbf{U}^* (\boldsymbol{\Sigma}^*)^{-2})_{m,:} \right\|_2 \lesssim \frac{1}{\lambda_r(\mathbf{G}^*)^2} \|\mathbf{G} - \mathbf{G}^*\| \|\mathbf{G}_{m,:} \mathbf{U}^*\|_2 + \frac{1}{\lambda_r(\mathbf{G}^*)} \|\mathbf{G}_{m,:} (\mathbf{U} \mathbf{H} - \mathbf{U}^*)\|_2 \quad (51)$$

holds, provided that  $\mathbf{G}$  and  $\mathbf{G}^*$  are sufficiently close.

- The first term on the right-hand side of (51) can already be controlled by Lemma 1 and Lemma 2.

- The second term on the right-hand side of (51), however, is considerably more difficult to analyze, due to the complicated statistical dependence between  $\mathbf{G}_{m,:}$  and  $\mathbf{U}\mathbf{H}$ . In order to decouple statistical dependency, we resort to the leave-one-out sequence  $\mathbf{U}^{(m)}$  introduced in Algorithm 5 and use the triangle inequality to bound

$$\|\mathbf{G}_{m,:}(\mathbf{U}\mathbf{H} - \mathbf{U}^*)\|_2 \leq \|\mathbf{G}_{m,:}(\mathbf{U}\mathbf{H} - \mathbf{U}^{(m)}\mathbf{H}^{(m)})\|_2 + \|\mathbf{G}_{m,:}(\mathbf{U}^{(m)}\mathbf{H}^{(m)} - \mathbf{U}^*)\|_2, \quad (52)$$

where  $\mathbf{H}^{(m)} := \mathbf{U}^{(m)\top}\mathbf{U}^*$ . As mentioned before, the leave-one-out estimate  $\mathbf{U}^{(m)}$  enjoys two nice properties.

- (1) The true estimate  $\mathbf{U}$  and the leave-one-out estimate  $\mathbf{U}^{(m)}$  are exceedingly close, as asserted by the following lemma (to be established in Appendix B.5).

**Lemma 4.** *Instate the assumptions of Theorem 1, and recall the definition of  $\mathbf{H}$  in (49). With probability at least  $1 - O(d^{-10})$ , the following holds simultaneously for all  $1 \leq m \leq d_1$ :*

$$\|\mathbf{U}^{(m)}\mathbf{U}^{(m)\top} - \mathbf{U}\mathbf{U}^\top\|_{\text{F}} \lesssim \frac{\zeta_{\text{op}}}{\sigma_r^{*2}} \left( \|\mathbf{U}\mathbf{H}\|_{2,\infty} + \sqrt{\frac{\mu r}{d_1}} \right),$$

where  $\zeta_{\text{op}}$  is defined in (48).

This result in turn allows us to control the first term on the right-hand side of (52).

- (2) Due to the statistical independence between  $\mathbf{A}_{m,:}$  and  $\mathbf{U}^{(m)}$ , the matrices  $\mathbf{G}_{m,:}$  and  $\mathbf{U}^{(m)}$  turn out to be nearly independent. This allows one to invoke simple concentration inequalities to develop tight bounds for the second term on the right-hand side of (52). The detailed proof can be found in Appendix B.4.

Finally, we make a remark on a technical issue encountered in the proof of Lemma 4. Recall that  $\mathbf{U}^{(m)}$  is obtained by simply replacing the  $m$ -th row of  $\mathbf{A}$  with its population version, which indicates the statistical dependency between  $\mathbf{U}^{(m)}$  and the  $m$ -th row of  $\mathbf{A}$ . However, there is still some delicate statistical dependency between  $\mathbf{U}^{(m)}$  and the columns of  $\mathbf{A}$  that need to be carefully coped with. Fortunately, the leave-two-out estimate  $\mathbf{U}^{(m,l)}$  — which is obtained by dropping not only the  $m$ -th row of  $\mathbf{A}$  but also the  $l$ -th of its columns — allows us to decouple the dependency between  $\mathbf{U}^{(m,l)}$  (and hence  $\mathbf{U}$  and  $\mathbf{U}^{(m)}$ ) and the  $l$ -th column of  $\mathbf{A}$ . This is precisely the main reason why we introduce additional leave-two-out estimates.

### 6.3 Proof of Theorem 1

We are now positioned to establish our main theorem. The proof is split into two parts.

#### 6.3.1 Statistical accuracy measured by $\|\cdot\|$

We begin by establishing the spectral norm bound (16c). Let  $\lambda_i$  and  $\lambda_i^*$  be the  $i$ -th largest eigenvalue of  $\mathbf{\Lambda}$  and  $\mathbf{\Lambda}^*$ , respectively. From Lemma 1 and Weyl's inequality, one finds that

$$\max_{1 \leq i \leq r} |\lambda_i - \lambda_i^*| = \|\mathbf{\Lambda} - \mathbf{\Lambda}^*\| \leq \|\mathbf{G} - \mathbf{G}^*\| \lesssim \zeta_{\text{op}} + \|\mathbf{A}^*\|_{2,\infty}^2 \leq \sigma_r^{*2} \cdot \mathcal{E}_{\text{general}}, \quad (53)$$

where  $\zeta_{\text{op}}$  and  $\mathcal{E}_{\text{general}}$  are defined in (48) and (17), respectively. Here, the last inequality arises from the simple fact that  $\|\mathbf{A}^*\|_{2,\infty}^2 \leq \frac{\mu_1 r \sigma_1^{*2}}{d_1}$  (cf. Lemma 11). By virtue of Fact 1, we know that  $\|\mathbf{\Lambda} - \mathbf{\Lambda}^*\| \ll \sigma_r^{*2}$ . Given that  $\mathbf{\Lambda}^* = \mathbf{\Sigma}^{*2}$  and  $\mathbf{\Lambda} = \mathbf{\Sigma}^2$ , this implies that for each  $1 \leq i \leq r$ ,

$$\frac{1}{4}\sigma_i^{*2} = \frac{1}{4}\lambda_i^* \leq \lambda_i^* - |\lambda_i - \lambda_i^*| \leq \lambda_i \leq \lambda_i^* + |\lambda_i - \lambda_i^*| \leq 4\lambda_i^* = 4\sigma_i^{*2},$$

thus indicating that

$$\frac{1}{2}\sigma_i^* \leq \sigma_i \leq 2\sigma_i^*. \quad (54)$$

In conclusion,

$$\|\Sigma - \Sigma^*\| = \max_{1 \leq i \leq r} |\sigma_i - \sigma_i^*| = \max_{1 \leq i \leq r} \frac{|\sigma_i^2 - \sigma_i^{*2}|}{\sigma_i + \sigma_i^*} \stackrel{(a)}{\leq} \max_{1 \leq i \leq r} \frac{\|\Lambda - \Lambda^*\|}{\frac{3}{2}\sigma_i^*} \stackrel{(b)}{\lesssim} \sigma_r^* \cdot \mathcal{E}_{\text{general}}$$

as claimed. Here, (a) comes from (54), whereas (b) follows from (53).

Next, we turn attention to (16a). First, it is seen that

$$\|\mathbf{U}\mathbf{R} - \mathbf{U}^*\| \leq \sqrt{2} \|\mathbf{U}\mathbf{U}^\top - \mathbf{U}^*\mathbf{U}^{*\top}\| = \|\sin \Theta(\mathbf{U}, \mathbf{U}^*)\| \quad (55)$$

where  $\mathbf{R}$  is defined in (14), and  $\Theta(\mathbf{U}, \mathbf{U}^*)$  denotes a diagonal matrix whose  $i$ -th diagonal entry is the  $i$ -th principal angle between the two subspaces represented by  $\mathbf{U}$  and  $\mathbf{U}^*$ . Here, the first inequality follows from a well-known inequality connecting two different subspace distance metrics [SS90, Chapter II], while the last identity follows from [SS90, Chapter II]. In addition, Lemma 1 and Fact 1 tell us that

$$\|\mathbf{G} - \mathbf{G}^*\| \lesssim \zeta_{\text{op}} + \|\mathbf{A}^*\|_{2,\infty}^2 \ll \sigma_r^{*2} \quad (56)$$

with probability at least  $1 - O(d^{-10})$ , which together with Weyl's inequality gives

$$\lambda_{r+1}(\mathbf{G}) \leq \lambda_{r+1}(\mathbf{G}^*) + \|\mathbf{G} - \mathbf{G}^*\| = \|\mathbf{G} - \mathbf{G}^*\| \leq \sigma_r^{*2}/2. \quad (57)$$

Therefore, [SS90, Chapter V, Theorem 3.6] (which is a version of the celebrated Davis-Kahan  $\sin \Theta$  Theorem [DK70]) reveals that

$$\|\sin \Theta(\mathbf{U}, \mathbf{U}^*)\| \leq \frac{\|\mathbf{G} - \mathbf{G}^*\|}{\lambda_r(\mathbf{G}^*) - \lambda_{r+1}(\mathbf{G})} \leq \frac{\|\mathbf{G} - \mathbf{G}^*\|}{\sigma_r^{*2} - \sigma_r^{*2}/2} = \frac{2\|\mathbf{G} - \mathbf{G}^*\|}{\sigma_r^{*2}},$$

where we have used the fact  $\lambda_r(\mathbf{G}^*) = \sigma_r^{*2}$ . The above bounds taken collectively imply that, with probability at least  $1 - O(d^{-10})$ ,

$$\|\mathbf{U}\mathbf{R} - \mathbf{U}^*\| \leq \frac{2\sqrt{2}\|\mathbf{G} - \mathbf{G}^*\|}{\sigma_r^{*2}} \lesssim \mathcal{E}_{\text{general}}. \quad (58)$$

### 6.3.2 Statistical accuracy measured by $\|\cdot\|_{2,\infty}$

Before continuing to the proof, we find it convenient to introduce a few more notations. In addition to the rotation matrix  $\mathbf{R}$  defined in (14) and the linear transformation  $\mathbf{H}$  defined in (49), we define

$$\text{sgn}(\mathbf{H}) := \tilde{\mathbf{U}}\tilde{\mathbf{V}}^\top \in \mathbb{R}^{d_1 \times d_1}, \quad (59)$$

where the columns of  $\tilde{\mathbf{U}} \in \mathbb{R}^{d_1 \times d_1}$  (resp.  $\tilde{\mathbf{V}} \in \mathbb{R}^{d_1 \times d_1}$ ) are the left (resp. right) singular vectors of  $\mathbf{H}$ . It is well-known that [TB77, Theorem 2]

$$\mathbf{R} = \text{sgn}(\mathbf{H}). \quad (60)$$

We now move on to establishing the advertised bound (16b).

1. To begin with, we claim that  $\mathbf{U}\mathbf{H}$  is extremely close to  $\mathbf{U}\mathbf{R}$ , provided that  $\|\mathbf{G} - \mathbf{G}^*\|$  is sufficiently small. To this end, recognizing that  $\|\mathbf{G} - \mathbf{G}^*\| \lesssim \zeta_{\text{op}} \ll \sigma_r^{*2}$  (according to Lemma 1 and Fact 1), we can apply [AFWZ17, Lemma 3] to show that

$$\|\mathbf{H}^{-1}\| \lesssim 1 \quad \text{and} \quad \sqrt{\|\mathbf{H} - \text{sgn}(\mathbf{H})\|} \leq \|\mathbf{U}\mathbf{U}^\top - \mathbf{U}^*\mathbf{U}^{*\top}\| \lesssim \frac{\zeta_{\text{op}}}{\sigma_r^{*2}},$$

where the last inequality follows from (58). Thus, invoke the identity (60) to arrive at

$$\begin{aligned} \|\mathbf{U}\mathbf{H} - \mathbf{U}\mathbf{R}\|_{2,\infty} &= \|\mathbf{U}\mathbf{H} - \mathbf{U}\text{sgn}(\mathbf{H})\|_{2,\infty} = \|\mathbf{U}\mathbf{H}\mathbf{H}^{-1}(\mathbf{H} - \text{sgn}(\mathbf{H}))\|_{2,\infty} \\ &\leq \|\mathbf{U}\mathbf{H}\|_{2,\infty} \|\mathbf{H}^{-1}\| \|\mathbf{H} - \text{sgn}(\mathbf{H})\| \\ &\lesssim \left(\frac{\zeta_{\text{op}}}{\sigma_r^{*2}}\right)^2 \|\mathbf{U}\mathbf{H}\|_{2,\infty} \ll \frac{\zeta_{\text{op}}}{\sigma_r^{*2}} \|\mathbf{U}\mathbf{H}\|_{2,\infty}. \end{aligned} \quad (61)$$

This in turn allows us to focus attention on bounding  $\|\mathbf{U}\mathbf{H} - \mathbf{U}^*\|_{2,\infty}$  (instead of  $\|\mathbf{U}\mathbf{R} - \mathbf{U}^*\|_{2,\infty}$ ).

2. Next, recall that  $\mathbf{G}^* = \mathbf{U}^* \boldsymbol{\Sigma}^{*2} \mathbf{U}^{*\top}$  and hence  $\mathbf{G}^* \mathbf{U}^* (\boldsymbol{\Sigma}^*)^{-2} = \mathbf{U}^*$ . Invoke the triangle inequality to reach

$$\begin{aligned} \|\mathbf{U}\mathbf{H} - \mathbf{U}^*\|_{2,\infty} &= \left\| \mathbf{U}\mathbf{H} - \mathbf{G}\mathbf{U}^* (\boldsymbol{\Sigma}^*)^{-2} + \mathbf{G}\mathbf{U}^* (\boldsymbol{\Sigma}^*)^{-2} - \mathbf{G}^* \mathbf{U}^* (\boldsymbol{\Sigma}^*)^{-2} \right\|_{2,\infty} \\ &\leq \left\| (\mathbf{G} - \mathbf{G}^*) \mathbf{U}^* (\boldsymbol{\Sigma}^*)^{-2} \right\|_{2,\infty} + \left\| \mathbf{U}\mathbf{H} - \mathbf{G}\mathbf{U}^* (\boldsymbol{\Sigma}^*)^{-2} \right\|_{2,\infty} \\ &\leq \left\| (\mathbf{G} - \mathbf{G}^*) \mathbf{U}^* \right\|_{2,\infty} \left\| (\boldsymbol{\Sigma}^*)^{-2} \right\| + \left\| \mathbf{U}\mathbf{H} - \mathbf{G}\mathbf{U}^* (\boldsymbol{\Sigma}^*)^{-2} \right\|_{2,\infty} \\ &\leq \frac{1}{\sigma_r^{*2}} \left\| (\mathbf{G} - \mathbf{G}^*) \mathbf{U}^* \right\|_{2,\infty} + \left\| \mathbf{U}\mathbf{H} - \mathbf{G}\mathbf{U}^* (\boldsymbol{\Sigma}^*)^{-2} \right\|_{2,\infty}. \end{aligned} \quad (62)$$

Regarding the first term of (62), Lemma 2 reveals that with probability at least  $1 - O(d^{-10})$ ,

$$\frac{1}{\sigma_r^{*2}} \left\| (\mathbf{G} - \mathbf{G}^*) \mathbf{U}^* \right\|_{2,\infty} \lesssim \frac{\zeta_{\text{op}} + \|\mathbf{A}^*\|_{2,\infty}^2}{\sigma_r^{*2}} \sqrt{\frac{\mu r}{d_1}}. \quad (63)$$

With regards to the second term of (62), Lemma 3 demonstrates that

$$\left\| \mathbf{U}\mathbf{H} - \mathbf{G}\mathbf{U}^* (\boldsymbol{\Sigma}^*)^{-2} \right\|_{2,\infty} \lesssim \frac{\kappa^2 (\zeta_{\text{op}} + \|\mathbf{A}^*\|_{2,\infty}^2)}{\sigma_r^{*2}} \left( \|\mathbf{U}\mathbf{H}\|_{2,\infty} + \sqrt{\frac{\mu r}{d_1}} \right) \quad (64)$$

with probability at least  $1 - O(d^{-10})$ . Combine (63) and (64) to arrive at

$$\|\mathbf{U}\mathbf{H} - \mathbf{U}^*\|_{2,\infty} \lesssim \frac{\kappa^2 (\zeta_{\text{op}} + \|\mathbf{A}^*\|_{2,\infty}^2)}{\sigma_r^{*2}} \left( \|\mathbf{U}\mathbf{H}\|_{2,\infty} + \sqrt{\frac{\mu r}{d_1}} \right). \quad (65)$$

3. As a byproduct of (65) and Fact 1, we see that

$$\|\mathbf{U}\mathbf{H} - \mathbf{U}^*\|_{2,\infty} \ll \|\mathbf{U}\mathbf{H}\|_{2,\infty} + \sqrt{\frac{\mu r}{d_1}}.$$

It then follows from the triangle inequality that

$$\|\mathbf{U}\mathbf{H}\|_{2,\infty} \leq \|\mathbf{U}\mathbf{H} - \mathbf{U}^*\|_{2,\infty} + \|\mathbf{U}^*\|_{2,\infty} \lesssim o(1) \|\mathbf{U}\mathbf{H}\|_{2,\infty} + \sqrt{\frac{\mu r}{d_1}},$$

thus indicating that

$$\|\mathbf{U}\mathbf{H}\|_{2,\infty} \leq 2 \sqrt{\frac{\mu r}{d_1}}. \quad (66)$$

Substitution into (61) and (65) gives

$$\|\mathbf{U}\mathbf{H} - \mathbf{U}\mathbf{R}\|_{2,\infty} \ll \frac{\zeta_{\text{op}}}{\sigma_r^{*2}} \sqrt{\frac{\mu r}{d_1}} \quad \text{and} \quad \|\mathbf{U}\mathbf{H} - \mathbf{U}^*\|_{2,\infty} \lesssim \frac{\kappa^2 (\zeta_{\text{op}} + \|\mathbf{A}^*\|_{2,\infty}^2)}{\sigma_r^{*2}} \sqrt{\frac{\mu r}{d_1}}.$$

Combining the above results yields

$$\begin{aligned} \|\mathbf{U}\mathbf{R} - \mathbf{U}^*\|_{2,\infty} &\leq \|\mathbf{U}\mathbf{H} - \mathbf{U}^*\|_{2,\infty} + \|\mathbf{U}\mathbf{H} - \mathbf{U}\mathbf{R}\|_{2,\infty} \\ &\lesssim \frac{\kappa^2 (\zeta_{\text{op}} + \|\mathbf{A}^*\|_{2,\infty}^2)}{\sigma_r^{*2}} \sqrt{\frac{\mu r}{d_1}}. \end{aligned}$$

Substituting the value of  $\zeta_{\text{op}}$  into the above inequality and using the upper bound  $\|\mathbf{A}^*\|_{2,\infty}^2 \leq \frac{\mu_1 r \sigma_1^{*2}}{d_1}$  (cf. Lemma 11), we conclude the proof.

## 7 Discussion

This paper has explored spectral methods tailored to subspace estimation for low-rank matrices with missing entries. In comparison to prior literature, our findings are particularly interesting when the column dimension  $d_2$  far exceeds the row dimension  $d_1$ . In many scenarios, even though the observed data are either too noisy or too incomplete to support reliable recovery of the entire matrix (so that prior matrix completion results often become inapplicable), they might still be informative enough if the purpose is merely to estimate the column subspace of the unknown matrix. In fact, this suggests a potentially useful paradigm for privacy-preserving estimation or learning: the inability to recover the entire matrix facilitates the protection of personal data, yet it is still possible to retrieve useful subspace information for inference and learning. Our main contribution lies in establishing  $\ell_{2,\infty}$  statistical guarantees for subspace estimation, therefore providing a stronger form of performance guarantees compared to the usual  $\ell_2$  perturbation bounds.

Moving forward, there are many directions that are worth pursuing. For example, our current theory is likely suboptimal with respect to the dependence on the rank  $r$  and the condition number  $\kappa$ . For instance, the conditions (15) and the risk bound (17) involve high-order polynomials of  $\kappa$  in multiple places, and the rank  $r$  in our current theory cannot exceed the order of  $d_1/\kappa^4$ ; all of these might be improvable via more refined analysis. In addition, it is natural to wonder whether we can extend our algorithm and theory to accommodate more general sampling patterns. Going beyond estimation, an important direction lies in statistical inference and uncertainty quantification for subspace estimation, namely, how to construct valid and hopefully optimal confidence regions that are likely to contain the unknown column subspace? It would also be interesting to investigate how to incorporate other structural prior (e.g. sparsity) to further reduce the sample complexity and/or improve the estimation accuracy. Finally, another interesting avenue for future exploration is the extension to distributed or decentralized settings.

## Acknowledgements

Y. Chen is supported in part by the grants AFOSR YIP award FA9550-19-1-0030, ONR N00014-19-1-2120, ARO YIP award W911NF-20-1-0097, ARO W911NF-18-1-0303, NSF DMS-2014279, CCF-1907661 and IIS-1900140, and the Princeton SEAS Innovation Award. Y. Chi is supported in part by the grants ONR N00014-18-1-2142 and N00014-19-1-2404, ARO W911NF-18-1-0303, and NSF CCF-1806154 and ECCS-1818571. H. V. Poor is supported in part by the NSF grant DMS-1736417, CCF-1908308, and a Princeton Schmidt Data-X Research Award. C. Cai is supported in part by Gordon Y. S. Wu Fellowships in Engineering.

## A Proofs for corollaries

### A.1 Proof of Corollary 1

Recall that the spectral algorithm considered herein (cf. Section 4.1) operates upon the noisy copy of the mode-1 matricization of the tensor  $\mathbf{T}^*$ , namely,

$$\mathbf{A}^* = \sum_{s=1}^r \mathbf{w}_s^* (\mathbf{w}_s^* \otimes \mathbf{w}_s^*)^\top. \quad (67)$$

Consequently, in order to apply Theorem 1, the main step boils down to estimating the spectrum and the incoherence parameters of  $\mathbf{A}^*$ . Specifically, we need to upper bound the condition number  $\kappa$ , as well as the incoherence parameters  $\mu_0, \mu_1$  and  $\mu_2$  as introduced in Definition 1.

Before proceeding, we introduce a few notations that simplify the presentation. Define

$$\lambda_i^* := \|\mathbf{w}_i^*\|_2^3, \quad 1 \leq i \leq r,$$

and let  $\lambda_{(i)}^*$  denote the  $i$ -th largest value in  $\{\lambda_i^*\}_{i=1}^r$ . We also recall that

$$\lambda_{\min}^* := \min_{1 \leq i \leq r} \|\mathbf{w}_i^*\|_2^3 \quad \text{and} \quad \lambda_{\max}^* := \max_{1 \leq i \leq r} \|\mathbf{w}_i^*\|_2^3.$$

In addition, we define two matrices of interest

$$\overline{\mathbf{W}}^* := [\overline{\mathbf{w}}_1^*, \dots, \overline{\mathbf{w}}_r^*] \in \mathbb{R}^{d \times r}, \quad \widetilde{\mathbf{W}}^* := [\overline{\mathbf{w}}_1^* \otimes \overline{\mathbf{w}}_1^*, \dots, \overline{\mathbf{w}}_r^* \otimes \overline{\mathbf{w}}_r^*] \in \mathbb{R}^{d^2 \times r},$$

where  $\overline{\mathbf{w}}_s^* := \mathbf{w}_s^* / \|\mathbf{w}_s^*\|_2$ , and  $\mathbf{a} \otimes \mathbf{b} := \begin{bmatrix} a_1 \mathbf{b} \\ \vdots \\ a_d \mathbf{b} \end{bmatrix}$ . In addition, let  $\mathbf{D}^* \in \mathbb{R}^{r \times r}$  be a diagonal matrix with diagonal entries

$$D_{s,s}^* = \|\mathbf{w}_s^*\|_2, \quad 1 \leq s \leq r.$$

These allow us to express

$$\mathbf{G}^* = \mathbf{A}^* \mathbf{A}^{*\top} = \overline{\mathbf{W}}^* \mathbf{D}^{*3} \widetilde{\mathbf{W}}^{*\top} \widetilde{\mathbf{W}}^* \mathbf{D}^{*3} \overline{\mathbf{W}}^{*\top}.$$

In the sequel, we begin by quantifying the spectrum of  $\mathbf{G}^*$ , which in turn allows us to understand the spectrum of  $\mathbf{A}^*$ .

- We first look at the eigenvalues of the matrices  $\overline{\mathbf{W}}^{*\top} \overline{\mathbf{W}}^*$  and  $\widetilde{\mathbf{W}}^{*\top} \widetilde{\mathbf{W}}^*$ . Towards this, let us write

$$\overline{\mathbf{W}}^{*\top} \overline{\mathbf{W}}^* = \mathbf{I}_r + \mathbf{C}, \quad \text{and} \quad \widetilde{\mathbf{W}}^{*\top} \widetilde{\mathbf{W}}^* = \mathbf{I}_r + \widetilde{\mathbf{C}} \quad (68)$$

for some matrices  $\mathbf{C}, \widetilde{\mathbf{C}} \in \mathbb{R}^{r \times r}$ . It follows immediately from the incoherence assumption (27) that

$$\|\mathbf{C}\|_\infty \leq \sqrt{\mu_5/d} \quad \text{and} \quad \|\widetilde{\mathbf{C}}\|_\infty \leq \mu_5/d,$$

thus leading to the simple bounds

$$\|\mathbf{C}\| \leq r \|\mathbf{C}\|_\infty \leq r \sqrt{\mu_5/d}, \quad \|\widetilde{\mathbf{C}}\| \leq r \|\widetilde{\mathbf{C}}\|_\infty \leq \mu_5 r/d. \quad (69)$$

These taken collectively with (68) and Weyl's inequality yield

$$\max_{i \in [r]} \left| \lambda_i(\overline{\mathbf{W}}^{*\top} \overline{\mathbf{W}}^*) - 1 \right| \leq \|\mathbf{C}\| \leq r \sqrt{\mu_5/d} \quad \text{and} \quad \max_{i \in [r]} \left| \lambda_i(\widetilde{\mathbf{W}}^{*\top} \widetilde{\mathbf{W}}^*) - 1 \right| \leq \|\widetilde{\mathbf{C}}\| \leq \mu_5 r/d,$$

which essentially tell us that

$$\|\overline{\mathbf{W}}^*\| = \sqrt{\lambda_1(\overline{\mathbf{W}}^{*\top} \overline{\mathbf{W}}^*)} \leq \sqrt{1 + r \sqrt{\mu_5/d}} \quad \text{and} \quad \|\widetilde{\mathbf{W}}^*\| = \sqrt{\lambda_1(\widetilde{\mathbf{W}}^{*\top} \widetilde{\mathbf{W}}^*)} \leq \sqrt{1 + \mu_5 r/d}. \quad (70)$$

- Returning to  $\mathbf{G}^*$ , one invokes the definition (68) to deduce that

$$\mathbf{G}^* = \overline{\mathbf{W}}^* \mathbf{D}^{*6} \overline{\mathbf{W}}^{*\top} + \overline{\mathbf{W}}^* \mathbf{D}^{*3} \widetilde{\mathbf{C}} \mathbf{D}^{*3} \overline{\mathbf{W}}^{*\top}.$$

Observe that the eigenvalues of  $\overline{\mathbf{W}}^* \mathbf{D}^{*3} (\overline{\mathbf{W}}^* \mathbf{D}^{*3})^\top$  are identical to those of  $(\overline{\mathbf{W}}^* \mathbf{D}^{*3})^\top \overline{\mathbf{W}}^* \mathbf{D}^{*3}$ , where the latter can be further decomposed as follows (in view of (68))

$$(\overline{\mathbf{W}}^* \mathbf{D}^{*3})^\top \overline{\mathbf{W}}^* \mathbf{D}^{*3} = \mathbf{D}^{*3} \overline{\mathbf{W}}^{*\top} \overline{\mathbf{W}}^* \mathbf{D}^{*3} = \mathbf{D}^{*6} + \mathbf{D}^{*3} \mathbf{C} \mathbf{D}^{*3}.$$

This taken together with Weyl's inequality, (69) and (70) shows that

$$\left| \lambda_i(\overline{\mathbf{W}}^* \mathbf{D}^{*6} \overline{\mathbf{W}}^{*\top}) - \lambda_i(\mathbf{D}^{*6}) \right| \leq \|\mathbf{D}^{*3} \mathbf{C} \mathbf{D}^{*3}\| \leq \|\mathbf{D}^*\|^6 \|\mathbf{C}\| \leq r \sqrt{\frac{\mu_5}{d}} \lambda_{\max}^{*2}$$

for each  $1 \leq i \leq r$ . In addition,

$$\left| \lambda_i(\mathbf{G}^*) - \lambda_i(\overline{\mathbf{W}}^* \mathbf{D}^{*6} \overline{\mathbf{W}}^{*\top}) \right| \leq \|\overline{\mathbf{W}}^* \mathbf{D}^{*3} \widetilde{\mathbf{C}} \mathbf{D}^{*3} \overline{\mathbf{W}}^{*\top}\| \leq \|\overline{\mathbf{W}}^*\|^2 \|\mathbf{D}^*\|^6 \|\widetilde{\mathbf{C}}\| \leq \frac{\mu_5 r}{d} \left(1 + r \sqrt{\frac{\mu_5}{d}}\right) \lambda_{\max}^{*2}.$$

As a result, invoke the triangle inequality to see that

$$\left| \lambda_i(\mathbf{G}^*) - \lambda_i(\mathbf{D}^{*6}) \right| \leq \left| \lambda_i(\mathbf{G}^*) - \lambda_i(\overline{\mathbf{W}}^* \mathbf{D}^{*6} \overline{\mathbf{W}}^{*\top}) \right| + \left| \lambda_i(\overline{\mathbf{W}}^* \mathbf{D}^{*6} \overline{\mathbf{W}}^{*\top}) - \lambda_i(\mathbf{D}^{*6}) \right|$$

$$\leq \frac{\mu_5 r}{d} \left(1 + r \sqrt{\frac{\mu_5}{d}}\right) \lambda_{\max}^{*2} + r \sqrt{\frac{\mu_5}{d}} \lambda_{\max}^{*2} \leq 3r \sqrt{\frac{\mu_5}{d}} \lambda_{\max}^{*2}$$

for each  $1 \leq i \leq r$ , where the last inequality holds under the assumption that  $r \sqrt{\mu_5/d} \leq 1$ . This means

$$\left| \lambda_i(\mathbf{G}^*) - \lambda_{(i)}^{*2} \right| \leq 3r \sqrt{\frac{\mu_5}{d}} \lambda_{\max}^{*2},$$

where  $\lambda_{(i)}^*$  denotes the  $i$ -th largest value in  $\{\lambda_i^*\}_{i=1}^r$ .

- Recalling that  $\mu_{\text{tc}} := \max\{\mu_3, \mu_4^2\}$ ,  $\kappa_{\text{tc}} := \lambda_{\max}^*/\lambda_{\min}^*$  and the rank assumption  $r \ll \kappa_{\text{tc}}^{-2} \sqrt{d/\mu_5}$ , we find that

$$\lambda_i(\mathbf{G}^*) = \lambda_{(i)}^{*2} + O\left(r \sqrt{\frac{\mu_5}{d}}\right) \lambda_{\max}^{*2} \quad \text{and} \quad \sigma_i(\mathbf{A}^*) = \lambda_{(i)}^* (1 + o(1)). \quad (71)$$

As a result, we immediately arrive at

$$\sigma_1(\mathbf{A}^*) = \lambda_{\max}^* (1 + o(1)), \quad \sigma_r(\mathbf{A}^*) = \lambda_{\min}^* (1 + o(1)), \quad \text{and} \quad \kappa = \frac{\sigma_1(\mathbf{A}^*)}{\sigma_r(\mathbf{A}^*)} \lesssim \kappa_{\text{tc}}.$$

Next, we turn attention to bounding the incoherence parameters of  $\mathbf{A}^*$ . Let  $\mathbf{A}^* = \mathbf{U}^* \mathbf{\Sigma}^* \mathbf{V}^{*\top}$  be the (compact) SVD of  $\mathbf{A}^*$ . It is seen from (67) that the column space of  $\mathbf{U}^*$  (resp.  $\mathbf{V}^*$ ) coincides with the column space of  $\overline{\mathbf{W}}^*$  (resp.  $\widetilde{\mathbf{W}}^*$ ). Therefore, there exist orthonormal matrices  $\mathbf{H}_1$  and  $\mathbf{H}_2$  such that

$$\mathbf{U}^* \mathbf{H}_1 = \overline{\mathbf{W}}^* (\overline{\mathbf{W}}^{*\top} \overline{\mathbf{W}}^*)^{-1/2} \quad \text{and} \quad \mathbf{V}^* \mathbf{H}_2 = \widetilde{\mathbf{W}}^* (\widetilde{\mathbf{W}}^{*\top} \widetilde{\mathbf{W}}^*)^{-1/2}.$$

These allow us to bound

$$\begin{aligned} \|\mathbf{U}^*\|_{2,\infty} &= \|\mathbf{U}^* \mathbf{H}_1\|_{2,\infty} \leq \|\overline{\mathbf{W}}^*\|_{2,\infty} \|(\overline{\mathbf{W}}^{*\top} \overline{\mathbf{W}}^*)^{-1/2}\| \leq \sqrt{\frac{\mu_4 r}{d}} \sqrt{\frac{1}{\lambda_r(\overline{\mathbf{W}}^{*\top} \overline{\mathbf{W}}^*)}} \lesssim \sqrt{\frac{\mu_4 r}{d}} \sqrt{\frac{1}{1-1/3}} \leq \sqrt{\frac{2\mu_4 r}{d}}, \\ \|\mathbf{V}^*\|_{2,\infty} &= \|\mathbf{V}^* \mathbf{H}_2\|_{2,\infty} \leq \|\widetilde{\mathbf{W}}^*\|_{2,\infty} \|(\widetilde{\mathbf{W}}^{*\top} \widetilde{\mathbf{W}}^*)^{-1/2}\| \leq \sqrt{\frac{\mu_4^2 r}{d^2}} \sqrt{\frac{1}{\lambda_r(\widetilde{\mathbf{W}}^{*\top} \widetilde{\mathbf{W}}^*)}} \leq \sqrt{\frac{\mu_4^2 r}{d^2}} \sqrt{\frac{1}{1-1/3}} \leq \sqrt{\frac{2\mu_4^2 r}{d^2}}, \end{aligned}$$

which follow from (70) and the assumption that  $r \ll \sqrt{d/\mu_5}$ . Moreover, the incoherence assumption (27) gives that

$$\mu_0 = \frac{d^3 \|\mathbf{A}^*\|_{\infty}^2}{\|\mathbf{A}^*\|_{\text{F}}^2} = \frac{d^3 \|\mathbf{T}^*\|_{\infty}^2}{\|\mathbf{T}^*\|_{\text{F}}^2} \leq \mu_3.$$

To conclude, the above analysis reveals that

$$\mu_0 \leq \mu_3, \quad \mu_1 \lesssim \mu_4, \quad \mu_2 \lesssim \mu_4^2, \quad \mu \lesssim \max\{\mu_3, \mu_4^2\} = \mu_{\text{tc}} \quad \text{and} \quad \kappa \lesssim \kappa_{\text{tc}},$$

where  $\mu = \max\{\mu_0, \mu_1, \mu_2\}$  and  $\kappa = \sigma_1(\mathbf{A}^*)/\sigma_r(\mathbf{A}^*)$ . With these estimates in place, Corollary 1 follows immediately from Theorem 1.

