Asymmetry Helps: Eigenvalue and Eigenvector Analyses of Asymmetrically Perturbed Low-Rank Matrices

Yuxin Chen^{*} Chen Cheng[†] Jianqing Fan[‡]

November 2018; Revised February 2020

Abstract

This paper is concerned with the interplay between statistical asymmetry and spectral methods. Suppose we are interested in estimating a rank-1 and symmetric matrix $M^* \in \mathbb{R}^{n \times n}$, yet only a randomly perturbed version M is observed. The noise matrix $M - M^*$ is composed of independent (but not necessarily homoscedastic) entries and is, therefore, not symmetric in general. This might arise if, for example, we have two independent samples for each entry of M^* and arrange them in an *asymmetric* fashion. The aim is to estimate the leading eigenvalue and the leading eigenvector of M^* .

We demonstrate that the leading eigenvalue of the data matrix M can be $O(\sqrt{n})$ times more accurate (up to some log factor) than its (unadjusted) leading singular value of M in eigenvalue estimation. Moreover, the eigen-decomposition approach is fully adaptive to heteroscedasticity of noise, without the need of any prior knowledge about the noise distributions. In a nutshell, this curious phenomenon arises since the statistical asymmetry automatically mitigates the bias of the eigenvalue approach, thus eliminating the need of careful bias correction. Additionally, we develop appealing non-asymptotic eigenvector perturbation bounds; in particular, we are able to bound the perterbation of any linear function of the leading eigenvector of M (e.g. entrywise eigenvector perturbation). We also provide partial theory for the more general rank-r case. The takeaway message is this: arranging the data samples in an asymmetric manner and performing eigen-decomposition could sometimes be quite beneficial.

Keywords: asymmetric matrices, eigenvalue perturbation, entrywise eigenvector perturbation, linear forms of eigenvectors, heteroscedasticity.

Contents

1	Introduction	2
2	Problem formulation 2.1 Models and assumptions 2.2 Our goal 2.3 Incoherence conditions 2.4 Netation	4 4 5
3	Preliminaries 3.1 Perturbation of eigenvalues of asymmetric matrices 3.2 The Neumann trick and eigenvector perturbation	5 6 6 6

Author names are sorted alphabetically. Corresponding author: Yuxin Chen.

*Department of Electrical Engineering, Princeton University, Princeton, NJ 08544, USA; Email: yuxin.chen@princeton.edu. †Department of Statistics, Stanford University, Stanford, CA 94305, USA; Email: chencheng@stanford.edu.

[‡]Department of Operations Research and Financial Engineering, Princeton University, Princeton, NJ 08544, USA; Email: jqfan@princeton.edu.

4	Perturbation analysis for the rank-1 case	7
	4.1 Main results: the rank-1 case	7
	4.1.1 A master bound	7
	4.1.2 Eigenvalue perturbation	8
	4.1.3 Perturbation of linear forms of eigenvectors	10
	4.1.4 Entrywise eigenvector perturbation	10
	4.2 Applications	11
	4.3 Why asymmetry helps?	12
	4.4 Proof outline of Theorem 3	13
F	Extension, porturbation analysis for the rank a case	14
9	5.1 Figure lue perturbation for the real <i>n</i> case	14
	5.1 Eigenvalue perturbation for the rank- τ case	14
	5.2 Application: spectral estimation when <i>M</i> is asymmetric and rank-1	10
	5.5 Proof of Theorem 4	10
6	Discussions	19
٨	Proofs for proliminary facts	24
A	A 1 Proof of Lemma 2	24
	A 2 Proof of Theorem 2	24
	A 3 Proof of Lemma 3	24
		20
В	Proof for the lower bound in Lemma 4	26
С	Proof of Lemma 5	26
D	Proof of Lemma 6	29
-		
Е	Proof of Lemma 8	29
	E.I Graphical tools	29
	E.2 Proof of Lemma 8	31
	E.3 Proof of Lemma 9	31
\mathbf{F}	Proof of Corollary 4	32
G	Proof of Corollary 5	32
н	Proof of Corollary 6	33
Ι	Proof of Theorem 5	33
	I.1 ℓ_2 eigenvector perturbation bounds	34
	I.2 Perturbation bounds for linear forms of eigenvectors	35
	I.3 Entrywise eigenvector perturbation bounds	36
J	Asymmetrization of data samples: two examples	36

1 Introduction

Consider an unknown symmetric and low-rank matrix $M^* \in \mathbb{R}^{n \times n}$. What we have observed is a corrupted version

$$\boldsymbol{M} = \boldsymbol{M}^{\star} + \boldsymbol{H},\tag{1}$$

with H denoting a noise matrix. A classical problem is concerned with estimating the leading eigenvalues and eigenspace of M^* given observation M.

The current paper concentrates on a scenario where the noise matrix H (and hence M) consists of independently generated random entries and is hence *asymmetric* in general. This might arise, for example, when we have available multiple (e.g. two) samples for each entry of M^* and arrange the samples in an asymmetric fashion. A natural approach that immediately comes to mind is based on singular value decomposition (SVD), which employs the leading singular values (resp. subspace) of M to approximately estimate the eigenvalues (resp. eigenspace) of M^* . By contrast, a much less popular alternative is based on eigen-decomposition of the asymmetric data matrix M, which attempts approximation using the leading eigenvalues and eigenspace of M. Given that eigen-decomposition of an asymmetric matrix is in general not as numerically stable as SVD, conventional wisdom often favors the SVD-based approach, unless certain symmetrization step is implemented prior to eigen-decomposition.

When comparing these two approaches numerically, however, a curious phenomenon arises, which largely motivates the research in this paper. Let us generate M^* as a random rank-1 matrix with leading eigenvalue $\lambda^* = 1$, and let H be a Gaussian random matrix whose entries are i.i.d. $\mathcal{N}(0, \sigma^2)$ with $\sigma = 1/\sqrt{n \log n}$. Fig. 1(a) compares the empirical accuracy of estimating the 1st eigenvalue of M^* via the leading eigenvalue (the blue line) and via the leading singular value of M (the red line). As it turns out, eigen-decomposition significantly outperforms vanilla SVD in estimating λ^* , and the advantage seems increasingly more remarkable as the dimensionality n grows. To facilitate comparison, we include an additional green line in Fig. 1(a), obtained by rescaling the red line by $2.5/\sqrt{n}$. Interestingly, this green line coincides almost perfectly with the blue line, thus suggesting orderwise gain of eigen-decomposition compared to SVD. What is more, this phenomenon does not merely happen under i.i.d. noise. Similar numerical behaviors are observed in the problem of matrix completion — as displayed in Fig. 1(b) — even though the components of the equivalent perturbation matrix are apparently far from identically distributed or homoscedastic.



(a) eigenvalue perturbation for matrix estimation (b) eigenvalue perturbation for matrix completion

Figure 1: Numerical error $|\lambda - \lambda^*|$ vs. the matrix dimension n, where λ is either the leading eigenvalue (the blue line) or the leading singular value (the red line) of M. Here, (a) is the case when $\{H_{ij}\}$ are i.i.d. $\mathcal{N}(0, \sigma^2)$ with $\sigma = 1/\sqrt{n \log n}$, and (b) is the matrix completion case with sampling rate $p = 3 \log n/n$, where $M_{i,j} = \frac{1}{p} M_{i,j}^*$ independently with probability p and 0 otherwise. The results are averaged over 100 independent trials. The green lines are obtained by rescaling the corresponding red lines by $2.5/\sqrt{n}$.

The goal of the current paper is thus to develop a systematic understanding of this phenomenon, that is, why statistical asymmetry empowers eigen-decomposition and how to exploit this feature in statistical estimation. Informally, our findings suggest that: when M^* is rank-1 and H is composed of zero-mean and independent (but not necessarily identically distributed or homoscedastic) entries,

- 1. the leading eigenvalue of M could be $O(\sqrt{n})$ times (up to some logarithmic factor) more accurate than the (unadjusted) leading singular value of M when estimating the 1st eigenvalue of M^{\star} ;¹
- 2. the perturbation of the leading eigenvector is well-controlled along an *arbitrary* deterministic direction; for example, the eigenvector perturbation is well-controlled in any coordinate, indicating that the eigenvector estimation error is spread out across all coordinates.

¹More precisely, this gain is possible when $\|H\|$ is nearly as large as $\|M^{\star}\|$ (up to some logarithmic factor).

We will further provide partial theory to accommodate the rank-r case. As an important application, such a theory allows us to estimate the leading singular value and singular vectors of an asymmetric rank-1 matrix via eigen-decomposition of a certain dilation matrix, which also outperforms the vanilla SVD approach.

We would like to immediately remark that: for some scenarios (e.g. the case with i.i.d. Gaussian noise), it is possible to adjust the leading singular value of M to obtain the same accuracy as the leading eigenvalue of M. As it turns out, the advantages of the eigen-decomposition approach may become more evident in the presence of heteroscedasticity — the case where the noise has location-varying and unknown variance. We shall elaborate on this point in Section 4.1.2.

All in all, when it comes to low-rank matrix estimation, arranging the observed matrix samples in an asymmetric manner and invoking eigen-decomposition properly could sometimes be statistically beneficial.

2 Problem formulation

2.1 Models and assumptions

In this section, we formally introduce our models and assumptions. Consider a symmetric and low-rank matrix $M^{\star} = [M_{ij}^{\star}]_{1 \leq i,j \leq n} \in \mathbb{R}^{n \times n}$. Suppose we are given a random copy of M^{\star} as follows

$$\boldsymbol{M} = \boldsymbol{M}^{\star} + \boldsymbol{H},\tag{2}$$

where $\boldsymbol{H} = [H_{ij}]_{1 \le i,j \le n}$ is a random noise matrix.

The current paper concentrates on independent — but not necessarily identically distributed or homoscedastic — noise. Specifically, we impose the following assumptions on H throughout this paper.

Assumption 1.

- 1. (Independent entries) The entries $\{H_{ij}\}_{1 \le i,j \le n}$ are independently generated;
- 2. (Zero mean) $\mathbb{E}[H_{ij}] = 0$ for all $1 \leq i, j \leq n$;
- 3. (Variance) $\operatorname{Var}(H_{ij}) = \mathbb{E} \left[H_{ij}^2 \right] \leq \sigma_n^2$ for all $1 \leq i, j \leq n$;
- 4. (Magnitude) Each H_{ij} $(1 \le i, j \le n)$ satisfies either of the following conditions:
 - (a) $|H_{ij}| \leq B_n;$
 - (b) H_{ij} has a symmetric distribution obeying $\mathbb{P}\{|H_{ij}| > B_n\} \le c_b n^{-12}$ for some universal constant $c_b > 0$.

Remark 1 (Notational convention). In what follows, the dependency of σ_n and B_n on n shall often be suppressed whenever it is clear from the context, so as to simplify notation.

Note that we do not enforce the constraint $H_{ij} = H_{ji}$, and hence \boldsymbol{H} and \boldsymbol{M} are in general asymmetric matrices. Also, Condition 3 does not require the H_{ij} 's to have equal variance across different locations; in fact, they can be heteroscedastic. In addition, while Condition 4(a) covers the class of bounded random variables, Condition 4(b) allows us to accommodate a large family of heavy-tailed distributions (e.g. sub-exponential distributions). An immediate consequence of Assumption 1 is the following bound on the spectral norm $\|\boldsymbol{H}\|$ of \boldsymbol{H} .

Lemma 1. Under Assumption 1, there exist some universal constants $c_0, C_0 > 0$ such that with probability exceeding $1 - C_0 n^{-10}$,

$$\|\boldsymbol{H}\| \le c_0 \sigma \sqrt{n \log n} + c_0 B \log n. \tag{3}$$

Proof. This is a standard *non-asymptotic* result that follows immediately from the matrix Bernstein inequality Tropp (2015) and the union bound (for Assumption 4(b)). We omit the details for conciseness. \Box

2.2 Our goal

The aim is to develop *non-asymptotic* eigenvalue and eigenvector perturbation bounds under this family of random and asymmetric noise matrices. Our theoretical development is divided into two parts. Below we introduce our goal as well as some notation used throughout.

Rank-1 symmetric case. For the rank-1 case, we assume the eigen-decomposition of M^{\star} to be

$$\boldsymbol{M}^{\star} = \lambda^{\star} \boldsymbol{u}^{\star} \boldsymbol{u}^{\star \top} \tag{4}$$

with λ^* and u^* being its leading eigenvalue and eigenvector, respectively. We also denote by λ and u the leading eigenvalue and eigenvector of M, respectively. The following quantities are the focal points of this paper (see Section 4):

- 1. Eigenvalue perturbation: $|\lambda \lambda^{\star}|$;
- 2. Perturbation of linear forms of eigenvectors: $\min\{|\boldsymbol{a}^{\top}(\boldsymbol{u}-\boldsymbol{u}^{\star})|, |\boldsymbol{a}^{\top}(\boldsymbol{u}+\boldsymbol{u}^{\star})|\}$ for any fixed unit vector $\boldsymbol{a} \in \mathbb{R}^{n}$;
- 3. Entrywise eigenvector perturbation: $\min\{\|\boldsymbol{u} \boldsymbol{u}^{\star}\|_{\infty}, \|\boldsymbol{u} + \boldsymbol{u}^{\star}\|_{\infty}\}.$

Rank-*r* symmetric case. For the general rank-*r* case, we let the eigen-decomposition of M^{\star} be

$$M^{\star} = U^{\star} \Sigma^{\star} U^{\star \top}, \tag{5}$$

where the columns of $U^* = [u_1^*, \dots, u_r^*] \in \mathbb{R}^{n \times r}$ are the eigenvectors, and $\Sigma^* = \text{diag}(\lambda_1^*, \dots, \lambda_r^*) \in \mathbb{R}^{r \times r}$ is a diagonal matrix with the eigenvalues arranged in descending order by their magnitude, i.e. $|\lambda_1^*| \geq \cdots \geq |\lambda_r^*|$. We let $\lambda_{\max}^* = |\lambda_1^*|$ and $\lambda_{\min}^* = |\lambda_r^*|$. In addition, we let the top-*r* eigenvalues (in magnitude) of M be $\lambda_1, \dots, \lambda_r$ (obeying $|\lambda_1| \geq \cdots \geq |\lambda_r|$) and their corresponding normalized eigenvectors be u_1, \dots, u_r . We will present partial eigenvalue perturbation results for this more general case, as detailed in Section 5.

As is well-known, eigen-decomposition can be applied to estimate the singular values and singular vectors of an asymmetric matrix M^* via the standard dilation trick Tropp (2015). As a consequence, our results are also applicable for singular value and singular vector estimation. See Section 5.2 for details.

2.3 Incoherence conditions

Finally, we single out an incoherence parameter that plays an important role in our theory, which captures how well the energy of the eigenvectors is spread out across all entries.

Definition 1 (Incoherence parameter). The incoherence parameter of a rank-*r* symmetric matrix M^* with eigen-decomposition $M^* = U^* \Sigma^* U^{*\top}$ is defined to be the smallest quantity μ obeying

$$\|\boldsymbol{U}^{\star}\|_{\infty} \le \sqrt{\frac{\mu}{n}},\tag{6}$$

where $\|\cdot\|_{\infty}$ denotes the entrywise ℓ_{∞} norm.

Remark 2. An alternative definition of the incoherence parameter (Candès and Recht, 2009; Keshavan et al., 2010; Chi et al., 2019; Chen et al., 2019a) is the smallest quantity μ_0 satisfying $\|\mathbf{U}^*\|_{2,\infty} \leq \sqrt{\mu_0 r/n}$. This is a weaker assumption than Definition 1, as it only requires the energy of \mathbf{U}^* to be spread out across all of its rows rather than all of its entries. Note that these two incoherent parameters are consistent in the rank-1 case; in the rank-r case one has $\mu_0 \leq \mu \leq \mu_0 r$.

2.4 Notation

The standard basis vectors in \mathbb{R}^n are denoted by e_1, \dots, e_n . For any vector z, we let $||z||_2$ and $||z||_{\infty}$ denote the ℓ_2 norm and the ℓ_{∞} norm of z, respectively. For any matrix M, denote by ||M||, $||M||_F$ and $||M||_{\infty}$ the spectral norm, the Frobenius norm and the entrywise ℓ_{∞} norm (the largest magnitude of all entries) of M, respectively. Let $[n] := \{1, \dots, n\}$. In addition, the notation f(n) = O(g(n)) or $f(n) \leq g(n)$ means that there is a constant c > 0 such that $|f(n)| \leq c|g(n)|$, $f(n) \geq g(n)$ means that there is a constant c > 0 such that $|f(n)| \geq c|g(n)|$, and $f(n) \approx g(n)$ means that there exist constants $c_1, c_2 > 0$ such that $c_1|g(n)| \leq |f(n)| \leq c_2|g(n)|$.

3 Preliminaries

Before continuing, we gather several preliminary facts that will be useful throughout. The readers familiar with matrix perturbation theory may proceed directly to the main theoretical development in Section 4.

3.1 Perturbation of eigenvalues of asymmetric matrices

We begin with a standard result concerning eigenvalue perturbation of a diagonalizable matrix Bauer and Fike (1960). Note that the matrices under study might be asymmetric.

Theorem 1 (Bauer-Fike Theorem). Consider a diagonalizable matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$ with eigen-decomposition $\mathbf{A} = \mathbf{V} \mathbf{\Lambda} \mathbf{V}^{-1}$, where $\mathbf{V} \in \mathbb{C}^{n \times n}$ is a non-singular eigenvector matrix and $\mathbf{\Lambda}$ is diagonal. Let $\tilde{\lambda}$ be an eigenvalue of $\mathbf{A} + \mathbf{H}$. Then there exists an eigenvalue λ of \mathbf{A} such that

$$|\lambda - \tilde{\lambda}| \le \|\boldsymbol{V}\| \, \|\boldsymbol{V}^{-1}\| \, \|\boldsymbol{H}\| \,. \tag{7}$$

In addition, if **A** is symmetric, then there exists an eigenvalue λ of **A** such that

$$|\lambda - \tilde{\lambda}| \le \|\boldsymbol{H}\|. \tag{8}$$

However, caution needs to be exercised as the Bauer-Fike Theorem does not specify which eigenvalue of A is close to an eigenvalue of A + H. Encouragingly, in the low-rank case of interest, the Bauer-Fike Theorem together with certain continuity of the spectrum allows one to localize the leading eigenvalues of the perturbed matrix.