## A.2 Proof of Corollary 2

In the problem of covariance estimation with missing data, the ground truth  $\mathbf{A}^*$  is effectively given by  $\mathbf{B}^* \mathbf{F}^*$ , which obeys

$$\mathbf{A}^* = \mathbf{B}^* \mathbf{F}^* = \mathbf{U}^* \mathbf{\Lambda}^{*1/2} \mathbf{F}^* \in \mathbb{R}^{d \times n}, \quad \mathbf{F}^* = [\mathbf{f}_1^*, \dots, \mathbf{f}_n^*] \in \mathbb{R}^{r \times n}$$

with  $\mathbf{f}_i^* \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(\mathbf{0}, \mathbf{I}_r)$ . We note that by our assumption on the sample size, one has  $n \gg \kappa_{\text{ce}}^2 (r + \log d)$ , where  $\kappa_{\text{ce}} = \lambda_1^*/\lambda_r^*$ . In addition, we note that under the assumption of Corollary 2, one has

$$\mathcal{E}_{\text{ce}} \ll \kappa_{\text{ce}}^{-1} \leq 1, \quad (72)$$

where  $\mathcal{E}_{\text{ce}}$  and  $\kappa_{\text{ce}}$  are defined in (38) and (35), respectively.



### A.2.1 Estimation error of the principle subspace

In this section, we will prove (37a) and (37b). To begin the proof, we verify the condition of the random noise (cf. (10)). From standard Gaussian concentration results, one is allowed to choose  $R \asymp \sigma \sqrt{\log(n+d)}$ , so that  $|\eta_{i,j}| \leq R$  for all  $i$  and  $j$  with probability  $1 - O((n+d)^{-12})$ . Under our sample size condition that  $n \gg \max\left\{\frac{\log^6(n+d)}{dp^2}, \frac{\log^3(n+d)}{p}\right\}$ , the requirement (10) is satisfied, namely,

$$\frac{R^2}{\sigma^2} \asymp \log(n+d) \lesssim \min\left\{\frac{p\sqrt{dn}}{\log(n+d)}, \frac{pn}{\log(n+d)}\right\}.$$

Next, we turn to the properties of  $\mathbf{B}^* \mathbf{F}^*$  and start by looking at its spectrum. Define

$$\mathbf{C} := \mathbf{F}^* \mathbf{F}^{*\top} - n\mathbf{I}_r,$$

which allows us to write

$$\mathbf{G}^* = \mathbf{B}^* \mathbf{F}^* (\mathbf{B}^* \mathbf{F}^*)^\top = \mathbf{U}^* \mathbf{\Lambda}^{*1/2} \mathbf{F}^* \mathbf{F}^{*\top} \mathbf{\Lambda}^{*1/2} \mathbf{U}^{*\top} = n \underbrace{\mathbf{U}^* \mathbf{\Lambda}^* \mathbf{U}^{*\top}}_{=\mathbf{S}^*} + \underbrace{\mathbf{U}^* \mathbf{\Lambda}^{*1/2} \mathbf{C} \mathbf{\Lambda}^{*1/2} \mathbf{U}^{*\top}}_{=\mathbf{\Delta}}. \quad (73)$$

Using standard results on Gaussian random matrices [Ver12], one obtains

$$\begin{aligned} \|\mathbf{C}\| &\lesssim \max\left\{\sqrt{n}(\sqrt{r} + \sqrt{\log(n+d)}), r + \log(n+d)\right\} \asymp \sqrt{n}(\sqrt{r} + \sqrt{\log(n+d)}), \\ |\sigma_i(\mathbf{F}^*) - \sqrt{n}| &\lesssim \sqrt{r} + \sqrt{\log(n+d)} \end{aligned} \quad (74)$$

with probability at least  $1 - O((n+d)^{-10})$ , provided that  $n \gg r + \log(n+d)$ . It then follows from Weyl's inequality that

$$|\lambda_i(\mathbf{G}^*) - \lambda_i(n\mathbf{S}^*)| = |\lambda_i(\mathbf{G}^*) - \lambda_i^* n| \leq \|\mathbf{\Delta}\| \leq \|\mathbf{C}\| \|\mathbf{U}^*\|^2 \|\mathbf{\Lambda}^*\| \lesssim \lambda_1^* \sqrt{n}(\sqrt{r} + \sqrt{\log(n+d)}). \quad (75)$$

Under the sample size assumption  $n \gg \kappa_{\text{ce}}^2(r + \log(n+d))$ , we conclude that

$$\lambda_i(\mathbf{G}^*) = \lambda_i^* n(1 + o(1)) \quad \text{and} \quad \sigma_i(\mathbf{B}^* \mathbf{F}^*) = \sqrt{\lambda_i^* n}(1 + o(1)), \quad (76)$$

and hence

$$\kappa(\mathbf{G}^*) = \frac{\lambda_1(\mathbf{G}^*)}{\lambda_r(\mathbf{G}^*)} \asymp \kappa_{\text{ce}} \quad \text{and} \quad \kappa(\mathbf{B}^* \mathbf{F}^*) = \frac{\sigma_1(\mathbf{B}^* \mathbf{F}^*)}{\sigma_r(\mathbf{B}^* \mathbf{F}^*)} \asymp \sqrt{\kappa_{\text{ce}}}. \quad (77)$$

Further, we look at the entrywise infinity norm of  $\mathbf{B}^* \mathbf{F}^*$ . From standard Gaussian concentration inequalities,

$$\begin{aligned} \|\mathbf{B}^* \mathbf{F}^*\|_\infty &= \max_{i,j} \left| \left\langle (\mathbf{U}^* \mathbf{\Lambda}^{*1/2})_{i,:}, \mathbf{f}_j^* \right\rangle \right| \lesssim \|\mathbf{U}^* \mathbf{\Lambda}^{*1/2}\|_{2,\infty} \sqrt{\log(n+d)} \\ &\leq \|\mathbf{U}^*\|_{2,\infty} \|\mathbf{\Lambda}^*\|^{1/2} \sqrt{\log(n+d)} \leq \sqrt{\frac{\lambda_1^* \mu_{\text{ce}} r \log(n+d)}{d}} \end{aligned}$$

holds with probability at least  $1 - O((n+d)^{-10})$ . Meanwhile, one has

$$\|\mathbf{B}^* \mathbf{F}^*\|_{\text{F}} \geq \|\mathbf{B}^*\|_{\text{F}} \sigma_r(\mathbf{F}^*) = \|\mathbf{\Lambda}^{*1/2}\|_{\text{F}} \sigma_r(\mathbf{F}^*) \gtrsim \sqrt{\lambda_r^* r n} = \sqrt{\frac{1}{\kappa_{\text{ce}}} \lambda_1^* r n},$$

where the last step follows from (76) and (77). As a result,

$$\|\mathbf{B}^* \mathbf{F}^*\|_\infty \leq \sqrt{\frac{\mu_{\text{ce}} \kappa_{\text{ce}} \log(n+d)}{nd}} \|\mathbf{B}^* \mathbf{F}^*\|_{\text{F}} \quad (78)$$

Recalling the definition of  $\mu_0$  in (5), one obtains

$$\mu_0 \lesssim \mu_{\text{ce}} \kappa_{\text{ce}} \log(n+d). \quad (79)$$

When it comes to the incoherence parameters  $\mu_1$  and  $\mu_2$  (cf. (6)), it can be easily verified that

$$\mu_1 = \frac{d}{r} \|\mathbf{U}^*\|_{2,\infty}^2 = \mu_{\text{ce}}.$$

In addition, recognizing the existence of an orthonormal matrix  $\mathbf{H}_2$  such that  $\mathbf{V}^* \mathbf{H}_2 = \mathbf{F}^{*\top} (\mathbf{F}^* \mathbf{F}^{*\top})^{-1/2}$ , we can bound

$$\begin{aligned} \|\mathbf{V}^*\|_{2,\infty} &= \|\mathbf{V}^* \mathbf{H}_2\|_{2,\infty} \leq \|\mathbf{F}^{*\top}\|_{2,\infty} \|(\mathbf{F}^* \mathbf{F}^{*\top})^{-1/2}\| \stackrel{(i)}{\leq} \frac{\sqrt{r} + \sqrt{\log(n+d)}}{\sigma_r(\mathbf{F}^*)} \\ &\stackrel{(ii)}{\lesssim} \frac{\sqrt{r} + \sqrt{\log(n+d)}}{\sqrt{n-d} - \sqrt{r} - \sqrt{\log(n+d)}} \stackrel{(iii)}{\lesssim} \frac{\sqrt{r \log(n+d)}}{\sqrt{n}}, \end{aligned}$$

where (i) follows from the standard Gaussian concentration result that  $\|\mathbf{F}^{*\top}\|_{2,\infty} - \sqrt{r} \lesssim \sqrt{\log(n+d)}$  with probability  $1 - O((n+d)^{-20})$ , (ii) arises from (74), and (iii) holds true under our sample size assumption. Consequently, we obtain

$$\mu_2 = \frac{n}{r} \|\mathbf{V}^*\|_{2,\infty}^2 \lesssim \log(n+d).$$

Thus far, we have shown that

$$\sigma_r^* \asymp \sqrt{\lambda_r^* n}, \quad \mu_0 \lesssim \mu_{\text{ce}} \kappa_{\text{ce}} \log(n+d), \quad \mu_1 = \mu_{\text{ce}}, \quad \mu_2 \lesssim \log(n+d), \quad \mu \lesssim \mu_{\text{ce}} \kappa_{\text{ce}} \log(n+d) \quad \text{and} \quad \kappa \asymp \sqrt{\kappa_{\text{ce}}},$$

where  $\sigma_r^* = \sigma_r^*(\mathbf{B}^* \mathbf{F}^*)$ ,  $\mu = \max\{\mu_0, \mu_1, \mu_2\}$  and  $\kappa = \kappa(\mathbf{B}^* \mathbf{F}^*)$ . Under the sample size assumption (36) and the rank condition  $r \ll \frac{d}{\mu_{\text{ce}} \kappa_{\text{ce}}^2}$ , it is straightforward to verify the condition (15) is satisfied, i.e.

$$\begin{aligned} p &\gg \max \left\{ \frac{\mu_{\text{ce}} \kappa_{\text{ce}}^3 r \log^3(n+d)}{\sqrt{dn}}, \frac{\mu_{\text{ce}} \kappa_{\text{ce}}^5 r \log^3(n+d)}{n} \right\} \gtrsim \max \left\{ \frac{\mu \kappa^4 r \log^2(n+d)}{\sqrt{dn}}, \frac{\mu \kappa^8 r \log^2(n+d)}{n} \right\}, \\ \frac{\sigma}{\sigma_r^*} &\asymp \frac{\sigma}{\sqrt{\lambda_r^* n}} \ll \min \left\{ \frac{\sqrt{p}}{\sqrt{\kappa_{\text{ce}}^4 \sqrt{dn} \sqrt{\log(n+d)}}}, \sqrt{\frac{p}{\kappa_{\text{ce}}^3 d \log(n+d)}} \right\} \lesssim \min \left\{ \frac{\sqrt{p}}{\kappa^4 \sqrt{d_1 d_2} \sqrt{\log d}}, \frac{1}{\kappa^3} \sqrt{\frac{p}{d_1 \log d}} \right\}, \\ r &\ll \frac{d}{\mu_{\text{ce}} \kappa_{\text{ce}}^2} \asymp \frac{d}{\mu_1 \kappa^4}. \end{aligned}$$

Applying Theorem 1 immediately establishes the claims (37a) and (37b) in Corollary 2. Along the way, we have also established the following upper bound (see Lemma 1), which will be useful in the sequel:

$$\|\mathbf{G} - \mathbf{G}^*\| \lesssim \lambda_r(\mathbf{G}^*) \cdot \mathcal{E}_{\text{ce}} \asymp \lambda_r^* n \cdot \mathcal{E}_{\text{ce}}. \quad (80)$$

Here, we recall that  $\mathbf{G} = \frac{1}{p^2} \mathcal{P}_{\text{off-diag}}(\mathcal{P}_\Omega(\mathbf{X}) \mathcal{P}_\Omega(\mathbf{X})^\top)$ .

### A.2.2 Estimation error of the covariance matrix

It remains to prove (37c) and (37d). Before proceeding, we first recall that  $\mathbf{U} \mathbf{\Lambda} \mathbf{U}^\top$  is the top- $r$  eigendecomposition of  $\mathbf{G}$ ,

$$\mathbf{\Sigma} = \mathbf{\Lambda}^{1/2}, \quad \mathbf{B} = \frac{1}{\sqrt{n}} \mathbf{U} \mathbf{\Sigma}, \quad \mathbf{B}^* = \mathbf{U}^* \mathbf{\Lambda}^{*1/2} \quad \text{and} \quad \mathbf{R} = \arg \min_{\mathbf{Q} \in \mathcal{O}^{r \times r}} \|\mathbf{U} \mathbf{Q} - \mathbf{U}^*\|_{\text{F}}. \quad (81)$$

Let us also define

$$\mathbf{K} := \arg \min_{\mathbf{Q} \in \mathcal{O}^{r \times r}} \|\mathbf{B} \mathbf{Q} - \mathbf{B}^*\|_{\text{F}}.$$

It is well known that the minimizer  $\mathbf{K}$  is given by [TB77]

$$\mathbf{K} = \text{sgn}(\mathbf{B}^\top \mathbf{B}^*),$$

where the  $\text{sgn}(\cdot)$  function is defined in (59). Since  $\mathbf{K}$  is an orthonormal matrix, one can express

$$\mathbf{S} - \mathbf{S}^* = (\mathbf{B} \mathbf{K}) (\mathbf{B} \mathbf{K})^\top - \mathbf{B}^* \mathbf{B}^{*\top} = (\mathbf{B} \mathbf{K} - \mathbf{B}^*) (\mathbf{B} \mathbf{K})^\top + \mathbf{B}^* (\mathbf{B} \mathbf{K} - \mathbf{B}^*)^\top. \quad (82)$$

As a result, everything boils down to controlling  $\|\mathbf{BK} - \mathbf{B}^*\|$  and  $\|\mathbf{BK} - \mathbf{B}^*\|_{2,\infty}$ . To this end, we use (81) to reach the following useful decomposition

$$\mathbf{BK} - \mathbf{B}^* = \frac{1}{\sqrt{n}}\mathbf{U}\mathbf{\Lambda}^{1/2}(\mathbf{K} - \mathbf{R}) + \mathbf{U} \left( \frac{1}{\sqrt{n}}\mathbf{\Lambda}^{1/2}\mathbf{R} - \mathbf{R}\mathbf{\Lambda}^{*1/2} \right) + (\mathbf{UR} - \mathbf{U}^*)\mathbf{\Lambda}^{*1/2}. \quad (83)$$

Given that  $\mathbf{U}\frac{1}{n}\mathbf{\Lambda}\mathbf{U}^\top$  is the top- $r$  eigendecomposition of  $\frac{1}{n}\mathbf{G}$ , an important step lies in controlling the difference between  $\frac{1}{n}\mathbf{G}$  and  $\mathbf{S}^*$ . Recalling the matrix  $\mathbf{\Delta}$  as defined in (73), one can use (75), (80) as well as the definition of  $\mathcal{E}_{ce}$  (cf. (38)) to obtain

$$\begin{aligned} \left\| \frac{1}{n}\mathbf{G} - \mathbf{S}^* \right\| &\leq \frac{1}{n} \|\mathbf{G} - \mathbf{G}^*\| + \left\| \frac{1}{n}\mathbf{G}^* - \mathbf{S}^* \right\| = \frac{1}{n} \|\mathbf{G} - \mathbf{G}^*\| + \frac{1}{n} \|\mathbf{\Delta}\| \\ &\lesssim \frac{1}{n} \lambda_r(\mathbf{G}^*) \cdot \mathcal{E}_{ce} + \frac{\lambda_1^*}{\sqrt{n}} (\sqrt{r} + \sqrt{\log(n+d)}) \asymp \lambda_r^* \cdot \mathcal{E}_{ce}, \end{aligned}$$

where the last inequality makes use of the identity  $\lambda_r(\mathbf{G}^*) \asymp n\lambda_r^*$ . Hence, apply [MWCC17, Lemma 46, Lemma 47] (with slight modification on  $\kappa$ ) and Weyl's inequality to show that

$$\left\| \frac{1}{\sqrt{n}}\mathbf{\Lambda}^{1/2}\mathbf{R} - \mathbf{R}\mathbf{\Lambda}^{*1/2} \right\| \lesssim \frac{\kappa(\mathbf{S}^*)}{\sqrt{\lambda_r(\mathbf{S}^*)}} \left\| \frac{1}{n}\mathbf{G} - \mathbf{S}^* \right\| \lesssim \kappa_{ce} \sqrt{\lambda_r^*} \cdot \mathcal{E}_{ce}; \quad (84)$$

$$\|\mathbf{K} - \mathbf{R}\| \lesssim \frac{\sqrt{\kappa(\mathbf{S}^*)}}{\lambda_r(\mathbf{S}^*)} \left\| \frac{1}{n}\mathbf{G} - \mathbf{S}^* \right\| \lesssim \sqrt{\kappa_{ce}} \cdot \mathcal{E}_{ce}. \quad (85)$$

In addition, it follows from Weyl's inequality that

$$\left\| \frac{1}{n}\mathbf{\Lambda} - \mathbf{\Lambda}^* \right\| \leq \left\| \frac{1}{n}\mathbf{G} - \mathbf{S}^* \right\| \lesssim \lambda_r^* \cdot \mathcal{E}_{ce},$$

which combined with (72) gives

$$\frac{1}{n} \|\mathbf{\Lambda}\| \leq \left\| \frac{1}{n}\mathbf{\Lambda} - \mathbf{\Lambda}^* \right\| + \|\mathbf{\Lambda}^*\| \lesssim \lambda_r^* \cdot \mathcal{E}_{ce} + \lambda_1^* \asymp \lambda_1^* \quad (86)$$

under our assumptions.

We are ready to upper bound the difference between  $\mathbf{BK} - \mathbf{B}^*$ . Plugging (37a), (84) (85) and (86) into (83) shows that

$$\begin{aligned} \|\mathbf{BK} - \mathbf{B}^*\| &\leq \frac{1}{\sqrt{n}} \|\mathbf{\Lambda}\|^{1/2} \|\mathbf{K} - \mathbf{R}\| + \left\| \frac{1}{\sqrt{n}}\mathbf{\Lambda}^{1/2}\mathbf{R} - \mathbf{R}\mathbf{\Lambda}^{*1/2} \right\| + \|\mathbf{UR} - \mathbf{U}^*\| \|\mathbf{\Lambda}^*\|^{1/2} \\ &\lesssim \sqrt{\kappa_{ce}\lambda_1^*} \cdot \mathcal{E}_{ce} + \kappa_{ce} \sqrt{\lambda_r^*} \cdot \mathcal{E}_{ce} + \sqrt{\lambda_1^*} \cdot \mathcal{E}_{ce} \\ &\lesssim \kappa_{ce} \sqrt{\lambda_1^*} \cdot \mathcal{E}_{ce}. \end{aligned} \quad (87)$$

Since  $\mathbf{K} \in \mathcal{O}^{r \times r}$ , this also implies that

$$\|\mathbf{B}\| = \|\mathbf{BK}\| \leq \|\mathbf{BK} - \mathbf{B}^*\| + \|\mathbf{B}^*\| \lesssim \kappa_{ce} \sqrt{\lambda_1^*} \cdot \mathcal{E}_{ce} + \sqrt{\lambda_1^*} \asymp \sqrt{\lambda_1^*}, \quad (88)$$

where the last step results from (72). In addition, (37b), (72) and the fact that  $\mathbf{R} \in \mathcal{O}^{r \times r}$  guarantees that

$$\|\mathbf{U}\|_{2,\infty} = \|\mathbf{UR}\|_{2,\infty} \leq \|\mathbf{UR} - \mathbf{U}^*\|_{2,\infty} + \|\mathbf{U}^*\|_{2,\infty} \lesssim \kappa_{ce}^{3/2} \mathcal{E}_{ce} \sqrt{\frac{\mu_{ce} r \log(n+d)}{d}} + \sqrt{\frac{\mu_{ce} r}{d}} \lesssim \sqrt{\frac{\kappa_{ce} \mu_{ce} r \log(n+d)}{d}}.$$

Consequently, it follows from the decomposition (83) that

$$\|\mathbf{BK} - \mathbf{B}^*\|_{2,\infty} \leq \|\mathbf{U}\|_{2,\infty} \frac{1}{\sqrt{n}} \|\mathbf{\Lambda}\|^{1/2} \|\mathbf{K} - \mathbf{R}\| + \|\mathbf{U}\|_{2,\infty} \left\| \frac{1}{\sqrt{n}}\mathbf{\Lambda}^{1/2}\mathbf{R} - \mathbf{R}\mathbf{\Lambda}^{*1/2} \right\| + \|\mathbf{UR} - \mathbf{U}^*\|_{2,\infty} \|\mathbf{\Lambda}^*\|^{1/2}$$

$$\begin{aligned}
&\lesssim \sqrt{\frac{\kappa_{\text{ce}}^2 \mu_{\text{ce}} r \lambda_1^* \log(n+d)}{d}} \mathcal{E}_{\text{ce}} + \sqrt{\frac{\kappa_{\text{ce}}^3 \mu_{\text{ce}} r \lambda_1^* \log(n+d)}{d}} \mathcal{E}_{\text{ce}} + \sqrt{\frac{\kappa_{\text{ce}}^3 \mu_{\text{ce}} r \lambda_1^* \log(n+d)}{d}} \mathcal{E}_{\text{ce}} \\
&\asymp \sqrt{\frac{\kappa_{\text{ce}}^3 \mu_{\text{ce}} r \lambda_1^* \log(n+d)}{d}} \mathcal{E}_{\text{ce}}.
\end{aligned} \tag{89}$$

Combining (89) and (72) gives that

$$\begin{aligned}
\|\mathbf{B}\|_{2,\infty} &= \|\mathbf{BK}\|_{2,\infty} \leq \|\mathbf{BK} - \mathbf{B}^*\|_{2,\infty} + \|\mathbf{B}^*\|_{2,\infty} \lesssim \sqrt{\frac{\kappa_{\text{ce}}^3 \mu_{\text{ce}} r \lambda_1^* \log(n+d)}{d}} \mathcal{E}_{\text{ce}} + \sqrt{\frac{\mu_{\text{ce}} r \lambda_1^*}{d}} \\
&\lesssim \sqrt{\frac{\mu_{\text{ce}} r \lambda_1^* \log(n+d)}{d}},
\end{aligned} \tag{90}$$

where we use the fact that  $\|\mathbf{B}^*\|_{2,\infty} = \|\mathbf{U}^* \mathbf{\Lambda}^{*1/2}\|_{2,\infty} \lesssim \|\mathbf{U}^*\|_{2,\infty} \|\mathbf{\Lambda}^*\|^{1/2} \lesssim \sqrt{\mu_{\text{ce}} r \lambda_1^* / d}$  and  $\mathcal{E}_{\text{ce}} \ll \kappa_{\text{ce}}^{-1}$ .

To finish up, we substitute (87) and (88) into (82) to find that

$$\begin{aligned}
\|\mathbf{S} - \mathbf{S}^*\| &\leq \|\mathbf{BK} - \mathbf{B}^*\| (\|\mathbf{B}^*\| + \|\mathbf{BK}\|) \leq \|\mathbf{BK} - \mathbf{B}^*\| (\|\mathbf{B}^*\| + \|\mathbf{B}\|) \\
&\lesssim \kappa_{\text{ce}} \lambda_1^* \cdot \mathcal{E}_{\text{ce}}.
\end{aligned}$$

Combining (89) and (90) reveals that

$$\begin{aligned}
\|\mathbf{S} - \mathbf{S}^*\|_{\infty} &\leq \|\mathbf{BK} - \mathbf{B}^*\|_{2,\infty} (\|\mathbf{BK}\|_{2,\infty} + \|\mathbf{B}^*\|_{2,\infty}) \leq \|\mathbf{BK} - \mathbf{B}^*\|_{2,\infty} (\|\mathbf{B}\|_{2,\infty} + \|\mathbf{B}^*\|_{2,\infty}) \\
&\lesssim \frac{\kappa_{\text{ce}}^2 \mu_{\text{ce}} r \lambda_1^* \log(n+d)}{d} \cdot \mathcal{E}_{\text{ce}}.
\end{aligned}$$

We have therefore established all claims.

### A.3 Proof of Corollary 3

Recall from our calculation (42) that

$$\mathbf{A}^* := \mathbb{E}[\mathbf{A}] = \frac{(q_{\text{in}} - q_{\text{out}}) \sqrt{n_u n_v}}{2} \mathbf{u}^* \mathbf{v}^{*\top}$$

is a rank-1 matrix, where

$$\mathbf{u}^* := \frac{1}{\sqrt{n_u}} \begin{bmatrix} \mathbf{1}_{n_u/2} \\ -\mathbf{1}_{n_u/2} \end{bmatrix} \quad \text{and} \quad \mathbf{v}^* := \frac{1}{\sqrt{n_v}} \begin{bmatrix} \mathbf{1}_{n_v/2} \\ -\mathbf{1}_{n_v/2} \end{bmatrix}.$$

Let  $\mathbf{u} \in \mathbb{R}^{n_u}$  be the leading eigenvector of  $\mathbf{G}$  (cf. (12) and Algorithm 4). To establish Corollary 3, the main step boils down to showing that, under the conditions of Corollary 3,

$$\min \{ \|\mathbf{u} - \mathbf{u}^*\|_{\infty}, \|\mathbf{u} + \mathbf{u}^*\|_{\infty} \} \lesssim \frac{1}{\sqrt{n_u}} \mathcal{E}_{\text{bsbm}}, \tag{91}$$

holds with probability exceeding  $1 - O(n^{-10})$ , where

$$\mathcal{E}_{\text{bsbm}} := \frac{q_{\text{in}}}{(q_{\text{in}} - q_{\text{out}})^2} \frac{\log n}{\sqrt{n_u n_v}} + \frac{\sqrt{q_{\text{in}}}}{q_{\text{in}} - q_{\text{out}}} \sqrt{\frac{\log n}{n_v}} + \frac{1}{\sqrt{n_u}}. \tag{92}$$

If this claim (91) holds, then under our condition (43) one has  $\mathcal{E}_{\text{bsbm}} \ll 1$ , and hence

$$\min \{ \|\mathbf{u} - \mathbf{u}^*\|_{\infty}, \|\mathbf{u} + \mathbf{u}^*\|_{\infty} \} \lesssim \frac{1}{\sqrt{n_u}} \mathcal{E}_{\text{bsbm}} < \frac{1}{\sqrt{n_u}}.$$

In other words, one has either  $\text{sign}(u_i) = \text{sign}(u_i^*)$  for all  $1 \leq i \leq n_u$ , or  $\text{sign}(u_i) = -\text{sign}(u_i^*)$  for all  $1 \leq i \leq n_u$ . This tells us that the entrywise rounding operation applied to  $\mathbf{u}$  is sufficient to recover exactly the community memberships of all nodes in  $\mathcal{U}$ .

The rest of the proof is devoted to establishing the claim (91). In order to apply Theorem 1, it suffices to estimate the spectrum and the incoherence parameters of  $\mathbf{A}^*$ , as well as some simple statistical properties of  $\mathbf{N} := \mathbf{A} - \mathbf{A}^*$ .

- We begin by looking at  $\mathbf{A}^*$ , which has rank 1 and satisfies

$$\sigma_1(\mathbf{A}^*) = \frac{(q_{\text{in}} - q_{\text{out}})\sqrt{n_u n_v}}{2}, \quad \|\mathbf{A}^*\|_\infty = \frac{q_{\text{in}} - q_{\text{out}}}{2}, \quad \|\mathbf{u}^*\|_\infty = \frac{1}{\sqrt{n_u}}, \quad \|\mathbf{v}^*\|_\infty = \frac{1}{\sqrt{n_v}}.$$

Recalling the definition of  $\mu_0, \mu_1, \mu_2$  in (5) and (6), we obtain

$$\mu_0 = \frac{n_u n_v}{\|\mathbf{A}^*\|_F^2} \|\mathbf{A}^*\|_\infty^2 = 1, \quad \mu_1 = n_u \|\mathbf{u}^*\|_{2,\infty}^2 = 1, \quad \mu_2 = n_v \|\mathbf{v}^*\|_{2,\infty}^2 = 1, \quad \kappa = 1.$$

- Next, we consider the maximum magnitude  $R$  and the maximum variance  $\sigma^2$  of all entries of  $\mathbf{N}$  (see Assumption 2). Clearly, one has

$$R = \max_{i,j} |N_{i,j}| \leq 1, \\ \sigma^2 = \max_{i,j} \text{Var}(N_{i,j}) = \max\{q_{\text{in}}(1 - q_{\text{in}}), q_{\text{out}}(1 - q_{\text{out}})\} \leq \max\{q_{\text{in}}, q_{\text{out}}\} \asymp q_{\text{in}},$$

which follows since  $N_{i,j}$  is a centered Bernoulli random variable with parameter either  $q_{\text{in}}$  or  $q_{\text{out}}$ . From the assumption (43) and the fact  $q_{\text{in}}^2 \geq (q_{\text{in}} - q_{\text{out}})^2$ , we know that

$$q_{\text{in}} \geq \frac{(q_{\text{in}} - q_{\text{out}})^2}{q_{\text{in}}} \gg \frac{\log n}{\sqrt{n_u n_v}} + \frac{\log n}{n_v}.$$

Putting the above estimates together, we can straightforwardly verify the random noise requirement (10), namely,

$$\frac{R^2}{\sigma^2} \lesssim \frac{1}{q_{\text{in}}} \lesssim \frac{\min\{\sqrt{n_u n_v}, n_v\}}{\log n}.$$

With the preceding bounds in place, Corollary 3 is an immediate consequence of Theorem 1.

## B Proofs for key lemmas

This section aims to establish the key lemmas listed in Section 6.2.

### B.1 Auxiliary quantities, notation, and preliminary facts

To simplify our treatment, the proofs shall consider the influence of missing data and that of noise altogether. Specifically, throughout this section, we shall define a rescaled version of  $\mathbf{A}$  as follows

$$\mathbf{A}^s := \frac{1}{p}\mathbf{A} = \mathbf{A}^* + \mathbf{E} \in \mathbb{R}^{d_1 \times d_2}, \quad (93)$$

where the matrix  $\mathbf{E}$  represents the aggregate perturbation

$$\mathbf{E} := \frac{1}{p}\mathcal{P}_\Omega(\mathbf{A}^*) - \mathbf{A}^* + \frac{1}{p}\mathcal{P}_\Omega(\mathbf{N}). \quad (94)$$

Clearly,  $\mathbf{E} \in \mathbb{R}^{d_1 \times d_2}$  is a random matrix with independent zero-mean entries and  $\mathbb{E}[\mathbf{A}^s] = \mathbf{A}^*$ . In addition, we define the corresponding leave-one-out and leave-two-out versions

$$\mathbf{A}^{s,(m)} := \frac{1}{p}\mathbf{A}^{(m)}, \quad (95)$$

$$\mathbf{A}^{s,(m,l)} := \frac{1}{p}\mathbf{A}^{(m,l)}, \quad (96)$$

for each  $1 \leq m \leq d_1, 1 \leq l \leq d_2$ .

As we shall see momentarily, it is convenient to introduce the following quantities regarding the above perturbation matrix  $\mathbf{E}$ : (1)  $\max_{i \in [d_1], j \in [d_2]} |E_{i,j}|$ ; (2)  $\max_{i \in [d_1], j \in [d_2]} \sqrt{\mathbb{E}[E_{i,j}^2]}$ ; (3)  $\max_{i \in [d_1]} \sqrt{\sum_{j \in [d_2]} \mathbb{E}[E_{i,j}^2]}$ ; (4)  $\max_{j \in [d_2]} \sqrt{\sum_{i \in [d_1]} \mathbb{E}[E_{i,j}^2]}$ . In our settings, it is easy to verify — using the definition of incoherence

parameters (cf. Definition 1), the assumptions of the random noise (cf. Assumption 2), and Lemma 11 — that the quantities defined above admit the following upper bounds

$$\max_{i \in [d_1], j \in [d_2]} |E_{i,j}| \leq \frac{\|\mathbf{A}^*\|_\infty + R}{p} \lesssim \frac{\sqrt{\mu r} \sigma_1^*}{\sqrt{d_1 d_2 p}} + \frac{\sigma \min\{\sqrt[4]{d_1 d_2}, \sqrt{d_2}\}}{\sqrt{p \log d}} =: B, \quad (97a)$$

$$\max_{i \in [d_1], j \in [d_2]} \sqrt{\mathbb{E}[E_{i,j}^2]} \leq \frac{\|\mathbf{A}^*\|_\infty + \sigma}{\sqrt{p}} \leq \sigma_1^* \sqrt{\frac{\mu r}{d_1 d_2 p}} + \frac{\sigma}{\sqrt{p}} =: \sigma_\infty, \quad (97b)$$

$$\max_{i \in [d_1]} \sqrt{\sum_{j \in [d_2]} \mathbb{E}[E_{i,j}^2]} \leq \frac{\|\mathbf{A}^*\|_{2,\infty} + \sigma \sqrt{d_2}}{\sqrt{p}} \leq \sigma_1^* \sqrt{\frac{\mu r}{d_1 p}} + \sigma \sqrt{\frac{d_2}{p}} =: \sigma_{\text{row}}, \quad (97c)$$

$$\max_{j \in [d_2]} \sqrt{\sum_{i \in [d_1]} \mathbb{E}[E_{i,j}^2]} \leq \frac{\|\mathbf{A}^{*\top}\|_{2,\infty} + \sigma \sqrt{d_1}}{\sqrt{p}} \leq \sigma_1^* \sqrt{\frac{\mu r}{d_2 p}} + \sigma \sqrt{\frac{d_1}{p}} =: \sigma_{\text{col}}, \quad (97d)$$

with probability exceeding  $1 - O(d^{-12})$ . Further, the following lemma singles out a few other useful properties about these quantities (to be established in Appendix B.6), which will be useful throughout the proof.

**Lemma 5.** *Instate the assumptions of Theorem 1. Then with probability at least  $1 - O(d^{-12})$ , we have*

$$B \lesssim \frac{\min\{\sqrt{\sigma_{\text{row}} \sigma_{\text{col}}}, \sigma_{\text{row}}\}}{\sqrt{\log d}}, \quad (98a)$$

$$\sigma_\infty^2 \lesssim B \log d \|\mathbf{A}^*\| \sqrt{\frac{\mu r}{d_2}} \lesssim \sigma_{\text{col}} \sqrt{\log d} \|\mathbf{A}^*\| \sqrt{\frac{\mu r}{d_1}}, \quad (98b)$$

$$\sqrt{\frac{\mu r}{d_1}} \gtrsim \frac{B \log^{3/2} d \|\mathbf{A}^*\|_\infty}{\|\mathbf{A}^*\|^2}, \quad (98c)$$

$$\sigma_r^{*2} \gg \max\left\{\kappa^2 \sigma_{\text{col}} \sigma_{\text{row}} \log d, \kappa^2 \sigma_{\text{col}} \sqrt{\log d} \|\mathbf{A}^*\|, \kappa^2 \|\mathbf{A}^*\|_{2,\infty}^2, \sigma_{\text{row}} \sqrt{\frac{d_1 \log d}{\mu r}} \|\mathbf{A}^{*\top}\|_{2,\infty}, B \log d \|\mathbf{A}^*\|_\infty\right\}. \quad (98d)$$

## B.2 Proof of Lemma 1

The main component of the proof is to demonstrate that

$$\|\mathbf{G} - \mathbf{G}^*\| \lesssim \underbrace{(\sigma_{\text{row}} + \sigma_{\text{col}}) (\sigma_{\text{col}} + \|\mathbf{A}^{*\top}\|_{2,\infty}) \log d + \sigma_{\text{col}} \sqrt{\log d} \|\mathbf{A}^*\| + \|\mathbf{A}^*\|_{2,\infty}^2}_{=: \delta_{\text{op}}}. \quad (99)$$

By substituting the values of  $\sigma_{\text{row}}$  and  $\sigma_{\text{col}}$  (cf. (97)) into the above expression, one derives

$$\delta_{\text{op}} \lesssim \zeta_{\text{op}} + \|\mathbf{A}^*\|_{2,\infty}^2, \quad (100)$$

where  $\zeta_{\text{op}}$  is defined in (48). Therefore, Lemma 1 is an immediate consequence of (99) and (100). The remainder of the proof amounts to justifying (99).

Recall the definitions of  $\mathbf{G}$  and  $\mathbf{G}^*$  in (13) and (47), respectively. Given that  $\mathbf{A}^s = \mathbf{A}^* + \mathbf{E}$ , we can expand

$$\begin{aligned} \mathbf{G} - \mathbf{G}^* &= \mathcal{P}_{\text{off-diag}}(\mathbf{A}^s \mathbf{A}^{s\top}) - \mathbf{A}^* \mathbf{A}^{*\top} = \mathcal{P}_{\text{off-diag}}(\mathbf{A}^s \mathbf{A}^{s\top} - \mathbf{A}^* \mathbf{A}^{*\top}) - \mathcal{P}_{\text{diag}}(\mathbf{A}^* \mathbf{A}^{*\top}) \\ &= \mathcal{P}_{\text{off-diag}}(\mathbf{E} \mathbf{E}^\top) + \mathcal{P}_{\text{off-diag}}(\mathbf{A}^* \mathbf{E}^\top + \mathbf{E} \mathbf{A}^{*\top}) - \mathcal{P}_{\text{diag}}(\mathbf{A}^* \mathbf{A}^{*\top}), \end{aligned} \quad (101)$$

where  $\mathcal{P}_{\text{off-diag}}$  and  $\mathcal{P}_{\text{diag}}$  are defined in Section 2.2. In what follows, we control these three terms separately.

### B.2.1 Step 1: bounding the term $\mathcal{P}_{\text{off-diag}}(\mathbf{E}\mathbf{E}^\top)$

We first consider the term  $\mathcal{P}_{\text{off-diag}}(\mathbf{E}\mathbf{E}^\top)$ . Since  $\{E_{i,j}\}_{i \in [d_1], j \in [d_2]}$  are independent zero-mean random variables, we can express

$$\mathcal{P}_{\text{off-diag}}(\mathbf{E}\mathbf{E}^\top) = \sum_{1 \leq l \leq d_2} (\mathbf{E}_{:,l} \mathbf{E}_{:,l}^\top - \mathbf{D}_l) \quad (102)$$

as a sum of independent zero-mean random matrices, where  $\mathbf{D}_l$  is a random diagonal matrix in  $\mathbb{R}^{d_1 \times d_1}$  with entries

$$(\mathbf{D}_l)_{i,i} = E_{i,l}^2. \quad (103)$$

We intend to invoke the truncated matrix Bernstein inequality [HSS16, Proposition A.7] to control the spectral norm of (102). To this end, we need to look at a few quantities.

- We first bound the spectral norm of the following covariance matrix

$$\mathbf{\Sigma}_{\text{ns}} := \sum_{1 \leq l \leq d_2} \mathbb{E} \left[ (\mathbf{E}_{:,l} \mathbf{E}_{:,l}^\top - \mathbf{D}_l)^2 \right] \in \mathbb{R}^{d_1 \times d_1}.$$

Straightforward computation reveals that  $\mathbf{\Sigma}_{\text{ns}}$  is a diagonal matrix with entries

$$(\mathbf{\Sigma}_{\text{ns}})_{i,i} = \sum_{1 \leq l \leq d_2} \mathbb{E} [E_{i,l}^2] \sum_{m:m \neq i} \mathbb{E} [E_{m,l}^2] \leq \sigma_{\text{row}}^2 \sigma_{\text{col}}^2$$

for each  $1 \leq i \leq d_1$ . This immediately reveals that

$$V_{\text{ns}} := \|\mathbf{\Sigma}_{\text{ns}}\| \leq \sigma_{\text{row}}^2 \sigma_{\text{col}}^2. \quad (104)$$

- Next, we turn to upper bounding the spectral norm of each summand  $\mathbf{E}_{:,l} \mathbf{E}_{:,l}^\top - \mathbf{D}_l$ . As shown in the proof of Lemma 12, one has

$$\mathbb{P} \left\{ \left| \|\mathbf{E}_{:,l}\|_2^2 - M_1 \right| \geq t \right\} \leq 2 \exp \left( -\frac{3}{8} \min \left\{ \frac{t^2}{V_1}, \frac{t}{L_1} \right\} \right), \quad t > 0,$$

where  $M_1, L_1$  and  $V_1$  are given respectively by

$$\begin{aligned} M_1 &:= \mathbb{E} \left[ \|\mathbf{E}_{:,l}\|_2^2 \right] \leq \sigma_{\text{col}}^2, \\ L_1 &:= \max_{1 \leq i \leq d_1} |E_{i,l}^2 - \mathbb{E} [E_{i,l}^2]| \leq 2B^2, \\ V_1 &:= \sum_{1 \leq i \leq d_1} \text{Var} (E_{i,l}^2) \leq B^2 \sigma_{\text{col}}^2. \end{aligned}$$

In addition, with probability exceeding  $1 - O(d^{-20})$ ,

$$\begin{aligned} \|\mathbf{E}_{:,l}\|_2^2 &\lesssim M_1 + L_1 \log d + \sqrt{V_1 \log d} \lesssim \sigma_{\text{col}}^2 + B^2 \log d + \sqrt{B^2 \sigma_{\text{col}}^2 \log d} \\ &\asymp B^2 \log d + \sigma_{\text{col}}^2, \end{aligned} \quad (105)$$

where the last line comes from the AM-GM inequality  $2\sqrt{B^2 \sigma_{\text{col}}^2 \log d} \leq B^2 \log d + \sigma_{\text{col}}^2$ . This together with the definition  $\mathbf{D}_l := \text{diag}(E_{1,l}^2, \dots, E_{d_1,l}^2)$  gives

$$\|\mathbf{E}_{:,l} \mathbf{E}_{:,l}^\top - \mathbf{D}_l\| \leq \|\mathbf{E}_{:,l}\|_2^2 + \|\mathbf{D}_l\| \leq 2\|\mathbf{E}_{:,l}\|_2^2 \lesssim B^2 \log d + \sigma_{\text{col}}^2.$$

Therefore, if we set

$$L_{\text{ns}} := C (B^2 \log d + \sigma_{\text{col}}^2) \quad (106)$$

for some sufficiently large constant  $C > 0$ , then the above argument reveals that

$$L_{\text{ns}} \geq \frac{C}{3} \left( M_1 + L_1 \log d + \sqrt{V_1 \log d} \right) \geq \frac{C}{3} \max \left\{ \sqrt{V_1 \log d}, L_1 \log d \right\}.$$

- In addition, one can easily bound that

$$\begin{aligned} \mathbb{E} \left[ \|\mathbf{E}_{:,l}\|_2^2 \mathbf{1} \left\{ \|\mathbf{E}_{:,l}\|_2^2 \geq L_{\text{ns}} \right\} \right] &\leq L_{\text{ns}} \mathbb{P} \left\{ \|\mathbf{E}_{:,l}\|_2^2 \geq L_{\text{ns}} \right\} + \int_{L_{\text{ns}}}^{\infty} \mathbb{P} \left\{ \|\mathbf{E}_{:,l}\|_2^2 \geq t \right\} dt \\ &\leq O(d^{-20}) L_{\text{ns}} + \int_{L_{\text{ns}}}^{\infty} \mathbb{P} \left\{ \|\mathbf{E}_{:,l}\|_2^2 \geq t \right\} dt. \end{aligned}$$

Moreover, we know that  $\min \{t^2/V_1, t/L_1\} \geq t/\max \{\sqrt{V_1/\log d}, L_1\}$  for any  $t \geq L_{\text{ns}}/2$ . As a result, for sufficiently large  $d$ , we have

$$\begin{aligned} \int_{L_{\text{ns}}}^{\infty} \mathbb{P} \left\{ \|\mathbf{E}_{:,l}\|_2^2 \geq t \right\} dt &\leq \int_{L_{\text{ns}}}^{\infty} \mathbb{P} \left\{ \left| \|\mathbf{E}_{:,l}\|_2^2 - M_1 \right| > \frac{t}{2} \right\} dt \\ &\leq 4 \int_{\frac{1}{2}L_{\text{ns}}}^{\infty} \exp \left( -\frac{3}{8} \min \left\{ \frac{t^2}{V_1}, \frac{t}{L_1} \right\} \right) dt \\ &\leq 4 \int_{\frac{1}{2}L_{\text{ns}}}^{\infty} \exp \left( -\frac{3}{8} \frac{t}{\max \{\sqrt{V_1/\log d}, L_1\}} \right) dt \\ &\lesssim \max \{\sqrt{V_1/\log d}, L_1\} \exp \left( -\frac{3}{16} \frac{L_{\text{ns}}}{\max \{\sqrt{V_1/\log d}, L_1\}} \right) \\ &\leq \max \{\sqrt{V_1/\log d}, L_1\} \exp \left( -\frac{3C}{32} \log d \right) \\ &\ll \frac{L_{\text{ns}}}{d^2}, \end{aligned}$$

provided that  $C > 0$  is sufficiently large. Consequently, we have

$$\begin{aligned} R_{\text{ns}} &:= \mathbb{E} \left[ \|\mathbf{E}_{:,l}\mathbf{E}_{:,l}^{\top} - \mathbf{D}_l\| \mathbf{1} \left\{ \|\mathbf{E}_{:,l}\mathbf{E}_{:,l}^{\top} - \mathbf{D}_l\| \geq L_{\text{ns}} \right\} \right] \\ &\leq \mathbb{E} \left[ 2 \|\mathbf{E}_{:,l}\|_2^2 \mathbf{1} \left\{ 2 \|\mathbf{E}_{:,l}\|_2^2 \geq L_{\text{ns}} \right\} \right] \ll \frac{L_{\text{ns}}}{d^2}. \end{aligned} \quad (107)$$

With estimates (104), (106) and (107) in place, we are ready to apply the truncated matrix Bernstein inequality [HSSS16, Proposition A.7] to obtain that, with probability at least  $1 - O(d^{-10})$ ,

$$\begin{aligned} \|\mathcal{P}_{\text{off-diag}}(\mathbf{E}\mathbf{E}^{\top})\| &= \left\| \sum_{1 \leq l \leq d_2} \mathbf{E}_{:,l}\mathbf{E}_{:,l}^{\top} - \mathbf{D}_l \right\| \lesssim d_2 R_{\text{ns}} + L_{\text{ns}} \log d + \sqrt{V_{\text{ns}} \log d} \\ &\asymp L_{\text{ns}} \log d + \sqrt{V_{\text{ns}} \log d} \\ &\lesssim B^2 \log^2 d + \sigma_{\text{col}}^2 \log d + \sigma_{\text{row}} \sigma_{\text{col}} \sqrt{\log d} \\ &\lesssim \sigma_{\text{col}} (\sigma_{\text{row}} + \sigma_{\text{col}}) \log d, \end{aligned} \quad (108)$$

where the last line results from the identity  $B^2 \log d \lesssim \sigma_{\text{row}} \sigma_{\text{col}}$  (See (98a)).

### B.2.2 Step 2: bounding the term $\mathcal{P}_{\text{off-diag}}(\mathbf{A}^* \mathbf{E}^{\top} + \mathbf{E} \mathbf{A}^{*\top})$

Next, we turn attention to  $\mathcal{P}_{\text{off-diag}}(\mathbf{A}^* \mathbf{E}^{\top} + \mathbf{E} \mathbf{A}^{*\top})$ . By symmetry, it suffices to control to the spectral norm of  $\mathcal{P}_{\text{off-diag}}(\mathbf{A}^* \mathbf{E}^{\top})$ . To this end, we first express

$$\mathcal{P}_{\text{off-diag}}(\mathbf{A}^* \mathbf{E}^{\top}) = \sum_{1 \leq l \leq d_2} (\mathbf{A}_{:,l}^* \mathbf{E}_{:,l}^{\top} - \widehat{\mathbf{D}}_l) \quad (109)$$

as a sum of independent zero-mean random matrices, where  $\widehat{\mathbf{D}}_l$  is a diagonal matrix obeying

$$(\widehat{\mathbf{D}}_l)_{i,i} = A_{i,l}^* E_{i,l}. \quad (110)$$



To control (109), we need to first look at two matrices defined as follows

$$\begin{aligned}\widehat{\Sigma}_{\text{crs}} &:= \sum_{1 \leq l \leq d_2} \mathbb{E} \left[ (\mathbf{A}_{:,l}^* \mathbf{E}_{:,l}^\top - \widehat{\mathbf{D}}_l) (\mathbf{A}_{:,l}^* \mathbf{E}_{:,l}^\top - \widehat{\mathbf{D}}_l)^\top \right]; \\ \widetilde{\Sigma}_{\text{crs}} &:= \sum_{1 \leq l \leq d_2} \mathbb{E} \left[ (\mathbf{A}_{:,l}^* \mathbf{E}_{:,l}^\top - \widehat{\mathbf{D}}_l)^\top (\mathbf{A}_{:,l}^* \mathbf{E}_{:,l}^\top - \widehat{\mathbf{D}}_l) \right].\end{aligned}$$

Straightforward computation shows that

$$\begin{aligned}(\widehat{\Sigma}_{\text{crs}})_{i,i} &= \sum_{1 \leq l \leq d_2} A_{i,l}^{*2} \mathbb{E} \left[ \|\mathbf{E}_{:,l}\|_2^2 - E_{i,l}^2 \right], \quad i \in [d_1], \\ (\widehat{\Sigma}_{\text{crs}})_{i,j} &= \sum_{1 \leq l \leq d_2} A_{i,l}^* A_{j,l}^* \mathbb{E} \left[ \|\mathbf{E}_{:,l}\|_2^2 - E_{i,l}^2 - E_{j,l}^2 \right], \quad i \neq j,\end{aligned}$$

and  $\widetilde{\Sigma}_{\text{crs}} \in \mathbb{R}^{d_1 \times d_1}$  is a diagonal matrix with entries

$$(\widetilde{\Sigma}_{\text{crs}})_{i,i} = \sum_{1 \leq l \leq d_2} \left( \|\mathbf{A}_{:,l}^*\|_2^2 - A_{i,l}^{*2} \right) \mathbb{E} \left[ E_{i,l}^2 \right], \quad i \in [d_1].$$

Hence we have

$$\|\widetilde{\Sigma}_{\text{crs}}\| \leq \max_{1 \leq i \leq d_1} |(\widetilde{\Sigma}_{\text{crs}})_{i,i}| \lesssim \sigma_{\text{row}}^2 \|\mathbf{A}^{*\top}\|_{2,\infty}^2. \quad (111)$$

To control the spectral norm of  $\widehat{\Sigma}_{\text{crs}}$ , we further decompose it as  $\widehat{\Sigma}_{\text{crs}} = \widehat{\Sigma}'_{\text{crs}} - \widehat{\Sigma}''_{\text{crs}}$ , where

$$\begin{aligned}(\widehat{\Sigma}'_{\text{crs}})_{i,i} &= \sum_{1 \leq l \leq d_2} A_{i,l}^{*2} \mathbb{E} \left[ \|\mathbf{E}_{:,l}\|_2^2 \right], \quad i \in [d_1], \\ (\widehat{\Sigma}'_{\text{crs}})_{i,j} &= \sum_{1 \leq l \leq d_2} A_{i,l}^* A_{j,l}^* \mathbb{E} \left[ \|\mathbf{E}_{:,l}\|_2^2 \right], \quad i \neq j,\end{aligned}$$

and

$$\begin{aligned}(\widehat{\Sigma}''_{\text{crs}})_{i,i} &= \sum_{1 \leq l \leq d_2} A_{i,l}^{*2} \mathbb{E} \left[ E_{i,l}^2 \right], \quad i \in [d_1], \\ (\widehat{\Sigma}''_{\text{crs}})_{i,j} &= \sum_{1 \leq l \leq d_2} A_{i,l}^* A_{j,l}^* \mathbb{E} \left[ E_{i,l}^2 + E_{j,l}^2 \right], \quad i \neq j.\end{aligned}$$

- The spectral norm of  $\widehat{\Sigma}'_{\text{crs}}$  can be easily upper bounded by

$$\|\widehat{\Sigma}'_{\text{crs}}\| \leq \max_{1 \leq l \leq d_2} \mathbb{E} \left[ \|\mathbf{E}_{:,l}\|_2^2 \right] \|\mathbf{A}^* \mathbf{A}^{*\top}\| \leq \sigma_{\text{col}}^2 \|\mathbf{A}^*\|^2. \quad (112)$$

- Regarding  $\widehat{\Sigma}''_{\text{crs}}$ , we first decompose  $\widehat{\Sigma}''_{\text{crs}} + \mathcal{P}_{\text{diag}}(\widehat{\Sigma}''_{\text{crs}}) = \mathbf{B}_1 + \mathbf{B}_2$ , where the diagonal entries of  $\mathbf{B}_1$  and  $\mathbf{B}_2$  are identical and equal to  $\sum_{1 \leq l \leq d_2} A_{i,l}^{*2} \mathbb{E} [E_{i,l}^2]$ , ( $1 \leq i \leq d_2$ ) while their off-diagonal parts are given by

$$(\mathbf{B}_1)_{i,j} = \sum_{1 \leq l \leq d_2} A_{i,l}^* \mathbb{E} [E_{i,l}^2] A_{j,l}^* \quad \text{and} \quad (\mathbf{B}_2)_{i,j} = \sum_{1 \leq l \leq d_2} A_{i,l}^* A_{j,l}^* \mathbb{E} [E_{j,l}^2], \quad i \neq j.$$

Let  $\mathbf{C}$  be a matrix in  $\mathbb{R}^{d_1 \times d_2}$  with entries  $C_{i,j} = A_{i,j}^* \mathbb{E} [E_{i,j}^2]$ . One can easily check that  $\mathbf{B}_1 = \sum_{1 \leq l \leq d_2} \mathbf{C}_{:,l} \mathbf{A}_{:,l}^{*\top} = \mathbf{C} \mathbf{A}^{*\top}$  and develop an upper bound

$$\|\mathbf{B}_1\| \leq \|\mathbf{C}\| \|\mathbf{A}^*\| \leq \|\mathbf{C}\|_{\text{F}} \|\mathbf{A}^*\| \leq \sigma_{\infty}^2 \|\mathbf{A}^*\|_{\text{F}} \|\mathbf{A}^*\|.$$

Note that the same bound also holds for  $\mathbf{B}_2$ . Therefore, we arrive at

$$\begin{aligned}\|\widehat{\Sigma}''_{\text{crs}}\| &\leq \|\mathcal{P}_{\text{diag}}(\widehat{\Sigma}''_{\text{crs}})\| + \|\mathcal{P}_{\text{diag}}(\widehat{\Sigma}''_{\text{crs}}) + \widehat{\Sigma}''_{\text{crs}}\| \leq \|\mathcal{P}_{\text{diag}}(\widehat{\Sigma}''_{\text{crs}})\| + \|\mathbf{B}_1\| + \|\mathbf{B}_2\| \\ &\leq \sigma_{\infty}^2 \|\mathbf{A}^*\|_{2,\infty}^2 + \sigma_{\infty}^2 \|\mathbf{A}^*\|_{\text{F}} \|\mathbf{A}^*\| \\ &\leq \sigma_{\infty}^2 \|\mathbf{A}^*\|_{2,\infty}^2 + \sigma_{\infty}^2 \sqrt{r} \|\mathbf{A}^*\|^2 \\ &\leq 2\sigma_{\infty}^2 \sqrt{r} \|\mathbf{A}^*\|^2,\end{aligned}$$

where we have used the facts that  $\|\mathbf{A}^*\|_{2,\infty} \leq \|\mathbf{A}^*\|$  and  $\|\mathbf{A}^*\|_F \leq \sqrt{r} \|\mathbf{A}^*\|$ . Consequently, the above bounds taken collectively yield

$$\|\widehat{\Sigma}_{\text{crs}}\| \leq \|\widehat{\Sigma}'_{\text{crs}}\| + \|\widehat{\Sigma}''_{\text{crs}}\| \lesssim (\sigma_{\text{col}}^2 + \sigma_{\infty}^2 \sqrt{r}) \|\mathbf{A}^*\|^2 \asymp \sigma_{\text{col}}^2 \|\mathbf{A}^*\|^2, \quad (113)$$

where the last step uses (97b) and (97d).

Putting (111), (112) and (113) together yields

$$V_{\text{crs}} := \max \left\{ \|\widehat{\Sigma}_{\text{crs}}\|, \|\widetilde{\Sigma}_{\text{crs}}\| \right\} \lesssim \sigma_{\text{col}}^2 \|\mathbf{A}^*\|^2 + \sigma_{\text{row}}^2 \|\mathbf{A}^{*\top}\|_{2,\infty}^2. \quad (114)$$

Second, we turn to the spectral norm of each summand  $\mathbf{A}_{:,l}^* \mathbf{E}_{:,l}^\top - \widehat{\mathbf{D}}_l$ . Recalling the definition that  $\widehat{\mathbf{D}}_l = \text{diag}(A_{1,l}^* E_{1,l}, \dots, A_{d_1,l}^* E_{d_1,l})$ , we can obtain

$$\|\mathbf{A}_{:,l}^* \mathbf{E}_{:,l}^\top - \widehat{\mathbf{D}}_l\| \leq \|\mathbf{A}_{:,l}^*\|_2 \|\mathbf{E}_{:,l}\|_2 + \|\widehat{\mathbf{D}}_l\| \leq 2 \|\mathbf{A}_{:,l}^*\|_2 \|\mathbf{E}_{:,l}\|_2 \leq 2 \|\mathbf{A}^{*\top}\|_{2,\infty} \|\mathbf{E}_{:,l}\|_2.$$

Set

$$L_{\text{crs}} := C \sqrt{L_{\text{ns}}} \|\mathbf{A}^{*\top}\|_{2,\infty} \asymp (\sigma_{\text{col}} + B \sqrt{\log d}) \|\mathbf{A}^{*\top}\|_{2,\infty}, \quad (115)$$

where  $L_{\text{ns}}$  is defined in (106) and  $C > 0$  is some sufficiently large universal constant. Then with probability at least  $1 - O(d^{-20})$ , one has

$$\|\mathbf{A}_{:,l}^* \mathbf{E}_{:,l}^\top - \widehat{\mathbf{D}}_l\| \leq 2 \|\mathbf{A}^{*\top}\|_{2,\infty} \|\mathbf{E}_{:,l}\|_2 \lesssim L_{\text{crs}},$$

where the last inequality comes from (105).

Third, we need to control

$$R_{\text{crs}} := \mathbb{E} \left[ \|\mathbf{A}_{:,l}^* \mathbf{E}_{:,l}^\top - \widehat{\mathbf{D}}_l\| \mathbf{1} \{ \|\mathbf{A}_{:,l}^* \mathbf{E}_{:,l}^\top - \widehat{\mathbf{D}}_l\| \geq L_{\text{crs}} \} \right].$$

From Jensen's inequality and (107), we know that

$$\mathbb{E} \left[ \|\mathbf{E}_{:,l}\|_2 \mathbf{1} \{ \|\mathbf{E}_{:,l}\|_2 \geq \sqrt{L_{\text{ns}}} \} \right] \leq \sqrt{\mathbb{E} \left[ \|\mathbf{E}_{:,l}\|_2^2 \mathbf{1} \{ \|\mathbf{E}_{:,l}\|_2^2 \geq L_{\text{ns}} \} \right]} \ll \frac{\sqrt{L_{\text{ns}}}}{d}.$$

By the definition of  $L_{\text{crs}}$  in (115) and the fact that

$$\left\{ \|\mathbf{A}_{:,l}^*\|_2 \|\mathbf{E}_{:,l}\|_2 \geq L_{\text{crs}} \right\} \subset \left\{ \|\mathbf{A}^{*\top}\|_{2,\infty} \|\mathbf{E}_{:,l}\|_2 \geq L_{\text{crs}} \right\} = \left\{ \|\mathbf{E}_{:,l}\|_2 \geq C \sqrt{L_{\text{ns}}} \right\},$$

one has

$$\begin{aligned} R_{\text{crs}} &\leq \mathbb{E} \left[ 2 \|\mathbf{A}_{:,l}^*\|_2 \|\mathbf{E}_{:,l}\|_2 \mathbf{1} \{ 2 \|\mathbf{A}_{:,l}^*\|_2 \|\mathbf{E}_{:,l}\|_2 \geq L_{\text{crs}} \} \right] \\ &\lesssim \|\mathbf{A}^{*\top}\|_{2,\infty} \mathbb{E} \left[ \|\mathbf{E}_{:,l}\|_2 \mathbf{1} \{ \|\mathbf{E}_{:,l}\|_2 \geq \frac{C}{2} \sqrt{L_{\text{ns}}} \} \right] \\ &\ll \frac{1}{d} \|\mathbf{A}^{*\top}\|_{2,\infty} \sqrt{L_{\text{ns}}} \asymp \frac{L_{\text{crs}}}{d}. \end{aligned} \quad (116)$$

With (114), (115) and (116) in place, we can apply the truncated matrix Bernstein inequality to obtain that, with with probability at least  $1 - O(d^{-10})$ ,

$$\begin{aligned} \|\mathcal{P}_{\text{off-diag}}(\mathbf{A}^* \mathbf{E}^\top)\| &\leq \left\| \sum_{1 \leq l \leq d_2} (\mathbf{A}_{:,l}^* \mathbf{E}_{:,l}^\top - \widehat{\mathbf{D}}_l) \right\| \lesssim d_2 R_{\text{crs}} + L_{\text{crs}} \log d + \sqrt{V_{\text{crs}} \log d} \\ &\asymp L_{\text{crs}} \log d + \sqrt{V_{\text{crs}} \log d} \\ &\lesssim (\sigma_{\text{col}} \log d + B \log^{3/2} d) \|\mathbf{A}^{*\top}\|_{2,\infty} + \sigma_{\text{row}} \sqrt{\log d} \|\mathbf{A}^{*\top}\|_{2,\infty} + \sigma_{\text{col}} \sqrt{\log d} \|\mathbf{A}^*\| \\ &\lesssim (\sigma_{\text{row}} + \sigma_{\text{col}}) \log d \|\mathbf{A}^{*\top}\|_{2,\infty} + \sigma_{\text{col}} \sqrt{\log d} \|\mathbf{A}^*\| \end{aligned} \quad (117)$$

under the condition (98a) that  $B \sqrt{\log d} \lesssim \sqrt{\sigma_{\text{row}} \sigma_{\text{col}}} \leq \max\{\sigma_{\text{row}}, \sigma_{\text{col}}\}$ .

### B.2.3 Step 3: combining Step 1 and Step 2

Taking together (108), (117) and (101), we conclude that

$$\begin{aligned}\|\mathbf{G} - \mathbf{G}^*\| &\lesssim (\sigma_{\text{row}} + \sigma_{\text{col}}) (\sigma_{\text{col}} + \|\mathbf{A}^{*\top}\|_{2,\infty}) \log d + \sigma_{\text{col}} \sqrt{\log d} \|\mathbf{A}^*\| + \|\mathcal{P}_{\text{diag}}(\mathbf{A}^* \mathbf{A}^{*\top})\| \\ &\lesssim (\sigma_{\text{row}} + \sigma_{\text{col}}) (\sigma_{\text{col}} + \|\mathbf{A}^{*\top}\|_{2,\infty}) \log d + \sigma_{\text{col}} \sqrt{\log d} \|\mathbf{A}^*\| + \|\mathbf{A}^*\|_{2,\infty}^2,\end{aligned}$$

where we have used the basic property  $\|\mathcal{P}_{\text{diag}}(\mathbf{A}^* \mathbf{A}^{*\top})\| = \|\mathbf{A}^*\|_{2,\infty}^2$ .

### B.3 Proof of Lemma 2

We first claim that, for any fixed matrix  $\mathbf{W}$ , with probability at least  $1 - O(d^{-10})$ , the following holds for any  $1 \leq m \leq d_1$ :

$$\begin{aligned}\|(\mathbf{G} - \mathbf{G}^*)_{m,:} \mathbf{W}\|_2 &\lesssim \left( \sigma_{\text{col}} (\sigma_{\text{row}} + \|\mathbf{A}^*\|_{2,\infty}) \sqrt{\log d} + B \log d \|\mathbf{A}^*\|_{\infty} + \|\mathbf{A}^*\|_{2,\infty}^2 \right) \|\mathbf{W}\|_{2,\infty} \\ &\quad + \sigma_{\text{row}} \sqrt{\log d} \|\mathbf{A}^{*\top}\|_{2,\infty} \|\mathbf{W}\|.\end{aligned}\tag{118}$$

In particular, taking  $\mathbf{W} = \mathbf{U}^*$  gives

$$\|(\mathbf{G} - \mathbf{G}^*)_{m,:} \mathbf{U}^*\|_2 \lesssim \delta_{\text{row}} \sqrt{\frac{\mu r}{d_1}},$$

where

$$\delta_{\text{row}} := \sigma_{\text{col}} (\sigma_{\text{row}} + \|\mathbf{A}^*\|_{2,\infty}) \sqrt{\log d} + B \log d \|\mathbf{A}^*\|_{\infty} + \sqrt{\frac{d_1 \log d}{\mu r}} \sigma_{\text{row}} \|\mathbf{A}^{*\top}\|_{2,\infty} + \|\mathbf{A}^*\|_{2,\infty}^2.\tag{119}$$

Using the values of  $B, \sigma_{\infty}, \sigma_{\text{row}}$  and  $\sigma_{\text{col}}$  specified in (97), one can easily verify that

$$\delta_{\text{row}} \lesssim \zeta_{\text{op}} + \|\mathbf{A}^*\|_{2,\infty}^2,\tag{120}$$

where  $\zeta_{\text{op}}$  is defined in (48). This leads to the advertised bound.

The rest of the proof is thus devoted to proving the claim (118). Recall the definitions of  $\mathbf{G}$  and  $\mathbf{G}^*$  in (13) and (47). For any  $m, i \in [d_1]$ , we can expand

$$\begin{aligned}(\mathbf{G} - \mathbf{G}^*)_{m,i} &= \langle \mathbf{A}_{m,:}^s, \mathbf{A}_{i,:}^s \rangle - \langle \mathbf{A}_{m,:}^*, \mathbf{A}_{i,:}^* \rangle = \langle \mathbf{E}_{m,:}, \mathbf{E}_{i,:} \rangle + \langle \mathbf{A}_{m,:}^*, \mathbf{E}_{i,:} \rangle + \langle \mathbf{E}_{m,:}, \mathbf{A}_{i,:}^* \rangle, \quad i \neq m; \\ (\mathbf{G} - \mathbf{G}^*)_{m,m} &= -G_{m,m}^* = -\|\mathbf{A}_{m,:}^*\|_2^2.\end{aligned}$$

This allows us to derive

$$\begin{aligned}\|(\mathbf{G} - \mathbf{G}^*)_{m,:} \mathbf{W}\|_2 &\leq \left\| \sum_{i:i \neq m} \langle \mathbf{E}_{m,:}, \mathbf{E}_{i,:} \rangle \mathbf{W}_{i,:} \right\|_2 + \left\| \sum_{i:i \neq m} \langle \mathbf{A}_{m,:}^*, \mathbf{E}_{i,:} \rangle \mathbf{W}_{i,:} \right\|_2 \\ &\quad + \left\| \sum_{i:i \neq m} \langle \mathbf{E}_{m,:}, \mathbf{A}_{i,:}^* \rangle \mathbf{W}_{i,:} \right\|_2 + \|G_{m,m}^* \mathbf{W}_{m,:}\|_2.\end{aligned}\tag{121}$$

We shall control each of these four terms separately.