Lemma 2. Suppose M^* is a rank-r symmetric matrix whose top-r eigenvalues obey $|\lambda_1^*| \ge \cdots \ge |\lambda_r^*| > 0$. If $||\mathbf{H}|| < |\lambda_r^*|/2$, then the top-r eigenvalues $\lambda_1, \cdots, \lambda_r$ of $\mathbf{M} = \mathbf{M}^* + \mathbf{H}$, sorted by modulus, obey that: for any $1 \le l \le r$,

$$|\lambda_l - \lambda_j^{\star}| \le \|\boldsymbol{H}\| \quad \text{for some } 1 \le j \le r.$$
(9)

In addition, if r = 1, then both the leading eigenvalue and the leading eigenvector of M are real-valued.

This result, which we establish in Appendix A.1, parallels Weyl's inequality for symmetric matrices. Note, however, that the above bound (9) might be quite loose for specific settings. We will establish much sharper perturbation bounds when H contains independent random entries (see, e.g. Corollary 1).

3.2 The Neumann trick and eigenvector perturbation

Next, we introduce a classical result dubbed as the "Neumann trick" Eldridge et al. (2018). This theorem, which is derived based on the Neumann series for a matrix inverse, has been applied to analyze eigenvectors in various settings Erdős et al. (2013); Jain and Netrapalli (2015); Eldridge et al. (2018).

Theorem 2 (Neumann trick). Consider the matrices M^* and M (see (5) and (2)). Suppose $||H|| < |\lambda_l|$ for some $1 \le l \le n$. Then

$$\boldsymbol{u}_{l} = \sum_{j=1}^{r} \frac{\lambda_{j}^{\star}}{\lambda_{l}} \left(\boldsymbol{u}_{j}^{\star \top} \boldsymbol{u}_{l} \right) \left\{ \sum_{s=0}^{\infty} \frac{1}{\lambda_{l}^{s}} \boldsymbol{H}^{s} \boldsymbol{u}_{j}^{\star} \right\}.$$
(10)

Proof. We supply the proof in Appendix A.2 for self-containedness.

Remark 3. In particular, if M^* is a rank-1 matrix and $||H|| < |\lambda_1|$, then

$$\boldsymbol{u}_{1} = \frac{\lambda_{1}^{\star}}{\lambda_{1}} \left(\boldsymbol{u}_{1}^{\star \top} \boldsymbol{u}_{1} \right) \left\{ \sum_{s=0}^{\infty} \frac{1}{\lambda_{1}^{s}} \boldsymbol{H}^{s} \boldsymbol{u}_{1}^{\star} \right\}.$$
(11)

An immediate consequence of the Neumann trick is the following lemma, which asserts that each of the top-r eigenvectors of M resides almost within the top-r eigen-subspace of M^* , provided that ||H|| is sufficiently small. The proof is deferred to Appendix A.3.

Lemma 3. Suppose M^* is a rank-r symmetric matrix with r non-zero eigenvalues obeying $1 = \lambda_{\max}^* = |\lambda_1^*| \ge \cdots \ge |\lambda_r^*| = \lambda_{\min}^* > 0$ and associated eigenvectors u_1^*, \cdots, u_r^* . Define $\kappa \triangleq \lambda_{\max}^* / \lambda_{\min}^*$. If $||H|| \le 1/(4\kappa)$, then the top-r eigenvectors u_1, \cdots, u_r of $M = M^* + H$ obey

$$\sum_{j=1}^{r} |\boldsymbol{u}_{j}^{\star \top} \boldsymbol{u}_{l}|^{2} \ge 1 - \frac{64\kappa^{4}}{9} \|\boldsymbol{H}\|^{2}, \qquad 1 \le l \le r.$$
(12)

In addition, if r = 1, then one further has

$$\min\{\|\boldsymbol{u}_1 - \boldsymbol{u}_1^{\star}\|_2, \|\boldsymbol{u}_1 + \boldsymbol{u}_1^{\star}\|_2\} \le \frac{8\sqrt{2}}{3}\|\boldsymbol{H}\|.$$
(13)

4 Perturbation analysis for the rank-1 case

4.1 Main results: the rank-1 case

This section presents perturbation analysis results when the truth M^* is a symmetric rank-1 matrix. We shall start by presenting a master bound which, as we will see, immediately leads to our main findings.

4.1.1 A master bound

Our master bound is concerned with the perturbation of linear forms of eigenvectors, as stated below.

Theorem 3 (Perturbation of linear forms of eigenvectors (rank-1)). Consider a rank-1 symmetric matrix $M^* = \lambda^* u^* u^{*\top} \in \mathbb{R}^{n \times n}$ with incoherence parameter μ (cf. Definition 1). Suppose the noise matrix H obeys Assumption 1, and assume the existence of some sufficiently small constant $c_1 > 0$ such that

$$\max\left\{\sigma\sqrt{n\log n}, B\log n\right\} \le c_1 |\lambda^\star|. \tag{14}$$

Then for any fixed vector $\mathbf{a} \in \mathbb{R}^n$ with $\|\mathbf{a}\|_2 = 1$, with probability at least $1 - O(n^{-10})$ one has

$$\left| \boldsymbol{a}^{\top} \left(\boldsymbol{u} - \frac{\boldsymbol{u}^{\star \top} \boldsymbol{u}}{\lambda/\lambda^{\star}} \boldsymbol{u}^{\star} \right) \right| \lesssim \frac{\max\left\{ \sigma \sqrt{n \log n}, B \log n \right\}}{\left| \lambda^{\star} \right|} \sqrt{\frac{\mu}{n}}.$$
(15)

Remark 4 (The noise size). We would like to remark on the range of the noise size covered by our theory. If the incoherence parameter of the truth M^* obeys $\mu \approx 1$, then even the magnitude of the largest entry of M^* cannot exceed the order of $|\lambda^*|/n$. One can thus interpret the condition (14) in this case as

$$\sigma \lesssim \sqrt{\frac{n}{\log n}} \| \boldsymbol{M}^{\star} \|_{\infty} \quad and \qquad B \lesssim \frac{n}{\log n} \| \boldsymbol{M}^{\star} \|_{\infty}.$$

In other words, the standard deviation σ of each noise component is allowed to be substantially larger (i.e. $\sqrt{n/\log n}$ times larger) than the magnitude of any of the true entries. In fact, this condition (14) matches, up to some log factor, the one required for spectral methods to perform noticeably better than random guessing.

In words, Theorem 3 tells us that: the quantity $\frac{\boldsymbol{u}^{\star\top}\boldsymbol{u}}{\lambda/\lambda^{\star}}\boldsymbol{a}^{\top}\boldsymbol{u}^{\star}$ serves as a remarkably accurate approximation of the linear form $\boldsymbol{a}^{\top}\boldsymbol{u}$. In particular, the approximation error is at most $O(1/\sqrt{n})$ under the condition (14) for incoherent matrices. Encouragingly, this approximation accuracy holds true for an arbitrary deterministic direction (reflected by \boldsymbol{a}). As a consequence, one can roughly interpret Theorem 3 as

$$\boldsymbol{u} \approx \frac{\lambda^{\star}}{\lambda} \boldsymbol{u}^{\star} \boldsymbol{u}^{\star \top} \boldsymbol{u} = \frac{1}{\lambda} \boldsymbol{M}^{\star} \boldsymbol{u}, \tag{16}$$

where such an approximation is fairly accurate along any fixed direction. Compared with the identity $u = \frac{1}{\lambda} M u = \frac{1}{\lambda} (M^* + H) u$, our results imply that H u is exceedingly small along any fixed direction, even though H and u are highly dependent. As we shall explain in Section 4.3, this observation usually cannot happen when H is a symmetric random matrix or when one uses the leading singular vector instead, due to the significant bias resulting from symmetry.

This master theorem has several interesting implications, as we shall elucidate momentarily.

4.1.2 Eigenvalue perturbation

To begin with, Theorem 3 immediately yields a much sharper non-asymptotic perturbation bound regarding the leading eigenvalue λ of M.

Corollary 1. Under the assumptions of Theorem 3, with probability at least $1 - O(n^{-10})$ we have

$$|\lambda - \lambda^{\star}| \lesssim \max\left\{\sigma\sqrt{n\log n}, B\log n\right\}\sqrt{\frac{\mu}{n}}.$$
 (17)

Proof. Without loss of generality, assume that $\lambda^* = 1$. Taking $a = u^*$ in Theorem 3, we get

$$\left|\boldsymbol{u}^{\star\top}\boldsymbol{u}\right|\frac{|\lambda-1|}{\lambda} = \left|\boldsymbol{u}^{\star\top}\boldsymbol{u} - \boldsymbol{u}^{\star\top}\boldsymbol{u}^{\star}\frac{\boldsymbol{u}^{\star\top}\boldsymbol{u}}{\lambda}\right| \lesssim \max\left\{\sigma\sqrt{n\log n}, B\log n\right\}\sqrt{\frac{\mu}{n}}.$$
(18)

From Lemma 1 and the condition (14), we know $||\boldsymbol{H}|| < 1/4$, which combines with Lemma 2 and Lemma 3 yields $\lambda \simeq |\boldsymbol{u}^{\star \top} \boldsymbol{u}| \simeq 1$. Substitution into (18) yields

$$|\lambda - 1| \lesssim \left| \frac{\lambda}{\boldsymbol{u}^{\star \top} \boldsymbol{u}} \right| \max\left\{ \sigma \sqrt{n \log n}, B \log n \right\} \sqrt{\frac{\mu}{n}} \lesssim \max\left\{ \sigma \sqrt{n \log n}, B \log n \right\} \sqrt{\frac{\mu}{n}}.$$

For the vast majority of applications we encounter, the maximum possible noise magnitude B (cf. Assumption 1) obeys $B \lesssim \sigma \sqrt{n/\log n}$, in which case the bound in Corollary 1 simplifies to

$$|\lambda - \lambda^*| \lesssim \sigma \sqrt{\mu \log n}.\tag{19}$$

This means that the eigenvalue estimation error is not much larger than the variability of each noise component. In addition, we remind the reader that for a fairly broad class of noise (see Remark 4), the leading eigenvalue λ of M is guaranteed to be real-valued, an observation that has been made in Lemma 2. In practice, however, one might still encounter some scenarios where λ is complex-valued. As a result, we recommend the practitioner to use the real part of λ as the eigenvalue estimate, which clearly enjoys the same statistical guarantee as in Corollary 1.

Comparison to the vanilla SVD-based approach. In order to facilitate comparison, we denote by λ_{svd} the largest singular value of M, and look at $|\lambda_{svd} - \lambda^*|$. Combining Weyl's inequality, Lemma 1 and the condition (14), we arrive at

$$|\lambda_{\mathsf{svd}} - \lambda^{\star}| \le \|\boldsymbol{H}\| \lesssim \max\left\{\sigma\sqrt{n\log n}, B\log n\right\}.$$
(20)

When $\mu \approx 1$, this error bound w.r.t. this (unadjusted) singular value could be \sqrt{n} times larger than the perturbation bound (17) derived for the leading eigenvalue. This corroborates our motivating experiments in Fig. 1.

Comparison to vanilla eigen-decomposition after symmetrization. The reader might naturally wonder what would happen if we symmetrize the data matrix before performing eigen-decomposition. Consider, for example, the i.i.d. Gaussian noise case where $H_{ij} \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, \sigma^2)$, and assume $\lambda^* > 0$ for simplicity. The leading eigenvalue λ_{sym} of the symmetrized matrix $(\mathbf{M} + \mathbf{M}^{\top})/2$ has been extensively studied in the literature Füredi and Komlós (1981); Yin et al. (1988); Péché (2006); Féral and Péché (2007); Benaych-Georges and Nadakuditi (2011); Renfrew and Soshnikov (2013); Knowles and Yin (2013). In particular, it has been shown (e.g. Capitaine et al. (2009)) that, with probability approaching one,

$$\lambda_{\mathsf{sym}} = \lambda^* + \frac{n\sigma^2}{2\lambda^*} + O(\sigma\sqrt{\log n}).$$
⁽²¹⁾

If $\sigma = |\lambda^{\star}| \sqrt{1/(n \log n)}$ (which is the setting in our numerical experiment), then this can be translated into

$$\frac{\lambda_{\mathsf{sym}} - \lambda^{\star}}{\lambda^{\star}} = \frac{1}{2\log n} + O\left(\frac{1}{\sqrt{n}}\right).$$

This implies that λ_{sym} suffers from a substantially larger bias than the leading eigenvalue λ obtained without symmetrization, since in this case we have (cf. Corollary 1)

$$\left|\frac{\lambda - \lambda^{\star}}{\lambda^{\star}}\right| \lesssim \frac{1}{\sqrt{n}}.$$
(22)

Comparison to properly adjusted eigen-decomposition and SVD-based methods. Armed with the approximation (21) in the i.i.d. Gaussian noise case, the careful reader might naturally suggest a properly corrected estimate $\lambda_{sym,c}$ as follows (again assuming $\lambda > 0$)

$$\lambda_{\text{sym,c}} = \frac{1}{2} \Big(\lambda_{\text{sym}} + \sqrt{\lambda_{\text{sym}}^2 - 2n\sigma^2} \Big), \tag{23}$$

which is a shrinkage-type estimate chosen to satisfy $\lambda_{sym} = \lambda_{sym,c} + \frac{n\sigma^2}{2\lambda_{sym,c}}$. A little algebra reveals that: if $\sigma = 1/\sqrt{n \log n}$, then

$$\left|\frac{\lambda_{\mathsf{sym},\mathsf{c}}-\lambda^{\star}}{\lambda^{\star}}\right|\lesssim\frac{1}{\sqrt{n}},$$

thus matching the estimation accuracy of λ (cf. (22)). In addition, some sort of universality results has been established as well in the literature (Capitaine et al., 2009), implying that the same approximation and correction are applicable to a broad family of zero-mean noise with identical variance. As we shall illustrate numerically in Section 4.2, this approach (i.e. $\lambda_{sym,c}$) performs almost identically to the one using vanilla eigen-decomposition without symmetrization. In addition, very similar observations have been made for the SVD-based approach (Silverstein, 1994; Yin et al., 1988; Péché, 2006; Féral and Péché, 2007; Benaych-Georges and Nadakuditi, 2012; Bryc and Silverstein, 2018); for the sake of brevity, we do not repeat the arguments here.

We would nevertheless like to single out a few statistical advantages of the eigen-decomposition approach without symmetrization. To begin with, λ is obtained via vanilla eigen-decomposition, and computing it does not rely on any kind of noise statistics. This is in stark contrast to the bias correction (23) in the presence of symmetric data, which requires prior knowledge about (or a very precise estimate of) the noise variance σ^2 . Leaving out this prior knowledge matter, a more important issue is that the approximation formula (21) assumes identical variance of noise components across all entries (i.e. homoscedasticity). While an approximation of this kind has been found for more general cases beyond homoscedastic noise (e.g. Bryc and Silverstein (2018)), the approximation formula (e.g. (Bryc and Silverstein, 2018, Theorem 1.1)) becomes fairly complicated, requires prior knowledge about all variance parameters, and is thus difficult to implement in practice. In comparison, the vanilla eigen-decomposition approach analyzed in Corollary 1 imposes no restriction on the noise statistics and is fully adaptive to heteroscedastic noise.

Lower bounds. To complete the picture, we provide a simple information-theoretic lower bound for the i.i.d. Gaussian noise case, which will be established in Appendix **B**.

Lemma 4. Fix any small constant $\varepsilon > 0$. Suppose that $H_{ij} \stackrel{i.i.d.}{\sim} \mathcal{N}(0, \sigma^2)$. Consider three matrices

$$\boldsymbol{M} = \lambda^{\star} \boldsymbol{u}^{\star \top} + \boldsymbol{H}, \qquad \widetilde{\boldsymbol{M}} = (\lambda^{\star} + \Delta) \, \boldsymbol{u}^{\star} \boldsymbol{u}^{\star \top} + \boldsymbol{H}, \qquad \widehat{\boldsymbol{M}} = (\lambda^{\star} - \Delta) \, \boldsymbol{u}^{\star} \boldsymbol{u}^{\star \top} + \boldsymbol{H}$$

with $\|\boldsymbol{u}^{\star}\|_{2} = 1$. If $\Delta \leq \sigma \sqrt{(\log_{2} 1.5 - \varepsilon) \log 2}$, then no algorithm can distinguish \boldsymbol{M} , $\widetilde{\boldsymbol{M}}$ and $\widehat{\boldsymbol{M}}$ with $p_{e} \leq \varepsilon$, where p_{e} is the minimax probability of error for testing three hypotheses (namely, the ones claiming that the true eigenvalues are λ^{\star} , $\lambda^{\star} + \Delta$, and $\lambda^{\star} - \Delta$, respectively).

In short, Lemma 4 asserts that one cannot possibly locate an eigenvalue to within a precision of Δ much better than σ , which reveals a fundamental limit that cannot be broken by any algorithm. In comparison, the vanilla eigen-decomposition method based on asymmetric data achieves an accuracy of $|\lambda - \lambda^*| \leq \sigma \sqrt{\log n}$ (cf. Corollary 1 and (19)) for the incoherent case, thus matching the information-theoretic lower bound up to some log factor. In fact, the extra $\sqrt{\log n}$ factor arises simply because we are aiming for a high-probability guarantee.

4.1.3 Perturbation of linear forms of eigenvectors

The master bound in Theorem 3 admits a more convenient form when controlling linear functions of the eigenvectors. The result is this:

Corollary 2. Under the same setting of Theorem 3, with probability at least $1 - O(n^{-10})$ we have

$$\min\left\{\left|\boldsymbol{a}^{\top}(\boldsymbol{u}-\boldsymbol{u}^{\star})\right|, \left|\boldsymbol{a}^{\top}(\boldsymbol{u}+\boldsymbol{u}^{\star})\right|\right\} \lesssim \left(\left|\boldsymbol{a}^{\top}\boldsymbol{u}^{\star}\right| + \sqrt{\frac{\mu}{n}}\right) \frac{\max\left\{\sigma\sqrt{n\log n}, B\log n\right\}}{\left|\lambda^{\star}\right|}.$$
(24)

Proof. Without loss of generality, assume that $u^{\star \top} u \ge 0$ and that $\lambda^{\star} = 1$. Then one has

$$\begin{split} |\boldsymbol{a}^{\top}(\boldsymbol{u}-\boldsymbol{u}^{\star})| &\leq \left|\boldsymbol{a}^{\top}\boldsymbol{u}-\boldsymbol{a}^{\top}\boldsymbol{u}^{\star}\frac{\boldsymbol{u}^{\star\top}\boldsymbol{u}}{\lambda}\right| + \left|\boldsymbol{a}^{\top}\boldsymbol{u}^{\star}\right| \left|\frac{\boldsymbol{u}^{\star\top}\boldsymbol{u}}{\lambda} - 1\right| \\ &\leq \max\left\{\sigma\sqrt{n\log n}, B\log n\right\}\sqrt{\frac{\mu}{n}} + \left|\boldsymbol{a}^{\top}\boldsymbol{u}^{\star}\right| \left|\frac{\boldsymbol{u}^{\star\top}\boldsymbol{u}}{\lambda} - 1\right|, \end{split}$$

where the last inequality arises from Theorem 3 as well as the definition of μ . In addition, apply Lemma 2 and Lemma 3 to obtain

$$\left|\frac{\boldsymbol{u}^{\star^{\top}}\boldsymbol{u}}{\lambda} - 1\right| \leq \frac{\boldsymbol{u}^{\star^{\top}}\boldsymbol{u}}{\lambda} \left|1 - \lambda\right| + \left|\boldsymbol{u}^{\star^{\top}}\boldsymbol{u} - 1\right| \lesssim \|\boldsymbol{H}\| \lesssim \max\left\{\sigma\sqrt{n\log n}, B\log n\right\}.$$

Putting the above bounds together concludes the proof.