- For the first term on the right-hand side of (121), we know that

$$\sum_{i:i \neq m} \langle \mathbf{E}_{m,:}, \mathbf{E}_{i,:} \rangle \mathbf{W}_{i,:} = \sum_{(i,j):i \neq m} E_{m,j} E_{i,j} \mathbf{W}_{i,:}$$

is a sum of independent zero-mean random vectors conditional on  $\{E_{m,j}\}_{j \in [d_2]}$ . In view of the matrix Bernstein inequality, it suffices to control the following two quantities

$$L_1 := \max_{(i,j):i \neq m} \|E_{m,j} E_{i,j} \mathbf{W}_{i,:}\|_2 \leq \max_{i,j} |E_{i,j}|^2 \|\mathbf{W}\|_{2,\infty} \leq B^2 \|\mathbf{W}\|_{2,\infty},$$

$$V_1 := \sum_{(i,j):i \neq m} E_{m,j}^2 \mathbb{E} [E_{i,j}^2] \|\mathbf{W}_{i,:}\|_2^2 \leq \|\mathbf{W}\|_{2,\infty}^2 \sum_j E_{m,j}^2 \sum_i \mathbb{E} [E_{i,j}^2] \leq \|\mathbf{W}\|_{2,\infty}^2 \sigma_{\text{col}}^2 \sum_j E_{m,j}^2,$$

where  $B$  and  $\sigma_{\text{col}}$  are defined in (97). According to Lemma 12, the following holds with probability at least  $1 - O(d^{-12})$ ,

$$V_1 \lesssim \sigma_{\text{col}}^2 (\sigma_{\text{row}}^2 + B^2 \log d) \|\mathbf{W}\|_{2,\infty}^2 \asymp \sigma_{\text{col}}^2 \sigma_{\text{row}}^2 \|\mathbf{W}\|_{2,\infty}^2,$$

where we use the condition (98a) (namely,  $B \lesssim \sigma_{\text{row}}/\sqrt{\log d}$ ). Apply the matrix Bernstein inequality to demonstrate that with probability exceeding  $1 - O(d^{-12})$ ,

$$\begin{aligned} \left\| \sum_{i:i \neq m} \langle \mathbf{E}_{m,:}, \mathbf{E}_{i,:} \rangle \mathbf{W}_{i,:} \right\|_2 &\lesssim L_1 \log d + \sqrt{V_1 \log d} \lesssim (B^2 \log d + \sigma_{\text{col}} \sigma_{\text{row}} \sqrt{\log d}) \|\mathbf{W}\|_{2,\infty} \\ &\asymp \sigma_{\text{col}} \sigma_{\text{row}} \sqrt{\log d} \|\mathbf{W}\|_{2,\infty}, \end{aligned} \quad (122)$$

where the last line follows from (98a) (i.e.  $B \lesssim \sqrt{\sigma_{\text{row}} \sigma_{\text{col}} / \log d}$ ).

- Regarding the second term on the right-hand side of (121), apply the same argument as above to show that

$$\left\| \sum_{i:i \neq m} \langle \mathbf{A}_{m,:}^*, \mathbf{E}_{i,:} \rangle \mathbf{W}_{i,:} \right\|_2 \lesssim (\sigma_{\text{col}} \sqrt{\log d} \|\mathbf{A}^*\|_{2,\infty} + B \log d \|\mathbf{A}^*\|_{\infty}) \|\mathbf{W}\|_{2,\infty} \quad (123)$$

holds with probability at least  $1 - O(d^{-12})$ .

- Turning to the third term on the right-hand side of (121), we have

$$\left\| \sum_{i:i \neq m} \langle \mathbf{E}_{m,:}, \mathbf{A}_{i,:}^* \rangle \mathbf{W}_{i,:} \right\|_2 \leq \left\| \sum_{1 \leq j \leq d_2} E_{m,j} (\mathbf{A}^{*\top} \mathbf{W})_{j,:} \right\|_2,$$

where the summands are independent zero-mean random vectors. Let us compute that

$$\begin{aligned} L_2 &:= \max_{j \in [d_2]} \left\| E_{m,j} (\mathbf{A}^{*\top} \mathbf{W})_{j,:} \right\|_2 \leq B \|\mathbf{A}^{*\top} \mathbf{W}\|_{2,\infty} \leq B \|\mathbf{A}^{*\top}\|_{2,\infty} \|\mathbf{W}\|; \\ V_2 &= \sum_j \mathbb{E} [E_{m,j}^2] \left\| (\mathbf{A}^{*\top} \mathbf{W})_{j,:} \right\|_2^2 \leq \sigma_{\text{row}}^2 \|\mathbf{A}^{*\top} \mathbf{W}\|_{2,\infty}^2 \leq \sigma_{\text{row}}^2 \|\mathbf{A}^{*\top}\|_{2,\infty}^2 \|\mathbf{W}\|^2. \end{aligned}$$

Then the matrix Bernstein inequality reveals that with probability exceeding  $1 - O(d^{-12})$ ,

$$\begin{aligned} \left\| \sum_j E_{m,j} (\mathbf{A}^{*\top} \mathbf{W})_{j,:} \right\|_2 &\lesssim L_2 \log d + \sqrt{V_2 \log d} \lesssim (B \log d + \sigma_{\text{row}} \sqrt{\log d}) \|\mathbf{A}^{*\top}\|_{2,\infty} \|\mathbf{W}\| \\ &\asymp \sigma_{\text{row}} \sqrt{\log d} \|\mathbf{A}^{*\top}\|_{2,\infty} \|\mathbf{W}\|, \end{aligned} \quad (124)$$

where the last line follows from the condition (98a) (i.e.  $B \lesssim \sigma_{\text{row}}/\sqrt{\log d}$ ).

- The last term on the right-hand side of (121) can simply be upper bounded by

$$\left\| \mathbf{G}_{m,m}^* \mathbf{W}_{m,:} \right\|_2 \leq \|\mathbf{A}^*\|_{2,\infty}^2 \|\mathbf{W}\|_{2,\infty}. \quad (125)$$

Putting (122), (123), (124) and (125) together yields

$$\begin{aligned} \left\| (\mathbf{G} - \mathbf{G}^*)_{m,:} \mathbf{W} \right\|_2 &\lesssim \left( \sigma_{\text{col}} (\sigma_{\text{row}} + \|\mathbf{A}^*\|_{2,\infty}) \sqrt{\log d} + B \log d \|\mathbf{A}^*\|_{\infty} + \|\mathbf{A}^*\|_{2,\infty}^2 \right) \|\mathbf{W}\|_{2,\infty} \\ &\quad + \sigma_{\text{row}} \sqrt{\log d} \|\mathbf{A}^{*\top}\|_{2,\infty} \|\mathbf{W}\| \end{aligned}$$

as claimed.

## B.4 Proof of Lemma 3

We claim for the moment that

$$\|\mathbf{U}\mathbf{H} - \mathbf{G}\mathbf{U}^*(\boldsymbol{\Sigma}^*)^{-2}\|_{2,\infty} \lesssim \frac{(\delta_{\text{op}} + \delta_{\text{loo}})\kappa^2}{\sigma_r^{*2}} \left( \|\mathbf{U}\mathbf{H}\|_{2,\infty} + \sqrt{\frac{\mu r}{d_1}} \right), \quad (126)$$

where  $\delta_{\text{op}}$  is defined in (99), and  $\delta_{\text{loo}}$  is defined as follows:

$$\delta_{\text{loo}} := \sigma_{\text{col}}\sigma_{\text{row}} \log d + \sigma_{\text{col}} \|\mathbf{A}^*\| \sqrt{\log d}. \quad (127)$$

Using the values of  $\sigma_{\text{row}}$  and  $\sigma_{\text{col}}$  specified in (97), one can easily see that  $\delta_{\text{loo}} \lesssim \zeta_{\text{op}}$ , where  $\zeta_{\text{op}}$  is defined in (48). In addition, recall that we have already shown that  $\delta_{\text{op}} \lesssim \zeta_{\text{op}} + \|\mathbf{A}^*\|_{2,\infty}^2$ . Putting these together establishes the lemma.

We now start to prove the claim (126). To this end, consider an arbitrary  $m \in [d_1]$ . In view of [AFWZ17, Lemma 1], we can decompose

$$\left\| (\mathbf{U}\mathbf{H} - \mathbf{G}\mathbf{U}^*(\boldsymbol{\Sigma}^*)^{-2})_{m,:} \right\|_2 \lesssim \frac{1}{\sigma_r^{*4}} \|\mathbf{G} - \mathbf{G}^*\| \|\mathbf{G}_{m,:}\mathbf{U}^*\|_2 + \frac{1}{\sigma_r^{*2}} \|\mathbf{G}_{m,:}\|_2 \|\mathbf{U}\mathbf{H} - \mathbf{U}^*\|_2. \quad (128)$$

- To bound the first term of (128), we apply Lemma 2, (120) and Fact 1 to reach

$$\|(\mathbf{G} - \mathbf{G}^*)\mathbf{U}^*\|_{2,\infty} \lesssim \delta_{\text{row}} \sqrt{\frac{\mu r}{d_1}} \ll \sigma_r^{*2} \sqrt{\frac{\mu r}{d_1}}.$$

The triangle inequality then gives

$$\|\mathbf{G}\mathbf{U}^*\|_{2,\infty} \leq \|(\mathbf{G} - \mathbf{G}^*)\mathbf{U}^*\|_{2,\infty} + \|\mathbf{G}^*\| \|\mathbf{U}^*\|_{2,\infty} \lesssim \sigma_1^{*2} \sqrt{\frac{\mu r}{d_1}}. \quad (129)$$

This taken collectively with the upper bound on  $\|\mathbf{G} - \mathbf{G}^*\|$  (cf. (99)) gives

$$\frac{1}{\sigma_r^{*4}} \|\mathbf{G} - \mathbf{G}^*\| \|\mathbf{G}_{m,:}\mathbf{U}^*\|_2 \leq \frac{1}{\sigma_r^{*4}} \|\mathbf{G} - \mathbf{G}^*\| \|\mathbf{G}\mathbf{U}^*\|_{2,\infty} \lesssim \frac{\delta_{\text{op}}\kappa^2}{\sigma_r^{*2}} \sqrt{\frac{\mu r}{d_1}}. \quad (130)$$

- Turning to the second term of (128), we start with the following bound

$$\|\mathbf{G}_{m,:}(\mathbf{U}\mathbf{H} - \mathbf{U}^*)\|_2 \leq \|\mathbf{G}_{m,:}(\mathbf{U}\mathbf{H} - \mathbf{U}^{(m)}\mathbf{H}^{(m)})\|_2 + \|\mathbf{G}_{m,:}(\mathbf{U}^{(m)}\mathbf{H}^{(m)} - \mathbf{U}^*)\|_2.$$

Lemma 2 tells us that

$$\|\mathbf{G}_{m,:}\|_2 \leq \|(\mathbf{G} - \mathbf{G}^*)_{m,:}\|_2 + \|\mathbf{G}^*\|_{2,\infty} \lesssim \sigma_1^{*2} \|\mathbf{U}^*\|_{2,\infty} \leq \sigma_1^{*2} \|\mathbf{U}^*\| = \sigma_1^{*2},$$

which makes use of the fact that  $\|\mathbf{G}^*\|_{2,\infty} \leq \|\mathbf{U}^*\|_{2,\infty} \|\boldsymbol{\Sigma}^*\|^2 \|\mathbf{U}^{*\top}\| = \sigma_1^{*2} \|\mathbf{U}^*\|_{2,\infty}$ . This combined with (133) (to be established shortly in the proof of Lemma 4) and the definitions of  $\mathbf{H}$  and  $\mathbf{H}^{(m)}$  gives

$$\begin{aligned} \|\mathbf{G}_{m,:}(\mathbf{U}\mathbf{H} - \mathbf{U}^{(m)}\mathbf{H}^{(m)})\|_2 &\leq \|\mathbf{G}_{m,:}\|_2 \|\mathbf{U}^{(m)}\mathbf{H}^{(m)} - \mathbf{U}\mathbf{H}\| \\ &= \|\mathbf{G}_{m,:}\|_2 \|(\mathbf{U}^{(m)}\mathbf{U}^{(m)\top} - \mathbf{U}\mathbf{U}^\top)\mathbf{U}^*\| \\ &\leq \|\mathbf{G}_{m,:}\|_2 \|\mathbf{U}^{(m)}\mathbf{U}^{(m)\top} - \mathbf{U}\mathbf{U}^\top\| \\ &\lesssim \frac{\sigma_1^{*2}}{\sigma_r^{*2}} \delta_{\text{loo}} \left( \|\mathbf{U}\mathbf{H}\|_{2,\infty} + \sqrt{\frac{\mu r}{d_1}} \right) \\ &\leq \delta_{\text{loo}}\kappa^2 \left( \|\mathbf{U}\mathbf{H}\|_{2,\infty} + \sqrt{\frac{\mu r}{d_1}} \right). \end{aligned} \quad (131)$$

In addition, Lemma 10 shows that with probability at least  $1 - O(d^{-11})$ ,

$$\|\mathbf{G}_{m,:}(\mathbf{U}^{(m)}\mathbf{H}^{(m)} - \mathbf{U}^*)\|_2 \lesssim \delta_{\text{loo}} \|\mathbf{U}^{(m)}\mathbf{H}^{(m)} - \mathbf{U}^*\|_{2,\infty} + \delta_{\text{op}}\kappa^2 \sqrt{\frac{\mu r}{d_1}}$$

$$\begin{aligned}
&\leq \delta_{\text{loo}} \|\mathbf{U}^{(m)} \mathbf{H}^{(m)}\|_{2,\infty} + \delta_{\text{loo}} \|\mathbf{U}^*\|_{2,\infty} + \delta_{\text{op}} \kappa^2 \sqrt{\frac{\mu r}{d_1}} \\
&\lesssim \delta_{\text{loo}} \|\mathbf{U} \mathbf{H}\|_{2,\infty} + (\delta_{\text{loo}} + \delta_{\text{op}} \kappa^2) \sqrt{\frac{\mu r}{d_1}}, \tag{132}
\end{aligned}$$

where the inequality (132) results from (143) (also established shortly in the proof of Lemma 4).

Then claim immediately follows from (130), (131), (132) and the union bound.

## B.5 Proof of Lemma 4

To begin with, recalling the definition of  $\delta_{\text{loo}}$  (cf. (127)), we claim that

$$\|\mathbf{U}^{(m)} \mathbf{U}^{(m)\top} - \mathbf{U} \mathbf{U}^\top\|_{\text{F}} = \frac{1}{\sigma_r^{*2}} \underbrace{(\sigma_{\text{col}} \sigma_{\text{row}} \log d + \sigma_{\text{col}} \|\mathbf{A}^*\| \sqrt{\log d})}_{=\delta_{\text{loo}}} \left( \|\mathbf{U} \mathbf{H}\|_{2,\infty} + \sqrt{\frac{\mu r}{d_1}} \right). \tag{133}$$

As mentioned before, one has  $\delta_{\text{loo}} \lesssim \zeta_{\text{op}}$ , from which the lemma follows immediately. The rest of the proof thus boils down to proving the claim (133).

We shall apply the Davis-Kahan sin  $\Theta$  theorem [DK70] to derive

$$\|\mathbf{U} \mathbf{U}^\top - \mathbf{U}^{(m)} \mathbf{U}^{(m)\top}\|_{\text{F}} \leq \frac{\|(\mathbf{G} - \mathbf{G}^{(m)}) \mathbf{U}^{(m)}\|_{\text{F}}}{\lambda_r(\mathbf{G}^{(m)}) - \lambda_{r+1}(\mathbf{G})} \leq \frac{2 \|(\mathbf{G} - \mathbf{G}^{(m)}) \mathbf{U}^{(m)}\|_{\text{F}}}{\sigma_r^{*2}}. \tag{134}$$

Here, the last inequality follows since, by Weyl's inequality,

$$\begin{aligned}
\lambda_r(\mathbf{G}^{(m)}) - \lambda_{r+1}(\mathbf{G}) &\geq \lambda_r(\mathbf{G}^*) - \|\mathbf{G}^{(m)} - \mathbf{G}^*\| - \lambda_{r+1}(\mathbf{G}^*) - \|\mathbf{G} - \mathbf{G}^*\| \\
&= \sigma_r^{*2} - \|\mathbf{G}^{(m)} - \mathbf{G}^*\| - \|\mathbf{G} - \mathbf{G}^*\| \\
&\geq \sigma_r^{*2}/2, \tag{135}
\end{aligned}$$

where the last line follows since  $\|\mathbf{G}^{(m)} - \mathbf{G}^*\| \lesssim \delta_{\text{op}} \ll \sigma_r^{*2}$  — an immediate consequence of Lemma 6 and Condition (98d). As a side note, the fact  $\|\mathbf{G}^{(m)} - \mathbf{G}^*\| \ll \sigma_r^{*2}$  also implies (according to [AFWZ17, Lemma 3])

$$\|(\mathbf{H}^{(m)})^{-1}\| \lesssim 1, \tag{136}$$

which will be useful later.

It remains to control the term  $\|(\mathbf{G} - \mathbf{G}^{(m)}) \mathbf{U}^{(m)}\|_{\text{F}}$  in (134). Recall the definitions of  $\mathbf{G}$  and  $\mathbf{G}^{(m)}$  in (13) and (45), respectively. It is straightforward to see that  $\mathbf{G} - \mathbf{G}^{(m)}$  is a rank-2 symmetric matrix with nonzero entries located only in the  $m$ -th row and the  $m$ -th column. Simple calculation reveals that

$$(\mathbf{G} - \mathbf{G}^{(m)})_{m,i} = \langle \mathbf{E}_{m,:}, \mathbf{A}_{i,:}^s \rangle, \quad i \neq m; \tag{137}$$

$$(\mathbf{G} - \mathbf{G}^{(m)})_{m,m} = 0. \tag{138}$$

We can then derive

$$\begin{aligned}
\|(\mathbf{G} - \mathbf{G}^{(m)}) \mathbf{U}^{(m)}\|_{\text{F}} &= \|(\mathbf{G} - \mathbf{G}^{(m)}) \mathbf{U}^{(m)} \mathbf{H}^{(m)} (\mathbf{H}^{(m)})^{-1}\|_{\text{F}} \leq \|(\mathbf{G} - \mathbf{G}^{(m)}) \mathbf{U}^{(m)} \mathbf{H}^{(m)}\|_{\text{F}} \|(\mathbf{H}^{(m)})^{-1}\| \\
&\lesssim \|(\mathbf{G} - \mathbf{G}^{(m)}) \mathbf{U}^{(m)} \mathbf{H}^{(m)}\|_{\text{F}} \\
&\leq \|\mathcal{P}_{m,:} (\mathbf{G} - \mathbf{G}^{(m)}) \mathbf{U}^{(m)} \mathbf{H}^{(m)}\|_{\text{F}} + \|\mathcal{P}_{:,m} (\mathbf{G} - \mathbf{G}^{(m)}) \mathbf{U}^{(m)} \mathbf{H}^{(m)}\|_{\text{F}}, \tag{139}
\end{aligned}$$

where the second line arises due to (136), and  $\mathcal{P}_{m,:}$  (resp.  $\mathcal{P}_{:,m}$ ) is the projection onto the subspace of matrix supported on  $\{m\} \times [d_2]$  (resp.  $[d_1] \times \{m\}$ ).

To bound the first term of (139), we make the observation (using (137) and (138)) that

$$\|\mathcal{P}_{m,:} (\mathbf{G} - \mathbf{G}^{(m)}) \mathbf{U}^{(m)} \mathbf{H}^{(m)}\|_{\text{F}} = \|(\mathbf{A}^s - \mathbf{A}^*)_{m,:} [\mathcal{P}_{-m,:} (\mathbf{A}^s)]^\top \mathbf{U}^{(m)} \mathbf{H}^{(m)}\|_2.$$

Controlling this quantity requires the assistance of leave-two-out matrices. Here, we only state our bound: with probability at least  $1 - O(d^{-11})$ , one has

$$\begin{aligned} & \|\mathcal{P}_{m,\cdot}(\mathbf{G} - \mathbf{G}^{(m)})\mathbf{U}^{(m)}\mathbf{H}^{(m)}\|_{\text{F}} \\ & \lesssim \sigma_{\text{col}}(\sigma_{\text{row}} \log d + \|\mathbf{A}^*\| \sqrt{\log d}) \left( \|\mathbf{U}^{(m)}\mathbf{H}^{(m)}\|_{2,\infty} + \sqrt{\frac{\mu r}{d_1}} \right). \end{aligned} \quad (140)$$

This bound will be restated in Lemma 7 and established in Appendix C. Turning to the second term of (139), we apply Lemma 6 (also established in Appendix C) to obtain

$$\begin{aligned} \|\mathcal{P}_{\cdot,m}(\mathbf{G} - \mathbf{G}^{(m)})\mathbf{U}^{(m)}\mathbf{H}^{(m)}\|_{\text{F}} &= \left( \sum_{1 \leq i \leq d_1} (\mathbf{G} - \mathbf{G}^{(m)})_{i,m}^2 \right)^{1/2} \left\| (\mathbf{U}^{(m)}\mathbf{H}^{(m)})_{m,\cdot} \right\|_2 \\ &\leq \|\mathbf{G} - \mathbf{G}^{(m)}\| \|\mathbf{U}^{(m)}\mathbf{H}^{(m)}\|_{2,\infty} \\ &\lesssim (\sigma_{\text{col}}(\sigma_{\text{row}} + \|\mathbf{A}^*\|_{2,\infty}) \sqrt{\log d}) \|\mathbf{U}^{(m)}\mathbf{H}^{(m)}\|_{2,\infty} \end{aligned} \quad (141)$$

with probability at least  $1 - O(d^{-11})$ . Hence, we can combine (140), (141) and (134) to yield

$$\begin{aligned} & \|\mathbf{U}^{(m)}\mathbf{U}^{(m)\top} - \mathbf{U}\mathbf{U}^\top\|_{\text{F}} \\ & \lesssim (\sigma_{\text{col}}\sigma_{\text{row}} \log d + \sigma_{\text{col}} \|\mathbf{A}^*\| \sqrt{\log d}) \left( \|\mathbf{U}^{(m)}\mathbf{H}^{(m)}\|_{2,\infty} + \sqrt{\frac{\mu r}{d_1}} \right) \\ & = \frac{\delta_{\text{lo0}}}{\sigma_r^{*2}} \left( \|\mathbf{U}^{(m)}\mathbf{H}^{(m)}\|_{2,\infty} + \sqrt{\frac{\mu r}{d_1}} \right), \end{aligned} \quad (142)$$

where  $\delta_{\text{lo0}}$  is defined in (127). As a result, the proof is complete as long as we can show that

$$\|\mathbf{U}^{(m)}\mathbf{H}^{(m)}\|_{2,\infty} \lesssim \|\mathbf{U}\mathbf{H}\|_{2,\infty} + \sqrt{\frac{\mu r}{d_1}}. \quad (143)$$

To finish up, it remains to justify this inequality (143). To this end, from the definitions of  $\mathbf{H}^{(m)}$  and  $\mathbf{H}$  we have

$$\begin{aligned} \|\mathbf{U}^{(m)}\mathbf{H}^{(m)}\|_{2,\infty} &\leq \|\mathbf{U}^{(m)}\mathbf{H}^{(m)} - \mathbf{U}\mathbf{H}\|_{2,\infty} + \|\mathbf{U}\mathbf{H}\|_{2,\infty} \\ &= \|(\mathbf{U}^{(m)}\mathbf{U}^{(m)\top} - \mathbf{U}\mathbf{U}^\top)\mathbf{U}^*\|_{2,\infty} + \|\mathbf{U}\mathbf{H}\|_{2,\infty} \\ &\leq \|\mathbf{U}^{(m)}\mathbf{U}^{(m)\top} - \mathbf{U}\mathbf{U}^\top\|_{\text{F}} \|\mathbf{U}^*\| + \|\mathbf{U}\mathbf{H}\|_{2,\infty} \\ &= \|\mathbf{U}^{(m)}\mathbf{U}^{(m)\top} - \mathbf{U}\mathbf{U}^\top\|_{\text{F}} + \|\mathbf{U}\mathbf{H}\|_{2,\infty}. \end{aligned} \quad (144)$$

Under the condition (98d), it is easily seen that  $\delta_{\text{lo0}} \ll \sigma_r^{*2}$ . This together with (142) gives

$$\|\mathbf{U}^{(m)}\mathbf{U}^{(m)\top} - \mathbf{U}\mathbf{U}^\top\|_{\text{F}} \leq 0.5 \|\mathbf{U}^{(m)}\mathbf{H}^{(m)}\|_{2,\infty} + 0.5 \sqrt{\frac{\mu r}{d_1}}, \quad (145)$$

which combined with (144) yields

$$\|\mathbf{U}^{(m)}\mathbf{H}^{(m)}\|_{2,\infty} \leq 2 \|\mathbf{U}\mathbf{H}\|_{2,\infty} + \sqrt{\frac{\mu r}{d_1}} \quad (146)$$

as claimed.

## B.6 Proof of Lemma 5

We start with (98a). In view of the definitions of  $B, \sigma_\infty, \sigma_{\text{row}}$  and  $\sigma_{\text{col}}$  in (97), we have with probability at least  $1 - O(d^{-12})$ ,

$$B^2 = \frac{\mu r \sigma_1^{*2}}{d_1 d_2 p^2} + \frac{\sigma^2 \min\{\sqrt{d_1 d_2}, d_2\}}{p \log d},$$



$$\begin{aligned}\sigma_{\text{row}}^2 &= \frac{\mu r \sigma_1^{*2}}{d_1 p} + \frac{\sigma^2 d_2}{p}, \\ \sigma_{\text{row}} \sigma_{\text{col}} &= \frac{\mu r \sigma_1^{*2}}{\sqrt{d_1 d_2} p} + \frac{2\sigma \sigma_1^* \sqrt{\mu r}}{p} + \frac{\sigma^2 \sqrt{d_1 d_2}}{p} \asymp \frac{\mu r \sigma_1^{*2}}{\sqrt{d_1 d_2} p} + \frac{\sigma^2 \sqrt{d_1 d_2}}{p},\end{aligned}$$

where we have used the AM-GM inequality (i.e.  $2\sigma \sigma_1^* \sqrt{\mu r} \leq \frac{\mu r \sigma_1^{*2}}{\sqrt{d_1 d_2}} + \sigma^2 \sqrt{d_1 d_2}$ ) in the last line. Therefore,

$$B^2 \log d \lesssim \sigma_{\text{row}} \sigma_{\text{col}} \quad \text{and} \quad B^2 \log d \lesssim \sigma_{\text{row}}^2$$

hold as long as  $p \gtrsim (d_1 d_2)^{-1/2} \log d$  and  $p \gtrsim d_2^{-1} \log d$ .

The next step is to establish (98b). Let us consider the first inequality. By (97), it is easily seen that

$$\begin{aligned}\sigma_\infty^2 &= \frac{\mu r \sigma_1^{*2}}{d_1 d_2 p} + \frac{\sigma^2}{p}; \\ B \log d \|\mathbf{A}^*\| \sqrt{\frac{\mu r}{d_2}} &= \frac{\mu r \sigma_1^{*2} \log d}{\sqrt{d_1} d_2 p} + \sigma \sigma_1^* \min \left\{ \sqrt[4]{d_1 d_2}, \sqrt{d_2} \right\} \sqrt{\frac{\mu r \log d}{d_2 p}}.\end{aligned}$$

As a consequence, the first inequality holds as long as  $\frac{\sigma}{\sigma_1^*} \lesssim \min \left\{ \sqrt[4]{d_1 d_2}, \sqrt{d_2} \right\} \sqrt{\frac{\mu r p \log d}{d_2}}$ , which is satisfied by our noise assumption that  $\frac{\sigma}{\sigma_1^*} \ll \frac{\sqrt{p}}{\kappa \sqrt[4]{d_1 d_2} \sqrt{\log d}}$ . To show the second inequality, we note that

$$B \sqrt{\log d} \|\mathbf{A}^*\| \sqrt{\frac{\mu r}{d_2}} \leq \left\{ \frac{\sigma_1^{*2} \sqrt{\mu r \log d}}{d_2 p} + \sigma \sigma_1^* \sqrt{\frac{d_1}{p}} \right\} \sqrt{\frac{\mu r}{d_1}}.$$

Recognizing that  $\sigma_{\text{col}} \|\mathbf{A}^*\| = \sigma_1^{*2} \sqrt{\frac{\mu r}{d_2 p}} + \sigma \sigma_1^* \sqrt{\frac{d_1}{p}}$ , we prove the second inequality provided that  $p \gtrsim d_2^{-1} \log d$ .

When it comes to (98c): by virtue of Lemma 11 and (97), one has

$$\frac{B \log^{3/2} d \|\mathbf{A}^*\|_\infty}{\|\mathbf{A}^*\|^2} \leq \frac{\mu r \log^{3/2} d}{d_1 d_2 p} + \frac{\sigma \sqrt{\mu r} \log d}{\sigma_1^* \sqrt[4]{d_1 d_2} \sqrt{p}}.$$

Consequently, (98c) holds provided  $p \gtrsim \frac{\sqrt{\mu r} \log^{3/2} d}{\sqrt{d_1} d_2}$  and  $\frac{\sigma}{\sigma_1^*} \lesssim \sqrt[4]{\frac{d_2}{d_1} \frac{\sqrt{p}}{\log d}}$ , which holds under our assumptions that  $p \gg \frac{\mu \kappa^4 r \log^2 d}{\sqrt{d_1 d_2}}$  and  $\frac{\sigma}{\sigma_1^*} \ll \frac{\sqrt{p}}{\kappa^3 \sqrt[4]{d_1 d_2} \sqrt{\log d}}$ .

Finally, using the definitions in (97) and Lemma 11, one obtains the following bounds:

$$\begin{aligned}\sigma_{\text{row}} \sigma_{\text{col}} \log d &\asymp \frac{\mu r \sigma_1^{*2} \log d}{\sqrt{d_1 d_2} p} + \frac{\sigma^2 \sqrt{d_1 d_2} \log d}{p}, \\ \sigma_{\text{col}} \sqrt{\log d} \|\mathbf{A}^*\| &= \sigma_1^{*2} \sqrt{\frac{\mu r \log d}{d_2 p}} + \sigma \sigma_1^* \sqrt{\frac{d_1 \log d}{p}}; \\ \|\mathbf{A}^*\|_{2,\infty}^2 &\lesssim \frac{\mu_1 r \sigma_1^{*2}}{d_1}; \\ \sigma_{\text{row}} \sqrt{\frac{d_1 \log d}{\mu r}} \|\mathbf{A}^{*\top}\|_{2,\infty} &\leq \sigma_1^{*2} \sqrt{\frac{\mu r \log d}{d_2 p}} + \sigma \sigma_1^* \sqrt{\frac{d_1 \log d}{p}}; \\ B \|\mathbf{A}^*\|_\infty \log d &\leq \frac{\mu r \sigma_1^{*2} \log d}{d_1 d_2 p} + \frac{\sigma \sigma_1^* \sqrt{\mu r \log d}}{\sqrt[4]{d_1 d_2} \sqrt{p}}.\end{aligned}$$

Therefore, it is easy to verify (98d) under the assumptions of Theorem 1.

## C Proofs for auxiliary lemmas

This section establishes several useful technical lemmas useful for proving our main theorems. Throughout this section, we shall frequently use the quantities  $B, \sigma_\infty, \sigma_{\text{row}}$  and  $\sigma_{\text{col}}$  defined in (97). In fact, it suffices to bear in mind the following bounds

$$\begin{aligned} B &\geq \max_{i,j} |E_{i,j}|; & \sigma_\infty &\geq \max_{i,j} \sqrt{\mathbb{E}[E_{i,j}^2]}; \\ \sigma_{\text{row}} &\geq \max_i \sqrt{\sum_j \mathbb{E}[E_{i,j}^2]}; & \sigma_{\text{col}} &\geq \max_j \sqrt{\sum_i \mathbb{E}[E_{i,j}^2]}. \end{aligned} \quad (147)$$

### C.1 Auxiliary technical lemmas

We first gather all technical lemmas to be established in this section, and begin with the following lemma, which shows that the leave-one-out sequence  $\mathbf{G}^{(m)}$  is close to  $\mathbf{G}$  and  $\mathbf{G}^*$  when measured by the spectral norm.

**Lemma 6.** *Instate the assumptions of Theorem 1. With probability at least  $1 - O(d^{-11})$ , one has*

$$\|\mathbf{G}^{(m)} - \mathbf{G}\| \lesssim \sigma_{\text{row}} (\sigma_{\text{col}} + \|\mathbf{A}^{*\top}\|_{2,\infty}) \sqrt{\log d}, \quad (148)$$

$$\|\mathbf{G}^{(m)} - \mathbf{G}^*\| \lesssim \delta_{\text{op}} = (\sigma_{\text{row}} + \sigma_{\text{col}}) (\sigma_{\text{col}} + \|\mathbf{A}^{*\top}\|_{2,\infty}) \log d + \sigma_{\text{col}} \sqrt{\log d} \|\mathbf{A}^*\| + \|\mathbf{A}^*\|_{2,\infty}^2, \quad (149)$$

where  $\sigma_{\text{row}}$  and  $\sigma_{\text{col}}$  are defined in (97).

*Proof.* See Appendix C.2. □

Similar to  $\mathbf{H}$  (defined in (49)), we also introduce the following matrices for each  $(m, l) \in [d_1] \times [d_2]$ :

$$\mathbf{H}^{(m)} := \mathbf{U}^{(m)\top} \mathbf{U}^*, \quad (150a)$$

$$\mathbf{H}^{(m,l)} := \mathbf{U}^{(m,l)\top} \mathbf{U}^*, \quad (150b)$$

where  $\mathbf{U}^{(m)}$  and  $\mathbf{U}^{(m,l)}$  are defined in Algorithm 5 and Algorithm 6, respectively.