The perturbation of linear forms of eigenvectors (or singular vectors) has not yet been well explored even for the symmetric case. One scenario that has been studied is linear forms of singular vectors under i.i.d. Gaussian noise (Koltchinskii and Xia, 2016; Xia, 2016). Our analysis — which is certainly different from Koltchinskii and Xia (2016) as our emphasis is eigen-decomposition — does not rely on the Gaussianality assumption, and accommodates a much broader class of random noise. Another work that has looked at linear forms of the leading singular vector is Ma et al. (2019) for phase retrieval and blind deconvolution, although the vector \boldsymbol{a} therein is specific to the problems (i.e. the design vectors) and cannot be made general.

Remark 5. The perturbation theory for linear forms of eigenvectors has been substantially extended in our follow-up work; the interested reader is referred to Cheng et al. (2020) for details.

4.1.4 Entrywise eigenvector perturbation

A straightforward consequence of Corollary 2 that is worth emphasizing is sharp entrywise control of the leading eigenvector as follows.

Corollary 3. Under the same setting of Theorem 3, with probability at least $1 - O(n^{-9})$ we have

$$\min\left\{\left\|\boldsymbol{u}-\boldsymbol{u}^{\star}\right\|_{\infty}, \left\|\boldsymbol{u}+\boldsymbol{u}^{\star}\right\|_{\infty}\right\} \lesssim \frac{\max\left\{\sigma\sqrt{n\log n}, B\log n\right\}}{\left|\lambda^{\star}\right|} \sqrt{\frac{\mu}{n}}.$$
(25)

Proof. Recognizing that $\|\boldsymbol{u} - \boldsymbol{u}^{\star}\|_{\infty} = \max_{i} |\boldsymbol{e}_{i}^{\top}\boldsymbol{u} - \boldsymbol{e}_{i}^{\top}\boldsymbol{u}^{\star}|$ and recalling our assumption $|\boldsymbol{e}_{i}^{\top}\boldsymbol{u}| \leq \sqrt{\mu/n}$, we can invoke Corollary 2 and the union bound to establish this entrywise bound.

We note that: while the ℓ_2 perturbation (or sin Θ distance) of eigenvectors or singular vectors has been extensively studied Davis and Kahan (1970); Wedin (1972); Vu (2011); Wang (2015); O'Rourke et al. (2013); Cai and Zhang (2018), the entrywise eigenvector behavior was much less explored. The prior literature

contains only a few entrywise eigenvector perturbation analysis results for settings very different from ours, e.g. the i.i.d. random matrix case Vu and Wang (2015); O'Rourke et al. (2016), the symmetric low-rank case Fan et al. (2018); Abbe et al. (2017); Eldridge et al. (2018), and the case with transition matrices for reversible Markov chains (Chen et al., 2019b). Our results add another instance to this body of works in providing entrywise eigenvector perburation bounds.

4.2 Applications

We apply our main results to two concrete matrix estimation problems and examine the effectiveness of these bounds. As before, M^* is a rank-1 matrix with incoherence parameter μ and leading eigenvalue λ^* .

Low-rank matrix estimation from Gaussian noise. Suppose that H is composed of i.i.d. Gaussian random variables $\mathcal{N}(0, \sigma^2)$.² If $\sigma \lesssim \frac{1}{\sqrt{n \log n}}$, applying Corollaries 1-3 reveals that with high probability,

$$|\lambda - \lambda^{\star}| \lesssim \sigma \sqrt{\mu \log n} \tag{26a}$$

$$\min\{\|\boldsymbol{u} - \boldsymbol{u}^{\star}\|_{\infty}, \|\boldsymbol{u} + \boldsymbol{u}^{\star}\|_{\infty}\} \lesssim \frac{\sigma\sqrt{\mu}\log n}{|\lambda^{\star}|}$$
(26b)

$$\min\{|\boldsymbol{a}^{\top}(\boldsymbol{u}-\boldsymbol{u}^{\star})|, |\boldsymbol{a}^{\top}(\boldsymbol{u}+\boldsymbol{u}^{\star})|\} \lesssim \left(|\boldsymbol{a}^{\top}\boldsymbol{u}^{\star}| + \sqrt{\frac{\mu}{n}}\right) \frac{\sigma\sqrt{n\log n}}{|\lambda^{\star}|}$$
(26c)

for any fixed unit vector $\mathbf{a} \in \mathbb{R}^n$. We have conducted additional numerical experiments in Fig. 2, which confirm our findings. It is also worth noting that empirically, eigen-decomposition and SVD applied to \mathbf{M} achieve nearly identical ℓ_2 and ℓ_{∞} errors when estimating the leading eigenvector of \mathbf{M}^* . In addition, we also include the numerical estimation error of the corrected eigenvalue $\lambda_{\text{sym,c}}$ (cf. (23)) of the symmetrized matrix $(\mathbf{M} + \mathbf{M}^{\top})/2$. As can be seen from Fig. 2, vanilla eigen-decomposition without symmetrization performs nearly identically to the one with symmetrization and proper correction.



Figure 2: Numerical simulation for rank-1 matrix estimation under i.i.d. Gaussian noise $\mathcal{N}(0, \sigma^2)$, where the rank-1 truth \mathbf{M}^* is generated randomly with leading eigenvalue 1. (a): $|\lambda - \lambda^*|$ vs. σ when n = 1000; (b) and (c): ℓ_{∞} and ℓ_2 eigenvector estimation errors vs. n with $\sigma = 1/\sqrt{n \log n}$, respectively. The blue (resp. red) lines represent the average errors over 100 independent trials using the vanilla eigen-decomposition (resp. SVD) approach applied to \mathbf{M} . The orange line in (a) represents the average errors over 100 independent trials using the corrected leading eigenvalue $\lambda_{\text{sym,c}}$ of the symmetrized matrix $(\mathbf{M} + \mathbf{M}^{\top})/2$ (cf. (23)).

Low-rank matrix completion. Suppose that M is generated using random partial entries of M^* as follows

$$M_{ij} = \begin{cases} \frac{1}{p} M_{ij}^{\star}, & \text{with probability } p, \\ 0, & \text{else,} \end{cases}$$
(27)

where p denotes the fraction of the entries of M^{\star} being revealed. It is straightforward to verify that $H = M - M^{\star}$ is zero-mean and obeys $|H_{ij}| \leq \frac{\mu}{np} := B$ and $\operatorname{Var}(H_{ij}) \leq \frac{\mu^2}{pn^2}$. Consequently, if $p \geq \frac{\mu^2 \log n}{n}$,

²In this case, one can take $B \simeq \sigma \sqrt{\log n}$, which clearly satisfies $B \log n \ll \sqrt{n\sigma^2 \log n}$.

then invoking Corollaries 1-3 yields

$$\frac{|\lambda - \lambda^{\star}|}{|\lambda^{\star}|} \lesssim \frac{1}{\sqrt{n}} \sqrt{\frac{\mu^3 \log n}{pn}}$$
(28a)

$$\min\{\|\boldsymbol{u} - \boldsymbol{u}^{\star}\|_{\infty}, \|\boldsymbol{u} + \boldsymbol{u}^{\star}\|_{\infty}\} \lesssim \frac{1}{\sqrt{n}} \sqrt{\frac{\mu^3 \log n}{pn}}$$
(28b)

$$\min\{|\boldsymbol{a}^{\top}(\boldsymbol{u}-\boldsymbol{u}^{\star})|, |\boldsymbol{a}^{\top}(\boldsymbol{u}+\boldsymbol{u}^{\star})|\} \lesssim \left(|\boldsymbol{a}^{\top}\boldsymbol{u}^{\star}| + \sqrt{\frac{\mu}{n}}\right) \sqrt{\frac{\mu^{2}\log n}{pn}}$$
(28c)

with high probability, where $a \in \mathbb{R}^n$ is any fixed unit vector. Additional numerical simulations have been carried out in Fig. 3 to verify these findings. Empirically, eigen-decomposition outperforms SVD in estimating both the leading eigenvalue and eigenvector of M^* .



Figure 3: Numerical simulation for rank-1 matrix completion, where the rank-1 truth M^* is randomly generated with leading eigenvalue 1 and sampling rate is $p = 3 \log n/n$. (a) $|\lambda - \lambda^*|$ vs. p when n = 1000; (b) and (c): ℓ_{∞} and ℓ_2 eigenvector estimation errors vs. n, respectively. The blue (resp. red) lines represent the average errors over 100 independent trials using the eigen-decomposition (resp. SVD) approach.

Finally, we remark that all the above applications assume the availability of an asymmetric data matrix M. One might naturally wonder whether there is anything useful we can say if only a symmetric matrix M is available. While this is in general difficult, our theory does have direct implications for both matrix completion and the case with i.i.d. Gaussian noise in the presence of symmetric data matrices; that is, it is possible to first asymmetrize the data matrix followed by eigen-decomposition. The interested reader is referred to Appendix J for details.

4.3 Why asymmetry helps?

We take a moment to develop some intuition underlying Theorem 3, focusing on the case with $\lambda^* = 1$ for simplicity. The key ingredient is the Neumann trick stated in Theorem 2. Specifically, in the rank-1 case we can expand

$$oldsymbol{u} = rac{1}{\lambda} \left(oldsymbol{u}^{\star op} oldsymbol{u}
ight) \sum_{s=0}^{\infty} rac{1}{\lambda^s} oldsymbol{H}^s oldsymbol{u}^{\star}.$$

A little algebra yields

$$\left| \boldsymbol{a}^{\top} \left(\boldsymbol{u} - \frac{\boldsymbol{u}^{\star \top} \boldsymbol{u}}{\lambda} \boldsymbol{u}^{\star} \right) \right| = \left| \frac{\boldsymbol{u}^{\star \top} \boldsymbol{u}}{\lambda} \sum_{s=1}^{\infty} \frac{\boldsymbol{a}^{\top} \boldsymbol{H}^{s} \boldsymbol{u}^{\star}}{\lambda^{s}} \right| \lesssim \sum_{s=1}^{\infty} \left| \frac{\boldsymbol{a}^{\top} \boldsymbol{H}^{s} \boldsymbol{u}^{\star}}{\lambda^{s}} \right|,$$
(29)

where the last inequality holds since (i) $|\boldsymbol{u}^{\star \top}\boldsymbol{u}| \leq 1$, and (ii) λ is real-valued and obeys $\lambda \approx 1$ if $||\boldsymbol{H}|| \ll 1$ (in view of Lemma 2). As a result, the perturbation can be well-controlled as long as $|\boldsymbol{a}^{\top}\boldsymbol{H}^{s}\boldsymbol{u}^{\star}|$ is small for every $s \geq 1$.

As it turns out, $\mathbf{a}^{\top} \mathbf{H}^s \mathbf{u}^*$ might be much better controlled when \mathbf{H} is random and asymmetric, in comparison to the case where \mathbf{H} is random and symmetric. To illustrate this point, it is perhaps the easiest to inspect the second-order term.

• Asymmetric case: when H is composed of independent zero-mean entries each with variance σ_n^2 , one has

$$\mathbb{E}\left[\boldsymbol{a}^{\top}\boldsymbol{H}^{2}\boldsymbol{u}^{\star}\right] = \boldsymbol{a}^{\top}\mathbb{E}\left[\boldsymbol{H}^{2}\right]\boldsymbol{u}^{\star} = \boldsymbol{a}^{\top}(\sigma^{2}\boldsymbol{I})\boldsymbol{u}^{\star} = \sigma^{2}\boldsymbol{a}^{\top}\boldsymbol{u}^{\star}.$$

• Symmetric case: when H is symmetric and its upper trangular part consists of independent zeromean entries with variance σ_n^2 , it holds that

$$\mathbb{E}\left[\boldsymbol{a}^{\top}\boldsymbol{H}^{2}\boldsymbol{u}^{\star}\right] = \boldsymbol{a}^{\top}\mathbb{E}\left[\boldsymbol{H}^{2}\right]\boldsymbol{u}^{\star} = \boldsymbol{a}^{\top}(n\sigma^{2}\boldsymbol{I})\boldsymbol{u}^{\star} = n\sigma^{2}\boldsymbol{a}^{\top}\boldsymbol{u}^{\star}.$$

In words, the term $a^{\top}H^2u^*$ in the symmetric case might have a significantly larger bias compared to the asymmetric case. This bias effect is substantial when $a^{\top}u^*$ is large (e.g. when $a = u^*$), which plays a crucial role in determining the size of eigenvalue perturbation.

The vanilla SVD-based approach can be interpreted in a similar manner. Specifically, we recognize that the leading singular value (resp. left singular vector) can be computed via the leading eigenvalue (resp. eigenvector) of the symmetric matrix MM^{\top} . Given that $MM^{\top}-M^*M^{*\top}$ is also symmetric, the aforementioned bias issue arises as well. This explains why vanilla eigen-decomposition might have an advantage over vanilla SVD when dealing with asymmetric matrices.

Finally, we remark that the aforementioned bias issue becomes less severe as $||\boldsymbol{H}||$ decreases. For example, when $||\boldsymbol{H}||$ is exceedingly small, the only dominant term on the right-hand side of (29) is $\boldsymbol{a}^{\top}\boldsymbol{H}\boldsymbol{u}^{\star}$, with all higher-order terms being vanishingly small. In this case, $\mathbb{E}[\boldsymbol{a}^{\top}\boldsymbol{H}\boldsymbol{u}^{\star}] = 0$ for both symmetric and asymmetric zero-mean noise matrices. As a consequence, the advantage of eigen-decomposition becomes negligible when dealing with nearly-zero noise. This observation is also confirmed in the numerical experiments reported in Fig. 2(a) and Fig. 3(a), where the two approaches achieve similar eigenvalue estimation accuracy when $\sigma \to 0$ (resp. $p \to 1$) in matrix estimation under Gaussian noise (resp. matrix completion). In fact, the case with very small $||\boldsymbol{H}||$ has been studied in the literature O'Rourke et al. (2013); Vu (2011); Eldridge et al. (2018). For example, it was shown in O'Rourke et al. (2013) that when $||\boldsymbol{H}|| \lesssim \frac{1}{\sqrt{n}} ||\boldsymbol{M}^{\star}||$, the singular value perturbation is also \sqrt{n} times smaller than the bound predicted by Weyl's theorem; similar improvement can be observed w.r.t. eigenvalue perturbation results achieve this gain even when $||\boldsymbol{H}||$ is nearly as large as $||\boldsymbol{M}^{\star}||$ (up to some logarithmic factor).

4.4 Proof outline of Theorem 3

This subsection outlines the main steps for establishing Theorem 3. To simplify presentation, we shall assume without loss of generality that

$$\lambda^{\star} = 1. \tag{30}$$

Throughout this paper, all the proofs are provided for the case when Conditions 1-3, 4(a) in Assumption 1 are valid. Otherwise, if Condition 4(b) is valid, then we can invoke the union bound to show that

$$\boldsymbol{M} = \boldsymbol{M}^{\star} + \tilde{\boldsymbol{H}} \tag{31}$$

with probability exceeding $1 - O(n^{-10})$, where $\tilde{H}_{ij} \triangleq H_{ij} \mathbb{1}_{\{|H_{ij}| \leq B\}}$ is the truncated noise and has magnitude bounded by *B*. Since H_{ij} has symmetric distribution, it is seen that $\mathbb{E}[\tilde{H}_{ij}] = 0$ and $\mathsf{Var}(\tilde{H}_{ij}) \leq \sigma^2$, which coincides with the case obeying Conditions 1-3, 4(a) in Assumption 1.

As already mentioned in Section 4.3, everything boils down to controlling $|a^{\top}H^{s}u^{\star}|$ for $s \geq 1$. This is accomplished via the following lemma.

Lemma 5 (Bounding higher-order terms). Consider any fixed unit vector $a \in \mathbb{R}^n$ and any positive integers s, k satisfying $Bsk \leq 2$ and $n\sigma^2 sk \leq 2$. Under the assumptions of Theorem 3,

$$\left| \mathbb{E}\left[\left(\boldsymbol{a}^{\top} \boldsymbol{H}^{s} \boldsymbol{u}^{\star} \right)^{k} \right] \right| \leq \frac{sk}{2} \max\left\{ \left(Bsk \right)^{sk}, \left(2n\sigma^{2}sk \right)^{sk/2} \right\} \left(\sqrt{\frac{\mu}{n}} \right)^{\kappa}.$$
(32)

Remark 6. A similar result in (Tao, 2013, Lemma 2.3) has studied the bilinear forms of the high order terms of an i.i.d. random matrix, with a few distinctions. First of all, Tao (2013) assumes that each entry of the noise matrix is i.i.d. and has finite fourth moment (if the noise variance is rescaled to be 1); these assumptions break in examples like matrix completion. Moreover, Tao (2013) focuses on the case with k = 2, and does not lead to high-probability bounds (which are crucial for, e.g. entrywise error control).

Using Markov's inequality and the union bound, we can translate Lemma 5 into a high probability bound as follows.

Corollary 4. Under the assumptions of Lemma 5, there exists some universal constant $c_2 > 0$ such that

$$|\boldsymbol{a}^{\top}\boldsymbol{H}^{s}\boldsymbol{u}^{\star}| \leq \left(c_{2}\max\left\{B\log n, \sqrt{n\sigma^{2}\log n}\right\}\right)^{s}\sqrt{\frac{\mu}{n}}, \quad \forall s \leq 20\log n$$

with probability $1 - O(n^{-10})$.

Proof. See Appendix \mathbf{F} .