Lemma 7 serves a crucial step towards proving Lemma 4.

**Lemma 7.** *Instate the assumptions of Theorem 1. For any fixed  $1 \leq m \leq d_1$ , with probability at least  $1 - O(d^{-11})$ , one has*

$$\begin{aligned} &\|(\mathbf{A}^s - \mathbf{A}^*)_{m,:} \mathcal{P}_{-m,:} (\mathbf{A}^s)^\top \mathbf{U}^{(m)} \mathbf{H}^{(m)}\|_2 \\ &\lesssim \sigma_{\text{col}} (\sigma_{\text{row}} \log d + \|\mathbf{A}^*\| \sqrt{\log d}) \left( \|\mathbf{U}^{(m)} \mathbf{H}^{(m)}\|_{2,\infty} + \sqrt{\frac{\mu r}{d_1}} \right), \end{aligned}$$

where  $\mathbf{A}^s$  is defined in (93), and  $\sigma_{\text{row}}$  and  $\sigma_{\text{col}}$  are both defined in (97).

*Proof.* See Appendix C.3. □

The proof of Lemma 7 relies on an upper bound on the  $\ell_{2,\infty}$  norm of  $\mathcal{P}_{-m,:} (\mathbf{A}^s)^\top \mathbf{U}^{(m)} \mathbf{H}^{(m)}$ , which is formalized below in Lemma 8. This is built upon a leave-two-out argument.

**Lemma 8.** *Instate the assumptions of Theorem 1. With probability at least  $1 - O(d^{-10})$ , the following holds simultaneously for all  $m \in [d_1]$ ,*

$$\|\mathcal{P}_{-m,:} (\mathbf{A}^s)^\top \mathbf{U}^{(m)} \mathbf{H}^{(m)}\|_{2,\infty} \lesssim (B \log d + \sigma_{\text{col}} \sqrt{\log d}) \|\mathbf{U}^{(m)} \mathbf{H}^{(m)}\|_{2,\infty} + \|\mathbf{A}^*\| \sqrt{\frac{\mu r}{d_2}},$$

where  $\mathbf{A}^s$  is defined in (93), and  $B$  and  $\sigma_{\text{col}}$  are defined in (97).

*Proof.* See Appendix C.4. □

The proof of Lemma 8 requires the proximity between  $\mathbf{U}^{(m)}$  and  $\mathbf{U}^{(m,l)}$ , which is demonstrated below in Lemma 9.

**Lemma 9.** *Instate the assumptions of Theorem 1. With probability at least  $1 - O(d^{-10})$ , the following holds simultaneously for any  $m \in [d_1]$  and  $l \in [d_2]$ ,*

$$\begin{aligned} \left\| \mathbf{U}^{(m)} \mathbf{U}^{(m)\top} - \mathbf{U}^{(m,l)} \mathbf{U}^{(m,l)\top} \right\| &\lesssim \frac{1}{\sigma_r^{*2}} (B \log d + \sigma_{\text{col}} \sqrt{\log d})^2 \left\| \mathbf{U}^{(m)} \mathbf{H}^{(m)} \right\|_{2,\infty} \\ &\quad + \frac{\sigma_\infty^2}{\sigma_r^{*2}} + \frac{1}{\sigma_r^{*2}} (B \log d + \sigma_{\text{col}} \sqrt{\log d}) \left\| \mathbf{A}^{*\top} \right\|_{2,\infty}, \end{aligned} \quad (151)$$

where  $B, \sigma_\infty$  and  $\sigma_{\text{col}}$  are defined in (97).

*Proof.* See Appendix C.5. □

Finally, Lemma 10 stated below constitutes the main part of Lemma 3 (recalling the decomposition in (51) and (52)).

**Lemma 10.** *Instate the assumptions of Theorem 1. For each fixed  $m \in [d_1]$ , the following holds with probability exceeding  $1 - O(d^{-11})$ ,*

$$\left\| \mathbf{G}_{m,:} (\mathbf{U}^{(m)} \mathbf{H}^{(m)} - \mathbf{U}^*) \right\|_2 \lesssim \delta_{\text{loo}} \left\| \mathbf{U}^{(m)} \mathbf{H}^{(m)} - \mathbf{U}^* \right\|_{2,\infty} + \delta_{\text{op}} \kappa^2 \sqrt{\frac{\mu r}{d_1}},$$

where  $\delta_{\text{op}}$  and  $\delta_{\text{loo}}$  are defined in (48) and (127), respectively.

*Proof.* See Appendix C.6. □

## C.2 Proof of Lemma 6

Recall the definitions of  $\mathbf{G}$  and  $\mathbf{G}^{(m)}$  in (13) and (45). As shown in (137) in the proof of Lemma 4 in Appendix B.5, we know that  $\mathbf{G} - \mathbf{G}^{(m)}$  is a rank-2 symmetric matrix with nonzero entries located only in the  $m$ -th row and the  $m$ -th column. In particular, one has

$$\begin{aligned} (\mathbf{G} - \mathbf{G}^{(m)})_{m,i} &= \langle \mathbf{E}_{m,:}, \mathbf{A}_{i,:}^s \rangle, \quad i \neq m, \\ (\mathbf{G} - \mathbf{G}^{(m)})_{m,m} &= 0, \end{aligned}$$

thus indicating that

$$(\mathbf{G} - \mathbf{G}^{(m)})_{m,:} = \mathbf{E}_{m,:} [\mathcal{P}_{-m,:} (\mathbf{A}^s)]^\top.$$

This allows us to upper bound

$$\begin{aligned} \left\| \mathbf{G} - \mathbf{G}^{(m)} \right\| &\leq \left\| \mathbf{G} - \mathbf{G}^{(m)} \right\|_{\text{F}} \lesssim \left\| (\mathbf{G} - \mathbf{G}^{(m)})_{m,:} \right\|_2 \\ &= \left\| \mathbf{E}_{m,:} [\mathcal{P}_{-m,:} (\mathbf{A}^s)]^\top \right\|_2 \\ &\lesssim \sigma_{\text{row}} (\sigma_{\text{col}} + \left\| \mathbf{A}^{*\top} \right\|_{2,\infty}) \sqrt{\log d}. \end{aligned}$$

Here, the last line follows the following. First, notice  $\mathbf{A}^s = \mathbf{A}^* + \mathbf{E}$  and

$$\mathbf{E}_{m,:} [\mathcal{P}_{-m,:} (\mathbf{A}^* + \mathbf{E})]^\top = \sum_{i:i \neq m} \langle \mathbf{E}_{m,:}, \mathbf{E}_{i,:} \rangle \mathbf{e}_i^\top + \sum_{i:i \neq m} \langle \mathbf{E}_{m,:}, \mathbf{A}_{i,:}^* \rangle \mathbf{e}_i^\top,$$

where  $\mathbf{e}_i$  is the  $i$ -th standard basis in  $\mathbb{R}^{d_1}$ . It follows from (122) and (124) shown in the proof of Lemma 2 (cf. Appendix B.3) that with probability at least  $1 - O(d^{-11})$ ,

$$\begin{aligned} \left\| \mathbf{E}_{m,:} [\mathcal{P}_{-m,:} (\mathbf{A}^* + \mathbf{E})]^\top \right\|_2 &\leq \left\| \sum_{i:i \neq m} \langle \mathbf{E}_{m,:}, \mathbf{E}_{i,:} \rangle \mathbf{e}_i^\top \right\|_2 + \left\| \sum_{i:i \neq m} \langle \mathbf{E}_{m,:}, \mathbf{A}_{i,:}^* \rangle \mathbf{e}_i^\top \right\|_2 \\ &\lesssim \sigma_{\text{col}} \sigma_{\text{row}} \sqrt{\log d} + \sigma_{\text{row}} \sqrt{\log d} \left\| \mathbf{A}^{*\top} \right\|_{2,\infty}. \end{aligned}$$

In addition, the above bound combined with Lemma 1 immediately yields (149). The proof is complete by taking the union bound over  $1 \leq m \leq d_1$ .

### C.3 Proof of Lemma 7

By construction, the  $m$ -th row of  $\mathbf{A}^s - \mathbf{A}^*$  is independent of  $[\mathcal{P}_{-m,:}(\mathbf{A}^s)]^\top \mathbf{U}^{(m)} \mathbf{H}^{(m)}$ . As a result,

$$(\mathbf{A}^s - \mathbf{A}^*)_{m,:} [\mathcal{P}_{-m,:}(\mathbf{A}^s)]^\top \mathbf{U}^{(m)} \mathbf{H}^{(m)} = \sum_{j \in [d_2]} E_{m,j} \left( [\mathcal{P}_{-m,:}(\mathbf{A}^s)]^\top \mathbf{U}^{(m)} \mathbf{H}^{(m)} \right)_{j,:}$$

can be viewed as a sum of independent zero-mean random vectors (where the randomness comes from  $\{E_{m,j}\}_{j \in [d_2]}$ ). It is straightforward to calculate that

$$\begin{aligned} L &:= \max_{j \in [d_2]} \left\| E_{m,j} \left( [\mathcal{P}_{-m,:}(\mathbf{A}^s)]^\top \mathbf{U}^{(m)} \mathbf{H}^{(m)} \right)_{j,:} \right\|_2 \leq B \left\| \mathcal{P}_{-m,:}(\mathbf{A}^s)^\top \mathbf{U}^{(m)} \mathbf{H}^{(m)} \right\|_{2,\infty}, \\ V &:= \sum_{j \in [d_2]} \mathbb{E} [E_{m,j}^2] \left\| \left( [\mathcal{P}_{-m,:}(\mathbf{A}^s)]^\top \mathbf{U}^{(m)} \mathbf{H}^{(m)} \right)_{j,:} \right\|_2^2 \leq \sigma_\infty^2 \left\| [\mathcal{P}_{-m,:}(\mathbf{A}^s)]^\top \mathbf{U}^{(m)} \mathbf{H}^{(m)} \right\|_{\mathbb{F}}^2. \end{aligned}$$

In view of the matrix Bernstein inequality, it boils down to controlling  $L$  and  $V$ . To this end, let us first bound  $L$ . From Lemma 8, one has that with probability at least  $1 - O(d^{-11})$ ,

$$L \lesssim B(B \log d + \sigma_{\text{col}} \sqrt{\log d}) \left\| \mathbf{U}^{(m)} \mathbf{H}^{(m)} \right\|_{2,\infty} + B \|\mathbf{A}^*\| \sqrt{\frac{\mu r}{d_2}}. \quad (152)$$

Regarding  $V$ , Lemma 13 guarantees the following upper bound with probability exceeding  $1 - O(d^{-11})$ ,

$$\begin{aligned} \left\| [\mathcal{P}_{-m,:}(\mathbf{A}^s)]^\top \mathbf{U}^{(m)} \mathbf{H}^{(m)} \right\|_{\mathbb{F}} &\leq \left\| \mathcal{P}_{-m,:}(\mathbf{A}^s) \right\| \left\| \mathbf{U}^{(m)} \mathbf{H}^{(m)} \right\|_{\mathbb{F}} \\ &\leq \sqrt{d_1} \|\mathbf{A}^s\| \left\| \mathbf{U}^{(m)} \mathbf{H}^{(m)} \right\|_{2,\infty} \\ &\lesssim \sqrt{d_1} (B \log d + (\sigma_{\text{row}} + \sigma_{\text{col}}) \sqrt{\log d} + \|\mathbf{A}^*\|) \left\| \mathbf{U}^{(m)} \mathbf{H}^{(m)} \right\|_{2,\infty} \\ &\lesssim \sqrt{d_1} (\sigma_{\text{row}} \sqrt{\log d} + \|\mathbf{A}^*\|) \left\| \mathbf{U}^{(m)} \mathbf{H}^{(m)} \right\|_{2,\infty}, \end{aligned} \quad (153)$$

where the last inequality follows from the condition (98d) that  $B \log d + \sigma_{\text{col}} \sqrt{\log d} \ll \sigma_r^*$ . Applying the matrix Bernstein inequality yields that with probability at least  $1 - O(d^{-11})$ : one has

$$\begin{aligned} \left\| (\mathbf{A}^s - \mathbf{A}^*)_{m,:} [\mathcal{P}_{-m,:}(\mathbf{A}^s)]^\top \mathbf{U}^{(m)} \mathbf{H}^{(m)} \right\|_2 &\lesssim L \log d + \sqrt{V \log d} \\ &\lesssim (B^2 \log^2 d + B \sigma_{\text{col}} \log^{3/2} d) \left\| \mathbf{U}^{(m)} \mathbf{H}^{(m)} \right\|_{2,\infty} + B \log d \|\mathbf{A}^*\| \sqrt{\frac{\mu r}{d_2}} \\ &\quad + \sqrt{d_1} \sigma_\infty (\sigma_{\text{row}} \log d + \|\mathbf{A}^*\| \sqrt{\log d}) \left\| \mathbf{U}^{(m)} \mathbf{H}^{(m)} \right\|_{2,\infty} \\ &\stackrel{(i)}{\lesssim} \sigma_{\text{col}} (\sigma_{\text{row}} \log d + \|\mathbf{A}^*\| \sqrt{\log d}) \left\| \mathbf{U}^{(m)} \mathbf{H}^{(m)} \right\|_{2,\infty} + \sigma_{\text{col}} \|\mathbf{A}^*\| \sqrt{\log d} \sqrt{\frac{\mu r}{d_1}} \\ &\lesssim \sigma_{\text{col}} (\sigma_{\text{row}} \log d + \|\mathbf{A}^*\| \sqrt{\log d}) \left( \left\| \mathbf{U}^{(m)} \mathbf{H}^{(m)} \right\|_{2,\infty} + \sqrt{\frac{\mu r}{d_1}} \right), \end{aligned} \quad (154)$$

where (i) uses (97) that  $\sigma_{\text{col}}^2 \asymp d_1 \sigma_\infty^2$  as well as the conditions (98a), (98d) and (98b) (namely,  $B \lesssim \sqrt{\sigma_{\text{row}} \sigma_{\text{col}} / \log d}$ ,  $B \log d \ll \|\mathbf{A}^*\|$  and  $B \log d \sqrt{\mu r / d_2} \lesssim \sigma_{\text{col}} \sqrt{\log d} \sqrt{\mu r / d_1}$ ).

### C.4 Proof of Lemma 8

For notational convenience, we denote

$$\mathbf{A}^{s,(m),0} := \mathcal{P}_{-m,:}(\mathbf{A}^s).$$

Fix an arbitrary  $l \in [d_2]$ , and we would like to upper bound  $\left\| \mathbf{A}_{:,l}^{s,(m),0 \top} \mathbf{U}^{(m)} \mathbf{H}^{(m)} \right\|_2$ . The main difficulty here lies in the complicated statistical dependence between  $\mathbf{A}_{:,l}^{s,(m),0}$  and  $\mathbf{U}^{(m)} \mathbf{H}^{(m)}$ . Recall the definitions of

the auxiliary matrices  $\mathbf{U}^{(m,l)}$  and  $\mathbf{H}^{(m,l)}$  in Algorithm 6 and (150b), respectively. By construction,  $\mathbf{A}_{:,l}^{\mathbf{s},(m),0}$  is independent of  $\mathbf{U}^{(m,l)}$  and  $\mathbf{H}^{(m,l)}$ . Moreover, Lemma 9 guarantees that  $\mathbf{U}^{(m)}\mathbf{H}^{(m)}$  is extremely close to  $\mathbf{U}^{(m,l)}\mathbf{H}^{(m,l)}$ . Thus, invoke the triangle inequality to upper bound

$$\begin{aligned} \left\| \mathbf{A}_{:,l}^{\mathbf{s},(m),0\top} \mathbf{U}^{(m)} \mathbf{H}^{(m)} \right\|_2 &\leq \underbrace{\left\| \left( \mathbf{A}_{:,l}^{\mathbf{s},(m),0} - \mathbb{E} \left[ \mathbf{A}_{:,l}^{\mathbf{s},(m),0} \right] \right)^\top \mathbf{U}^{(m,l)} \mathbf{H}^{(m,l)} \right\|_2}_{=:\alpha_1} + \underbrace{\left\| \mathbb{E} \left[ \mathbf{A}_{:,l}^{\mathbf{s},(m),0} \right]^\top \mathbf{U}^{(m,l)} \mathbf{H}^{(m,l)} \right\|_2}_{=:\alpha_2} \\ &\quad + \underbrace{\left\| \mathbf{A}_{:,l}^{\mathbf{s},(m),0} \right\|_2 \left\| \mathbf{U}^{(m)} \mathbf{H}^{(m)} - \mathbf{U}^{(m,l)} \mathbf{H}^{(m,l)} \right\|_2}_{=:\alpha_3}. \end{aligned}$$

Before moving on, we make note of the following two useful upper bounds on  $\left\| \mathbf{U}^{(m)} \mathbf{U}^{(m)\top} - \mathbf{U}^{(m,l)} \mathbf{U}^{(m,l)\top} \right\|$  based on Lemma 9. On the one hand, one has with probability at least  $1 - O(d^{-13})$ ,

$$\begin{aligned} \left\| \mathbf{U}^{(m)} \mathbf{U}^{(m)\top} - \mathbf{U}^{(m,l)} \mathbf{U}^{(m,l)\top} \right\| &\lesssim \frac{1}{\sigma_r^{*2}} (B \log d + \sigma_{\text{col}} \sqrt{\log d})^2 \left\| \mathbf{U}^{(m)} \mathbf{H}^{(m)} \right\|_{2,\infty} \\ &\quad + \frac{\sigma_\infty^2}{\sigma_r^{*2}} + \frac{1}{\sigma_r^{*2}} (B \log d + \sigma_{\text{col}} \sqrt{\log d}) \left\| \mathbf{A}^{*\top} \right\|_{2,\infty} \\ &\stackrel{(i)}{\lesssim} \frac{1}{\sigma_r^{*2}} \left( B^2 \log^2 d + \sigma_{\text{col}}^2 \log d + B \log d \left\| \mathbf{A}^{*\top} \right\|_{2,\infty} + \sigma_{\text{col}} \sqrt{\log d} \left\| \mathbf{A}^* \right\| \right) \\ &\lesssim \frac{\delta_{\text{op}}}{\sigma_r^{*2}} \ll 1. \end{aligned} \tag{155}$$

where (i) follows from the facts that  $\sigma_\infty \leq \sigma_{\text{col}}$ ,  $\left\| \mathbf{U}^{(m)} \mathbf{H}^{(m)} \right\|_{2,\infty} \leq \left\| \mathbf{U}^{(m)} \mathbf{H}^{(m)} \right\| \leq 1$  and  $\left\| \mathbf{A}^{*\top} \right\|_{2,\infty} \leq \left\| \mathbf{A}^* \right\|$ ; (155) is due to the definition of  $\sigma_{\text{col}}$  in (97d), the definition of  $\delta_{\text{op}}$  (cf. (48)) as well as conditions (98a) and (98d). On the other hand, we can also bound

$$\begin{aligned} \left\| \mathbf{U}^{(m)} \mathbf{U}^{(m)\top} - \mathbf{U}^{(m,l)} \mathbf{U}^{(m,l)\top} \right\| &\stackrel{(i)}{\lesssim} \frac{1}{\sigma_r^{*2}} (B \log d + \sigma_{\text{col}} \sqrt{\log d})^2 \left\| \mathbf{U}^{(m)} \mathbf{H}^{(m)} \right\|_{2,\infty} \\ &\quad + \frac{1}{\sigma_r^{*2}} (B \log d + \sigma_{\text{col}} \sqrt{\log d}) \left\| \mathbf{A}^* \right\| \sqrt{\frac{\mu r}{d_2}} \\ &= o(1) \left\| \mathbf{U}^{(m)} \mathbf{H}^{(m)} \right\|_{2,\infty} + \frac{1}{\sigma_r^{*2}} (B \log d + \sigma_{\text{col}} \sqrt{\log d}) \left\| \mathbf{A}^* \right\| \sqrt{\frac{\mu r}{d_2}} \end{aligned} \tag{156}$$

where (i) arises from (98b) and the inequality that  $\left\| \mathbf{A}^{*\top} \right\|_{2,\infty} \leq \left\| \mathbf{A}^* \right\| \left\| \mathbf{V}^* \right\|_{2,\infty} \leq \left\| \mathbf{A}^* \right\| \sqrt{\mu r / d_2}$ , and (156) is due to conditions (98a) and (98d). In the sequel, we control the  $\alpha_i$ 's separately.

- For  $\alpha_1$ , it is easy to see that

$$\left( \mathbf{A}_{:,l}^{\mathbf{s},(m),0} - \mathbb{E} \left[ \mathbf{A}_{:,l}^{\mathbf{s},(m),0} \right] \right)^\top \mathbf{U}^{(m,l)} \mathbf{H}^{(m,l)} = \sum_{i:i \neq m} E_{i,l} \left( \mathbf{U}^{(m,l)} \mathbf{H}^{(m,l)} \right)_{i,:}$$

is a sum of independent zero-mean random vectors conditional on  $\{E_{i,j}\}_{i \in [d_1] \setminus \{m\}, j \in [d_2] \setminus \{l\}}$ . Straightforward calculation gives that

$$\begin{aligned} L &:= \max_{i:i \neq m} \left\| E_{i,l} \left( \mathbf{U}^{(m,l)} \mathbf{H}^{(m,l)} \right)_{i,:} \right\|_2 \leq B \left\| \mathbf{U}^{(m,l)} \mathbf{H}^{(m,l)} \right\|_{2,\infty}, \\ V &:= \sum_{i:i \neq m} \mathbb{E} \left[ E_{i,l}^2 \right] \left\| \left( \mathbf{U}^{(m,l)} \mathbf{H}^{(m,l)} \right)_{i,:} \right\|_2^2 \leq \sigma_{\text{col}}^2 \left\| \mathbf{U}^{(m,l)} \mathbf{H}^{(m,l)} \right\|_{2,\infty}^2. \end{aligned}$$

Then we apply the matrix Bernstein inequality to obtain that with probability at least  $1 - O(d^{-13})$ ,

$$\left\| \left( \mathbf{A}_{:,l}^{\mathbf{s},(m),0} - \mathbb{E} \left[ \mathbf{A}_{:,l}^{\mathbf{s},(m),0} \right] \right)^\top \mathbf{U}^{(m,l)} \mathbf{H}^{(m,l)} \right\|_2 \lesssim L \log d + \sqrt{V \log d}$$

$$\begin{aligned}
&\leq (B \log d + \sigma_{\text{col}} \sqrt{\log d}) \|\mathbf{U}^{(m,l)} \mathbf{H}^{(m,l)}\|_{2,\infty} \\
&\leq (B \log d + \sigma_{\text{col}} \sqrt{\log d}) \left( \|\mathbf{U}^{(m)} \mathbf{H}^{(m)}\|_{2,\infty} + \|\mathbf{U}^{(m)} \mathbf{U}^{(m)\top} - \mathbf{U}^{(m,l)} \mathbf{U}^{(m,l)\top}\| \right), \tag{157}
\end{aligned}$$

where the last line results from the following observation:

$$\begin{aligned}
\|\mathbf{U}^{(m,l)} \mathbf{H}^{(m,l)}\|_{2,\infty} &\leq \|\mathbf{U}^{(m)} \mathbf{H}^{(m)}\|_{2,\infty} + \|\mathbf{U}^{(m)} \mathbf{U}^{(m)\top} - \mathbf{U}^{(m,l)} \mathbf{U}^{(m,l)\top}\|_{2,\infty} \\
&\leq \|\mathbf{U}^{(m)} \mathbf{H}^{(m)}\|_{2,\infty} + \|\mathbf{U}^{(m)} \mathbf{U}^{(m)\top} - \mathbf{U}^{(m,l)} \mathbf{U}^{(m,l)\top}\|. \tag{158}
\end{aligned}$$

- Turning to  $\alpha_2$ , we obtain the simple upper bound

$$\left\| \mathbb{E} \left[ \mathbf{A}_{:,l}^{\text{s},(m),0} \right]^\top \mathbf{U}^{(m,l)} \mathbf{H}^{(m,l)} \right\|_2 \leq \|\mathbf{A}_{:,l}^*\|_2 \|\mathbf{U}^{(m,l)} \mathbf{H}^{(m,l)}\| \leq \|\mathbf{A}^{*\top}\|_{2,\infty}. \tag{159}$$

- With regards to  $\alpha_3$ , Lemma 12 reveals that with probability at least  $1 - O(d^{-13})$ ,

$$\begin{aligned}
\|\mathbf{A}_{:,l}^{\text{s},(m),0}\|_2 \|\mathbf{U}^{(m)} \mathbf{H}^{(m)} - \mathbf{U}^{(m,l)} \mathbf{H}^{(m,l)}\| &\leq \|\mathbf{A}_{:,l}\|_2 \|\mathbf{U}^{(m)} \mathbf{U}^{(m)\top} - \mathbf{U}^{(m,l)} \mathbf{U}^{(m,l)\top}\| \\
&\lesssim (\|\mathbf{A}^{*\top}\|_{2,\infty} + B \sqrt{\log d} + \sigma_{\text{col}}) \|\mathbf{U}^{(m)} \mathbf{U}^{(m)\top} - \mathbf{U}^{(m,l)} \mathbf{U}^{(m,l)\top}\| \\
&\lesssim \|\mathbf{A}^{*\top}\|_{2,\infty} + (B \sqrt{\log d} + \sigma_{\text{col}}) \|\mathbf{U}^{(m)} \mathbf{U}^{(m)\top} - \mathbf{U}^{(m,l)} \mathbf{U}^{(m,l)\top}\|, \tag{160}
\end{aligned}$$

where we use (155) in the last step.

Combining (157), (159), (160) implies that with probability greater than  $1 - O(d^{-13})$ ,

$$\begin{aligned}
&\left\| \mathbf{A}_{:,l}^{\text{s},(m),0\top} \mathbf{U}^{(m)} \mathbf{H}^{(m)} \right\|_2 \\
&\lesssim (B \log d + \sigma_{\text{col}} \sqrt{\log d}) \left( \|\mathbf{U}^{(m)} \mathbf{H}^{(m)}\|_{2,\infty} + \|\mathbf{U}^{(m)} \mathbf{U}^{(m)\top} - \mathbf{U}^{(m,l)} \mathbf{U}^{(m,l)\top}\| \right) + \|\mathbf{A}^{*\top}\|_{2,\infty} \\
&\stackrel{(i)}{\lesssim} (B \log d + \sigma_{\text{col}} \sqrt{\log d}) \|\mathbf{U}^{(m)} \mathbf{H}^{(m)}\|_{2,\infty} + \frac{1}{\sigma_r^2} (B \log d + \sigma_{\text{col}} \sqrt{\log d})^2 \|\mathbf{A}^*\| \sqrt{\frac{\mu r}{d_2}} + \|\mathbf{A}^*\| \sqrt{\frac{\mu r}{d_2}} \\
&\stackrel{(ii)}{\lesssim} (B \log d + \sigma_{\text{col}} \sqrt{\log d}) \|\mathbf{U}^{(m)} \mathbf{H}^{(m)}\|_{2,\infty} + \|\mathbf{A}^*\| \sqrt{\frac{\mu r}{d_2}},
\end{aligned}$$

where (i) is by (156) and  $\|\mathbf{A}^{*\top}\|_{2,\infty} \leq \|\mathbf{A}^*\| \sqrt{\mu r / d_2}$ , and (ii) follows from conditions (98a) and (98d). The proof is complete by taking the union bound over  $1 \leq l \leq d_2$ .

## C.5 Proof of Lemma 9

Fix arbitrary  $m \in [d_1]$  and  $l \in [d_2]$ . Recalling the definitions of  $\mathbf{G}^{(m)}$  and  $\mathbf{G}^{(m,l)}$  in (45) and (46b), we see that  $\mathbf{G}^{(m)} - \mathbf{G}^{(m,l)}$  is symmetric with entries

$$\begin{aligned}
(\mathbf{G}^{(m)} - \mathbf{G}^{(m,l)})_{i,j} &= E_{i,l} E_{j,l} + A_{i,l}^* E_{j,l} + E_{i,l} A_{j,l}^*, \quad i \neq m, j \neq m, i \neq j, \\
(\mathbf{G}^{(m)} - \mathbf{G}^{(m,l)})_{i,m} &= A_{m,l}^* E_{i,l}, \quad i \neq m, \\
(\mathbf{G}^{(m)} - \mathbf{G}^{(m,l)})_{i,i} &= 0, \quad 1 \leq i \leq d_1.
\end{aligned}$$

Note that  $\mathbf{G}^{(m)} - \mathbf{G}^{(m,l)}$  depends only on  $\{E_{i,l}\}_{i \in [d_1] \setminus \{m\}}$  and is hence statistically independent of  $\mathbf{U}^{(m,l)}$  and  $\mathbf{H}^{(m,l)}$ . In particular, we can express

$$\mathcal{P}_{-m}(\mathbf{G}^{(m)} - \mathbf{G}^{(m,l)}) = \mathcal{P}_{\text{off-diag}} \mathcal{P}_{-m}(\mathbf{E}_{:,l} \mathbf{E}_{:,l}^\top + \mathbf{E}_{:,l} \mathbf{A}_{:,l}^{*\top} + \mathbf{A}_{:,l}^* \mathbf{E}_{:,l}^\top),$$

where  $\mathcal{P}_{-m}$  is the projection onto the subspace of matrices supported on  $\{(i, j) \in [d_1] \times [d_2] : i \neq m \text{ and } j \neq m\}$  and  $\mathcal{P}_{\text{off-diag}}$  extracts the off-diagonal part. In addition,

$$(\mathbf{G}^{(m)} - \mathbf{G}^{(m,l)})_{m,:} = \mathbf{A}_{m,l}^* \mathbf{E}_{:,l}^\top - \mathbf{A}_{m,l}^* \mathbf{E}_{m,l} \mathbf{e}_m^\top,$$

where  $\mathbf{e}_m$  stands for the  $m$ -th standard basis in  $\mathbb{R}^{d_1}$ .

In the sequel, we shall apply the Davis-Kahan sin  $\Theta$  theorem to prove the claim. Towards this end, we need to control  $\|\mathbf{G}^{(m)} - \mathbf{G}^{(m,l)}\|$  and  $\|(\mathbf{G}^{(m)} - \mathbf{G}^{(m,l)})\mathbf{U}^{(m,l)}\|$ .

### C.5.1 Step 1: controlling $\|\mathbf{G}^{(m)} - \mathbf{G}^{(m,l)}\|$

Recall that we have already dealt with  $\mathcal{P}_{-m}(\mathbf{G}^{(m)} - \mathbf{G}^{(m,l)})$  in the proof of Lemma 1 in Appendix B.2. Straightforward computation gives

$$\begin{aligned} \|\mathcal{P}_{-m}(\mathbf{G}^{(m)} - \mathbf{G}^{(m,l)})\| &\lesssim \|\mathbf{E}_{:,l}\|_2^2 + \|\mathbf{E}_{:,l}\|_2 \|\mathbf{A}^{*\top}\|_{2,\infty} \leq \|\mathbf{E}_{:,l}\|_2^2 + \|\mathbf{E}_{:,l}\|_2 \|\mathbf{A}^*\|, \\ \|\mathcal{P}_m(\mathbf{G}^{(m)} - \mathbf{G}^{(m,l)})\| &\lesssim \left\| (\mathbf{G}^{(m)} - \mathbf{G}^{(m,l)})_{m,:} \right\|_2 \leq \|\mathbf{E}_{:,l}\|_2 \|\mathbf{A}^*\|_\infty, \end{aligned}$$

where  $\mathcal{P}_m$  is the projection onto the subspace of matrices supported on  $\{(i, j) \in [d_1] \times [d_2] : i = m \text{ or } j = m\}$ . In view of (105) (shown in the proof of Lemma 1 in Appendix B.2), we know that

$$\begin{aligned} \|\mathbf{G}^{(m)} - \mathbf{G}^{(m,l)}\| &\leq \|\mathcal{P}_{-m}(\mathbf{G}^{(m)} - \mathbf{G}^{(m,l)})\| + \|\mathcal{P}_m(\mathbf{G}^{(m)} - \mathbf{G}^{(m,l)})\| \\ &\lesssim \|\mathbf{E}_{:,l}\|_2^2 + \|\mathbf{E}_{:,l}\|_2 \|\mathbf{A}^*\| \\ &\lesssim B^2 \log d + \sigma_{\text{col}}^2 + (B\sqrt{\log d} + \sigma_{\text{col}}) \|\mathbf{A}^*\| \\ &\ll \sigma_r^{*2}, \end{aligned}$$

where the last step results from the conditions (98a) and (98d). Since  $\|\mathbf{G}^{(m)} - \mathbf{G}^*\| \lesssim \delta_{\text{op}} \ll \sigma_r^{*2}$  by Lemma 6 and the condition (98d), this also implies

$$\|\mathbf{G}^{(m,l)} - \mathbf{G}^*\| \ll \sigma_r^{*2} \quad \text{and} \quad \|(\mathbf{H}^{(m,l)})^{-1}\| \lesssim 1, \quad (161)$$

according to [AFWZ17, Lemma 2]. Moreover, it follows from Weyl's inequality that

$$\begin{aligned} \lambda_r(\mathbf{G}^{(m)}) - \lambda_{r+1}(\mathbf{G}^{(m)}) - \|\mathbf{G}^{(m)} - \mathbf{G}^{(m,l)}\| &\geq \lambda_r(\mathbf{G}^*) - \lambda_{r+1}(\mathbf{G}^*) - 2\|\mathbf{G}^{(m)} - \mathbf{G}^*\| - \|\mathbf{G}^{(m)} - \mathbf{G}^{(m,l)}\| \\ &\gtrsim \sigma_r^{*2}. \end{aligned} \quad (162)$$

### C.5.2 Step 2: controlling $\|(\mathbf{G}^{(m)} - \mathbf{G}^{(m,l)})\mathbf{U}^{(m,l)}\|$

In view of (161), we can obtain

$$\begin{aligned} \|(\mathbf{G}^{(m)} - \mathbf{G}^{(m,l)})\mathbf{U}^{(m,l)}\| &\leq \|(\mathbf{G}^{(m)} - \mathbf{G}^{(m,l)})\mathbf{U}^{(m,l)}\mathbf{H}^{(m,l)}\| \|(\mathbf{H}^{(m,l)})^{-1}\| \\ &\lesssim \|(\mathbf{G}^{(m)} - \mathbf{G}^{(m,l)})\mathbf{U}^{(m,l)}\mathbf{H}^{(m,l)}\| \\ &\leq \underbrace{\|\mathcal{P}_m(\mathbf{G}^{(m)} - \mathbf{G}^{(m,l)})\mathbf{U}^{(m,l)}\mathbf{H}^{(m,l)}\|_{\text{F}}}_{=:\alpha_1} + \underbrace{\|\mathcal{P}_{-m}(\mathbf{G}^{(m)} - \mathbf{G}^{(m,l)})\mathbf{U}^{(m,l)}\mathbf{H}^{(m,l)}\|}_{=:\alpha_2}. \end{aligned} \quad (163)$$

Therefore, it suffices to control  $\alpha_1$  and  $\alpha_2$  separately.

- Regarding  $\alpha_1$ , Lemma 12 reveals that, with probability at least  $1 - O(d^{-13})$ ,

$$\begin{aligned} &\|\mathcal{P}_m(\mathbf{G}^{(m)} - \mathbf{G}^{(m,l)})\mathbf{U}^{(m,l)}\mathbf{H}^{(m,l)}\|_{\text{F}} \\ &\leq \left\| (\mathbf{G}^{(m)} - \mathbf{G}^{(m,l)})_{m,:} \mathbf{U}^{(m,l)} \mathbf{H}^{(m,l)} \right\|_2 + \left\| (\mathbf{G}^{(m)} - \mathbf{G}^{(m,l)})_{m,:} \right\|_2 \|\mathbf{U}^{(m,l)} \mathbf{H}^{(m,l)}\|_{2,\infty} \\ &\leq \|\mathbf{A}^*\|_\infty \|\mathbf{E}_{:,l}\|_2 \|\mathbf{U}^{(m,l)} \mathbf{H}^{(m,l)}\|_2 + \|\mathbf{A}^*\|_\infty \|\mathbf{E}_{:,l}\|_2 \|\mathbf{U}^{(m,l)} \mathbf{H}^{(m,l)}\|_{2,\infty} \\ &\lesssim (B \log d + \sigma_{\text{col}} \sqrt{\log d}) \|\mathbf{A}^*\|_\infty \|\mathbf{U}^{(m,l)} \mathbf{H}^{(m,l)}\|_{2,\infty}. \end{aligned} \quad (164)$$



- When it comes to  $\alpha_2$ , since the spectral norm of a submatrix is always less than that of its original matrix, we can further upper bound

$$\begin{aligned} \|\mathcal{P}_{-m}(\mathbf{G}^{(m)} - \mathbf{G}^{(m,l)})\mathbf{U}^{(m,l)}\mathbf{H}^{(m,l)}\| &\leq \underbrace{\left\| (\mathbf{E}_{:,l}\mathbf{E}_{:,l}^\top - \mathbf{D}_l) \mathbf{U}^{(m,l)} \mathbf{H}^{(m,l)} \right\|}_{=:\beta_1} \\ &\quad + \underbrace{\left\| (\mathbf{A}_{:,l}^* \mathbf{E}_{:,l}^\top + \mathbf{E}_{:,l} \mathbf{A}_{:,l}^{*\top} - 2\widehat{\mathbf{D}}_l) \mathbf{U}^{(m,l)} \mathbf{H}^{(m,l)} \right\|}_{=:\beta_2}, \end{aligned}$$

where  $\mathbf{D}_l$  and  $\widehat{\mathbf{D}}_l$  are defined in (103) and (110) in Appendix B.2. In what follows, let us bound  $\beta_1$  and  $\beta_2$ .

- To bound  $\beta_1$ , we have

$$\left\| (\mathbf{E}_{:,l}\mathbf{E}_{:,l}^\top - \mathbf{D}_l) \mathbf{U}^{(m,l)} \mathbf{H}^{(m,l)} \right\| \leq \|\mathbf{E}_{:,l}\|_2 \left\| (\mathbf{E}_{:,l})^\top \mathbf{U}^{(m,l)} \mathbf{H}^{(m,l)} \right\|_2 + \|\mathbf{D}_l \mathbf{U}^{(m,l)} \mathbf{H}^{(m,l)}\|.$$

By Lemma 12, one has with probability at least  $1 - O(d^{-13})$ ,

$$\begin{aligned} \|\mathbf{E}_{:,l}\|_2 \left\| (\mathbf{E}_{:,l})^\top \mathbf{U}^{(m,l)} \mathbf{H}^{(m,l)} \right\|_2 &\lesssim (B\sqrt{\log d} + \sigma_{\text{col}}) (B \log d + \sigma_{\text{col}} \sqrt{\log d}) \|\mathbf{U}^{(m,l)} \mathbf{H}^{(m,l)}\|_{2,\infty} \\ &\leq (B \log d + \sigma_{\text{col}} \sqrt{\log d})^2 \|\mathbf{U}^{(m,l)} \mathbf{H}^{(m,l)}\|_{2,\infty}. \end{aligned}$$

As for the second term  $\|\mathbf{D}_l \mathbf{U}^{(m,l)} \mathbf{H}^{(m,l)}\|$ , we first observe that

$$\|\mathbb{E}[\mathbf{D}_l] \mathbf{U}^{(m,l)} \mathbf{H}^{(m,l)}\| \leq \|\mathbb{E}[\mathbf{D}_l]\| \|\mathbf{U}^{(m,l)} \mathbf{H}^{(m,l)}\| \leq \max_{i \in [d_1]} \mathbb{E}[E_{i,l}^2] \|\mathbf{U}^{(m,l)} \mathbf{H}^{(m,l)}\| \leq \sigma_\infty^2.$$

Additionally, the deviation  $(\mathbf{D}_l - \mathbb{E}[\mathbf{D}_l])\mathbf{U}^{(m,l)}\mathbf{H}^{(m,l)} = \sum_{i \in [d_1]} (E_{i,l}^2 - \mathbb{E}[E_{i,l}^2]) \mathbf{e}_i \mathbf{e}_i^\top \mathbf{U}^{(m,l)} \mathbf{H}^{(m,l)}$  is a sum of independent zero-mean random matrices. By the matrix Bernstein inequality, we have with probability at least  $1 - O(d^{-13})$ ,

$$\begin{aligned} &\left\| (\mathbf{D}_l - \mathbb{E}[\mathbf{D}_l]) \mathbf{U}^{(m,l)} \mathbf{H}^{(m,l)} \right\| \\ &\lesssim \left( \max_{i \in [d_1]} |E_{i,l}^2 - \mathbb{E}[E_{i,l}^2]| \log d + \sqrt{\sum_{i \in [d_1]} \mathbb{E}[E_{i,l}^4] \log d} \right) \max_{i \in [d_1]} \|\mathbf{e}_i \mathbf{e}_i^\top \mathbf{U}^{(m,l)} \mathbf{H}^{(m,l)}\|_{2,\infty} \\ &\lesssim \left( \max_{i \in [d_1]} |E_{i,l}^2 - \mathbb{E}[E_{i,l}^2]| \log d + \sqrt{\sum_{i \in [d_1]} \mathbb{E}[E_{i,l}^4] \log d} \right) \|\mathbf{U}^{(m,l)} \mathbf{H}^{(m,l)}\|_{2,\infty} \\ &\lesssim \left( B^2 \log d + \sqrt{(B^2 \sigma_{\text{col}}^2) \log d} \right) \|\mathbf{U}^{(m,l)} \mathbf{H}^{(m,l)}\|_{2,\infty} \\ &\stackrel{(i)}{\lesssim} \left( B^2 \log d + (B^2 + \sigma_{\text{col}}^2) \sqrt{\log d} \right) \|\mathbf{U}^{(m,l)} \mathbf{H}^{(m,l)}\|_{2,\infty} \\ &\asymp (B^2 \log d + \sigma_{\text{col}}^2 \sqrt{\log d}) \|\mathbf{U}^{(m,l)} \mathbf{H}^{(m,l)}\|_{2,\infty}, \end{aligned}$$

where we have used the AM-GM inequality in (i). Combining the estimates above yields

$$\beta_1 = \left\| (\mathbf{E}_{:,l}\mathbf{E}_{:,l}^\top - \mathbf{D}_l) \mathbf{U}^{(m,l)} \mathbf{H}^{(m,l)} \right\| \lesssim (B \log d + \sigma_{\text{col}} \sqrt{\log d})^2 \|\mathbf{U}^{(m,l)} \mathbf{H}^{(m,l)}\|_{2,\infty} + \sigma_\infty^2. \quad (165)$$

- Turning to  $\beta_2$ , we see from Lemma 12 that with probability at least  $1 - O(d^{-13})$ , one has

$$\begin{aligned} \left\| \mathbf{A}_{:,l}^* \mathbf{E}_{:,l}^\top \mathbf{U}^{(m,l)} \mathbf{H}^{(m,l)} \right\| &\leq \|\mathbf{A}^{*\top}\|_{2,\infty} \|\mathbf{E}_{:,l}^\top \mathbf{U}^{(m,l)} \mathbf{H}^{(m,l)}\| \\ &\lesssim (B \log d + \sigma_{\text{col}} \sqrt{\log d}) \|\mathbf{A}^{*\top}\|_{2,\infty} \|\mathbf{U}^{(m,l)} \mathbf{H}^{(m,l)}\|_{2,\infty} \end{aligned}$$

and

$$\left\| \mathbf{E}_{:,l} \mathbf{A}_{:,l}^{*\top} \mathbf{U}^{(m,l)} \mathbf{H}^{(m,l)} \right\| \leq \|\mathbf{E}_{:,l}\|_2 \|\mathbf{A}^{*\top}\|_{2,\infty} \|\mathbf{U}^{(m,l)} \mathbf{H}^{(m,l)}\|$$

$$\lesssim (B\sqrt{\log d} + \sigma_{\text{col}}) \|\mathbf{A}^{\star\top}\|_{2,\infty}.$$

In addition,  $\widehat{\mathbf{D}}_l \mathbf{U}^{(m,l)} \mathbf{H}^{(m,l)} = \sum_{i \in [d_1]} \mathbf{A}_{i,l}^* E_{i,l} \mathbf{e}_i \mathbf{e}_i^\top \mathbf{U}^{(m,l)} \mathbf{H}^{(m,l)}$  is a sum of independent zero-mean random matrices. It then follows from the matrix Bernstein inequality that with probability at least  $1 - O(d^{-13})$ ,

$$\begin{aligned} \|\widehat{\mathbf{D}}_l \mathbf{U}^{(m,l)} \mathbf{H}^{(m,l)}\| &\lesssim \left( \max_{i \in [d_1]} |A_{i,l}^* E_{i,l}| \log d + \sqrt{\sum_{i \in [d_1]} A_{i,l}^{*2} \mathbb{E}[E_{i,l}^2] \log d} \right) \|\mathbf{U}^{(m,l)} \mathbf{H}^{(m,l)}\|_{2,\infty} \\ &\lesssim (B \log d + \sigma_{\text{col}} \sqrt{\log d}) \|\mathbf{A}^*\|_\infty \|\mathbf{U}^{(m,l)} \mathbf{H}^{(m,l)}\|_{2,\infty}. \end{aligned}$$

Hence, we know that

$$\begin{aligned} \beta_2 &= \left\| (\mathbf{A}_{:,l}^* \mathbf{E}_{:,l}^\top + \mathbf{E}_{:,l} \mathbf{A}_{:,l}^{*\top} - 2\widehat{\mathbf{D}}_l) \mathbf{U}^{(m,l)} \mathbf{H}^{(m,l)} \right\| \\ &\lesssim (B \log d + \sigma_{\text{col}} \sqrt{\log d}) \|\mathbf{A}^{\star\top}\|_{2,\infty} \|\mathbf{U}^{(m,l)} \mathbf{H}^{(m,l)}\|_{2,\infty} + (B\sqrt{\log d} + \sigma_{\text{col}}) \|\mathbf{A}^{\star\top}\|_{2,\infty} \\ &\lesssim (B \log d + \sigma_{\text{col}} \sqrt{\log d}) \|\mathbf{A}^{\star\top}\|_{2,\infty}, \end{aligned} \quad (166)$$

which results from the facts that  $\|\mathbf{A}^*\|_\infty \leq \|\mathbf{A}^{\star\top}\|_{2,\infty}$  and  $\|\mathbf{U}^{(m,l)} \mathbf{H}^{(m,l)}\|_{2,\infty} \leq \|\mathbf{U}^{(m,l)} \mathbf{H}^{(m,l)}\| \leq 1$ .

Putting (165) and (166) together yields that

$$\begin{aligned} \|\mathcal{P}_{-m}(\mathbf{G}^{(m)} - \mathbf{G}^{(m,l)}) \mathbf{U}^{(m,l)} \mathbf{H}^{(m,l)}\| &\lesssim (B \log d + \sigma_{\text{col}} \sqrt{\log d})^2 \|\mathbf{U}^{(m,l)} \mathbf{H}^{(m,l)}\|_{2,\infty} \\ &\quad + \sigma_\infty^2 + (B \log d + \sigma_{\text{col}} \sqrt{\log d}) \|\mathbf{A}^{\star\top}\|_{2,\infty}. \end{aligned} \quad (167)$$

This combined with (164) and (163) implies

$$\begin{aligned} \|(\mathbf{G}^{(m)} - \mathbf{G}^{(m,l)}) \mathbf{U}^{(m,l)}\| &\lesssim (B \log d + \sigma_{\text{col}} \sqrt{\log d})^2 \|\mathbf{U}^{(m,l)} \mathbf{H}^{(m,l)}\|_{2,\infty} + \sigma_\infty^2 \\ &\quad + (B \log d + \sigma_{\text{col}} \sqrt{\log d}) \|\mathbf{A}^{\star\top}\|_{2,\infty} + (B \log d + \sigma_{\text{col}} \sqrt{\log d}) \|\mathbf{A}^*\|_\infty \|\mathbf{U}^{(m,l)} \mathbf{H}^{(m,l)}\|_{2,\infty} \\ &\stackrel{(i)}{\leq} (B \log d + \sigma_{\text{col}} \sqrt{\log d})^2 \|\mathbf{U}^{(m,l)} \mathbf{H}^{(m,l)}\|_{2,\infty} + \sigma_\infty^2 \\ &\quad + (B \log d + \sigma_{\text{col}} \sqrt{\log d}) \|\mathbf{A}^{\star\top}\|_{2,\infty} + (B \log d + \sigma_{\text{col}} \sqrt{\log d}) \|\mathbf{A}^*\|_\infty \\ &\stackrel{(ii)}{\lesssim} (B \log d + \sigma_{\text{col}} \sqrt{\log d})^2 \|\mathbf{U}^{(m,l)} \mathbf{H}^{(m,l)}\|_{2,\infty} + \sigma_\infty^2 \\ &\quad + (B \log d + \sigma_{\text{col}} \sqrt{\log d}) \|\mathbf{A}^{\star\top}\|_{2,\infty}, \end{aligned} \quad (168)$$

where (i) is due to the facts that  $\|\mathbf{U}^{(m,l)} \mathbf{H}^{(m,l)}\|_{2,\infty} \leq \|\mathbf{U}^{(m,l)} \mathbf{H}^{(m,l)}\| \leq 1$ , and (168) arises from the inequality  $\|\mathbf{A}^*\|_\infty \leq \|\mathbf{A}^{\star\top}\|_{2,\infty}$ .

### C.5.3 Step 3: combining Step 1 and Step 2

From (162) and (168), we apply the Davis-Kahan sin  $\Theta$  theorem to obtain that with probability exceeding  $1 - O(d^{-13})$ ,

$$\begin{aligned} \|\mathbf{U}^{(m)} \mathbf{U}^{(m)\top} - \mathbf{U}^{(m,l)} \mathbf{U}^{(m,l)\top}\| &\leq \frac{\|(\mathbf{G}^{(m)} - \mathbf{G}^{(m,l)}) \mathbf{U}^{(m,l)}\|}{\lambda_r(\mathbf{G}^{(m)}) - \lambda_{r+1}(\mathbf{G}^{(m)}) - \|\mathbf{G}^{(m)} - \mathbf{G}^{(m,l)}\|} \\ &\lesssim \frac{1}{\sigma_r^{*2}} \|(\mathbf{G}^{(m)} - \mathbf{G}^{(m,l)}) \mathbf{U}^{(m,l)}\| \\ &\stackrel{(i)}{\lesssim} \frac{1}{\sigma_r^{*2}} (B \log d + \sigma_{\text{col}} \sqrt{\log d})^2 \left( \|\mathbf{U}^{(m)} \mathbf{H}^{(m)}\|_{2,\infty} + \|\mathbf{U}^{(m)} \mathbf{U}^{(m)\top} - \mathbf{U}^{(m,l)} \mathbf{U}^{(m,l)\top}\| \right) \end{aligned}$$

$$\begin{aligned}
& + \frac{\sigma_\infty^2}{\sigma_r^{*2}} + \frac{1}{\sigma_r^{*2}} (B \log d + \sigma_{\text{col}} \sqrt{\log d}) \|\mathbf{A}^{*\top}\|_{2,\infty} \\
& \stackrel{\text{(ii)}}{\lesssim} \frac{1}{\sigma_r^{*2}} (B \log d + \sigma_{\text{col}} \sqrt{\log d})^2 \|\mathbf{U}^{(m)} \mathbf{H}^{(m)}\|_{2,\infty} + o(1) \|\mathbf{U}^{(m)} \mathbf{U}^{(m)\top} - \mathbf{U}^{(m,l)} \mathbf{U}^{(m,l)\top}\| \\
& + \frac{\sigma_\infty^2}{\sigma_r^{*2}} + \frac{1}{\sigma_r^{*2}} (B \log d + \sigma_{\text{col}} \sqrt{\log d}) \|\mathbf{A}^{*\top}\|_{2,\infty}.
\end{aligned}$$

Here, we have used (158) in (i), and the condition (98d) (i.e.  $\max\{B \log d, \sigma_{\text{col}} \sqrt{\log d}\} \ll \sigma_r^*$ ) in (ii). Rearrange the inequalities and taking the union bound over  $m \in [d_1]$  and  $l \in [d_2]$  complete the proof.

## C.6 Proof of Lemma 10

Recall the definition of  $\mathbf{G}$  in (12). We can express

$$\mathbf{G}_{m,:} (\mathbf{U}^{(m)} \mathbf{H}^{(m)} - \mathbf{U}^*) = \mathbf{A}_{m,:}^s [\mathcal{P}_{-m,:} (\mathbf{A}^s)]^\top (\mathbf{U}^{(m)} \mathbf{H}^{(m)} - \mathbf{U}^*).$$

Consequently, one can upper bound

$$\begin{aligned}
\left\| \mathbf{G}_{m,:} (\mathbf{U}^{(m)} \mathbf{H}^{(m)} - \mathbf{U}^*) \right\|_2 & \leq \underbrace{\left\| \mathbf{A}_{m,:}^s [\mathcal{P}_{-m,:} (\mathbf{A}^s)]^\top (\mathbf{U}^{(m)} \mathbf{H}^{(m)} - \mathbf{U}^*) \right\|_2}_{=:\beta_1} \\
& + \underbrace{\left\| (\mathbf{A}^s - \mathbf{A}^*) [\mathcal{P}_{-m,:} (\mathbf{A}^s)]^\top (\mathbf{U}^{(m)} \mathbf{H}^{(m)} - \mathbf{U}^*) \right\|_2}_{=:\beta_2}.
\end{aligned}$$

In what follows, we shall control  $\beta_1$  and  $\beta_2$  separately.

- To upper bound  $\beta_1$ , we have

$$\left\| \mathbf{A}_{m,:}^s [\mathcal{P}_{-m,:} (\mathbf{A}^s)]^\top (\mathbf{U}^{(m)} \mathbf{H}^{(m)} - \mathbf{U}^*) \right\|_2 \leq \left\| \mathbf{A}_{m,:}^s [\mathcal{P}_{-m,:} (\mathbf{A}^s)]^\top \right\|_2 \|\mathbf{U}^{(m)} \mathbf{H}^{(m)} - \mathbf{U}^*\|.$$

It is straightforward to derive

$$\left\| \mathbf{A}_{m,:}^s [\mathcal{P}_{-m,:} (\mathbf{A}^s)]^\top \right\|_2 \leq \left\| \mathbf{A}_{m,:}^s \mathbf{A}^{s\top} \right\|_2 \leq \left\| \mathbf{A}_{m,:}^s \mathbf{A}^{*\top} \right\|_2 + \left\| \mathbf{A}_{m,:}^s \mathbf{E}^\top \right\|_2,$$

whose first term can be bounded by

$$\left\| \mathbf{A}_{m,:}^s \mathbf{A}^{*\top} \right\|_2 \leq \left\| \mathbf{A}_{m,:}^* \right\|_2 \|\mathbf{A}^*\| \leq \sigma_1^* \|\mathbf{A}^*\|_{2,\infty}.$$

In addition, Lemma 14 indicates that

$$\begin{aligned}
\left\| \mathbf{A}_{m,:}^s \mathbf{E}^\top \right\|_2^2 & = \sum_i \left( \sum_j A_{m,j}^* E_{i,j} \right)^2 \lesssim (\sigma_{\text{col}}^2 + \sigma_\infty^2 \log^2 d) \|\mathbf{A}^*\|_{2,\infty}^2 + B^2 \|\mathbf{A}^*\|_\infty^2 \log^3 d \\
& \leq (\sigma_{\text{col}}^2 + B^2 \log^2 d) \|\mathbf{A}^*\|_{2,\infty}^2 + B^2 \|\mathbf{A}^*\|_\infty^2 \log^3 d
\end{aligned}$$

holds with probability at least  $1 - O(d^{-11})$ . Hence, we have

$$\begin{aligned}
\left\| \mathbf{A}_{m,:}^s [\mathcal{P}_{-m,:} (\mathbf{A}^s)]^\top \right\|_2 & \leq \left\| \mathbf{A}_{m,:}^s \mathbf{A}^{*\top} \right\|_2 + \left\| \mathbf{A}_{m,:}^s \mathbf{E}^\top \right\|_2 \\
& \lesssim \sigma_1^* \|\mathbf{A}^*\|_{2,\infty} + (\sigma_{\text{col}} + B \log d) \|\mathbf{A}^*\|_{2,\infty} + B \log^{3/2} d \|\mathbf{A}^*\|_\infty \\
& \lesssim \sigma_1^* \|\mathbf{A}^*\|_{2,\infty} + \sigma_1^{*2} \sqrt{\frac{\mu r}{d_1}} \lesssim \sigma_1^{*2} \sqrt{\frac{\mu r}{d_1}},
\end{aligned} \tag{169}$$

using conditions (98a), (98c) and (98d). Moreover, from Lemma 1 and Lemma 6, we know that

$$\|\mathbf{U}^{(m)} \mathbf{H}^{(m)} - \mathbf{U}^*\| \lesssim \|\mathbf{U}^{(m)} \mathbf{U}^{(m)\top} - \mathbf{U}^* \mathbf{U}^{*\top}\| \leq \frac{\|\mathbf{G}^{(m)} - \mathbf{G}^*\|}{\lambda_r(\mathbf{G}^*) - \lambda_{r+1}(\mathbf{G}^{(m)})}$$

$$\begin{aligned}
&\leq \frac{\|\mathbf{G}^{(m)} - \mathbf{G}^*\|}{\lambda_r(\mathbf{G}^*) - \lambda_{r+1}(\mathbf{G}^*) - \|\mathbf{G}^{(m)} - \mathbf{G}^*\|} \\
&\lesssim \frac{1}{\sigma_r^{*2}} \|\mathbf{G}^{(m)} - \mathbf{G}^*\| \lesssim \frac{\delta_{\text{op}}}{\sigma_r^{*2}},
\end{aligned} \tag{170}$$

where  $\delta_{\text{op}}$  is defined in (48). Combining (169) and (170) yields

$$\left\| \mathbf{A}_{m,:}^* \mathcal{P}_{-m,:} (\mathbf{A}^s)^\top (\mathbf{U}^{(m)} \mathbf{H}^{(m)} - \mathbf{U}^*) \right\|_2 \lesssim \left\| \mathbf{A}_{m,:}^* \mathcal{P}_{-m,:} (\mathbf{A}^s)^\top \right\|_2 \left\| \mathbf{U}^{(m)} \mathbf{H}^{(m)} - \mathbf{U}^* \right\| \lesssim \delta_{\text{op}} \kappa^2 \sqrt{\frac{\mu r}{d_1}}. \tag{171}$$

- Next, we look at  $\beta_2$ . Before we start, we pause to note that by (156), one has

$$\begin{aligned}
\left\| \mathbf{U}^{(m)} \mathbf{U}^{(m)\top} - \mathbf{U}^{(m,l)} \mathbf{U}^{(m,l)\top} \right\| &\lesssim \frac{1}{\sigma_r^{*2}} (B \log d + \sigma_{\text{col}} \sqrt{\log d})^2 (\left\| \mathbf{U}^{(m)} \mathbf{H}^{(m)} - \mathbf{U}^* \right\|_{2,\infty} + \left\| \mathbf{U}^* \right\|_{2,\infty}) \\
&\quad + \frac{1}{\sigma_r^{*2}} (B \log d + \sigma_{\text{col}} \sqrt{\log d}) \left\| \mathbf{A}^* \right\| \sqrt{\frac{\mu r}{d_2}}.
\end{aligned} \tag{172}$$

We now ready to control  $(\mathbf{A}^s - \mathbf{A}^*)_{m,:} [\mathcal{P}_{-m,:} (\mathbf{A}^s)]^\top (\mathbf{U}^{(m)} \mathbf{H}^{(m)} - \mathbf{U}^*)$ , which can be accomplished in the same way as in the proof of Lemma 7 in Appendix C.3. We omit the proof details for conciseness here and only give the proof sketch. First, we can use  $\mathbf{U}^{(m,l)} \mathbf{H}^{(m,l)} - \mathbf{U}^*$  as the surrogate for  $\mathbf{U}^{(m)} \mathbf{H}^{(m)} - \mathbf{U}^*$  to deal with the statistical dependence issue, and apply the Bernstein inequality to show that with probability at least  $1 - O(d^{-11})$ ,

$$\begin{aligned}
&\left\| [\mathcal{P}_{-m,:} (\mathbf{A}^s)]^\top (\mathbf{U}^{(m)} \mathbf{H}^{(m)} - \mathbf{U}^*) \right\|_{2,\infty} \\
&\lesssim (B \log d + \sigma_{\text{col}} \sqrt{\log d}) \left\| \mathbf{U}^{(m)} \mathbf{H}^{(m)} - \mathbf{U}^* \right\|_{2,\infty} \\
&\quad + \left\| \mathbf{A}^{*\top} \right\|_{2,\infty} \left\| \mathbf{U}^{(m)} \mathbf{H}^{(m)} - \mathbf{U}^* \right\| \\
&\quad + (\left\| \mathbf{A}^{*\top} \right\|_{2,\infty} + B \sqrt{\log d} + \sigma_{\text{col}}) \left\| \mathbf{U}^{(m)} \mathbf{U}^{(m)\top} - \mathbf{U}^{(m,l)} \mathbf{U}^{(m,l)\top} \right\| \\
&\stackrel{(i)}{\lesssim} (B \log d + \sigma_{\text{col}} \sqrt{\log d}) \left\| \mathbf{U}^{(m)} \mathbf{H}^{(m)} - \mathbf{U}^* \right\|_{2,\infty} + \frac{\delta_{\text{op}}}{\sigma_r^{*2}} \left\| \mathbf{A}^* \right\| \sqrt{\frac{\mu r}{d_2}} \\
&\quad + \frac{1}{\sigma_r^{*2}} (B \log d + \sigma_{\text{col}} \sqrt{\log d})^3 (\left\| \mathbf{U}^{(m)} \mathbf{H}^{(m)} - \mathbf{U}^* \right\|_{2,\infty} + \left\| \mathbf{U}^* \right\|_{2,\infty}) \\
&\quad + \frac{1}{\sigma_r^{*2}} (B \log d + \sigma_{\text{col}} \sqrt{\log d})^2 \left\| \mathbf{A}^* \right\| \sqrt{\frac{\mu r}{d_2}} \\
&\lesssim (B \log d + \sigma_{\text{col}} \sqrt{\log d}) \left\| \mathbf{U}^{(m)} \mathbf{H}^{(m)} - \mathbf{U}^* \right\|_{2,\infty} + \frac{\delta_{\text{op}}}{\sigma_r^*} \left( \left\| \mathbf{U}^* \right\|_{2,\infty} + \kappa \sqrt{\frac{\mu r}{d_2}} \right),
\end{aligned} \tag{173}$$

where (i) follows from (155), (170) and (172) and the inequality  $\left\| \mathbf{A}^{*\top} \right\|_{2,\infty} \leq \left\| \mathbf{A}^* \right\| \sqrt{\mu r / d_2}$ ; (173) arises from the definition of  $\delta_{\text{op}}$  in (99) and conditions (98a) and (98d) (namely,  $B \log d + \sigma_{\text{col}} \sqrt{\log d} \ll \sigma_r^* / \kappa$  and  $(B \log d + \sigma_{\text{col}} \sqrt{\log d})^2 \lesssim \delta_{\text{op}} \ll \sigma_r^{*2}$ ). Applying the matrix Bernstein inequality yields that with probability at least  $1 - O(d^{-11})$ ,

$$\begin{aligned}
&\left\| (\mathbf{A}^s - \mathbf{A}^*)_{m,:} [\mathcal{P}_{-m,:} (\mathbf{A}^s)]^\top (\mathbf{U}^{(m)} \mathbf{H}^{(m)} - \mathbf{U}^*) \right\|_2 \\
&\lesssim B \log d \left\| [\mathcal{P}_{-m,:} (\mathbf{A}^s)]^\top (\mathbf{U}^{(m)} \mathbf{H}^{(m)} - \mathbf{U}^*) \right\|_{2,\infty} \\
&\quad + \sqrt{d_1} \sigma_\infty (\sigma_{\text{row}} \log d + \left\| \mathbf{A}^* \right\| \sqrt{\log d}) \left\| \mathbf{U}^{(m)} \mathbf{H}^{(m)} - \mathbf{U}^* \right\|_{2,\infty} \\
&\stackrel{(i)}{\lesssim} B \log d (B \log d + \sigma_{\text{col}} \sqrt{\log d}) \left\| \mathbf{U}^{(m)} \mathbf{H}^{(m)} - \mathbf{U}^* \right\|_{2,\infty} \\
&\quad + \delta_{\text{op}} \frac{B \log d}{\sigma_r^*} \left( \left\| \mathbf{U}^* \right\|_{2,\infty} + \kappa \sqrt{\frac{\mu r}{d_2}} \right)
\end{aligned}$$

$$\begin{aligned}
& + \sigma_{\text{col}}(\sigma_{\text{row}} \log d + \|\mathbf{A}^*\| \sqrt{\log d}) \|\mathbf{U}^{(m)} \mathbf{H}^{(m)} - \mathbf{U}^*\|_{2,\infty} \\
& \stackrel{\text{(ii)}}{\lesssim} \sigma_{\text{col}}(\sigma_{\text{row}} \log d + \|\mathbf{A}^*\| \sqrt{\log d}) \|\mathbf{U}^{(m)} \mathbf{H}^{(m)} - \mathbf{U}^*\|_{2,\infty} \\
& \quad + \delta_{\text{op}} \frac{B \log d}{\sigma_r^*} \left( \|\mathbf{U}^*\|_{2,\infty} + \kappa \sqrt{\frac{\mu r}{d_2}} \right) \\
& \stackrel{\text{(iii)}}{\lesssim} \sigma_{\text{col}}(\sigma_{\text{row}} \log d + \|\mathbf{A}^*\| \sqrt{\log d}) \|\mathbf{U}^{(m)} \mathbf{H}^{(m)} - \mathbf{U}^*\|_{2,\infty} + o(1) \delta_{\text{op}} \|\mathbf{U}^*\|_{2,\infty} + \delta_{\text{op}} \frac{\kappa \sigma_{\text{col}} \sqrt{\log d}}{\sigma_r^*} \sqrt{\frac{\mu r}{d_1}} \\
& \stackrel{\text{(iv)}}{\lesssim} \sigma_{\text{col}}(\sigma_{\text{row}} \log d + \|\mathbf{A}^*\| \sqrt{\log d}) \|\mathbf{U}^{(m)} \mathbf{H}^{(m)} - \mathbf{U}^*\|_{2,\infty} + o(1) \delta_{\text{op}} \sqrt{\frac{\mu r}{d_1}}.
\end{aligned}$$

Here, (i) follows from (97) and (173); (ii) is due to conditions (98a) and (98d) that  $B^2 \log d \lesssim \sigma_{\text{col}} \sigma_{\text{row}}$ ,  $B \log d \ll \sigma_r^*$  and  $B \log d (B \log d + \sigma_{\text{col}} \sqrt{\log d}) \lesssim B \log^2 d + \sigma_{\text{col}}^2 \log d \leq \delta_{\text{op}}$ ; (iii) holds true because of (98b) and (98d) that  $B \log d \ll \sigma_r^*$ ; and (iv) arises from (98d) that  $\sigma_{\text{col}} \sqrt{\log d} \ll \sigma_r^* / \kappa$ . Recalling the definition of  $\delta_{\text{lo0}}$  in (127), we obtain that

$$\left\| (\mathbf{A}^s - \mathbf{A}^*)_{m,:} [\mathcal{P}_{-m,:}(\mathbf{A}^s)]^\top (\mathbf{U}^{(m)} \mathbf{H}^{(m)} - \mathbf{U}^*) \right\|_2 \lesssim \delta_{\text{lo0}} \|\mathbf{U}^{(m)} \mathbf{H}^{(m)} - \mathbf{U}^*\|_{2,\infty} + o(1) \delta_{\text{op}} \sqrt{\frac{\mu r}{d_1}}. \quad (174)$$

Putting (171) and (174) together, we arrive at the advertised bound.

## D Proofs for lower bounds

### D.1 Proof of Theorem 2

Without loss of generality, it suffices to focus on the set of matrices with  $\sigma_r(\mathbf{A}^*) \in [0.9, 1.1]$ ; otherwise one can always rescale the matrices  $\mathbf{A}^*$  and  $\mathbf{N}$  by the same factor  $1/\sigma_r(\mathbf{A}^*)$  simultaneously.