In addition, in view of Lemma 1 and the condition (14), one has

$$\|\boldsymbol{H}\| \lesssim \max\left\{B\log n, \sqrt{n\sigma^2 \log n}\right\} < 1/10$$
(33)

with probability $1 - O(n^{-10})$, which together with Lemma 2 implies $\lambda \ge 3 \|\boldsymbol{H}\|$. This further leads to

$$\sum_{s:s\geq 20\log n} \left(\frac{\|\boldsymbol{H}\|}{\lambda}\right)^s \leq \frac{\|\boldsymbol{H}\|}{\lambda} \sum_{s:s\geq 20\log n-1} \left(\frac{\|\boldsymbol{H}\|}{\lambda}\right)^s \leq \frac{\|\boldsymbol{H}\|}{\lambda} \sum_{s:s\geq 20\log n-1} \frac{1}{3^s}$$
$$\lesssim \max\left\{B\log n, \sqrt{n\sigma^2\log n}\right\} \cdot n^{-10}.$$

Putting the above bounds together and using the fact that λ is real-valued and $\lambda \geq 1/2$ (cf. Lemma 2), we have

$$\begin{split} \left| \boldsymbol{a}^{\top} \left(\boldsymbol{u} - \frac{\boldsymbol{u}^{\star^{\top}} \boldsymbol{u}}{\lambda} \boldsymbol{u}^{\star} \right) \right| &= \left| \frac{\boldsymbol{u}^{\star^{\top}} \boldsymbol{u}}{\lambda} \sum_{s=1}^{+\infty} \frac{\boldsymbol{a}^{\top} \boldsymbol{H}^{s} \boldsymbol{u}^{\star}}{\lambda^{s}} \right| \\ &\lesssim \sum_{s=1}^{20 \log n} \frac{1}{\lambda^{s}} \left| \boldsymbol{a}^{\top} \boldsymbol{H}^{s} \boldsymbol{u}^{\star} \right| + \sum_{s=20 \log n}^{+\infty} \left(\frac{\|\boldsymbol{H}\|}{\lambda} \right)^{s} \\ &\leq \sqrt{\frac{\mu}{n}} \sum_{s=1}^{20 \log n} \left(2c_{2} \max\left\{ B \log n, \sqrt{n\sigma^{2} \log n} \right\} \right)^{s} + \frac{\max\left\{ B \log n, \sqrt{n\sigma^{2} \log n} \right\}}{n^{10}} \\ &\lesssim \max\left\{ B \log n, \sqrt{n\sigma^{2} \log n} \right\} \sqrt{\frac{\mu}{n}}, \end{split}$$

as long as max $\{B \log n, \sqrt{n\sigma^2 \log n}\}$ is sufficiently small. Here, the last line also uses the fact that $\mu \ge 1$ (and hence $\sqrt{\mu/n} \gg n^{-10}$). This concludes the proof.

5 Extension: perturbation analysis for the rank-r case

5.1 Eigenvalue perturbation for the rank-*r* case

The eigenvalue perturbation analysis in Section 4 can be extended to accommodate the case where M^{\star} is symmetric and rank-*r*, as detailed in this section. As before, assume that the *r* non-zero eigenvalues of M^{\star} obey $\lambda_{\max}^{\star} = |\lambda_1^{\star}| \geq \cdots \geq |\lambda_r^{\star}| = \lambda_{\min}^{\star}$. Once again, we start with a master bound.

Theorem 4 (Perturbation of linear forms of eigenvectors (rank-r)). Consider a rank-r symmetric matrix $M^* \in \mathbb{R}^{n \times n}$ with incoherence parameter μ . Define $\kappa \triangleq \lambda^*_{\max}/\lambda^*_{\min}$. Suppose that

$$\frac{\max\left\{\sigma\sqrt{n\log n}, B\log n\right\}}{\lambda_{\max}^{\star}} \le \frac{c_1}{\kappa} \tag{34}$$

for some sufficiently small constants $c_1 > 0$. Then for any fixed unit vector $\mathbf{a} \in \mathbb{R}^n$ and any $1 \le l \le r$, with probability at least $1 - O(n^{-10})$ one has

$$\left| \boldsymbol{a}^{\top} \left(\boldsymbol{u}_{l} - \sum_{j=1}^{r} \frac{\lambda_{j}^{\star} \boldsymbol{u}_{j}^{\star \top} \boldsymbol{u}_{l}}{\lambda_{l}} \boldsymbol{u}_{j}^{\star} \right) \right| \lesssim \max \left\{ \sigma \sqrt{n \log n}, B \log n \right\} \frac{\kappa}{|\lambda_{l}|} \sqrt{\frac{\mu r}{n}}$$
(35)

$$\lesssim \frac{\max\left\{\sigma\sqrt{n\log n}, B\log n\right\}}{\lambda_{\max}^{\star}} \kappa^2 \sqrt{\frac{\mu r}{n}}.$$
(36)

This result allows us to control the perturbation of the linear form of eigenvectors. The perturbation upper bound grows as either the rank r or the condition number κ increases.

One of the most important consequences of Theorem 4 is a refinement of the Bauer-Fike theorem concerning eigenvalue perturbations as follows.

Corollary 5. Consider the lth $(1 \le l \le r)$ eigenvalue λ_l of M. Under the assumptions of Theorem 4, with probability at least $1 - O(n^{-10})$, there exists $1 \le j \le r$ such that

$$\left|\lambda_{l} - \lambda_{j}^{\star}\right| \lesssim \max\left\{\sigma\sqrt{n\log n}, B\log n\right\} \kappa r \sqrt{\frac{\mu}{n}},\tag{37}$$

provided that

$$\frac{\max\left\{\sigma\sqrt{n\log n}, B\log n\right\}}{\lambda_{\max}^{\star}} \le c_1/\kappa^2 \tag{38}$$

for some sufficiently small constant $c_1 > 0$.

Proof. See Appendix G.

In comparison, the Bauer-Fike theorem (Lemma 2) together with Lemma 1 gives a perturbation bound

$$\left|\lambda_{l} - \lambda_{j}^{\star}\right| \leq \|\boldsymbol{H}\| \lesssim \max\left\{\sigma\sqrt{n\log n}, B\log n\right\} \quad \text{for some } 1 \leq j \leq r.$$
(39)

For the low-rank case where $r \ll \sqrt{n}$, the eigenvalue perturbation bound derived in Corollary 5 can be much sharper than the Bauer-Fike theorem.

Another result that comes from Theorem 4 is the following bound that concerns linear forms of the eigen-subspace.

Corollary 6. Under the same setting of Theorem 4, with probability $1 - O(n^{-9})$ we have

$$\left\|\boldsymbol{a}^{\top}\boldsymbol{U}\right\|_{2} \lesssim \kappa\sqrt{r} \left\|\boldsymbol{a}^{\top}\boldsymbol{U}^{\star}\right\|_{2} + \frac{\max\left\{\sigma\sqrt{n\log n}, B\log n\right\}}{\lambda_{\max}^{\star}}\kappa^{2}r\sqrt{\frac{\mu}{n}}.$$
(40)

Proof. See Appendix H.

Consequently, by taking $a = e_i$ $(1 \le i \le n)$ in Corollary 6, we arrive at the following statement regarding the alternative definition of the incoherence of the eigenvector matrix U (see Remark 2).

Corollary 7. Under the same setting of Theorem 4, with probability $1 - O(n^{-8})$ we have

$$\|\boldsymbol{U}\|_{2,\infty} \lesssim \kappa r \sqrt{\frac{\mu}{n}}.$$
(41)

Proof. Given that $\|\boldsymbol{U}\|_{2,\infty} = \max_{1 \le i \le n} \|\boldsymbol{e}_i^\top \boldsymbol{U}\|_2$ and recalling our assumption implies $\|\boldsymbol{U}^\star\|_{2,\infty} \le \sqrt{\mu r/n}$, we can invoke Corollary 6 and the union bound to derive the advertised entrywise bounds.

Remark 7. The eigenvector matrix is often employed to form a reasonably good initial guess for several nonconvex statistical estimation problems Keshavan et al. (2010), and the above kind of incoherence property is crucial in guaranteeing fast convergence of the subsequent nonconvex iterative refinement procedures Ma et al. (2019).

Unfortunately, these results fall short of providing simple perturbation bounds for the eigenvectors; in other words, the above-mentioned bounds do not imply the size of the difference between U and U^* . The challenge arises in part due to the lack of orthonormality of the eigenvectors when dealing with asymmetric matrices. Analyzing the eigenspace perturbation for the general rank-r case will likely require new analysis techniques, which we leave for future work. There is, however, some special case in which we can develop eigenvector perturbation theory, as detailed in the next subsection.

Remark 8. The theory for the rank-r case has recently been significantly improved; see our follow-up work Cheng et al. (2020) for details.

5.2 Application: spectral estimation when M^* is asymmetric and rank-1

In some scenarios, the above general rank results allow us to improve spectral estimation when M^* is asymmetric. Consider the case where $M^* = \lambda^* u^* v^{*\top} \in \mathbb{R}^{n_1 \times n_2}$ is an asymmetric rank-1 matrix with leading singular value λ^* . Suppose that we observe two independent noisy copies of M^* , namely,

$$M_1 = M^* + H_1, \qquad M_2 = M^* + H_2,$$
 (42)

where H_1 and H_2 are independent noise matrices. The goal is to estimate the singular value and singular vectors of M^* from M_1 and M_2 .

We attempt estimation via the standard dilation trick (e.g. Tao (2012)). This consists of embedding the matrices of interest within a larger block matrix

$$\boldsymbol{M}_{\mathsf{d}}^{\star} \triangleq \left[\begin{array}{cc} \boldsymbol{0} & \boldsymbol{M}^{\star} \\ \boldsymbol{M}^{\star \top} & \boldsymbol{0} \end{array} \right], \qquad \boldsymbol{M}_{\mathsf{d}} \triangleq \left[\begin{array}{cc} \boldsymbol{0} & \boldsymbol{M}_{1} \\ \boldsymbol{M}_{2}^{\top} & \boldsymbol{0} \end{array} \right].$$
(43)

Here, we place M_1 and M_2 in two different subblocks, in order to "asymmetrize" the dilation matrix. The rationale is that M_d^{\star} is a rank-2 symmetric matrix with exactly two nonzero eigenvalues

$$\lambda_1(\boldsymbol{M}_{\mathsf{d}}^{\star}) = \lambda^{\star} \qquad and \qquad \lambda_2(\boldsymbol{M}_{\mathsf{d}}^{\star}) = -\lambda^{\star},$$

whose corresponding eigenvectors are given by

$$\frac{1}{\sqrt{2}} \begin{pmatrix} \boldsymbol{u}^{\star} \\ \boldsymbol{v}^{\star} \end{pmatrix} \quad and \quad \frac{1}{\sqrt{2}} \begin{pmatrix} \boldsymbol{u}^{\star} \\ -\boldsymbol{v}^{\star} \end{pmatrix},$$

respectively. This motivates us to perform eigen-decomposition of M_d , and use the top-2 eigenvalues and eigenvectors to estimate λ^* , u^* and v^* , respectively.

Eigenvalue perturbation analysis. As an immediate consequence of Corollary 5, the two leading eigenvalues of M_d provide fairly accurate estimates of the leading singular value λ^* of M^* , as stated below.

Corollary 8. Assume $\mathbf{M}^{\star} \in \mathbb{R}^{n_1 \times n_2}$ is a rank-1 matrix with leading singular value λ^{\star} and incoherence parameter μ . Define $n \triangleq n_1 + n_2$. Suppose that $\lambda_1^{\mathsf{d}} \ge \lambda_2^{\mathsf{d}}$ are the two leading eigenvalues of \mathbf{M}_{d} (cf. (43)), and that \mathbf{H}_1 and \mathbf{H}_2 are independent and satisfy Assumption 1. Then with probability at least $1 - O(n^{-10})$,

$$\max\left\{\left|\lambda_{1}^{\mathsf{d}}-\lambda^{\star}\right|,\left|\lambda_{2}^{\mathsf{d}}+\lambda^{\star}\right|\right\} \lesssim \max\left\{\sigma\sqrt{n\log n},B\log n\right\}\sqrt{\frac{\mu}{n}},\tag{44}$$

provided that

$$\frac{\max\left\{\sigma\sqrt{n\log n}, B\log n\right\}}{\lambda^{\star}} \le c_1 \tag{45}$$

for some sufficiently small constant $c_1 > 0$.

Proof. To begin with, it follows from Corollary 5 that both λ_1^d and λ_2^d are close to either λ^* or $-\lambda^*$. Repeating similar arguments as in the proof of Lemma 2 (which we omit here), we can immediately show the separation between these two eigenvalues, namely, λ_1^d (resp. λ_2^d) is close to λ^* (resp. $-\lambda^*$).

Eigenvector perturbation analysis. We then move on to studying the eigenvector perturbation bounds. Specifically, denote by u_1^d and u_2^d the eigenvectors of M_d associated with its two leading eigenvalues λ_1^d and λ_2^d , respectively. Without loss of generality, we assume that $\lambda_1^d \ge \lambda_2^d$. If we write

$$oldsymbol{u}_1^{\mathsf{dilation}} = egin{pmatrix} oldsymbol{u}_{1,1}^{\mathsf{d}} \ oldsymbol{u}_{1,2}^{\mathsf{d}} \end{pmatrix} \qquad ext{with} \ oldsymbol{u}_{1,1}^{\mathsf{d}} \in \mathbb{R}^{n_1}, oldsymbol{u}_{1,2}^{\mathsf{d}} \in \mathbb{R}^{n_2},$$

then we can employ $u_{1,1}^{\mathsf{d}}$ and $u_{1,2}^{\mathsf{d}}$ to estimate u^{\star} and v^{\star} after proper normalization, namely,

$$\boldsymbol{u} \triangleq \frac{\boldsymbol{u}_{1,1}^{\mathsf{d}}}{\|\boldsymbol{u}_{1,1}^{\mathsf{d}}\|_{2}}, \qquad \boldsymbol{v} \triangleq \frac{\boldsymbol{u}_{1,2}^{\mathsf{d}}}{\|\boldsymbol{u}_{1,2}^{\mathsf{d}}\|_{2}}.$$
(46)

The following theorem develops error bounds for both \boldsymbol{u} and \boldsymbol{v} , which we establish in Appendix I. Here, we denote min $\|\boldsymbol{x} \pm \boldsymbol{y}\|_2 = \min\{\|\boldsymbol{x} - \boldsymbol{y}\|_2, \|\boldsymbol{x} + \boldsymbol{y}\|_2\}$, and min $\|\boldsymbol{x} \pm \boldsymbol{y}\|_{\infty} = \min\{\|\boldsymbol{x} - \boldsymbol{y}\|_{\infty}, \|\boldsymbol{x} + \boldsymbol{y}\|_{\infty}\}$.

Theorem 5. Suppose $\mathbf{M}^{\star} = \lambda^{\star} \mathbf{u}^{\star} \mathbf{v}^{\star \top} \in \mathbb{R}^{n_1 \times n_2}$ is a rank-1 matrix with leading singular value λ^{\star} and incoherence parameter μ , where $\|\mathbf{u}^{\star}\|_2 = \|\mathbf{v}^{\star}\|_2 = 1$. Define $n \triangleq n_1 + n_2$, and fix any unit vectors $\mathbf{a} \in \mathbb{R}^{n_1}$ and $\mathbf{b} \in \mathbb{R}^{n_2}$. Then with probability at least $1 - O(n^{-10})$, the estimates \mathbf{u} and \mathbf{v} (cf. (46)) obey

$$\max\left\{\min\|\boldsymbol{u} \pm \boldsymbol{u}^{\star}\|_{2}, \ \min\|\boldsymbol{v} \pm \boldsymbol{v}^{\star}\|_{2}\right\} \lesssim \frac{\max\left\{\sigma\sqrt{n\log n}, B\log n\right\}}{\lambda^{\star}},\tag{47a}$$

$$\max\left\{\min\left\|\boldsymbol{u} \pm \boldsymbol{u}^{\star}\right\|_{\infty}, \ \min\left\|\boldsymbol{v} \pm \boldsymbol{v}^{\star}\right\|_{\infty}\right\} \lesssim \frac{\max\left\{\sigma\sqrt{n\log n}, B\log n\right\}}{\lambda^{\star}} \sqrt{\frac{\mu}{n}}, \tag{47b}$$

$$\min\left\{\left|\boldsymbol{a}^{\top}(\boldsymbol{u}-\boldsymbol{u}^{\star})\right|, \left|\boldsymbol{a}^{\top}(\boldsymbol{u}+\boldsymbol{u}^{\star})\right|\right\} \lesssim \left(\left|\boldsymbol{a}^{\top}\boldsymbol{u}^{\star}\right| + \sqrt{\frac{\mu}{n}}\right) \frac{\max\left\{\sigma\sqrt{n\log n}, B\log n\right\}}{\lambda^{\star}}, \quad (47c)$$

$$\min\left\{\left|\boldsymbol{b}^{\top}(\boldsymbol{v}-\boldsymbol{v}^{\star})\right|,\left|\boldsymbol{b}^{\top}(\boldsymbol{v}+\boldsymbol{v}^{\star})\right|\right\} \lesssim \left(\left|\boldsymbol{b}^{\top}\boldsymbol{v}^{\star}\right| + \sqrt{\frac{\mu}{n}}\right) \frac{\max\left\{\sigma\sqrt{n\log n}, B\log n\right\}}{\lambda^{\star}},\tag{47d}$$

provided that there exists some some sufficiently small constant $c_1 > 0$ such that

$$\frac{\max\left\{\sigma\sqrt{n\log n}, B\log n\right\}}{\lambda^{\star}} \le c_1.$$
(48)

Similar to the symmetric rank-1 case, the estimation errors of the estimates u and v are well-controlled in any deterministic direction (e.g. the entrywise errors are well-controlled). This allows us to complete the theory for the case when M^* is a real-valued and rank-1 matrix.

Further, we conduct numerical experiments for matrix completion when M^* is a rank-1 and asymmetric matrix in Fig. 4. Here, we suppose that at most 1 sample is observed for each entry, and we estimate the singular value and singular vectors of M^* via the above-mentioned dilation trick, coupled with the asymmetrization procedure discussed in Section J. The numerical performance confirms that the proposed technique outperforms vanilla SVD in spectral estimation.