Let us start with the minimax spectral norm bound (22a). Recognizing the elementary fact that

$$\|\mathbf{U}\mathbf{U}^\top - \mathbf{U}^*\mathbf{U}^{*\top}\| \asymp \min_{\mathbf{R} \in \mathcal{O}^{r \times r}} \|\mathbf{U}\mathbf{R} - \mathbf{U}^*\|,$$

we have

$$\inf_{\widehat{\mathbf{U}}} \sup_{\mathbf{A}^* \in \mathcal{M}^*} \mathbb{E} \left[ \min_{\mathbf{R} \in \mathcal{O}^{r \times r}} \|\widehat{\mathbf{U}}\mathbf{R} - \mathbf{U}(\mathbf{A}^*)\| \right] \asymp \inf_{\widehat{\mathbf{U}}} \sup_{\mathbf{A}^* \in \mathcal{M}^*} \mathbb{E} \left[ \|\widehat{\mathbf{U}}\widehat{\mathbf{U}}^\top - \mathbf{U}(\mathbf{A}^*)(\mathbf{U}(\mathbf{A}^*))^\top\| \right].$$

In light of this, we shall focus attention on bounding  $\|\mathbf{U}\mathbf{U}^\top - \mathbf{U}(\mathbf{A}^*)(\mathbf{U}(\mathbf{A}^*))^\top\|$  in the remainder of the proof. In addition, it can be easily seen (which we omit for brevity) that it is sufficient to establish the lower bounds for the rank-1 case (i.e.  $r = 1$ ).<sup>6</sup> In what follows, we assume that

$$\mathbf{A} = \mathcal{P}_\Omega \left( \underbrace{\mathbf{u}^* \mathbf{v}^{*\top}}_{=\mathbf{A}^*} + \mathbf{N} \right),$$

where  $\mathbf{v}^* \sim \mathcal{N}(\mathbf{0}, \frac{1}{d_2} \mathbf{I}_{d_2})$  and  $N_{i,j} \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, \sigma^2)$ . Without loss of generality, we assume throughout that  $d_1/2$  is an integer.

**Step 1: constructing a collection of hypotheses.** Let us begin by constructing a family of well-separated unit vectors  $\{\mathbf{u}^i\}_{1 \leq i \leq M} \subseteq \mathbb{R}^{d_1}$ . In view of the celebrated Varshamov-Gilbert bound [Mas07, Lemma 4.7], one can find a set of vectors  $\{\mathbf{w}^i\}_{i=1}^M \subseteq \{-1, 1\}^{d_1/2}$  obeying

$$\log M \geq d_1/32 \quad \text{and} \quad \min \{\|\mathbf{w}^i \pm \mathbf{w}^j\|_2\} \geq \sqrt{d_1}/2, \quad \forall i \neq j, \quad (175)$$

<sup>6</sup>Suppose we wish to estimate  $\mathbf{U}^* \in \mathbb{R}^{d_1 \times r}$  where  $\mathbf{U}^*$  takes the form  $\mathbf{U}^* = \begin{bmatrix} \mathbf{u}^* & \mathbf{0} \\ \mathbf{0} & \mathbf{Q}^* \end{bmatrix}$ ,  $\mathbf{u}^* \in \mathbb{R}^{d_1/2}$ ,  $\|\mathbf{u}^*\|_2 = 1$  and  $\mathbf{Q}^* \in \mathbb{R}^{d_1/2 \times (r-1)}$  consists of orthonormal columns. If there is an oracle informing us of  $\mathbf{Q}^*$ , then the problem of estimating  $\mathbf{U}^*$  is reduced to the rank-1 case. This suggests that we can focus on the rank-1 case to derive the lower bound.

where we denote  $\min \|\mathbf{a} \pm \mathbf{b}\|_2 = \min\{\|\mathbf{a} - \mathbf{b}\|_2, \|\mathbf{a} + \mathbf{b}\|_2\}$ . For some  $\delta \in (0, 1)$  to be chosen later, we generate the  $d_1$ -dimensional vectors

$$\mathbf{u}^i := \frac{\delta}{\sqrt{d_1/2}} \begin{bmatrix} \mathbf{w}^i \\ \mathbf{0} \end{bmatrix} + \sqrt{\frac{1-\delta^2}{d_1/2}} \begin{bmatrix} \mathbf{0} \\ \mathbf{1} \end{bmatrix} \in \mathbb{R}^{d_1}, \quad 1 \leq i \leq M, \quad (176)$$

where  $\mathbf{0}$  (resp.  $\mathbf{1}$ ) denotes the all-zero (resp. all-one) vector. By construction, it is easily seen that  $\|\mathbf{u}^i\|_2 = 1$  for all  $1 \leq i \leq M$ , and that

$$\begin{aligned} \|\mathbf{u}^i \mathbf{u}^{i\top} - \mathbf{u}^j \mathbf{u}^{j\top}\| &\geq \frac{1}{\sqrt{2}} \|\mathbf{u}^i \mathbf{u}^{i\top} - \mathbf{u}^j \mathbf{u}^{j\top}\|_{\text{F}} = \frac{1}{\sqrt{2}} \sqrt{\text{tr}(\mathbf{u}^i \mathbf{u}^{i\top} - \mathbf{u}^j \mathbf{u}^{j\top})(\mathbf{u}^i \mathbf{u}^{i\top} - \mathbf{u}^j \mathbf{u}^{j\top})} \\ &= \frac{1}{\sqrt{2}} \sqrt{2 - 2\langle \mathbf{u}^i, \mathbf{u}^j \rangle} \geq \frac{1}{\sqrt{2}} \sqrt{2 - 2|\langle \mathbf{u}^i, \mathbf{u}^j \rangle|} \\ &= \frac{1}{\sqrt{2}} \sqrt{\|\mathbf{u}^i\|_2^2 + \|\mathbf{u}^j\|_2^2 - 2|\langle \mathbf{u}^i, \mathbf{u}^j \rangle|} = \frac{1}{\sqrt{2}} \min \|\mathbf{u}^i \pm \mathbf{u}^j\|_2 \\ &= \frac{1}{\sqrt{2}} \cdot \frac{\delta}{\sqrt{d_1/2}} \min \{\|\mathbf{w}^i \pm \mathbf{w}^j\|_2\} \geq \frac{\delta}{2}, \end{aligned} \quad (177)$$

where the last inequality arises from (175). We shall then associate each vector  $\mathbf{u}^i$  ( $1 \leq i \leq M$ ) with a hypothesis as follows:

$$\mathcal{H}_i: \quad \mathbf{A} = \mathcal{P}_\Omega(\mathbf{u}^i \mathbf{v}^{*\top} + \mathbf{N}), \quad 1 \leq i \leq M.$$

In the sequel, for each  $1 \leq i \leq M$  and  $1 \leq k \leq d_2$ , we denote

- $\mathbb{P}^i$ : the distribution of  $\mathbf{A}$  under the hypothesis  $\mathcal{H}_i$ ;
- $\mathbb{P}_\Omega^i$ : the distribution of  $\mathbf{A}$  under the hypothesis  $\mathcal{H}_i$ , conditional on  $\Omega$ ;
- $\mathbb{P}_{\Omega,k}^i$ : the distribution of the  $k$ -th column of  $\mathbf{A}$  under the hypothesis  $\mathcal{H}_i$ , conditional on  $\Omega$ .

Additionally, standard Gaussian concentration inequalities imply that: with high probability, one has

$$\|\mathbf{u}^i \mathbf{v}^{*\top}\| = \|\mathbf{u}^i\|_2 \|\mathbf{v}^*\|_2 = 1 + o(1),$$

and hence  $\mathbf{u}^i \mathbf{v}^{*\top} \in \mathcal{M}^*$  for all  $1 \leq i \leq M$ .

**Step 2: bounding the KL divergence between each pair of hypotheses.** Fix any  $1 \leq i \neq j \leq M$ . The next step lies in upper bounding the KL divergence of  $\mathbb{P}^j$  from  $\mathbb{P}^i$ . Towards this, we observe that

$$\begin{aligned} \text{KL}(\mathbb{P}^i \parallel \mathbb{P}^j) &= \text{KL}(\mathbb{E}_\Omega[\mathbb{P}_\Omega^i] \parallel \mathbb{E}_\Omega[\mathbb{P}_\Omega^j]) \leq \mathbb{E}_\Omega[\text{KL}(\mathbb{P}_\Omega^i \parallel \mathbb{P}_\Omega^j)] \\ &= \mathbb{E}_\Omega \left[ \sum_{1 \leq k \leq d_2} \text{KL}(\mathbb{P}_{\Omega,k}^i \parallel \mathbb{P}_{\Omega,k}^j) \right]. \end{aligned} \quad (178)$$

Here, the penultimate inequality arises from the convexity of KL divergence and Jensen's inequality, whereas the last line follows since the noise components are independently generated and KL divergence is additive for independent distributions.

Before moving on, we find it convenient to introduce additional notation to simplify presentation. For any vector  $\mathbf{u} := [u_i]_{1 \leq i \leq d_1}$  and any index set  $\mathcal{A} \subseteq [d_1]$ , we define

$$\mathbf{u}_{\mathcal{A}} := [u_i]_{i \in \mathcal{A}} \in \mathbb{R}^{|\mathcal{A}|},$$

which is obtained by maintaining only those entries of  $\mathbf{u}$  lying within  $\mathcal{A}$ . In addition, define

$$\begin{aligned} \Omega_k &:= \{m \in [d_1]: (m, k) \in \Omega\}; \\ \widehat{\Omega}_k &:= \{m \in [d_1/2]: (m, k) \in \Omega\}; \\ \widetilde{\Omega}_k &:= \{d_1/2 < m \leq d_1: (m, k) \in \Omega\}; \end{aligned}$$

$$\begin{aligned}\widehat{\Omega}_k^{(i,j),\text{diff}} &:= \{m \in [d_1/2]: (m, k) \in \Omega \text{ and } u_m^i \neq u_m^j\}; \\ \widehat{\Omega}_k^{(i,j),\text{same}} &:= \{m \in [d_1/2]: (m, k) \in \Omega \text{ and } u_m^i = u_m^j\}.\end{aligned}$$

By construction, one clearly has  $\Omega_k = \widehat{\Omega}_k \cup \widetilde{\Omega}_k = \widehat{\Omega}_k^{(i,j),\text{diff}} \cup \widehat{\Omega}_k^{(i,j),\text{same}} \cup \widetilde{\Omega}_k$ , and

$$\|\mathbf{u}_{\Omega_k}^i\|_2^2 = \frac{2\delta^2}{d_1} |\widehat{\Omega}_k| + \frac{2(1-\delta^2)}{d_1} |\widetilde{\Omega}_k|, \quad 1 \leq i \leq M. \quad (180)$$

With these in place, we are in a position to control the KL divergence. We first make the observation that: conditional on the sampling set  $\Omega_k$  and under the hypothesis  $\mathcal{H}_i$ , the entries of  $\mathbf{A}_{:,k}$  within  $\Omega_k$  follow a multivariate Gaussian distribution  $\mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_{\Omega_k}^i)$ , where

$$\boldsymbol{\Sigma}_{\Omega_k}^i := \sigma^2 \mathbf{I}_{|\Omega_k|} + \frac{1}{d_2} \mathbf{u}_{\Omega_k}^i \mathbf{u}_{\Omega_k}^{i\top}.$$

As a result, invoking the KL divergence for multivariate Gaussians, we can deduce that

$$\begin{aligned}\text{KL}(\mathbb{P}_{\Omega,k}^i \parallel \mathbb{P}_{\Omega,k}^j) &= \text{KL}(\mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_{\Omega_k}^i) \parallel \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_{\Omega_k}^j)) = \frac{1}{2} \left( \text{tr}((\boldsymbol{\Sigma}_{\Omega_k}^j)^{-1} \boldsymbol{\Sigma}_{\Omega_k}^i) - |\Omega_k| \right). \\ &= \frac{1}{2} \text{tr} \left( \left( \mathbf{I}_{|\Omega_k|} + \frac{1}{\sigma^2 d_2} \mathbf{u}_{\Omega_k}^j \mathbf{u}_{\Omega_k}^{j\top} \right)^{-1} \left( \mathbf{I}_{|\Omega_k|} + \frac{1}{\sigma^2 d_2} \mathbf{u}_{\Omega_k}^i \mathbf{u}_{\Omega_k}^{i\top} \right) \right) - \frac{1}{2} |\Omega_k| \\ &\stackrel{(i)}{=} \frac{1}{2} \text{tr} \left( \left( \mathbf{I}_{|\Omega_k|} - \frac{1}{\sigma^2 d_2 + \|\mathbf{u}_{\Omega_k}^j\|_2^2} \mathbf{u}_{\Omega_k}^j \mathbf{u}_{\Omega_k}^{j\top} \right) \left( \mathbf{I}_{|\Omega_k|} + \frac{1}{\sigma^2 d_2} \mathbf{u}_{\Omega_k}^i \mathbf{u}_{\Omega_k}^{i\top} \right) \right) - \frac{1}{2} |\Omega_k| \\ &= \frac{1}{2\sigma^2 d_2} \|\mathbf{u}_{\Omega_k}^i\|_2^2 - \frac{1}{2(\sigma^2 d_2 + \|\mathbf{u}_{\Omega_k}^j\|_2^2)} \|\mathbf{u}_{\Omega_k}^j\|_2^2 - \frac{\langle \mathbf{u}_{\Omega_k}^i, \mathbf{u}_{\Omega_k}^j \rangle^2}{2\sigma^2 d_2 (\sigma^2 d_2 + \|\mathbf{u}_{\Omega_k}^j\|_2^2)} \\ &\stackrel{(ii)}{=} \frac{\|\mathbf{u}_{\Omega_k}^i\|_2^4 - \langle \mathbf{u}_{\Omega_k}^i, \mathbf{u}_{\Omega_k}^j \rangle^2}{2\sigma^2 d_2 (\sigma^2 d_2 + \|\mathbf{u}_{\Omega_k}^i\|_2^2)} = \frac{\langle \mathbf{u}_{\Omega_k}^i, \mathbf{u}_{\Omega_k}^i + \mathbf{u}_{\Omega_k}^j \rangle \langle \mathbf{u}_{\Omega_k}^i, \mathbf{u}_{\Omega_k}^i - \mathbf{u}_{\Omega_k}^j \rangle}{2\sigma^2 d_2 (\sigma^2 d_2 + \|\mathbf{u}_{\Omega_k}^i\|_2^2)}, \quad (181)\end{aligned}$$

where (i) follows from the Woodbury matrix identity, (ii) arises since  $\|\mathbf{u}_{\Omega_k}^i\|_2 = \|\mathbf{u}_{\Omega_k}^j\|_2$  (cf. (180)). Next, straightforward calculations yield

$$\begin{aligned}\langle \mathbf{u}_{\Omega_k}^i, \mathbf{u}_{\Omega_k}^i - \mathbf{u}_{\Omega_k}^j \rangle &= \frac{4\delta^2}{d_1} |\widehat{\Omega}_k^{(i,j),\text{diff}}|, \\ \langle \mathbf{u}_{\Omega_k}^i, \mathbf{u}_{\Omega_k}^i + \mathbf{u}_{\Omega_k}^j \rangle &= \frac{4\delta^2}{d_1} |\widehat{\Omega}_k^{(i,j),\text{same}}| + \frac{4(1-\delta^2)}{d_1} |\widetilde{\Omega}_k|.\end{aligned}$$

Substituting the above identities and the identity (180) into (181) gives

$$\begin{aligned}\text{KL}(\mathbb{P}_{\Omega,k}^i \parallel \mathbb{P}_{\Omega,k}^j) &\leq \frac{\langle \mathbf{u}_{\Omega_k}^i, \mathbf{u}_{\Omega_k}^i + \mathbf{u}_{\Omega_k}^j \rangle \langle \mathbf{u}_{\Omega_k}^i, \mathbf{u}_{\Omega_k}^i - \mathbf{u}_{\Omega_k}^j \rangle}{2\sigma^2 d_2 \cdot \sigma^2 d_2} \\ &\leq \frac{16\delta^2 |\widehat{\Omega}_k^{(i,j),\text{diff}}| \cdot (\delta^2 |\widehat{\Omega}_k^{(i,j),\text{same}}| + (1-\delta^2) |\widetilde{\Omega}_k|)}{\sigma^4 d_2^2 d_1^2} \\ &\leq \frac{16\delta^2 |\widehat{\Omega}_k^{(i,j),\text{diff}}| \cdot (|\widehat{\Omega}_k^{(i,j),\text{same}}| + |\widetilde{\Omega}_k|)}{\sigma^4 d_2^2 d_1^2},\end{aligned}$$

where we have used the fact that  $\delta \in (0, 1)$ . In addition, the elementary inequality  $2\|\mathbf{u}_{\Omega_k}^i\|_2^2 \geq \langle \mathbf{u}_{\Omega_k}^i, \mathbf{u}_{\Omega_k}^i + \mathbf{u}_{\Omega_k}^j \rangle$  together with the preceding identities yields

$$\begin{aligned}\text{KL}(\mathbb{P}_{\Omega,k}^i \parallel \mathbb{P}_{\Omega,k}^j) &= \frac{\langle \mathbf{u}_{\Omega_k}^i, \mathbf{u}_{\Omega_k}^i + \mathbf{u}_{\Omega_k}^j \rangle \langle \mathbf{u}_{\Omega_k}^i, \mathbf{u}_{\Omega_k}^i - \mathbf{u}_{\Omega_k}^j \rangle}{2\sigma^2 d_2 \|\mathbf{u}_{\Omega_k}^i\|_2^2} \leq \frac{\langle \mathbf{u}_{\Omega_k}^i, \mathbf{u}_{\Omega_k}^i - \mathbf{u}_{\Omega_k}^j \rangle}{\sigma^2 d_2} \\ &= \frac{4\delta^2 |\widehat{\Omega}_k^{(i,j),\text{diff}}|}{\sigma^2 d_2 d_1}.\end{aligned}$$

Putting the above bounds and the inequality (178) together leads to

$$\text{KL}(\mathbb{P}^i \parallel \mathbb{P}^j) \leq \mathbb{E} \left[ \sum_{k=1}^{d_2} \frac{16\delta^2 |\widehat{\Omega}_k^{(i,j),\text{diff}}| \cdot (|\widehat{\Omega}_k^{(i,j),\text{same}}| + |\widetilde{\Omega}_k|)}{\sigma^4 d_2^2 d_1^2} \right] \lesssim \frac{\delta^2 p^2}{\sigma^4 d_2}; \quad (182a)$$

$$\text{KL}(\mathbb{P}^i \parallel \mathbb{P}^j) \leq \sum_{k=1}^{d_2} \frac{4\delta^2 |\widehat{\Omega}_k^{(i,j),\text{diff}}|}{\sigma^2 d_2 d_1} \lesssim \frac{\delta^2 p}{\sigma^2}. \quad (182b)$$

**Step 3: invoking Fano's inequality.** Fano's inequality [Tsy08, Corollary 2.6] asserts that if

$$\frac{1}{M} \sum_{i=2}^M \text{KL}(\mathbb{P}^i \parallel \mathbb{P}^1) \leq \frac{\log M}{32}, \quad (183)$$

then the minimax probability of testing error necessarily obeys

$$p_{e,M} := \inf_{\psi} \max_{1 \leq j \leq M} \mathbb{P} \{ \psi \neq j \mid \mathcal{H}_j \} \geq 0.2,$$

where the infimum is taken over all tests. In view of (175) and the upper bounds (182), we observe that the bound (183) would hold by taking

$$\delta = c_1 \min \left\{ \frac{\sigma^2 \sqrt{d_1 d_2}}{p} + \sigma \sqrt{\frac{d_1}{p}}, 1 \right\} \quad (184)$$

for some sufficiently small constant  $c_1 > 0$ . Therefore, adopting the standard reduction scheme as introduced in [Tsy08, Chapter 2.2], we arrive at

$$\begin{aligned} \inf_{\widehat{U}} \sup_{\mathbf{A}^* \in \mathcal{M}^*} \mathbb{E} \left[ \left\| \widehat{U} \widehat{U}^\top - \mathbf{U}(\mathbf{A}^*) (\mathbf{U}(\mathbf{A}^*))^\top \right\| \right] &\gtrsim \min_{i \neq j} \left\| \mathbf{u}^i \mathbf{u}^{i\top} - \mathbf{u}^j \mathbf{u}^{j\top} \right\| \gtrsim \delta \\ &\asymp \min \left\{ \frac{\sigma^2 \sqrt{d_1 d_2}}{p} + \sigma \sqrt{\frac{d_1}{p}}, 1 \right\}, \end{aligned}$$

where the penultimate inequality comes from (177), and the last line makes use of our choice (184). Combined with the high-probability fact that  $\|\mathbf{u}^i \mathbf{v}^{*\top}\| \in [0.9, 1.1]$ , we establish the minimax spectral norm bound (22a).

Given that  $\|\mathbf{Z}\|_{2,\infty} \geq \frac{1}{\sqrt{d_1}} \|\mathbf{Z}\|$  holds for any  $\mathbf{Z} \in \mathbb{R}^{d_1 \times r}$ , the advertised  $\ell_{2,\infty}$  lower bound (22b) follows immediately from the spectral norm lower bound (22a).

## D.2 Proof of Theorem 3

To begin with, the sampling set  $\Omega$  can be equivalently viewed as the edge set of a random bipartite graph  $\mathcal{G}(d_1, d_2, p)$ . Here, we recall that  $\mathcal{G}(d_1, d_2, p)$  is generated by (i) taking the complete bipartite graph connecting two disjoint vertex sets  $\mathcal{U}$  and  $\mathcal{V}$ , where  $|\mathcal{U}| = d_1$  and  $|\mathcal{V}| = d_2$ , and (ii) removing each edge independently with probability  $1 - p$ . As shown in [Joh12, Theorem 6], if  $p < \frac{1-\epsilon}{\sqrt{d_1 d_2}}$  for some constant  $0 < \epsilon < 1$  and if  $d_1 \leq d_2$ , then with probability  $1 - o(1)$ , there is no connected component in  $\mathcal{G}(d_1, d_2, p)$  containing more than  $O(\log d_1)$  (resp.  $O(\sqrt{d_1 d_2} \log d_1)$ ) vertices in  $\mathcal{U}$  (resp.  $\mathcal{V}$ ). In what follows, we let  $\mathcal{C}_1, \dots, \mathcal{C}_K$  denote the collection of connected components in  $\mathcal{G}(d_1, d_2, p)$ , and denote by  $\mathcal{U}_i$  (resp.  $\mathcal{V}_i$ ) the set of vertices in  $\mathcal{U}$  (resp.  $\mathcal{V}$ ) that reside within  $\mathcal{C}_i$ .

Generate  $\mathbf{u}^*$  and  $\mathbf{v}^*$  such that

$$\mathbf{u}_i^* = \begin{cases} 1/\sqrt{d_1}, & \text{with prob. } 0.5 \\ -1/\sqrt{d_1}, & \text{else} \end{cases} \quad \text{and} \quad \mathbf{v}_j^* = \begin{cases} 1/\sqrt{d_2}, & \text{with prob. } 0.5 \\ -1/\sqrt{d_2}, & \text{else} \end{cases}$$

for each  $1 \leq i \leq d_1$  and  $1 \leq j \leq d_2$ . Letting  $\mathbf{u}_{\mathcal{S}} \in \mathbb{R}^{|\mathcal{S}|}$  represent a vector comprising the entries of  $\mathbf{u}$  whose indices come from  $\mathcal{S}$ , we generate

$$\tilde{\mathbf{u}}_{\mathcal{C}_i}^* = z_i \mathbf{u}_{\mathcal{C}_i}^* \quad \text{and} \quad \tilde{\mathbf{v}}_{\mathcal{C}_i}^* = z_i \mathbf{v}_{\mathcal{C}_i}^*;$$



here,  $z_i$  is a set of independent Bernoulli variables with  $z_i = 1$  with probability 0.5 and  $z_i = -1$  otherwise. As one can easily verify (which we omit for brevity),

- $\mathbf{u}_{\mathcal{C}_i}^* \mathbf{v}_{\mathcal{C}_i}^{*\top} = \tilde{\mathbf{u}}_{\mathcal{C}_i}^* \tilde{\mathbf{v}}_{\mathcal{C}_i}^{*\top}$  for each  $i$ , and hence  $\mathcal{P}_\Omega(\mathbf{u}^* \mathbf{v}^{*\top}) = \mathcal{P}_\Omega(\tilde{\mathbf{u}}^* \tilde{\mathbf{v}}^{*\top})$ ;
- with probability  $1 - o(1)$ , one has  $\min \|\mathbf{u}^* \pm \tilde{\mathbf{u}}^*\|_2 \asymp 1$  and  $\|\mathbf{u}^* \mathbf{v}^{*\top} - \tilde{\mathbf{u}}^* \tilde{\mathbf{v}}^{*\top}\|_F \asymp 1$ .

This concludes the proof.

## E A few more auxiliary lemmas

In this section, we establish a few auxiliary facts that are useful throughout the proof of the main theorem. We begin with some basic properties about the truth  $\mathbf{A}^*$  and  $\mathbf{G}^*$ .

**Lemma 11.** *Recall the definition of the incoherence parameters in Definition 1. Then one has*

$$\begin{aligned} \|\mathbf{A}^*\|_{2,\infty} &\leq \sqrt{\frac{\mu_1 r \sigma_1^{*2}}{d_1}}, & \|\mathbf{A}^{*\top}\|_{2,\infty} &\leq \sqrt{\frac{\mu_2 r \sigma_1^{*2}}{d_2}}, \\ \|\mathbf{G}^*\|_{2,\infty} &\leq \sqrt{\frac{\mu_1 r \sigma_1^{*4}}{d_1}}, & \|\mathbf{A}^*\|_\infty &\leq \min \left\{ \sqrt{\frac{\mu_1 \mu_2 r^2}{d_1 d_2}}, \sigma_1^* \|\mathbf{U}^*\|_{2,\infty}, \sigma_1^* \|\mathbf{V}^*\|_{2,\infty} \right\}. \end{aligned}$$

Next, we summarize several facts related to the matrix  $\mathbf{E}$  defined in (94), which contains independent zero-mean entries.

**Lemma 12.** *Fix any matrices  $\mathbf{W}_1$  and  $\mathbf{W}_2$ . With probability greater than  $1 - O(d^{-20})$ , the following holds*

$$\begin{aligned} \max_{i \in [d_1]} \sum_{j \in [d_2]} E_{i,j}^2 &\lesssim B^2 \log d + \sigma_{\text{row}}^2, \\ \max_{i \in [d_2]} \|\mathbf{E}_{i,:} \mathbf{W}_1\|_2 &\lesssim (B \log d + \sigma_{\text{row}} \sqrt{\log d}) \|\mathbf{W}_1\|_{2,\infty}, \\ \max_{j \in [d_2]} \sum_{i \in [d_1]} E_{i,j}^2 &\lesssim B^2 \log d + \sigma_{\text{col}}^2, \\ \max_{j \in [d_2]} \left\| (\mathbf{E}_{:,j})^\top \mathbf{W}_2 \right\|_2 &\lesssim (B \log d + \sigma_{\text{col}} \sqrt{\log d}) \|\mathbf{W}_2\|_{2,\infty}, \end{aligned}$$

where  $\sigma_{\text{row}}$ ,  $\sigma_{\text{col}}$ , and  $B$  are respectively upper bounds on  $\max_{i \in [d_1]} \sqrt{\sum_{j \in [d_2]} \mathbb{E}[E_{i,j}^2]}$ ,  $\max_{j \in [d_2]} \sqrt{\sum_{i \in [d_1]} \mathbb{E}[E_{i,j}^2]}$ , and  $\max_{i \in [d_1], j \in [d_2]} |E_{i,j}|$ ; see (97) for precise definitions. As a result, one has

$$\begin{aligned} \|\mathbf{E}\|_{2,\infty} &\lesssim B \sqrt{\log d} + \sigma_{\text{row}}, \\ \|\mathbf{A}^s\|_{2,\infty} &\lesssim \|\mathbf{A}^*\|_{2,\infty} + B \sqrt{\log d} + \sigma_{\text{row}}, \\ \|\mathbf{E}^\top\|_{2,\infty} &\lesssim B \sqrt{\log d} + \sigma_{\text{col}}, \\ \|\mathbf{A}^{s\top}\|_{2,\infty} &\lesssim \|\mathbf{A}^{*\top}\|_{2,\infty} + B \sqrt{\log d} + \sigma_{\text{col}}, \end{aligned}$$

where  $\mathbf{A}^s = \mathbf{A}^* + \mathbf{E}$  is defined in (93).

**Lemma 13.** *With probability greater than  $1 - O(d^{-20})$ , one has*

$$\|\mathbf{E}\| \lesssim B \log d + (\sigma_{\text{row}} + \sigma_{\text{col}}) \sqrt{\log d},$$

where  $B$ ,  $\sigma_{\text{row}}$  and  $\sigma_{\text{col}}$  are defined in (97).

**Lemma 14.** *Fix any vector  $\mathbf{w} \in \mathbb{R}^{d_2}$ . With probability at least  $1 - O(d^{-20})$ , one has*

$$\sum_{i \in [d_1]} \left( \sum_{j \in [d_2]} w_j E_{i,j} \right)^2 \lesssim \|\mathbf{w}\|_2^2 (\sigma_{\text{col}}^2 + \sigma_\infty^2 \log^2 d) + \|\mathbf{w}\|_\infty^2 B^2 \log^3 d,$$

where  $B$ ,  $\sigma_\infty$  and  $\sigma_{\text{col}}$  are defined in (97).

## E.1 Proof of Lemma 11

Given the SVD of  $\mathbf{A}^* = \mathbf{U}^* \boldsymbol{\Sigma}^* \mathbf{V}^{*\top}$ , one has  $\mathbf{G}^* = \mathbf{A}^* \mathbf{A}^{*\top} = \mathbf{U}^* \boldsymbol{\Sigma}^{*2} \mathbf{U}^{*\top}$ . Using the definition of the incoherence parameters, one can derive

$$\begin{aligned}\|\mathbf{A}^*\|_{2,\infty} &= \max_{i \in [d_1]} \|\mathbf{U}_{i,:}^* \boldsymbol{\Sigma}^* \mathbf{V}^{*\top}\|_2 \leq \max_{i \in [d_1]} \|\mathbf{U}_{i,:}^*\|_2 \|\boldsymbol{\Sigma}^*\| \|\mathbf{V}^*\| \leq \sigma_1^* \|\mathbf{U}^*\|_{2,\infty} \leq \sqrt{\frac{\mu_1 r \sigma_1^{*2}}{d_1}}; \\ \|\mathbf{A}^{*\top}\|_{2,\infty} &= \max_{j \in [d_2]} \|\mathbf{V}_{j,:}^* \boldsymbol{\Sigma}^* \mathbf{U}^{*\top}\|_2 \leq \max_{j \in [d_2]} \|\mathbf{V}_{j,:}^*\|_2 \|\boldsymbol{\Sigma}^*\| \|\mathbf{U}^*\| \leq \sigma_1^* \|\mathbf{V}^*\|_{2,\infty} \leq \sqrt{\frac{\mu_2 r \sigma_1^{*2}}{d_2}}; \\ \|\mathbf{G}^*\|_{2,\infty} &= \max_{i \in [d_1]} \|\mathbf{U}_{i,:}^* \boldsymbol{\Sigma}^{*2} \mathbf{U}^{*\top}\|_2 \leq \max_{i \in [d_1]} \|\mathbf{U}_{i,:}^*\|_2 \|\boldsymbol{\Sigma}^{*2}\| \|\mathbf{U}^*\| \leq \sigma_1^{*2} \|\mathbf{U}^*\|_{2,\infty} \leq \sqrt{\frac{\mu_1 r \sigma_1^{*4}}{d_1}}.\end{aligned}$$

Moreover, the Cauchy-Schwartz inequality allows one to upper bound

$$\|\mathbf{A}^*\|_\infty = \max_{(i,j) \in [d_1] \times [d_2]} \left| \mathbf{U}_{i,:}^* \boldsymbol{\Sigma}^* (\mathbf{V}_{j,:}^*)^\top \right| \leq \|\mathbf{U}^*\|_{2,\infty} \|\boldsymbol{\Sigma}^*\| \|\mathbf{V}^*\|_{2,\infty} \leq \sigma_1^* \|\mathbf{U}^*\|_{2,\infty} \|\mathbf{V}^*\|_{2,\infty}.$$

In view of the simple bounds  $\|\mathbf{U}^*\|_{2,\infty} \leq \|\mathbf{U}^*\| \leq 1$  and  $\|\mathbf{V}^*\|_{2,\infty} \leq \|\mathbf{V}^*\| \leq 1$ , we conclude that

$$\|\mathbf{A}^*\|_\infty \leq \sigma_1^* \|\mathbf{U}^*\|_{2,\infty} \quad \text{and} \quad \|\mathbf{A}^*\|_\infty \leq \sigma_1^* \|\mathbf{V}^*\|_{2,\infty}.$$

## E.2 Proof of Lemma 12

We shall only prove the results concerning  $\sigma_{\text{col}}$ ; the results concerning  $\sigma_{\text{row}}$  follow immediately via nearly identical arguments.

In view of the Bernstein inequality, we have

$$\mathbb{P} \left\{ \left| \sum_{i \in [d_1]} E_{i,j}^2 - M_1 \right| \geq t \right\} \leq 2 \exp \left( -\frac{3}{8} \min \left\{ \frac{t^2}{V_1}, \frac{t}{L_1} \right\} \right), \quad t > 0,$$

where  $M_1, L_1$  and  $S_1$  are given respectively by

$$\begin{aligned}M_1 &:= \sum_{i \in [d_1]} \mathbb{E} [E_{i,l}^2] \leq \sigma_{\text{col}}^2, \\ L_1 &:= \max_{i \in [d_1]} |E_{i,l}^2 - \mathbb{E} [E_{i,l}^2]| \leq B^2 + \sigma_\infty^2 \leq 2B^2, \\ V_1 &:= \sum_{i \in [d_1]} \text{Var} (E_{i,l}^2) \leq \sum_{i \in [d_1]} \mathbb{E} [E_{i,l}^4] \leq B^2 \sigma_{\text{col}}^2.\end{aligned}$$

Here, we have made use of the fact that  $\sigma_\infty \leq B$ . As a result, one has

$$\begin{aligned}\sum_{i \in [d_1]} E_{i,j}^2 &\lesssim M_1 + L_1 \log d + \sqrt{V_1 \log d} \lesssim \sigma_{\text{col}}^2 + B^2 \log d + B \sigma_{\text{col}} \sqrt{\log d} \\ &\asymp \sigma_{\text{col}}^2 + B^2 \log d\end{aligned}$$

with probability exceeding  $1 - O(d^{-20})$ , where the last line arises from the AM-GM inequality (namely,  $2B \sigma_{\text{col}} \sqrt{\log d} \leq \sigma_{\text{col}}^2 + B^2 \log d$ ). As an immediate consequence, with probability at least  $1 - O(d^{-20})$ ,

$$\begin{aligned}\|\mathbf{E}_{:,j}\|_2 &= \sqrt{\sum_{i \in [d_1]} E_{i,j}^2} \lesssim \sigma_{\text{col}} + B \sqrt{\log d}, \\ \|\mathbf{A}_{:,j}^s\|_2 &\leq \|\mathbf{A}_{:,j}^*\|_2 + \|\mathbf{E}_{:,j}\|_2 \lesssim \|\mathbf{A}^{*\top}\|_{2,\infty} + \sigma_{\text{col}} + B \sqrt{\log d}.\end{aligned}$$

Next, we turn to the claim concerning a fixed matrix  $\mathbf{W}_2$ . Observe that  $(\mathbf{E}_{:,l})^\top \mathbf{W}_2 = \sum_{i \in [d_1]} E_{i,l}(\mathbf{W}_2)_{i,:}$  is a sum of independent zero-mean random vectors. In order to invoke standard concentration inequalities, we compute

$$L_2 := \max_{i \in [d_1]} \|E_{i,l}(\mathbf{W}_2)_{i,:}\|_2 \leq B \|\mathbf{W}_2\|_{2,\infty},$$

$$V_2 := \sum_{i \in [d_1]} \mathbb{E} [E_{i,l}^2] \|(\mathbf{W}_2)_{i,:}\|_2^2 \leq \sigma_{\text{col}}^2 \|\mathbf{W}_2\|_{2,\infty}^2.$$

Invoking the matrix Bernstein inequality yields that with probability exceeding  $1 - O(d^{-20})$ ,

$$\|(\mathbf{E}_{:,l})^\top \mathbf{W}_2\|_2 \lesssim L_2 \log d + \sqrt{V_2 \log d} \lesssim (B \log d + \sigma_{\text{col}} \sqrt{\log d}) \|\mathbf{W}_2\|_{2,\infty}.$$

### E.3 Proof of Lemma 13

First, we can write

$$\mathbf{E} = \sum_{i \in [d_1], j \in [d_2]} E_{i,j} \mathbf{e}_i \mathbf{e}_j^\top$$

as a sum of independent zero-mean random matrices (since  $\mathbb{E}[E_{i,j}] = 0$ ). We make the observation that

$$L := \max_{i \in [d_1], j \in [d_2]} \|E_{i,j} \mathbf{e}_i \mathbf{e}_j^\top\| \leq B;$$

$$V := \max \left\{ \left\| \sum_{i \in [d_1], j \in [d_2]} \mathbb{E} [E_{i,j}^2] \mathbf{e}_i \mathbf{e}_i^\top \right\|, \left\| \sum_{i \in [d_1], j \in [d_2]} \mathbb{E} [E_{i,j}^2] \mathbf{e}_j \mathbf{e}_j^\top \right\| \right\} \leq \sigma_{\text{row}}^2 + \sigma_{\text{col}}^2.$$

It then follows from the matrix Bernstein inequality that, with probability at least  $1 - O(d^{-10})$ ,

$$\|\mathbf{E}\| \lesssim L \log d + \sqrt{V \log d} \lesssim B \log d + (\sigma_{\text{row}} + \sigma_{\text{col}}) \sqrt{\log d}.$$

### E.4 Proof of Lemma 14

Let us define a sequence of independent zero-mean random variables  $\{X_i\}_{1 \leq i \leq d_1}$  as follows

$$X_i := \sum_{j \in [d_2]} w_j E_{i,j}.$$

It is easily seen that

$$\max_{j \in [d_2]} |w_j E_{i,j}| \leq \|\mathbf{w}\|_\infty B; \quad \mathbb{E} [X_i^2] = \sum_{j \in [d_2]} w_j^2 \sigma_{i,j}^2 \leq \|\mathbf{w}\|_2^2 \sigma_\infty^2.$$

We can therefore apply the Bernstein inequality to show that, with probability at least  $1 - O(d^{-11})$ ,

$$|X_i| \lesssim (\|\mathbf{w}\|_\infty B) \log d + \sqrt{(\|\mathbf{w}\|_2^2 \sigma_\infty^2) \log d} =: R. \quad (185)$$

Next, let us introduce a sequence of independent random variables  $\{Y_i\}_{1 \leq i \leq d_1}$ , obtained by truncating  $X_i$

$$Y_i \triangleq X_i \mathbf{1} \{|X_i| \leq CR\}$$

for some sufficiently large absolute constant  $C > 0$ . From (185) and the union bound, we know that  $Y_i = X_i$  holds simultaneously for all  $1 \leq i \leq d_1$  with probability at least  $1 - O(d^{-10})$ .

Further, it is straightforward to compute that

$$M_2 := \sum_{i \in [d_1]} \mathbb{E} [Y_i^2] \leq \sum_{i \in [d_1]} \mathbb{E} [X_i^2] \leq \sum_{i \in [d_1], j \in [d_2]} w_j^2 \sigma_{i,j}^2 \leq \|\mathbf{w}\|_2^2 \sigma_{\text{col}}^2;$$

$$L_2 := \max_{i \in [d_1]} |Y_i^2 - \mathbb{E} [Y_i^2]| \lesssim R^2 \lesssim \|\mathbf{w}\|_\infty^2 B^2 \log^2 d + \|\mathbf{w}\|_2^2 \sigma_\infty^2 \log d;$$

$$V_2 := \sum_{i \in [d_1]} \text{Var} (Y_i^2) \leq \sum_{i \in [d_1]} \mathbb{E} [Y_i^4] \leq \sum_{i \in [d_1]} \mathbb{E} [X_i^4] \lesssim \sum_{i \in [d_1]} \sum_{j \in [d_2]} w_j^4 \mathbb{E} [E_{i,j}^4] + \sum_{i \in [d_1]} \sum_{j_1 \neq j_2} w_{j_1}^2 w_{j_2}^2 \mathbb{E} [E_{i,j_1}^2] \mathbb{E} [E_{i,j_2}^2]$$

$$\lesssim \|\mathbf{w}\|_\infty^2 \|\mathbf{w}\|_2^2 B^2 \sigma_{\text{col}}^2 + \|\mathbf{w}\|_2^4 \sigma_\infty^2 \sigma_{\text{col}}^2.$$

We then apply the Bernstein inequality to conclude that with probability at least  $1 - O(d^{-10})$ :

$$\begin{aligned} \sum_{i \in [d_1]} Y_i^2 &\lesssim M_2 + L_2 \log d + \sqrt{V_2 \log d} \\ &\lesssim \|\mathbf{w}\|_2^2 \sigma_{\text{col}}^2 + \|\mathbf{w}\|_\infty^2 B^2 \log^3 d + \|\mathbf{w}\|_2^2 \sigma_\infty^2 \log^2 d + \left( \|\mathbf{w}\|_\infty \|\mathbf{w}\|_2 B \sigma_{\text{col}} + \|\mathbf{w}\|_2^2 \sigma_\infty \sigma_{\text{col}} \right) \sqrt{\log d} \\ &\asymp \|\mathbf{w}\|_2^2 (\sigma_{\text{col}}^2 + \sigma_\infty^2 \log^2 d) + \|\mathbf{w}\|_\infty^2 B^2 \log^3 d, \end{aligned}$$

where the last line arises from the AM-GM inequality (namely,  $2 \|\mathbf{w}\|_\infty \|\mathbf{w}\|_2 B \sigma_{\text{col}} \sqrt{\log d} \leq \|\mathbf{w}\|_\infty^2 B^2 \log d + \|\mathbf{w}\|_2^2 \sigma_{\text{col}}^2$  and  $2 \|\mathbf{w}\|_2^2 \sigma_\infty \sigma_{\text{col}} \sqrt{\log d} \leq \|\mathbf{w}\|_2^2 \sigma_\infty^2 \log d + \|\mathbf{w}\|_2^2 \sigma_{\text{col}}^2$ ).

## References

- [Abb18] E. Abbe. Community detection and stochastic block models. *Foundations and Trends in Communications and Information Theory*, 14(1-2):1–162, 2018.
- [ABH15] E. Abbe, A. S. Bandeira, and G. Hall. Exact recovery in the stochastic block model. *IEEE Transactions on Information Theory*, 62(1):471–487, 2015.
- [ABKK17] N. Agarwal, A. S. Bandeira, K. Koiliaris, and A. Kolla. Multisection in the stochastic block model using semidefinite programming. In *Compressed Sensing and its Applications*, pages 125–162. Springer, 2017.
- [AFW19] E. Abbe, J. Fan, and K. Wang. An  $\ell_p$  analysis of kernel principal component analysis with application to spectral clustering. *working paper*, 2019.
- [AFWZ17] E. Abbe, J. Fan, K. Wang, and Y. Zhong. Entrywise eigenvector analysis of random matrices with low expected rank. *arXiv:1709.09565*, 2017.
- [AHB14] T. Alzahrani, K. J. Horadam, and S. Boztas. Community detection in bipartite networks using random walks. In *Complex Networks V*, pages 157–165. Springer, 2014.
- [AL18] A. A. Amini and E. Levina. On semidefinite relaxations for the block model. *The Annals of Statistics*, 46(1):149–179, 2018.
- [ASSS19] A. Agarwal, D. Shah, D. Shen, and D. Song. On robustness of principal component regression. In *Advances in Neural Information Processing Systems*, pages 9889–9900, 2019.
- [Ban18] A. S. Bandeira. Random laplacian matrices and convex relaxations. *Foundations of Computational Mathematics*, 18(2):345–379, 2018.
- [BL08] P. J. Bickel and E. Levina. Covariance regularization by thresholding. *The Annals of Statistics*, 36(6):2577–2604, 2008.
- [BM16] B. Barak and A. Moitra. Noisy tensor completion via the sum-of-squares hierarchy. In *Proceedings of the Conference on Learning Theory*, pages 417–445, 2016.
- [BY12] Z. Bai and J. Yao. On sample eigenvalues in a generalized spiked population model. *Journal of Multivariate Analysis*, 106:167–177, 2012.
- [CC18a] Y. Chen and E. Candès. The projected power method: An efficient algorithm for joint alignment from pairwise differences. *Communications on Pure and Applied Mathematics*, 71(8):1648–1714, 2018.
- [CC18b] Y. Chen and Y. Chi. Harnessing structures in big data via guaranteed low-rank matrix estimation: Recent theory and fast algorithms via convex and nonconvex optimization. *IEEE Signal Processing Magazine*, 35(4):14–31, July 2018.

- [CCF18] Y. Chen, C. Cheng, and J. Fan. Asymmetry helps: Eigenvalue and eigenvector analyses of asymmetrically perturbed low-rank matrices. *arXiv preprint arXiv:1811.12804*, *The Annals of Statistics*, to appear, 2018.
- [CCF+19] Y. Chen, Y. Chi, J. Fan, C. Ma, and Y. Yan. Noisy matrix completion: understanding statistical guarantees for convex relaxation via nonconvex optimization. *arXiv:1902.07698*, 2019.
- [CCFM19] Y. Chen, Y. Chi, J. Fan, and C. Ma. Gradient descent with random initialization: Fast global convergence for nonconvex phase retrieval. *Mathematical Programming*, 176(1-2):5–37, July 2019.
- [CFMW19] Y. Chen, J. Fan, C. Ma, and K. Wang. Spectral method and regularized mle are both optimal for top- $k$  ranking. *The Annals of Statistics*, 47(4):2204–2235, 2019.
- [CFMY19] Y. Chen, J. Fan, C. Ma, and Y. Yan. Inference and uncertainty quantification for noisy matrix completion. *Proceedings of the National Academy of Sciences of the U.S.A.*, 116(46):22931–22937, 2019.
- [Cha15] S. Chatterjee. Matrix estimation by universal singular value thresholding. *The Annals of Statistics*, 43(1):177–214, 2015.
- [CJSX14] Y. Chen, A. Jalali, S. Sanghavi, and H. Xu. Clustering partially observed graphs via convex optimization. *The Journal of Machine Learning Research*, 15(1):2213–2238, 2014.
- [CKR17] J. Cho, D. Kim, and K. Rohe. Asymptotic theory for estimating the singular vectors and values of a partially-observed low rank matrix with noise. *Statistica Sinica*, pages 1921–1948, 2017.
- [CKST16] Y. Chen, G. Kamath, C. Suh, and D. Tse. Community recovery in graphs with locality. In *Proceedings of the International Conference on Machine Learning*, pages 689–698, 2016.
- [CL15] T. T. Cai and X. Li. Robust and computationally feasible community detection in the presence of arbitrary outlier nodes. *The Annals of Statistics*, 43(3):1027–1059, 2015.
- [CLC19] Y. Chi, Y. M. Lu, and Y. Chen. Nonconvex optimization meets low-rank matrix factorization: An overview. *IEEE Transactions on Signal Processing*, 67(20):5239 – 5269, October 2019.
- [CLL19] J. Chen, D. Liu, and X. Li. Nonconvex rectangular matrix completion via gradient descent without  $\ell_{2,\infty}$  regularization. *arXiv preprint arXiv:1901.06116*, 2019.
- [CLPC19] C. Cai, G. Li, H. V. Poor, and Y. Chen. Nonconvex low-rank tensor completion from noisy data. In *Advances in Neural Information Processing Systems*, pages 1861–1872, 2019.
- [CLX18] Y. Chen, X. Li, and J. Xu. Convexified modularity maximization for degree-corrected stochastic block models. *The Annals of Statistics*, 46(4):1573–1602, 2018.
- [CMW13] T. T. Cai, Z. Ma, and Y. Wu. Sparse PCA: Optimal rates and adaptive estimation. *The Annals of Statistics*, 41(6):3074–3110, 2013.
- [CMW15] T. Cai, Z. Ma, and Y. Wu. Optimal estimation and rank detection for sparse spiked covariance matrices. *Probability Theory and Related Fields*, 161(3-4):781–815, 2015.
- [CO06] A. Coja-Oghlan. A spectral heuristic for bisecting random graphs. *Random Structures & Algorithms*, 29(3):351–398, 2006.
- [CO10] A. Coja-Oghlan. Graph partitioning via adaptive spectral techniques. *Combinatorics, Probability and Computing*, 19(2):227–284, 2010.
- [CPC20] C. Cai, H. V. Poor, and Y. Chen. Uncertainty quantification for nonconvex tensor completion: Confidence intervals, heteroscedasticity and optimality. *Proceedings of the International Conference on Machine Learning*, to appear, 2020.

- [CR09] E. J. Candès and B. Recht. Exact matrix completion via convex optimization. *Foundations of Computational Mathematics*, 9(6):717, 2009.
- [CRV15] P. Chin, A. Rao, and V. Vu. Stochastic block model and community detection in sparse graphs: A spectral algorithm with optimal rate of recovery. In *Proceedings of the Conference on Learning Theory*, pages 391–423, 2015.
- [CSG16] Y. Chen, C. Suh, and A. J. Goldsmith. Information recovery from pairwise measurements. *IEEE Transactions on Information Theory*, 62(10):5881–5905, 2016.
- [CT10] E. Candès and T. Tao. The power of convex relaxation: Near-optimal matrix completion. *IEEE Transactions on Information Theory*, 56(5):2053–2080, May 2010.
- [CTP19a] J. Cape, M. Tang, and C. E. Priebe. Signal-plus-noise matrix models: eigenvector deviations and fluctuations. *Biometrika*, 106(1):243–250, 2019.
- [CTP19b] J. Cape, M. Tang, and C. E. Priebe. The two-to-infinity norm and singular subspace geometry with applications to high-dimensional statistics. *The Annals of Statistics*, 47(5):2405–2439, 2019.
- [CW15] Y. Chen and M. J. Wainwright. Fast low-rank estimation by projected gradient descent: General statistical and algorithmic guarantees. *arXiv:1509.03025*, 2015.
- [CY12] T. T. Cai and M. Yuan. Adaptive covariance matrix estimation through block thresholding. *The Annals of Statistics*, 40(4):2014–2042, 2012.
- [CZ16] T. T. Cai and A. Zhang. Minimax rate-optimal estimation of high-dimensional covariance matrices with incomplete data. *Journal of Multivariate Analysis*, 150:55–74, 2016.
- [CZ18] T. T. Cai and A. Zhang. Rate-optimal perturbation bounds for singular subspaces with applications to high-dimensional statistics. *The Annals of Statistics*, 46(1):60–89, 2018.
- [Dhi01] I. S. Dhillon. Co-clustering documents and words using bipartite spectral graph partitioning. In *Proceedings of the Seventh ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 269–274. ACM, 2001.
- [DK70] C. Davis and W. M. Kahan. The rotation of eigenvectors by a perturbation. iii. *SIAM Journal on Numerical Analysis*, 7(1):1–46, 1970.
- [EBW17] J. Eldridge, M. Belkin, and Y. Wang. Unperturbed: spectral analysis beyond Davis-Kahan. *arXiv preprint arXiv:1706.06516*, 2017.
- [EK15] N. El Karoui. On the impact of predictor geometry on the performance on high-dimensional ridge-regularized generalized robust regression estimators. *Probability Theory and Related Fields*, pages 1–81, 2015.
- [EKBB<sup>+</sup>13] N. El Karoui, D. Bean, P. J. Bickel, C. Lim, and B. Yu. On robust regression with high-dimensional predictors. *Proceedings of the National Academy of Sciences of the U.S.A.*, 110(36):14557–14562, 2013.
- [EvdG19] A. Elsener and S. van de Geer. Sparse spectral estimation with missing and corrupted measurements. *Stat*, 8(1):e229, 2019.
- [FP16] L. Florescu and W. Perkins. Spectral thresholds in the bipartite stochastic block model. In *Proceedings of the Conference on Learning Theory*, pages 943–959, 2016.
- [FPV15] V. Feldman, W. Perkins, and S. Vempala. Subsampled power iteration: a unified algorithm for block models and planted csp’s. In *Advances in Neural Information Processing Systems*, pages 2836–2844, 2015.

- [FWZ16] J. Fan, W. Wang, and Y. Zhong. An  $\ell_\infty$  eigenvector perturbation bound and its application to robust covariance estimation. *arXiv preprint arXiv:1603.03516*, 2016.
- [FWZZ18] J. Fan, K. Wang, Y. Zhong, and Z. Zhu. Robust high dimensional factor models with applications to statistical machine learning. *arXiv preprint arXiv:1808.03889*, 2018.
- [GLMZ16] C. Gao, Y. Lu, Z. Ma, and H. H. Zhou. Optimal estimation and completion of matrices with biclustering structures. *The Journal of Machine Learning Research*, 17(1):5602–5630, 2016.
- [GMZZ17] C. Gao, Z. Ma, A. Y. Zhang, and H. H. Zhou. Achieving optimal misclassification proportion in stochastic block models. *The Journal of Machine Learning Research*, 18(1):1980–2024, 2017.
- [GRESS16] A. Gonen, D. Rosenbaum, Y. C. Eldar, and S. Shalev-Shwartz. Subspace learning with partial information. *The Journal of Machine Learning Research*, 17(1):1821–1841, 2016.
- [Gro11] D. Gross. Recovering low-rank matrices from few coefficients in any basis. *IEEE Transactions on Information Theory*, 57(3):1548–1566, March 2011.
- [GRY11] S. Gandy, B. Recht, and I. Yamada. Tensor completion and low-n-rank tensor recovery via convex optimization. *Inverse Problems*, 27(2):025010, 2011.
- [GSSB17] M. Ghassemi, Z. Shakeri, A. D. Sarwate, and W. U. Bajwa. STARK: Structured dictionary learning through rank-one tensor recovery. In *Proceedings of the 2017 IEEE 7th International Workshop on Computational Advances in Multi-Sensor Adaptive Processing*, pages 1–5. IEEE, 2017.
- [GV16] O. Guédon and R. Vershynin. Community detection in sparse networks via grothendieck’s inequality. *Probability Theory and Related Fields*, 165(3-4):1025–1049, 2016.
- [HSSS16] S. B. Hopkins, T. Schramm, J. Shi, and D. Steurer. Fast spectral algorithms from sum-of-squares proofs: Tensor decomposition and planted sparse vectors. In *Proceedings of the Forty-eighth Annual ACM Symposium on Theory of Computing*, pages 178–191. ACM, 2016.
- [HWX16] B. Hajek, Y. Wu, and J. Xu. Achieving exact cluster recovery threshold via semidefinite programming: Extensions. *IEEE Transactions on Information Theory*, 62(10):5918–5937, 2016.
- [HZC18] B. Hao, A. Zhang, and G. Cheng. Sparse and low-rank tensor estimation via cubic sketchings. *arXiv preprint arXiv:1801.09326*, 2018.
- [JH12] J. Josse and F. Husson. Handling missing values in exploratory multivariate data analysis methods. *Journal de la Société Française de Statistique*, 153(2):79–99, 2012.
- [JL09] I. M. Johnstone and A. Y. Lu. On consistency and sparsity for principal components analysis in high dimensions. *Journal of the American Statistical Association*, 104(486):682–693, 2009.
- [JMRT16] A. Javanmard, A. Montanari, and F. Ricci-Tersenghi. Phase transitions in semidefinite relaxations. *Proceedings of the National Academy of Sciences of the U.S.A.*, 113(16):E2218–E2223, 2016.
- [JNS13] P. Jain, P. Netrapalli, and S. Sanghavi. Low-rank matrix completion using alternating minimization. In *Proceedings of the Forty-fifth Annual ACM Symposium on Theory of Computing*, pages 665–674, 2013.
- [JO14] P. Jain and S. Oh. Provable tensor factorization with missing data. In *Advances in Neural Information Processing Systems*, pages 1431–1439, 2014.
- [Joh01] I. M. Johnstone. On the distribution of the largest eigenvalue in principal components analysis. *The Annals of statistics*, 29(2):295–327, 2001.
- [Joh12] T. Johansson. The giant component of the random bipartite graph, 2012.

- [Kie97] H. A. Kiers. Weighted least squares fitting using ordinary least squares algorithms. *Psychometrika*, 62(2):251–266, 1997.
- [KL16] V. Koltchinskii and K. Lounici. Asymptotics and concentration bounds for bilinear forms of spectral projectors of sample covariance. In *Annales de l’Institut Henri Poincaré, Probabilités et Statistiques*, volume 52, pages 1976–2013. Institut Henri Poincaré, 2016.
- [KMO10a] R. H. Keshavan, A. Montanari, and S. Oh. Matrix completion from a few entries. *IEEE Transactions on Information Theory*, 56(6):2980–2998, June 2010.
- [KMO10b] R. H. Keshavan, A. Montanari, and S. Oh. Matrix completion from noisy entries. *Journal of Machine Learning Research*, 11(Jul):2057–2078, 2010.
- [KOKC13] H.-J. Kim, E. Ollila, V. Koivunen, and C. Croux. Robust and sparse estimation of tensor decompositions. In *Proceedings of the 2013 IEEE Global Conference on Signal and Information Processing*, pages 965–968. IEEE, 2013.
- [KX16] V. Koltchinskii and D. Xia. Perturbation of linear forms of singular vectors under gaussian noise. In *High Dimensional Probability VII*, pages 397–423. Springer, 2016.
- [LBEK18] L. Lei, P. J. Bickel, and N. El Karoui. Asymptotics for high dimensional regression m-estimates: fixed design results. *Probability Theory and Related Fields*, 172(3-4):983–1079, 2018.
- [LCJ14] D. B. Larremore, A. Clauset, and A. Z. Jacobs. Efficiently inferring community structure in bipartite networks. *Physical Review E*, 90(1):012805, 2014.
- [Lei19] L. Lei. Unified  $\ell_{2 \rightarrow \infty}$  eigenspace perturbation theory for symmetric random matrices. *arXiv preprint arXiv:1909.04798*, 2019.
- [LM62] D. N. Lawley and A. E. Maxwell. Factor analysis as a statistical method. *Journal of the Royal Statistical Society. Series D (The Statistician)*, 12(3):209–229, 1962.
- [LMWY13] J. Liu, P. Musialski, P. Wonka, and J. Ye. Tensor completion for estimating missing values in visual data. *IEEE transactions on Pattern Analysis and Machine Intelligence*, 35(1):208–220, 2013.
- [LMX15] M. Lelarge, L. Massoulié, and J. Xu. Reconstruction in the labelled stochastic block model. *IEEE Transactions on Network Science and Engineering*, 2(4):152–163, 2015.
- [Lou13] K. Lounici. Sparse principal component analysis with missing observations. In *High Dimensional Probability VI*, pages 327–356. Springer, 2013.
- [Lou14] K. Lounici. High-dimensional covariance matrix estimation with missing observations. *Bernoulli*, 20(3):1029–1058, 2014.
- [LR15] J. Lei and A. Rinaldo. Consistency of spectral clustering in stochastic block models. *The Annals of Statistics*, 43(1):215–237, 2015.
- [LW12] P.-L. Loh and M. J. Wainwright. High-dimensional regression with noisy and missing data: Provable guarantees with nonconvexity. *The Annals of Statistics*, 40(3):1637–1664, 06 2012.
- [Ma13] Z. Ma. Sparse principal component analysis and iterative thresholding. *The Annals of Statistics*, 41(2):772–801, 2013.
- [Mas07] P. Massart. *Concentration inequalities and model selection*, volume 6. Springer, 2007.
- [Mas14] L. Massoulié. Community detection thresholds and the weak ramanujan property. In *Proceedings of the Forty-sixth Annual ACM Symposium on Theory of Computing*, pages 694–703. ACM, 2014.



- [MHWG14] C. Mu, B. Huang, J. Wright, and D. Goldfarb. Square deal: Lower bounds and improved relaxations for tensor recovery. In *Proceedings of the International Conference on Machine Learning*, pages 73–81, 2014.
- [MNS14] E. Mossel, J. Neeman, and A. Sly. Consistency thresholds for binary symmetric block models. *arXiv preprint arXiv:1407.1591*, 3(5), 2014.
- [MNS15] E. Mossel, J. Neeman, and A. Sly. Reconstruction and estimation in the planted partition model. *Probability Theory and Related Fields*, 162(3-4):431–461, 2015.
- [MS18] A. Montanari and N. Sun. Spectral algorithms for tensor completion. *Communications on Pure and Applied Mathematics*, 71(11):2381–2425, 2018.
- [MSC17] X. Mao, P. Sarkar, and D. Chakrabarti. Estimating mixed memberships with sharp eigenvector deviations. *arXiv preprint arXiv:1709.00407*, 2017.
- [MWCC17] C. Ma, K. Wang, Y. Chi, and Y. Chen. Implicit regularization in nonconvex statistical estimation: gradient descent converges linearly for phase retrieval, matrix completion and blind deconvolution. *arXiv preprint arXiv:1711.10467*, 2017.
- [Nad08] B. Nadler. Finite sample approximation results for principal component analysis: A matrix perturbation approach. *The Annals of Statistics*, 36(6):2791–2817, 2008.
- [OVW18] S. O’Rourke, V. Vu, and K. Wang. Random perturbation of low rank matrices: Improving classical bounds. *Linear Algebra and its Applications*, 540:26–59, 2018.
- [Pau07] D. Paul. Asymptotics of sample eigenstructure for a large dimensional spiked covariance model. *Statistica Sinica*, pages 1617–1642, 2007.
- [PO19] E. Pavez and A. Ortega. Covariance matrix estimation with non uniform and data dependent missing observations. *arXiv preprint arXiv:1910.00667*, 2019.
- [PS17] A. Potechin and D. Steurer. Exact tensor completion with sum-of-squares. In *Proceedings of the Conference on Learning Theory*, pages 1619–1673, 2017.
- [PW19] A. Pananjady and M. J. Wainwright. Value function estimation in Markov reward processes: Instance-dependent  $\ell_\infty$ -bounds for policy evaluation. *arXiv preprint arXiv:1909.08749*, 2019.
- [RCY11] K. Rohe, S. Chatterjee, and B. Yu. Spectral clustering and the high-dimensional stochastic blockmodel. *The Annals of Statistics*, 39(4):1878–1915, 2011.
- [RM14] E. Richard and A. Montanari. A statistical model for tensor PCA. In *Advances in Neural Information Processing Systems*, pages 2897–2905, 2014.
- [RPP13] B. Romera-Paredes and M. Pontil. A new convex relaxation for tensor completion. In *Advances in Neural Information Processing Systems*, pages 2967–2975, 2013.
- [RV15] M. Rudelson and R. Vershynin. Delocalization of eigenvectors of random matrices with independent entries. *Duke Mathematical Journal*, 164(13):2507–2538, 2015.
- [SCC17] P. Sur, Y. Chen, and E. J. Candès. The likelihood ratio test in high-dimensional logistic regression is asymptotically a rescaled chi-square. *Probability Theory and Related Fields*, pages 1–72, 2017.
- [SHKM14] O. Semerci, N. Hao, M. E. Kilmer, and E. L. Miller. Tensor-based formulation and nuclear norm regularization for multienergy computed tomography. *IEEE Transactions on Image Processing*, 23(4):1678–1693, 2014.
- [SHSS16] Y. Shen, Q. Huang, N. Srebro, and S. Sanghavi. Normalized spectral map synchronization. In *Advances in Neural Information Processing Systems*, pages 4925–4933, 2016.

- [Sin11] A. Singer. Angular synchronization by eigenvectors and semidefinite programming. *Applied and Computational Harmonic Analysis*, 30(1):20–36, 2011.
- [SL16] R. Sun and Z.-Q. Luo. Guaranteed matrix completion via non-convex factorization. *IEEE Transactions on Information Theory*, 62(11):6535–6579, 2016.
- [SRK09] J. Salmi, A. Richter, and V. Koivunen. Sequential unfolding SVD for tensors with applications in array signal processing. *IEEE Transactions on Signal Processing*, 57(12):4719–4733, 2009.
- [SS90] G. W. Stewart and J.-g. Sun. Matrix perturbation theory. 1990.
- [STFP12] D. L. Sussman, M. Tang, D. E. Fishkind, and C. E. Priebe. A consistent adjacency spectral embedding for stochastic blockmodel graphs. *Journal of the American Statistical Association*, 107(499):1119–1128, 2012.
- [TB77] J. M. Ten Berge. Orthogonal procrustes rotation for two or more matrices. *Psychometrika*, 42(2):267–276, 1977.
- [Tsy08] A. B. Tsybakov. *Introduction to nonparametric estimation*. Springer Science & Business Media, 2008.
- [Ver12] R. Vershynin. Introduction to the non-asymptotic analysis of random matrices. *Compressed Sensing, Theory and Applications*, pages 210 – 268, 2012.
- [Vu11] V. Vu. Singular vectors under random perturbation. *Random Structures & Algorithms*, 39(4):526–538, 2011.
- [Vu18] V. Vu. A simple SVD algorithm for finding hidden partitions. *Combinatorics, Probability and Computing*, 27(1):124–140, 2018.
- [Wed72] P.-Å. Wedin. Perturbation bounds in connection with singular value decomposition. *BIT Numerical Mathematics*, 12(1):99–111, 1972.
- [Wed73] P.-Å. Wedin. Perturbation theory for pseudo-inverses. *BIT Numerical Mathematics*, 13(2):217–232, 1973.
- [XY17] D. Xia and M. Yuan. On polynomial time methods for exact low-rank tensor completion. *Foundations of Computational Mathematics*, pages 1–49, 2017.
- [XYZ17] D. Xia, M. Yuan, and C.-H. Zhang. Statistically optimal and computationally efficient low rank tensor completion from noisy entries. *arXiv preprint arXiv:1711.04934*, 2017.
- [XZ19] D. Xia and F. Zhou. The sup-norm perturbation of HOSVD and low rank tensor denoising. *Journal of Machine Learning Research*, 20(61):1–42, 2019.
- [YP14] S.-Y. Yun and A. Proutiere. Accurate community detection in the stochastic block model via spectral algorithms. *arXiv preprint arXiv:1412.7335*, 2014.
- [YP16] S.-Y. Yun and A. Proutiere. Optimal cluster recovery in the labeled stochastic block model. In *Advances in Neural Information Processing Systems*, pages 965–973, 2016.
- [YWS14] Y. Yu, T. Wang, and R. J. Samworth. A useful variant of the davis–kahan theorem for statisticians. *Biometrika*, 102(2):315–323, 2014.
- [YZ16] M. Yuan and C.-H. Zhang. On tensor completion via nuclear norm minimization. *Foundations of Computational Mathematics*, 16(4):1031–1068, 2016.
- [YZ17] M. Yuan and C.-H. Zhang. Incoherent tensor norms and their applications in higher order tensor completion. *IEEE Transactions on Information Theory*, 63(10):6753–6766, 2017.
- [ZA18] Z. Zhou and A. A. Amini. Optimal bipartite network clustering. *arXiv preprint arXiv:1803.06031*, 2018.

- [ZB18] Y. Zhong and N. Boumal. Near-optimal bounds for phase synchronization. *SIAM Journal on Optimization*, 28(2):989–1016, 2018.
- [ZCW18] A. Zhang, T. T. Cai, and Y. Wu. Heteroskedastic PCA: Algorithm, optimality, and applications. *arXiv preprint arXiv:1810.08316v2*, 2018.
- [ZWS19] Z. Zhu, T. Wang, and R. J. Samworth. High-dimensional principal component analysis with heterogeneous missingness. *arXiv preprint arXiv:1906.12125*, 2019.
- [ZX18] A. Zhang and D. Xia. Tensor SVD: Statistical and computational limits. *IEEE Transactions on Information Theory*, 64(11):7311–7338, 2018.