Finally, we remark that the asymptotic behavior of the eigenvalues of asymmetric random matrices has been extensively explored in the physics literature (e.g. (Sommers et al., 1988; Khoruzhenko, 1996; Brezin and Zee, 1998; Chalker and Mehlig, 1998; Feinberg and Zee, 1997; Lytova and Tikhomirov, 2018)). Their focus, however, has largely been to pin down the asymptotic density of the eigenvalues, similar to the semicircle law in the symmetric case. Nevertheless, a sharp perturbation bound for the leading eigenvalue particularly for the low-rank case — is beyond their reach. A few recent papers began to explore the locations of eigenvalue outliers that fall outside the bulk predicted by the circular law Tao (2013); Rajagopalan (2015); Benaych-Georges and Rochet (2016); Bordenave and Capitaine (2016). The results reported therein either do not focus on obtaining the right convergence rate (e.g. providing only a bound like $|\lambda - \lambda^*| = o(|\lambda^*|)$) or are restricted to a special family of ground truth (e.g. the one with a diagonal block equal to identity) or i.i.d. noise. As a result, these prior results are insufficient to demonstrate the power and benefits of the eigen-decomposition method in the presence of data asymmetry.



(a) eigenvalue perturbation vs. n (b) eigenvalue perturbation vs. p

(c) ℓ_{∞} eigenvector perturbation

Figure 4: Numerical experiments for rank-1 matrix completion, where $\mathbf{M}^{\star} = \mathbf{u}^{\star} \mathbf{v}^{\star \top} \in \mathbb{R}^{n_1 \times n_2}$ is randomly generated with leading singular value $\lambda^{\star} = 1$. Let $n = n_1 = 2n_2$. Each entry is observed independently with probability p. (a) $|\lambda - \lambda^{\star}|$ vs. n with $p = 3 \log n / n$; (b) $|\lambda - \lambda^{\star}|$ vs. p with n = 1000; (c) ℓ_{∞} eigenvector estimation error vs. n with $p = 3 \log n / n$. The blue (resp. red) lines represent the average errors over 100 independent trials using using the eigen-decomposition (resp. SVD) approach applied to $\mathbf{M}_{\text{dilation}}$ (resp. \mathbf{M}).

5.3 Proof of Theorem 4

Without loss of generality, we shall assume $\lambda_{\max}^{\star} = \lambda_1^{\star} = 1$ throughout the proof. To begin with, Lemma 2 implies that for all $1 \le l \le r$,

$$\lambda_l \geq |\lambda_{\min}^{\star}| - \|\boldsymbol{H}\| > 1/(2\kappa) > \|\boldsymbol{H}\|$$
(49)

as long as $\|H\| < 1/(2\kappa)$. In view of the Neumann trick (Theorem 2), we can derive

$$\begin{vmatrix} \boldsymbol{a}^{\top}\boldsymbol{u}_{l} - \sum_{j=1}^{r} \frac{\lambda_{j}^{*}\boldsymbol{u}_{j}^{*\top}\boldsymbol{u}_{l}}{\lambda_{l}} \boldsymbol{a}^{\top}\boldsymbol{u}_{j}^{*} \end{vmatrix} = \left| \sum_{j=1}^{r} \frac{\lambda_{j}^{*}}{\lambda_{l}} (\boldsymbol{u}_{j}^{*\top}\boldsymbol{u}_{l}) \left\{ \sum_{s=1}^{\infty} \frac{1}{\lambda_{l}^{*}} \boldsymbol{a}^{\top} \boldsymbol{H}^{s} \boldsymbol{u}_{j}^{*} \right\} \right| \\
\leq \left(\sum_{j=1}^{r} \frac{|\lambda_{j}^{*}|}{|\lambda_{l}|} |\boldsymbol{u}_{j}^{*\top}\boldsymbol{u}_{l}| \right) \left\{ \max_{1 \leq j \leq r} \sum_{s=1}^{\infty} \frac{1}{|\lambda_{l}|^{s}} |\boldsymbol{a}^{\top} \boldsymbol{H}^{s} \boldsymbol{u}_{j}^{*}| \right\} \\
\leq \sqrt{r} \sum_{j=1}^{r} |\boldsymbol{u}_{j}^{*\top}\boldsymbol{u}_{l}|^{2} \left\{ \max_{1 \leq j \leq r} \frac{|\lambda_{j}^{*}|}{|\lambda_{l}|} \right\} \left\{ \max_{1 \leq j \leq r} \sum_{s=1}^{\infty} \frac{1}{|\lambda_{l}|^{s}} |\boldsymbol{a}^{\top} \boldsymbol{H}^{s} \boldsymbol{u}_{j}^{*}| \right\} \\
\leq \sqrt{r} \cdot \frac{1}{|\lambda_{l}|} \cdot \left\{ \max_{1 \leq j \leq r} \sum_{s=1}^{\infty} \frac{1}{|\lambda_{l}|^{s}} |\boldsymbol{a}^{\top} \boldsymbol{H}^{s} \boldsymbol{u}_{j}^{*}| \right\},$$
(50)

where the third line follows since $\sum_{j=1}^{r} |\boldsymbol{u}_{j}^{\star \top} \boldsymbol{u}_{l}|^{2} \leq ||\boldsymbol{u}_{l}||_{2}^{2} = 1$, and the last inequality makes use of (49). Apply Corollary 4 to reach

$$(50) \leq \frac{\sqrt{r}}{|\lambda_l|} \sum_{s=1}^{\infty} \left(2c_2 \kappa \max\left\{ B \log n, \sqrt{n\sigma^2 \log n} \right\} \right)^s \sqrt{\frac{\mu}{n}}$$
$$\lesssim \frac{\kappa}{|\lambda_l|} \max\left\{ B \log n, \sqrt{n\sigma^2 \log n} \right\} \sqrt{\frac{\mu r}{n}}$$
$$\lesssim \kappa^2 \max\left\{ B \log n, \sqrt{n\sigma^2 \log n} \right\} \sqrt{\frac{\mu r}{n}},$$

with the proviso that $|\lambda_l| > 1/(2\kappa)$ and $\max \left\{ B \log n, \sqrt{n\sigma^2 \log n} \right\} \le c_1/\kappa$ for some sufficiently small constant $c_1 > 0$. The condition $|\lambda_l| > 1/(2\kappa)$ follows immediately by combining Lemma 2, Lemma 1 and the condition (34).



Figure 5: Numerical experiments for the spiked covariance model, where the sample vectors are zero-mean Gaussian vectors with covariance matrix Σ^* (cf. (51)). We plot $|\lambda - \lambda^*|$ vs. n with d = n/10. The blue (resp. red) lines represent the average errors over 100 independent trials when λ is the leading eigenvalue of $\widehat{\Sigma}_{asym}$ (resp. $\hat{\Sigma}$).

6 Discussions

In this paper, we demonstrate the remarkable advantage of eigen-decomposition over SVD in the presence of asymmetric noise matrices. This is in stark contrast to conventional wisdom, which is generally not in favor of eigen-decomposition for asymmetric matrices. Our results only reflect the tip of an iceberg, and there are many outstanding issues left answered. We conclude the paper with a few future directions.

Sharper eigenvalue perturbation bounds for the rank-r case. Our current results in Section 5 provide an eigenvalue perturbation bound on the order of r/\sqrt{n} , assuming the truth is rank-r. However, numerical experiments suggest that the dependency on r might be improvable. It would be interesting to see whether further theoretical refinement is possible, e.g. whether it is possible to improve it to $O(\sqrt{r/n})$.

Eigenvector perturbation bounds for the rank-r **case.** As mentioned before, the current theory falls short of providing eigenvector perturbation bounds for the general rank-r case. The main difficulty lies in the lack of orthogonality of the eigenvectors of the observed matrix M. Nevertheless, when the size of the noise is not too large, it is possible to establish certain near-orthogonality of the eigenvectors, which might in turn lead to sharp control of eigenvector perturbation.

A challenging signal-to-noise ratio regime. Take the rank-1 case for example: the present work focuses on the regime where $\|H\| \leq \|M^*\|/\sqrt{\log n}$, and it is known that spectral methods fail to yield reliable estimation if $\|H\| \gg \|M^*\|$. There is, however, a "gray" region (which includes, for example, the case with $\|H\| \approx \|M^*\|$) that has not been addressed. Developing non-asymptotic yet informative perturbation bounds for this regime is likely very challenging and requires new analysis techniques, which we leave for future investigation.

Correlated noise. The current theoretical development relies heavily on the assumption that the noise matrix H contains independent random entries. There is no shortage of examples where the noise matrix is asymmetric but is not composed of independent entries. For instance, in blind deconvolution Li et al. (2018), the noise matrix is a sum of independent asymmetric matrices. Can we develop eigenvalue perturbation theory for this class of noise?

Statistical inference of eigenvalues and eigenvectors. In various applications like network analysis and inference, one might be interested in determining the (asymptotic) eigenvalue and eigenvector distributions of a random data matrix, in order to produce valid confidence intervals Johnstone (2001); Bai and Yao (2008); Cai et al. (2017); Cape et al. (2018); Xia (2018); Bao et al. (2018); Chen et al. (2019c). Can we use the current framework to characterize the distributions of the leading eigenvalues as well as certain linear forms of the eigenvectors of M when the noise matrix is non-symmetric?

Asymmetrization for other applications. Given the abundant applications of spectral estimation, our findings are likely to be useful for other matrix eigenvalue problems and might extend to the tensor case Zhang

and Xia (2018); Cai et al. (2019). Here, we conclude the paper with an example in *covariance estimation* (Baik and Silverstein, 2006; Fan et al., 2018). Imagine that we observe a collection of n independent Gaussian vectors $X_1, \dots, X_n \in \mathbb{R}^d$, which have mean zero and covariance matrix

$$\boldsymbol{\Sigma}^{\star} = \boldsymbol{v}\boldsymbol{v}^{\top} + \boldsymbol{I}_d \tag{51}$$

with \boldsymbol{v} being a unit vector. This falls under the category of the spiked covariance model (Johnstone and Lu, 2009). One strategy to estimate the spectral norm $\lambda^* = 2$ of $\boldsymbol{\Sigma}^*$ is to look at the spectrum of the sample covariance matrix $\hat{\boldsymbol{\Sigma}} = \frac{1}{n} \sum_{i=1}^{n} \boldsymbol{X}_i \boldsymbol{X}_i^{\top}$. Motivated by the results of this paper, we propose an alternative strategy by looking at the following asymmetrized sample covariance matrix

$$\widehat{\boldsymbol{\Sigma}}_{\mathsf{asym}} = \frac{2}{n} \Big(\sum_{i=1}^{n/2} \mathsf{Upper}\left(\boldsymbol{X}_{i} \boldsymbol{X}_{i}^{\mathsf{T}}\right) + \sum_{i=n/2+1}^{n} \mathsf{Lower}\left(\boldsymbol{X}_{i} \boldsymbol{X}_{i}^{\mathsf{T}}\right) \Big), \tag{52}$$

where $\mathsf{Upper}(\cdot)$ (resp. $\mathsf{Lower}(\cdot)$) extracts out the upper (resp. lower) triangular part of the matrix, including (resp. excluding) the diagonal entries. As can be seen from Fig. 5, the largest eigenvalue of the asymmetrized $\hat{\Sigma}_{\mathsf{asym}}$ is much closer to the true spectral norm of Σ^{\star} , compared to the largest singular value of the sample covariance matrix $\hat{\Sigma}$. We leave the theoretical understanding of such findings to future investigation.

Acknowledgment

Y. Chen is supported in part by the AFOSR YIP award FA9550-19-1-0030, by the ARO grants W911NF-20-1-0097 and W911NF-18-1-0303, by the ONR grant N00014-19-1-2120, by the NSF grants CCF-1907661 and IIS-1900140, and by the Princeton SEAS innovation award. J. Fan is supported in part by the NSF grants DMS-1662139 and DMS-1712591, by the ONR grant N00014-19-1-2120, and by the NIH grant 2R01-GM072611-13. C. Cheng is supported in part by the Elite Undergraduate Training Program of School of Mathematical Sciences in Peking University, and by the William R. Hewlett Stanford graduate fellowship. We thank Cong Ma for helpful discussions, and Zhou Fan for telling us an example of asymmetrizing the Gaussian matrix.

References

- ABBE, E., FAN, J., WANG, K. and ZHONG, Y. (2017). Entrywise eigenvector analysis of random matrices with low expected rank. *under revision, Annals of Statistics*.
- BAI, Z. and YAO, J.-F. (2008). Central limit theorems for eigenvalues in a spiked population model. In Annales de l'Institut Henri Poincaré, Probabilités et Statistiques, vol. 44. Institut Henri Poincaré.
- BAIK, J. and SILVERSTEIN, J. W. (2006). Eigenvalues of large sample covariance matrices of spiked population models. *Journal of multivariate analysis* **97** 1382–1408.
- BAO, Z., DING, X. and WANG, K. (2018). Singular vector and singular subspace distribution for the matrix denoising model. arXiv preprint arXiv:1809.10476.
- BAUER, F. L. and FIKE, C. T. (1960). Norms and exclusion theorems. Numerische Mathematik 2 137-141.
- BENAYCH-GEORGES, F. and NADAKUDITI, R. R. (2011). The eigenvalues and eigenvectors of finite, low rank perturbations of large random matrices. *Advances in Mathematics* **227** 494–521.
- BENAYCH-GEORGES, F. and NADAKUDITI, R. R. (2012). The singular values and vectors of low rank perturbations of large rectangular random matrices. *Journal of Multivariate Analysis* **111** 120–135.
- BENAYCH-GEORGES, F. and ROCHET, J. (2016). Outliers in the single ring theorem. *Probability Theory* and Related Fields 165 313–363.
- BORDENAVE, C. and CAPITAINE, M. (2016). Outlier eigenvalues for deformed iid random matrices. *Com*munications on Pure and Applied Mathematics **69** 2131–2194.

- BREZIN, E. and ZEE, A. (1998). Non-hermitean delocalization: Multiple scattering and bounds. Nuclear Physics B 509 599–614.
- BRYC, W. and SILVERSTEIN, J. W. (2018). Singular values of large non-central random matrices. arXiv preprint arXiv:1802.02960.
- CAI, C., LI, G., CHI, Y., POOR, H. V. and CHEN, Y. (2019). Subspace estimation from unbalanced and incomplete data matrices: $\ell_{2,\infty}$ statistical guarantees. arXiv preprint arXiv:1910.04267.
- CAI, T., HAN, X. and PAN, G. (2017). Limiting laws for divergent spiked eigenvalues and largest non-spiked eigenvalue of sample covariance matrices. arXiv preprint arXiv:1711.00217.
- CAI, T. T. and ZHANG, A. (2018). Rate-optimal perturbation bounds for singular subspaces with applications to high-dimensional statistics. *The Annals of Statistics* **46** 60–89.
- CANDÈS, E. J. and RECHT, B. (2009). Exact matrix completion via convex optimization. Foundations of Computational Mathematics 9 717–772.
- CAPE, J., TANG, M. and PRIEBE, C. E. (2018). Signal-plus-noise matrix models: eigenvector deviations and fluctuations. arXiv preprint arXiv:1802.00381.
- CAPITAINE, M., DONATI-MARTIN, C. and FÉRAL, D. (2009). The largest eigenvalues of finite rank deformation of large wigner matrices: convergence and nonuniversality of the fluctuations. *The Annals of Probability* 37 1–47.
- CHALKER, J. T. and MEHLIG, B. (1998). Eigenvector statistics in non-hermitian random matrix ensembles. *Physical review letters* **81** 3367.
- CHEN, Y., CHI, Y., FAN, J., MA, C. and YAN, Y. (2019a). Noisy matrix completion: Understanding statistical guarantees for convex relaxation via nonconvex optimization. arXiv:1902.07698.
- CHEN, Y., FAN, J., MA, C. and WANG, K. (2019b). Spectral method and regularized MLE are both optimal for top-K ranking. Annals of Statistics 47 2204–2235.
- CHEN, Y., FAN, J., MA, C. and YAN, Y. (2019c). Inference and uncertainty quantification for noisy matrix completion. *Proceedings of the National Academy of Sciences* **116** 22931–22937.
- CHENG, C., WEI, Y. and CHEN, Y. (2020). Tackling small eigen-gaps: Fine-grained eigenvector estimation and inference under heteroscedastic noise. arXiv preprint arXiv:2001.04620.
- CHI, Y., LU, Y. M. and CHEN, Y. (2019). Nonconvex optimization meets low-rank matrix factorization: An overview. *IEEE Transactions on Signal Processing* 67 5239 – 5269.
- DAVIS, C. and KAHAN, W. M. (1970). The rotation of eigenvectors by a perturbation. iii. SIAM Journal on Numerical Analysis 7 1–46.
- ELDRIDGE, J., BELKIN, M. and WANG, Y. (2018). Unperturbed: spectral analysis beyond Davis-Kahan. In *Proceedings of Algorithmic Learning Theory*.
- EMBREE, M. and TREFETHEN, L. N. (2001). Generalizing eigenvalue theorems to pseudospectra theorems. SIAM Journal on Scientific Computing 23 583–590.
- ERDŐS, L., KNOWLES, A., YAU, H.-T., YIN, J. ET AL. (2013). Spectral statistics of Erdős–Rényi graphs I: local semicircle law. *The Annals of Probability* **41** 2279–2375.
- FAN, J., WANG, W. and ZHONG, Y. (2018). An ℓ_{∞} eigenvector perturbation bound and its application to robust covariance estimation. *Journal of Machine Learning* 18 1–42.
- FEINBERG, J. and ZEE, A. (1997). Non-hermitian random matrix theory: Method of hermitian reduction. Nuclear Physics B 504 579–608.

- FÉRAL, D. and PÉCHÉ, S. (2007). The largest eigenvalue of rank one deformation of large wigner matrices. Communications in mathematical physics 272 185–228.
- FÜREDI, Z. and KOMLÓS, J. (1981). The eigenvalues of random symmetric matrices. Combinatorica 1 233–241.
- JAIN, P. and NETRAPALLI, P. (2015). Fast exact matrix completion with finite samples. In *Conference on Learning Theory*.
- JOHNSTONE, I. M. (2001). On the distribution of the largest eigenvalue in principal components analysis. Annals of statistics 295–327.
- JOHNSTONE, I. M. and LU, A. Y. (2009). On consistency and sparsity for principal components analysis in high dimensions. *Journal of the American Statistical Association* **104** 682–693.
- KESHAVAN, R. H., MONTANARI, A. and OH, S. (2010). Matrix completion from a few entries. *IEEE Transactions on Information Theory* 56 2980 –2998.
- KHORUZHENKO, B. (1996). Large-*n* eigenvalue distribution of randomly perturbed asymmetric matrices. Journal of Physics A: Mathematical and General **29** L165.
- KNOWLES, A. and YIN, J. (2013). The isotropic semicircle law and deformation of wigner matrices. Communications on Pure and Applied Mathematics 66 1663–1749.
- KOLTCHINSKII, V. and XIA, D. (2016). Perturbation of linear forms of singular vectors under Gaussian noise. In *High Dimensional Probability VII*. Springer, 397–423.
- LI, X., LING, S., STROHMER, T. and WEI, K. (2018). Rapid, robust, and reliable blind deconvolution via nonconvex optimization. Applied and Computational Harmonic Analysis.
- LYTOVA, A. and TIKHOMIROV, K. (2018). On delocalization of eigenvectors of random non-hermitian matrices. arXiv preprint arXiv:1810.01590.
- MA, C., WANG, K., CHI, Y. and CHEN, Y. (2019). Implicit regularization in nonconvex statistical estimation: Gradient descent converges linearly for phase retrieval, matrix completion and blind deconvolution. accepted to Foundations of Computational Mathematics.
- O'ROURKE, S., VU, V. and WANG, K. (2013). Random perturbation of low rank matrices: Improving classical bounds. arXiv preprint arXiv:1311.2657.
- O'ROURKE, S., VU, V. and WANG, K. (2016). Eigenvectors of random matrices: a survey. Journal of Combinatorial Theory, Series A 144 361–442.
- PÉCHÉ, S. (2006). The largest eigenvalue of small rank perturbations of hermitian random matrices. Probability Theory and Related Fields 134 127–173.
- RAJAGOPALAN, A. B. (2015). Outlier eigenvalue fluctuations of perturbed iid matrices. arXiv preprint arXiv:1507.01441.
- RENFREW, D. and SOSHNIKOV, A. (2013). On finite rank deformations of wigner matrices ii: Delocalized perturbations. *Random Matrices: Theory and Applications* **2** 1250015.
- SILVERSTEIN, J. W. (1994). The spectral radii and norms of large dimensional non-central random atrices matrices. Stochastic Models 10 525–532.
- SOMMERS, H., CRISANTI, A., SOMPOLINSKY, H. and STEIN, Y. (1988). Spectrum of large random asymmetric matrices. *Physical review letters* **60** 1895.
- TAO, T. (2012). Topics in Random Matrix Theory. Graduate Studies in Mathematics, American Mathematical Society, Providence, Rhode Island.

- TAO, T. (2013). Outliers in the spectrum of iid matrices with bounded rank perturbations. Probability Theory and Related Fields 155 231–263.
- TROPP, J. A. (2015). An introduction to matrix concentration inequalities. Found. Trends Mach. Learn. 8 1–230.
- TSYBAKOV, A. B. (2009). Introduction to nonparametric estimation.
- Vu, V. (2011). Singular vectors under random perturbation. Random Structures & Algorithms 39 526–538.
- VU, V. and WANG, K. (2015). Random weighted projections, random quadratic forms and random eigenvectors. Random Structures and Algorithms 47 792–821.
- WANG, R. (2015). Singular vector perturbation under gaussian noise. SIAM Journal on Matrix Analysis and Applications 36 158–177.
- WEDIN, P. (1972). Perturbation bounds in connection with singular value decomposition. BIT Numerical Mathematics 12 99–111.
- XIA, D. (2016). Statistical inference for large matrices. Ph.D. thesis, Georgia Institute of Technology.
- XIA, D. (2018). Confidence interval of singular vectors for high-dimensional and low-rank matrix regression. arXiv preprint arXiv:1805.09871.
- YIN, Y.-Q., BAI, Z.-D. and KRISHNAIAH, P. R. (1988). On the limit of the largest eigenvalue of the large dimensional sample covariance matrix. *Probability theory and related fields* 78 509–521.
- ZHANG, A. and XIA, D. (2018). Tensor SVD: Statistical and computational limits. IEEE Transactions on Information Theory 64 7311 – 7338.

A Proofs for preliminary facts

A.1 Proof of Lemma 2

Without loss of generality, suppose the leading eigenvalue of M^{\star} is 1.

(1) We start by proving the rank-1 case, in which the eigenvalues of M^* are either 0 or 1. Theorem 1 immediately implies that

$$\Lambda(\boldsymbol{M}) \subseteq \mathcal{B}(1, \|\boldsymbol{H}\|) \cup \mathcal{B}(0, \|\boldsymbol{H}\|), \qquad (53)$$

where $\Lambda(\mathbf{M})$ is the set of eigenvalues of \mathbf{M} , and $\mathcal{B}(z,r) \triangleq \{x \in \mathbb{C} : |x-z| \leq r\}$ denotes a disk of radius r centered at z.

To establish the lemma, our aim is to show that there is exactly one eigenvalue of M — denoted by λ — lying in the disk $\mathcal{B}(1, ||\mathbf{H}||)$, and it has multiplicity 1. If this is true, then the assumption $||\mathbf{H}|| < 1/2$ indicates that $|\lambda| \ge 1 - ||\mathbf{H}|| > ||\mathbf{H}|| \ge z$ for any $z \in \mathcal{B}(0, ||\mathbf{H}||)$, and hence λ must be the leading eigenvalue. Furthermore, since M is a real-valued matrix, both λ and its complex conjugate $\overline{\lambda}$ are eigenvalues of M. However, if $\lambda \neq \overline{\lambda}$, then both of these two eigenvalues fall within $\mathcal{B}(1, ||\mathbf{H}||)$, resulting in contradiction. As a result, one necessarily has $\lambda = \overline{\lambda}$ and both of them are real-valued. A similar argument demonstrates that the eigenvector u associated with λ is also real-valued.

We then justify the existence and uniqueness of an eigenvalue in $\mathcal{B}(1, ||\mathbf{H}||)$. Denote $\Lambda(\mathbf{M}) = \{\lambda_1, \dots, \lambda_n\}$, and define a set of auxiliary matrices

$$\boldsymbol{M}_t = \boldsymbol{M}^\star + t \boldsymbol{H}, \qquad 0 \le t \le 1.$$

Similar to (53), one has

$$\Lambda(\boldsymbol{M}_t) \subseteq \mathcal{B}(1, t \|\boldsymbol{H}\|) \cup \mathcal{B}(0, t \|\boldsymbol{H}\|).$$
(54)

Recognizing that the set of eigenvalues of M_t depends continuously on t (e.g. (Embree and Trefethen, 2001, Theorem 6)), we can write

$$\Lambda(\boldsymbol{M}_t) = \left\{ \lambda_1(t), \lambda_2(t), \cdots, \lambda_n(t) \right\},\tag{55}$$

with each $\lambda_j(t), 1 \leq j \leq n$ being a continuous function in t. Meanwhile, as long as $\|\mathbf{H}\| < 1/2$ and $0 \leq t \leq 1$, the two disks $\mathcal{B}(1, t \|\mathbf{H}\|)$ and $\mathcal{B}(0, t \|\mathbf{H}\|)$ are always disjoint. Thus, the continuity of the spectrum (w.r.t. t) requires $\lambda_j(t)$ to always stay within the same disk where $\lambda_j(0) \in \{0, 1\}$ lies, namely,

$$\lambda_j(t) \in \mathcal{B}(\lambda_j(0), t \, \|\boldsymbol{H}\|). \tag{56}$$

Given that M^{\star} (or M_0) has n-1 eigenvalues equal to 0 and one eigenvalue equal to 1, we establish the lemma for the rank-1 case.

(2) We now turn to the rank-*r* case. Repeating the above argument for the rank-1 case, we can immediately show that: if $\|\boldsymbol{H}\| < \lambda_r^*/2$, then (i) there are exactly n - r eigenvalues lying within $\mathcal{B}(0, \|\boldsymbol{H}\|)$; (ii) all other eigenvalues lie within $\bigcup_{1 \le j \le r} \mathcal{B}(\lambda_j^*, \|\boldsymbol{H}\|)$, which are exactly the top-*r* leading eigenvalues of \boldsymbol{M} . This concludes the proof.

A.2 Proof of Theorem 2

From the definition of eigenvectors,

$$(\boldsymbol{M}^{\star} + \boldsymbol{H}) \boldsymbol{u}_{l} = \lambda_{l} \boldsymbol{u}_{l}, \quad \text{or equivalently,} \quad \frac{1}{\lambda_{l}} \boldsymbol{M}^{\star} \boldsymbol{u}_{l} = \left(\boldsymbol{I} - \frac{1}{\lambda_{l}} \boldsymbol{H}\right) \boldsymbol{u}_{l}.$$

When $\|H\|_2 < |\lambda_l|$, one can invert $I - \frac{1}{\lambda_l}H$ and use the assumption (5) to reach

$$egin{aligned} oldsymbol{u}_l &= \left(oldsymbol{I} - rac{1}{\lambda_l}oldsymbol{H}
ight)^{-1}rac{1}{\lambda_l}oldsymbol{M}^\staroldsymbol{u}_l \ &= rac{1}{\lambda_l}\left(oldsymbol{I} - rac{1}{\lambda_l}oldsymbol{H}
ight)^{-1}\left(\sum_{j=1}^r\lambda_j^\staroldsymbol{u}_j^\staroldsymbol{u}_j^{ op}
ight)oldsymbol{u}_l \end{aligned}$$

$$=\sum_{j=1}^{r}\frac{\lambda_{j}^{\star}}{\lambda_{l}}\left(\boldsymbol{u}_{j}^{\star\top}\boldsymbol{u}_{l}\right)\left(\boldsymbol{I}-\frac{1}{\lambda_{l}}\boldsymbol{H}\right)^{-1}\boldsymbol{u}_{j}^{\star},$$

where the last line follows by rearranging terms. Finally, replacing $\left(I - \frac{1}{\lambda_l}H\right)^{-1}$ with the Neumann series $\sum_{s=0}^{\infty} \frac{1}{\lambda_{i}^{s}} \boldsymbol{H}^{s}$, we establish the theorem.

A.3 Proof of Lemma 3

(1) We start with the rank-1 case. Towards this, we resort to the Neumann trick in Theorem 2, which in the rank-1 case reads

$$\boldsymbol{u} = \frac{1}{\lambda} \left(\boldsymbol{u}^{\star \top} \boldsymbol{u} \right) \sum_{s=0}^{\infty} \left(\frac{1}{\lambda} \boldsymbol{H} \right)^{s} \boldsymbol{u}^{\star}.$$
(57)

From Lemma 2, we know that λ is real-valued and that $\lambda > 1 - ||\mathbf{H}|| \ge 3/4 > ||\mathbf{H}||$ under our assumption. This together with (57) yields

$$\left\| \boldsymbol{u} - \frac{\boldsymbol{u}^{\star \top} \boldsymbol{u}}{\lambda} \boldsymbol{u}^{\star} \right\|_{2} \leq \frac{1}{\lambda} \sum_{s=1}^{\infty} \left\| \frac{1}{\lambda} \boldsymbol{H} \right\|^{s} \| \boldsymbol{u}^{\star} \|_{2} = \frac{1}{\lambda} \sum_{s=1}^{\infty} \left\| \frac{1}{\lambda} \boldsymbol{H} \right\|^{s} = \frac{\| \boldsymbol{H} \|}{\lambda \left(\lambda - \| \boldsymbol{H} \| \right)} \leq \frac{8}{3} \| \boldsymbol{H} \|,$$
(58)

where the last inequality holds since $\lambda \geq 3/4$ and $\lambda - \|\boldsymbol{H}\| \geq 1 - 2\|\boldsymbol{H}\| \geq 1/2$. Next, by decomposing \boldsymbol{u} into two orthogonal components $\boldsymbol{u} = (\boldsymbol{u}^{\star \top} \boldsymbol{u}) \boldsymbol{u}^{\star} + (\boldsymbol{u} - (\boldsymbol{u}^{\star \top} \boldsymbol{u}) \boldsymbol{u}^{\star})$, we obtain

$$\begin{aligned} \left| \boldsymbol{u}^{\star \top} \boldsymbol{u} \right| &= \left\| \left(\boldsymbol{u}^{\star \top} \boldsymbol{u} \right) \boldsymbol{u}^{\star} \right\|_{2} = \sqrt{1 - \left\| \boldsymbol{u} - \left(\boldsymbol{u}^{\star \top} \boldsymbol{u} \right) \boldsymbol{u}^{\star} \right\|_{2}^{2}} \\ &\geq 1 - \left\| \boldsymbol{u} - \left(\boldsymbol{u}^{\star \top} \boldsymbol{u} \right) \boldsymbol{u}^{\star} \right\|_{2}^{2} \\ &\geq 1 - \left\| \boldsymbol{u} - \frac{\boldsymbol{u}^{\star \top} \boldsymbol{u}}{\lambda} \boldsymbol{u}^{\star} \right\|_{2}^{2} \end{aligned}$$

$$(59)$$

$$\geq 1 - \frac{64}{9} \left\| \boldsymbol{H} \right\|^2.$$
(60)

The inequality (59) holds since $(\boldsymbol{u}^{\star\top}\boldsymbol{u}) \boldsymbol{u}^{\star}$ is orthogonal projection of \boldsymbol{u} onto the subspace spanned by \boldsymbol{u}^{\star} , and hence $\|\boldsymbol{u} - (\boldsymbol{u}^{\star\top}\boldsymbol{u}) \boldsymbol{u}^{\star}\|_2 \leq \|\boldsymbol{u} - \frac{1}{\lambda} (\boldsymbol{u}^{\star\top}\boldsymbol{u}) \boldsymbol{u}^{\star}\|_2$. Finally, (60) together with the fact that \boldsymbol{u} is real-valued (cf. Lemma 2) leads to the advertised bound:

$$\min\{\|\boldsymbol{u} - \boldsymbol{u}^{\star}\|_{2}, \|\boldsymbol{u} + \boldsymbol{u}^{\star}\|_{2}\} = \sqrt{\|\boldsymbol{u}\|_{2}^{2} + \|\boldsymbol{u}^{\star}\|_{2}^{2} - 2|\boldsymbol{u}^{\star\top}\boldsymbol{u}|} \leq \frac{8\sqrt{2}}{3}\|\boldsymbol{H}\|.$$

(2) For the rank-r case, it is seen that for any $1 \le l \le r$,

$$\sum_{j=1}^{r} |\boldsymbol{u}_{j}^{\star \top} \boldsymbol{u}_{l}|^{2} = \left\| \sum_{j=1}^{r} (\boldsymbol{u}_{j}^{\star \top} \boldsymbol{u}_{l}) \boldsymbol{u}_{j}^{\star} \right\|_{2}^{2} = 1 - \left\| \boldsymbol{u}_{l} - \sum_{j=1}^{r} (\boldsymbol{u}_{j}^{\star \top} \boldsymbol{u}_{l}) \boldsymbol{u}_{j}^{\star} \right\|_{2}^{2}$$
$$\geq 1 - \left\| \boldsymbol{u}_{l} - \sum_{j=1}^{r} \frac{\lambda_{j}^{\star} \boldsymbol{u}_{j}^{\star \top} \boldsymbol{u}_{l}}{\lambda_{l}} \boldsymbol{u}_{j}^{\star} \right\|_{2}^{2}, \tag{61}$$

where the inequality arises since $\sum_{j=1}^{r} (\boldsymbol{u}_{j}^{\star \top} \boldsymbol{u}_{l}) \boldsymbol{u}_{j}^{\star}$ is the Euclidean projection of \boldsymbol{u}_{l} onto the span of $\{\boldsymbol{u}_{1}^{\star}, \cdots, \boldsymbol{u}_{r}^{\star}\}$. In addition, observe that

$$\left\|\sum_{j=1}^{r}\lambda_{j}^{\star}(\boldsymbol{u}_{j}^{\star\top}\boldsymbol{u}_{l})\boldsymbol{u}_{j}^{\star}\right\|_{2} \leq \sqrt{\sum_{j}\left(\lambda_{j}^{\star}\right)^{2}\left|\boldsymbol{u}_{j}^{\star\top}\boldsymbol{u}_{l}\right|^{2}} \leq \lambda_{\max}^{\star}\sqrt{\sum_{j}\left|\boldsymbol{u}_{j}^{\star\top}\boldsymbol{u}_{l}\right|^{2}} = \lambda_{\max}^{\star}\|\boldsymbol{u}_{l}\|_{2} = 1.$$

This taken collectively with Theorem 2 leads to

$$\left\| oldsymbol{u}_l - \sum_{j=1}^r rac{\lambda_j^\star oldsymbol{u}_j^{\star op} oldsymbol{u}_l}{\lambda_l} oldsymbol{u}_j^\star
ight\|_2 = \left\| \sum_{j=1}^r rac{\lambda_j^\star}{\lambda_l} (oldsymbol{u}_j^{\star op} oldsymbol{u}_l) \left\{ \sum_{s=1}^\infty rac{1}{\lambda_l^s} oldsymbol{H}^s oldsymbol{u}_j^\star
ight\}
ight\|_2$$

$$= \left\| \frac{1}{\lambda_l} \sum_{s=1}^{\infty} \frac{1}{\lambda_l^s} \boldsymbol{H}^s \left\{ \sum_{j=1}^r \lambda_j^* (\boldsymbol{u}_j^{\star \top} \boldsymbol{u}_l) \boldsymbol{u}_j^{\star} \right\} \right\|_2$$

$$\leq \frac{1}{|\lambda_l|} |\sum_{s=1}^{\infty} \frac{1}{|\lambda_l|^s} \|\boldsymbol{H}\|^s \left\| \sum_{j=1}^r \lambda_j^* (\boldsymbol{u}_j^{\star \top} \boldsymbol{u}_l) \boldsymbol{u}_j^{\star} \right\|_2$$

$$\leq \frac{1}{|\lambda_l|} \cdot \frac{\|\boldsymbol{H}\|}{|\lambda_l| - \|\boldsymbol{H}\|}$$

$$\leq \frac{8\kappa^2}{3} \|\boldsymbol{H}\|.$$

The last two lines follow since, when $\|\boldsymbol{H}\| < 1/(4\kappa)$, one can invoke Lemma 2 to show (i) $|\lambda_l| > |\lambda_{\min}^{\star}| - \|\boldsymbol{H}\| \ge 3/(4\kappa)$ and (ii) $|\lambda_l| - \|\boldsymbol{H}\| \ge |\lambda_{\min}^{\star}| - 2\|\boldsymbol{H}\| \ge 1/(2\kappa)$. Substitution into (61) yields

$$(61) \ge 1 - \frac{64\kappa^4}{9} \|\boldsymbol{H}\|^2$$

as claimed.

B Proof for the lower bound in Lemma 4

Here, we establish Lemma 4, towards which we intend to invoke Fano's inequality. Let P_{ij} and \tilde{P}_{ij} represent the distributions of M_{ij} and \tilde{M}_{ij} , respectively, and denote by P, \tilde{P} , \hat{P} the distributions of M, \tilde{M} and \hat{M} , respectively. Then the KL divergence between P and \tilde{P} can be calculated as

$$\mathsf{KL}\big(\tilde{P} \,\|\, P\big) = \sum_{1 \le i,j \le n} \mathsf{KL}\big(\tilde{P}_{ij} \,\|\, P_{ij}\big) = \sum_{1 \le i,j \le n} \frac{(\Delta u_i^\star u_j^\star)^2}{2\sigma^2} = \frac{\Delta^2}{\sigma^2}.$$

where the first identity comes from the property of KL divergence for product measures, the second identity follows since $\mathsf{KL}(\mathcal{N}(\mu_1, \sigma^2) \parallel \mathcal{N}(\mu_2, \sigma^2)) = \frac{(\mu_1 - \mu_2)^2}{2\sigma^2}$, and the third one holds since $\Vert \boldsymbol{u}^\star \Vert_2 = 1$. The same argument yields $\mathsf{KL}(\hat{P} \parallel P) = \Delta^2/\sigma^2$. In view of (Tsybakov, 2009, Corollary 2.6), if

$$\alpha \log 2 \ge \frac{1}{2} \left\{ \mathsf{KL}\big(\tilde{P} \,\|\, P\big) + \mathsf{KL}\big(\hat{P} \,\|\, P\big) \right\} = \frac{\Delta^2}{\sigma^2}$$

then one has $p_{\rm e} \geq \frac{\log 3 - \log 2}{\log 2} - \alpha$. Taking $\alpha = \log_2 1.5 - \varepsilon$ immediately establishes the proof.

C Proof of Lemma 5

To establish this lemma, we exploit entrywise independence of H and develop a combinatorial trick.

To begin with, we expand the quantity of interest as

$$\left(\boldsymbol{a}^{\top}\boldsymbol{H}^{s}\boldsymbol{u}^{\star}\right)^{k} = \sum_{1 \leq i_{t}^{(b)} \leq n, \ 0 \leq t \leq s, \ 1 \leq b \leq k} \prod_{b=1}^{k} a_{i_{0}^{(b)}} \left(\prod_{t=1}^{s} H_{i_{t-1}^{(b)}i_{t}^{(b)}}\right) u_{i_{s}^{(b)}}^{\star}.$$

Throughout this section, we will use

$$\mathcal{I} := \left\{ i_t^{(b)} \mid 0 \le t \le s, 1 \le b \le k \right\} \in [n]^{(s+1)k}$$
(62)

to denote such a collection of (s+1)k indices. Thus, one can write

$$\mathbb{E}\left[\left(\boldsymbol{a}^{\top}\boldsymbol{H}^{s}\boldsymbol{u}^{\star}\right)^{k}\right] = \sum_{\mathcal{I}\in[n]^{(s+1)k}} \mathbb{E}\left[\prod_{b=1}^{k} a_{i_{0}^{(b)}}\left(\prod_{t=1}^{s} H_{i_{t-1}^{(b)}i_{t}^{(b)}}\right) u_{i_{s}^{(b)}}^{\star}\right].$$



Figure 6: (Left) An illustration of $\mathcal{I} := \{i_t^{(b)} \mid 0 \leq t \leq s, 1 \leq b \leq k\}$ and the corresponding edges $\{e_t^{(b)} \mid 1 \le t \le s, 1 \le b \le k\}$ when s = 4 and k = 2. (Right) An illustration of a type (or a partition of sk edges), where the edges of the same color belong to the same subset of the partition.

We shall often think of this sum-product graphically by viewing each index pair $(i_{t-1}^{(b)}, i_t^{(b)})$ as a directed edge over the vertex set $[n]^{(s+1)k}$. As a result, this gives us a set of sk edges in total $\{e_t^{(b)} \mid 1 \le t \le s, 1 \le b \le k\}$, where $e_t^{(b)}$ represents $(i_{t-1}^{(b)}, i_t^{(b)})$. See Fig. 6 for an illustration. In what follows, two edges $e_{t_1}^{(b_1)}$ and $e_{t_2}^{(b_2)}$ are said to be equivalent in \mathcal{I} , denoted by

$$e_{t_1}^{(b_1)} = e_{t_2}^{(b_2)}$$

if the values of the corresponding vertices are identical (namely, $i_{t_1-1}^{(b_1)} = i_{t_2-1}^{(b_2)}$ and $i_{t_1}^{(b_1)} = i_{t_2}^{(b_2)}$). When $(s+1)k \ll n$, most of the summands in the above expansion vanish. In fact, for any summand associated with a given \mathcal{I} : as long as there exists a distinct edge $(i_{t^*-1}^{(b^*)}, i_{t^*}^{(b^*)})$ (i.e. not equal to any other edge associated with \mathcal{I}), then the contribution of this summand is zero, namely,

$$\mathbb{E}\left[\prod_{b=1}^{k} a_{i_{0}^{(b)}}\left(\prod_{t=1}^{s} H_{i_{t-1}^{(b)}i_{t}^{(b)}}\right) u_{i_{s}^{(b)}}^{\star}\right] = \left(\prod_{b=1}^{k} a_{i_{0}^{(b)}}u_{i_{s}^{(b)}}^{\star}\right) \mathbb{E}\left[H_{i_{t^{*}-1}^{(b^{*})}i_{t^{*}}^{(b^{*})}}\right] \mathbb{E}\left[\prod_{(t,b)\in[s]\times[k], (t,b)\neq(t^{*},b^{*})} H_{i_{t-1}^{(b)}i_{t}^{(b)}}\right] = 0.$$
(63)

As a result, for any term with non-zero contribution, every edge $(i_{t-1}^{(b)}, i_t^{(b)})$ must appear at least twice.

To enable simple yet effective upper bounds on the non-vanishing terms, we group those terms of the same "type" and look at each "type" separately. Specifically:

- We represent a collection of edges $\{(i_{t-1}^{(b)}, i_t^{(b)}) \mid 1 \le t \le s, 1 \le b \le k\}$ as $\{e_t^{(b)} \mid 1 \le t \le s, 1 \le b \le k\}$. See Fig. 6 for an illustration.
- Each type encodes a set of constraints across edges, namely, which edges correspond to the same pair of vertices. To be precise, we define each type to be a partition of $\{e_t^{(b)} \mid 1 \le t \le s, 1 \le b \le k\}$ into a disjoint union of subsets, so that all edges falling within the same subset are equivalent. For instance, when s = k = 2, one possible type is $\{\{e_1^{(1)}, e_2^{(2)}\}, \{e_1^{(2)}, e_2^{(1)}\}\}$, which encodes the constraints $e_1^{(1)} = e_2^{(2)}$ and $e_1^{(2)} = e_2^{(1)}$.
- For each index collection \mathcal{I} (defined in (62)), we write type(\mathcal{I}) = \mathcal{T} if the associated edge set of \mathcal{I} satisfies the constraints encoded by a type \mathcal{T} .
- Define the set Γ^+ as

 Γ^+ : the set of all types s.t. each subset of the associated partition has cardinality at least 2. (64)

With this grouping strategy and (63) in mind, we can derive

$$\left| \mathbb{E} \left[\left(\boldsymbol{a}^{\top} \boldsymbol{H}^{s} \boldsymbol{u}^{\star} \right)^{k} \right] \right| = \left| \sum_{\mathcal{I} \in [n]^{(s+1)k}} \mathbb{E} \left[\prod_{b=1}^{k} a_{i_{0}^{(b)}} \left(\prod_{t=1}^{s} H_{i_{t-1}^{(b)} i_{t}^{(b)}} \right) u_{i_{s}^{(b)}}^{\star} \right] \right|$$

$$= \left| \sum_{\mathcal{T} \in \Gamma^+} \sum_{\mathcal{I} \in [n]^{(s+1)k}: \, \mathsf{type}(\mathcal{I}) = \mathcal{T}} \mathbb{E} \left[\prod_{b=1}^k a_{i_0^{(b)}} \left(\prod_{t=1}^s H_{i_{t-1}^{(b)}i_t^{(b)}} \right) u_{i_s^{(b)}}^{\star} \right] \right|, \tag{65}$$

where the last line follows since all types outside Γ^+ have vanishing contributions (cf. (63)). To bound the right-hand side of (65), we need the following lemma. Here and throughout, $|\mathcal{T}|$ represents the number of non-empty subsets in the partition associated with \mathcal{T} .

Lemma 6. For any fixed unit vector $\mathbf{a} \in \mathbb{R}^{n \times 1}$ and any type $\mathcal{T} \in \Gamma^+$, one has

$$\left|\sum_{\mathcal{I}\in[n]^{(s+1)k}:\,\mathsf{type}(\mathcal{I})=\mathcal{T}}\mathbb{E}\left[\prod_{b=1}^{k}a_{i_{0}^{(b)}}\left(\prod_{t=1}^{s}H_{i_{t-1}^{(b)}i_{t}^{(b)}}\right)u_{i_{s}^{(b)}}^{\star}\right]\right| \leq B^{sk}\left(\frac{\sigma^{2}}{B^{2}}\right)^{|\mathcal{T}|}\left(\sqrt{\frac{\mu}{n}}\right)^{k}n^{|\mathcal{T}|}.$$
(66)

Armed with this lemma, we can further obtain

$$(65) \leq \sum_{\mathcal{T}\in\Gamma^{+}} B^{sk} \left(\frac{\sigma^{2}}{B^{2}}\right)^{|\mathcal{T}|} \left(\sqrt{\frac{\mu}{n}}\right)^{k} n^{|\mathcal{T}|} = \sum_{\mathcal{T}\in\Gamma^{+}} B^{sk-2|\mathcal{T}|} \left(n\sigma^{2}\right)^{|\mathcal{T}|} \left(\sqrt{\frac{\mu}{n}}\right)^{k}$$
$$\leq \sum_{\mathcal{T}\in\Gamma^{+}} \sum_{l=1}^{sk/2} \mathbb{1}_{\{|\mathcal{T}|=l\}} B^{sk-2|\mathcal{T}|} \left(n\sigma^{2}\right)^{|\mathcal{T}|} \left(\sqrt{\frac{\mu}{n}}\right)^{k}$$
$$= \sum_{l=1}^{sk/2} \left\{\sum_{\mathcal{T}\in\Gamma^{+}} \mathbb{1}_{\{|\mathcal{T}|=l\}}\right\} B^{sk-2l} \left(n\sigma^{2}\right)^{l} \left(\sqrt{\frac{\mu}{n}}\right)^{k}, \tag{67}$$

where we have grouped the types based on their cardinality. The last inequality results from $|\mathcal{T}| \leq sk/2$ in Lemma 9. The following lemma bounds the number of distinct types having the same cardinality.

Lemma 7.
$$\sum_{\mathcal{T}\in\Gamma^+} \mathbb{1}_{\{|\mathcal{T}|=l\}} \leq {\binom{sk}{l}} l^{sk-l} \leq 2^{sk} l^{sk-l}.$$

Proof. The quantity of interest is the number of ways to partition sk edges into l disjoint subsets, where each subset contains at least 2 edges. To bound this quantity, we first pick l edges and assign each of them to a distinct subset; there are $\binom{sk}{l}$ different ways to achieve it. We still have sk - l edges left, and the number of ways to assign them to l subsets is clearly upper bounded by l^{sk-l} . This concludes the proof. \Box

Consequently, Lemma 7 yields

$$(67) \leq 2^{sk} \sum_{l=1}^{sk/2} (Bl)^{sk-2l} (n\sigma^2 l)^l \left(\sqrt{\frac{\mu}{n}}\right)^k$$

$$\stackrel{(i)}{\leq} 2^{sk} \sum_{l=1}^{sk/2} \left(\frac{Bsk}{2}\right)^{sk-2l} \left(\frac{n\sigma^2 sk}{2}\right)^l \left(\sqrt{\frac{\mu}{n}}\right)^k$$

$$\stackrel{(ii)}{\leq} 2^{sk} \sum_{l=1}^{sk/2} \max\left\{\left(\frac{Bsk}{2}\right)^{sk}, \left(\frac{n\sigma^2 sk}{2}\right)^{sk/2}\right\} \left(\sqrt{\frac{\mu}{n}}\right)^k$$

$$= \frac{sk}{2} \max\left\{(Bsk)^{sk}, (2n\sigma^2 sk)^{sk/2}\right\} \left(\sqrt{\frac{\mu}{n}}\right)^k,$$

where (i) uses the condition $l \leq sk/2$, and (ii) relies on the fact that $a^{sk-2l}b^l = a^{sk-2l}(\sqrt{b})^{2l} \leq \max\{a^{sk}, (\sqrt{b})^{sk}\}$ for any a, b > 0.

Taken collectively, the preceding bounds conclude the proof of Lemma 5, provided that Lemma 6 can be established.

D Proof of Lemma 6

To begin with,

$$\left| \mathbb{E} \left[\prod_{b=1}^{k} a_{i_{0}^{(b)}} \left(\prod_{t=1}^{s} H_{i_{t-1}^{(b)} i_{t}^{(b)}} \right) u_{i_{s}^{(b)}}^{\star} \right] \right| \leq \mathbb{E} \left[\prod_{b=1}^{k} \left| a_{i_{0}^{(b)}} \right| \left(\prod_{t=1}^{s} \left| H_{i_{t-1}^{(b)} i_{t}^{(b)}} \right| \right) \left| u_{i_{s}^{(b)}}^{\star} \right| \right] \\ \leq \left(\sqrt{\frac{\mu}{n}} \right)^{k} \left(\prod_{b=1}^{k} \left| a_{i_{0}^{(b)}} \right| \right) \mathbb{E} \left[\prod_{b=1}^{k} \prod_{t=1}^{s} \left| H_{i_{t-1}^{(b)} i_{t}^{(b)}} \right| \right], \tag{68}$$

where the last inequality holds since each entry of u^{\star} is bounded in magnitude by $\sqrt{\mu/n}$ (see Definition 1).

According to the definition of Γ^+ (see (64)), for any type $\mathcal{T} \in \Gamma^+$, each edge $(i_{t-1}^{(b)}, i_t^{(b)})$ is repeated at least twice. The total number of distinct edges is exactly $|\mathcal{T}|$. For notational simplicity, suppose the distinct edges are $e_1, \dots, e_{|\mathcal{T}|}$, where e_i is repeated by $l_i \geq 2$ times. Then we can write

$$\mathbb{E}\left[\prod_{b=1}^{k}\prod_{t=1}^{s}|H_{i_{t-1}^{(b)}i_{t}^{(b)}}|\right] = \prod_{i=1}^{|\mathcal{T}|}\mathbb{E}\left[|H_{\boldsymbol{e}_{i}}|^{l_{i}}\right] \leq \prod_{i=1}^{|\mathcal{T}|}\sigma^{2}B^{l_{i}-2} = \left(\frac{\sigma^{2}}{B^{2}}\right)^{|\mathcal{T}|}B^{\sum_{i}l_{i}} = \left(\frac{\sigma^{2}}{B^{2}}\right)^{|\mathcal{T}|}B^{sk},$$

where the inequality follows since

$$\mathbb{E}\left[\left|H_{ij}\right|^{l}\right] \leq B^{l-2}\mathbb{E}\left[\left|H_{ij}\right|^{2}\right] \leq \sigma^{2}B^{l-2}, \qquad l \geq 2.$$

Substitution into (68) yields

$$\left| \sum_{\mathcal{I} \in [n]^{(s+1)k}: \, \mathsf{type}(\mathcal{I}) = \mathcal{T}} \mathbb{E} \left[\prod_{b=1}^{k} a_{i_{0}^{(b)}} \left(\prod_{t=1}^{s} H_{i_{t-1}^{(b)} i_{t}^{(b)}} \right) u_{i_{s}^{(b)}}^{\star} \right] \right| \\ \leq \left(\sqrt{\frac{\mu}{n}} \right)^{k} \left(\frac{\sigma^{2}}{B^{2}} \right)^{|\mathcal{T}|} B^{sk} \left\{ \sum_{\mathcal{I} \in [n]^{(s+1)k}: \, \mathsf{type}(\mathcal{I}) = \mathcal{T}} \prod_{b=1}^{k} \left| a_{i_{0}^{(b)}} \right| \right\}.$$

$$(69)$$

It then boils down to controlling $\sum_{\mathcal{I} \in [n]^{(s+1)k}: \text{type}(\mathcal{I}) = \mathcal{T}} \prod_{b=1}^{k} |a_{i_0^{(b)}}|$. This is achieved by the following lemma, which we establish in Appendix E.

Lemma 8. For any type $\mathcal{T} \in \Gamma^+$ and any unit vector $\boldsymbol{a} \in \mathbb{R}^n$, one has

$$\sum_{\mathcal{I} \in [n]^{(s+1)k}: \, \mathsf{type}(\mathcal{I}) = \mathcal{T}} \prod_{b=1}^{k} \left| a_{i_{0}^{(b)}} \right| \le n^{|\mathcal{T}|}.$$
(70)

This lemma together with (69) concludes the proof.

E Proof of Lemma 8

E.1 Graphical tools

The proof of Lemma 8 is combinatorial in nature. Before proceeding, we introduce a few graphical notions that will be useful.

To begin with, recall that the index collection we have used so far is

$$\mathcal{I} := \left\{ i_t^{(b)} \mid 0 \le t \le s, 1 \le b \le k \right\} \in [n]^{(s+1)k},$$
(71)

which involves (s + 1)k vertices. As it turns out, it is helpful to further augment it into 2(s + 1)k vertices via simple duplication. Specifically, introduce the following set of 2(s + 1)k vertices

$$\mathcal{V} = \left\{ q_t^{(b)} \mid 0 \le t \le s, 1 \le b \le k \right\} \bigcup \left\{ p_t^{(b)} \mid 1 \le t \le s+1, 1 \le b \le k \right\}.$$
(72)

Throughout this paper, we shall



Figure 7: An example of the induced graph, where s = 4, k = 2. The left plot shows the partition, while the right plot displays the induced graph and the vertex set Q_0 (the yellow region).

• view $\{q_t^{(b)} \mid 0 \le t \le s, 1 \le b \le k\}$ as a copy of $\{i_t^{(b)} \mid 0 \le t \le s, 1 \le b \le k\};$

• view
$$\{p_t^{(b)} \mid 1 \le t \le s+1, 1 \le b \le k\}$$
 as another copy of $\{i_{t-1}^{(b)} \mid 1 \le t \le s+1, 1 \le b \le k\}$.

The main incentive is that it allows us to reparametrize

$$\prod_{b=1}^{k} a_{i_{0}^{(b)}} \left(\prod_{t=1}^{s} H_{i_{t-1}^{(b)} i_{t}^{(b)}} \right) u_{i_{s}^{(b)}}^{\star} \longrightarrow \prod_{b=1}^{k} a_{q_{0}^{(b)}} \left(\prod_{t=1}^{s} H_{p_{t}^{(b)} q_{t}^{(b)}} \right) u_{p_{s+1}^{\star}}^{\star}$$
(73)

Recall that we have categorized \mathcal{I} into different "types". Given that \mathcal{V} is simply a redundant representation of \mathcal{I} , we can also associate a type with each \mathcal{V} by enforcing proper constraints. These constraints can be encoded via the following graphical notion.

Definition 2 (Induced graph). For any given type \mathcal{T} , we define the induced graph $\mathcal{G}(\mathcal{T})$ with the vertex set \mathcal{V} and the (undirected) edge set $\mathcal{E}(\mathcal{T})$ as $\mathcal{E}(\mathcal{T}) = \mathcal{E}_1 \cup \mathcal{E}_2(\mathcal{T})$, where

$$\mathcal{E}_{1} \triangleq \left\{ \left(p_{t}^{(b)}, q_{t-1}^{(b)} \right) \mid 1 \leq t \leq s+1, 1 \leq b \leq k \right\};$$

$$\mathcal{E}_{2}(\mathcal{T}) \triangleq \left\{ \left(p_{t_{1}}^{(b_{1})}, p_{t_{2}}^{(b_{2})} \right), \left(q_{t_{1}}^{(b_{1})}, q_{t_{2}}^{(b_{2})} \right) \mid \forall t_{1}, b_{1}, t_{2}, b_{2} : \left(i_{t_{1}-1}^{(b_{1})}, i_{t_{1}}^{(b_{1})} \right) = \left(i_{t_{2}-1}^{(b_{2})}, i_{t_{2}}^{(b_{2})} \right) \right\}.$$
(74)

In words,

- (1) we connect the vertices $p_t^{(b)}$ (which is a copy of $i_t^{(b)}$) and $q_{t-1}^{(b)}$ (which represents $i_{t-1}^{(r)}$) by an edge;
- (2) whenever two edges $(i_{t_1-1}^{(b_1)}, i_{t_1}^{(b_1)})$ and $(i_{t_2-1}^{(b_2)}, i_{t_2}^{(b_2)})$ are equivalent in \mathcal{I} (or \mathcal{T}), we draw two edges connecting the corresponding vertices in the induced graph.

As an illustration, Fig. 7 displays the induced graph for a simple example with s = 4 and k = 2, where

$$\mathcal{T} = \left\{ \{\underbrace{(i_0^{(1)}, i_1^{(1)})}_{\boldsymbol{e}_1^{(1)}}, \underbrace{(i_0^{(2)}, i_1^{(2)})}_{\boldsymbol{e}_1^{(2)}}, \{\underbrace{(i_1^{(1)}, i_2^{(1)})}_{\boldsymbol{e}_2^{(1)}}, \underbrace{(i_3^{(1)}, i_4^{(1)})}_{\boldsymbol{e}_4^{(1)}}, \{\underbrace{(i_1^{(2)}, i_2^{(2)})}_{\boldsymbol{e}_2^{(2)}}, \underbrace{(i_3^{(2)}, i_4^{(2)})}_{\boldsymbol{e}_4^{(2)}}, \{\underbrace{(i_2^{(1)}, i_3^{(1)})}_{\boldsymbol{e}_3^{(1)}}, \{\underbrace{(i_2^{(2)}, i_3^{(2)})}_{\boldsymbol{e}_3^{(2)}}, \underbrace{(i_2^{(2)}, i_3^{(2)})}_{\boldsymbol{e}_3^{(2)}}, \underbrace{(i_2^{(2)},$$

Here, the 45-degree lines correspond to the edges in \mathcal{E}_1 , and the remaining edges come from $\mathcal{E}_2(\mathcal{T})$. Throughout the rest of the paper, we will abuse the notation and write $type(\mathcal{V}) = \mathcal{T}$ if \mathcal{V} is induced by an index set of type \mathcal{T} .

One useful feature of the above induced graph is that: each edge subset in the partition associated with \mathcal{T} corresponds to a connected component in $\mathcal{G}(\mathcal{T})$. For our purpose, it is convenient to divide all connected components of $\mathcal{G}(\mathcal{T})$ into two classes. To this end, we first define $\mathcal{Q}_0 \triangleq \{q_0^{(b)} \mid 1 \leq b \leq k\}$ (as illustrated in Fig. 7).

- Class 1: a connected component \mathcal{C} belongs to Class 1 if $|\mathcal{C} \cap \mathcal{Q}_0| \leq 1$;
- Class 2: a connected component C belongs to Class 2 if $|C \cap Q_0| \ge 2$.

We denote by $m_1(\mathcal{T})$ (resp. $m_2(\mathcal{T})$) the total number of Class 1 (resp. 2) connected components in $\mathcal{G}(\mathcal{T})$.

E.2 Proof of Lemma 8

Making use of the augmented set \mathcal{V} (defined in (72)) and the way of reparametrization (73), we can write

$$\sum_{\mathcal{I} \in [n]^{(s+1)k}: \, \mathsf{type}(\mathcal{I}) = \mathcal{T}} \mathbb{E}\left[\prod_{b=1}^{k} \left| a_{i_{0}^{(b)}} \right|\right] \leq \sum_{\mathcal{V} \in [n]^{2(s+1)k}: \, \mathsf{type}(\mathcal{V}) = \mathcal{T}} \mathbb{E}\left[\prod_{b=1}^{k} \left| a_{q_{0}^{(b)}} \right|\right].$$
(75)

As we shall see, the benefit of this bound is to allow us sum over the connected components of $\mathcal{G}(\mathcal{T})$ in a separate manner, since all indices in the same connected component must be identical. Specifically,

- Let $\mathcal{X}_1, \mathcal{X}_2, \cdots$ (resp. $\mathcal{Y}_1, \mathcal{Y}_2, \cdots$) denote the collection of Class 1 (resp. Class 2) connected components in $\mathcal{G}(\mathcal{T})$;
- Denote by x_i (resp. y_i) the value assigned to all indices in \mathcal{X}_i (resp. \mathcal{Y}_i);

With these notations in mind, we can decompose

$$\begin{aligned} (75) &\leq \sum_{x_1=1}^n \sum_{x_2=1}^n \cdots \sum_{y_1=1}^n \sum_{y_2=1}^n \cdots \mathbb{E}\left[\left| a_{x_1} \right|^{|\mathcal{X}_1 \cap \mathcal{Q}_0|} \cdot \left| a_{x_2} \right|^{|\mathcal{X}_2 \cap \mathcal{Q}_0|} \cdots \left| a_{y_1} \right|^{|\mathcal{Y}_1 \cap \mathcal{Q}_0|} \cdot \left| a_{y_2} \right|^{|\mathcal{Y}_2 \cap \mathcal{Q}_0|} \cdots \right] \\ &= \mathbb{E}\left[\left(\sum_{x_1=1}^n \left| a_{x_1} \right|^{|\mathcal{X}_1 \cap \mathcal{Q}_0|} \right) \left(\sum_{x_2=1}^n \left| a_{x_2} \right|^{|\mathcal{X}_2 \cap \mathcal{Q}_0|} \right) \cdots \left(\sum_{y_1=1}^n \left| a_{y_1} \right|^{|\mathcal{Y}_1 \cap \mathcal{Q}_0|} \right) \left(\sum_{y_2=1}^n \left| a_{y_2} \right|^{|\mathcal{Y}_2 \cap \mathcal{Q}_0|} \right) \cdots \right] \\ &\stackrel{(i)}{\leq} \mathbb{E}\left[\left(\sum_{x_1=1}^n 1 \right) \left(\sum_{x_2=1}^n 1 \right) \cdots \left(\sum_{y_1=1}^n \left| a_{y_1} \right|^2 \right) \left(\sum_{y_2=1}^n \left| a_{y_2} \right|^2 \right) \cdots \right] \\ &\stackrel{(ii)}{=} \left(\sum_{x_1=1}^n 1 \right) \left(\sum_{x_2=1}^n 1 \right) \cdots 1 \cdot 1 \cdots = n^{m_1(\mathcal{T})}, \end{aligned}$$

where $m_1(\mathcal{T})$ is the total number of Class 1 connected components in the induced graph $\mathcal{G}(\mathcal{T})$. Here, (i) comes from the definitions of Class 1 and Class 2 connected components, and (ii) follows since $\|\boldsymbol{a}\|_2^2 = 1$. We can thus finish the proof by observing that $m_1(\mathcal{T}) \leq |\mathcal{T}|$, as claimed in the following lemma.

Lemma 9. For any type $\mathcal{T} \in \Gamma^+$, one has $m_1(\mathcal{T}) \leq |\mathcal{T}| \leq sk/2$.

E.3 Proof of Lemma 9

Introduce two disjoint vertex sets:

$$\mathcal{Q}_{\setminus 0} \triangleq \left\{ q_t^{(b)} \mid 1 \le t \le s, 1 \le b \le k \right\} \subset \mathcal{V}; \tag{76}$$

$$\mathcal{Q}_0 \triangleq \left\{ q_0^{(b)} \mid 1 \le b \le k \right\} \subset \mathcal{V}.$$
(77)

Clearly, one can define a partition of $\mathcal{Q}_{\backslash 0}$ — denoted by $\mathcal{T}_{\mathcal{Q}_{\backslash 0}}$ — induced by \mathcal{T} . Specifically, we say that $q_{t_1}^{(b_1)}$ and $q_{t_2}^{(b_2)}$ belong to the same connected subgraph of $\mathcal{Q}_{\backslash 0}$ if $(i_{t_1-1}^{(b_1)}, i_{t_1}^{(b_1)}) = (i_{t_2-1}^{(b_2)}, i_{t_2}^{(b_2)})$ in \mathcal{T} . Clearly, $|\mathcal{T}| = |\mathcal{T}_{\mathcal{Q}_{\backslash 0}}|$.

In addition, for any connected subgraph C_q in $\mathcal{Q}_{\setminus 0}$, we denote by C the corresponding connected component in $\mathcal{G}(\mathcal{T})$. We can thus define a mapping ψ that maps C_q to C, whose domain is the set of all connected subgraphs in $\mathcal{Q}_{\setminus 0}$. We claim that the collection of Class 1 connected components — denoted by $\{\mathcal{X}_j\}$ obeys

$$\{\mathcal{X}_j \mid 1 \le j \le m_1(\mathcal{T})\} \subseteq \operatorname{Im}\psi,\tag{78}$$

where $\text{Im}\psi$ denotes the image of ψ , and $m_1(\mathcal{T})$ is the total number of Class 1 connected components. If this holds true, then

$$\mathcal{T}| = \left|\mathcal{T}_{\mathcal{Q}_{\setminus 0}}\right| \ge \left|\mathrm{Im}\psi\right| \ge \left|\left\{\mathcal{X}_{j} \mid 1 \le j \le m_{1}(\mathcal{T})\right\}\right| = m_{1}(\mathcal{T}).$$

To justify the claim (78), it suffices to show that $\mathcal{X}_j \cap \mathcal{Q}_{\setminus 0} \neq \emptyset$. Given that $\mathcal{T} \in \Gamma^+$ (so that each subset of the partition associated with \mathcal{T} has cardinality at least 2) and that \mathcal{X}_j is defined over the induced graph $\mathcal{G}(\mathcal{T})$ (which is a redundant representation of the original index collection), one must have

$$\max\left\{ \left| \mathcal{X}_{j} \cap \left\{ p_{t}^{(b)} \mid 1 \le t \le s+1, 1 \le b \le k \right\} \right|, \left| \mathcal{X}_{j} \cap \left\{ q_{t}^{(b)} \mid 0 \le t \le s, 1 \le b \le k \right\} \right| \right\} \ge 2.$$
(79)

Further, from the construction of the induced graph, it is easily seen that

$$\left| \mathcal{X}_{j} \cap \left\{ p_{t}^{(b)} \mid 1 \le t \le s+1, 1 \le b \le k \right\} \right| = \left| \mathcal{X}_{j} \cap \left\{ q_{t}^{(b)} \mid 0 \le t \le s, 1 \le b \le k \right\} \right|,$$

thus resulting in

$$\left|\mathcal{X}_{j} \cap \left\{q_{t}^{(b)} \mid 0 \leq t \leq s, 1 \leq b \leq k\right\}\right| = \left|\mathcal{X}_{j} \cap \left(\mathcal{Q}_{0} \cup \mathcal{Q}_{\setminus 0}\right)\right| \geq 2.$$

However, by definition of Class 1 connected component, we have $|\mathcal{X}_j \cap \mathcal{Q}_0| \leq 1$, implying that $\mathcal{X}_j \cap \mathcal{Q}_{\setminus 0} \neq \emptyset$. This establishes the claim (78).

Finally, when $\mathcal{T} \in \Gamma^+$, each connected subgraph of $\mathcal{T}_{\mathcal{Q}_{\setminus 0}}$ also contains at least two nodes. As a consequence,

$$|\mathcal{T}| = |\mathcal{T}_{\mathcal{Q}_{\setminus 0}}| \le |\mathcal{Q}_{\setminus 0}|/2 = sk/2.$$

F Proof of Corollary 4

In the sequel, we assume that $20 \log n$ is an integer to avoid the clumsy notation $\lfloor 20 \log n \rfloor$. But it is straightforward to extend it to the case where $20 \log n$ is not an integer.

It follows from Markov's inequality that for any even integer k,

$$\mathbb{P}\left(\left|\boldsymbol{a}^{\top}\boldsymbol{H}^{s}\boldsymbol{u}^{\star}\right| \geq \tau\right) \leq \frac{\mathbb{E}\left[\left|\boldsymbol{a}^{\top}\boldsymbol{H}^{s}\boldsymbol{u}^{\star}\right|^{k}\right]}{\tau^{k}}.$$

This combined with Lemma 5 gives

$$\mathbb{P}\left(\left|\boldsymbol{a}^{\top}\boldsymbol{H}^{s}\boldsymbol{u}^{\star}\right| \leq \left(c_{2}\max\left\{B\log n, \sqrt{n\sigma^{2}\log n}\right\}\right)^{s}\sqrt{\frac{\mu}{n}}\right) \leq \frac{\left|\mathbb{E}\left[\left(\boldsymbol{a}^{\top}\boldsymbol{H}^{s}\boldsymbol{u}^{\star}\right)^{k}\right]\right|}{\left(c_{2}\max\left\{B\log n, \sqrt{n\sigma^{2}\log n}\right\}\right)^{sk}\left(\sqrt{\frac{\mu}{n}}\right)^{k}} \\ \leq \frac{sk}{2}\max\left\{\left(\frac{sk}{c_{2}\log n}\right)^{sk}, \left(\frac{2sk}{c_{2}^{2}\log n}\right)^{sk/2}\right\}.$$

For any $s \leq 20 \log n$, choose k such that $sk \in [20 \log n, 40 \log n]$. It is straightforward to show — using the union bound — that with probability at least $1 - O(n^{-10})$,

$$|\boldsymbol{a}^{\top}\boldsymbol{H}^{s}\boldsymbol{u}^{\star}| \leq \left(c_{2}\max\left\{B\log n, \sqrt{n\sigma^{2}\log n}\right\}\right)^{s}\sqrt{\mu/n}, \quad \forall s \leq 20\log n,$$

as long as $c_2 > 0$ is some sufficiently large constant.

G Proof of Corollary 5

Taking $\boldsymbol{a} = \boldsymbol{u}_i^{\star}$ for some $1 \leq i \leq r$ in (35), we see that with high probability,

$$\left| \boldsymbol{u}_{l}^{\star \top} \left(\boldsymbol{u}_{l} - \sum_{j=1}^{r} \frac{\lambda_{j}^{\star} \boldsymbol{u}_{j}^{\star \top} \boldsymbol{u}_{l}}{\lambda_{l}} \boldsymbol{u}_{j}^{\star} \right) \right| \lesssim \max \left\{ B \log n, \sqrt{n\sigma^{2} \log n} \right\} \frac{\kappa}{|\lambda_{l}|} \sqrt{\frac{\mu r}{n}}$$

In addition, since $u_1^{\star}, \dots, u_r^{\star}$ are orthogonal to each other, we have that: for any $1 \leq i \leq r$,

$$\left| \boldsymbol{u}_{i}^{\star \top} \left(\boldsymbol{u}_{l} - \sum_{j=1}^{r} \frac{\lambda_{j}^{\star} \boldsymbol{u}_{j}^{\star \top} \boldsymbol{u}_{l}}{\lambda_{l}} \boldsymbol{u}_{j}^{\star} \right) \right| = \left| \boldsymbol{u}_{i}^{\star \top} \boldsymbol{u}_{l} - \frac{\lambda_{i}^{\star} \boldsymbol{u}_{i}^{\star \top} \boldsymbol{u}_{l}}{\lambda_{l}} \right| = \left| \boldsymbol{u}_{i}^{\star \top} \boldsymbol{u}_{l} \right| \left| \frac{\lambda_{l} - \lambda_{i}^{\star}}{\lambda_{l}} \right| \ge \left| \boldsymbol{u}_{i}^{\star \top} \boldsymbol{u}_{l} \right| \min_{1 \le j \le r} \left| \frac{\lambda_{l} - \lambda_{j}^{\star}}{\lambda_{l}} \right|$$

Therefore, combining the above two bounds yields

$$\min_{1 \le j \le r} \left| \frac{\lambda_l - \lambda_j^*}{\lambda_l} \right| \max_{1 \le i \le r} \left| \boldsymbol{u}_i^{\star \top} \boldsymbol{u}_l \right| \lesssim \max\left\{ B \log n, \sqrt{n\sigma^2 \log n} \right\} \frac{\kappa}{|\lambda_l|} \sqrt{\frac{\mu r}{n}}.$$
(80)

Finally, it comes from Lemma 3 that if $\|\boldsymbol{H}\| \ll 1/\kappa^2$, then $\sum_{1 \le i \le r} |\boldsymbol{u}_i^{\star \top} \boldsymbol{u}_l|^2 \gtrsim 1$, and hence

$$\max_{1 \le j \le r} \left| \boldsymbol{u}_i^{\star \top} \boldsymbol{u}_l \right| \ge \sqrt{\frac{1}{r} \sum_{1 \le i \le r} \left| \boldsymbol{u}_i^{\star \top} \boldsymbol{u}_l \right|^2} \gtrsim 1/\sqrt{r}.$$
(81)

This combined with (80) establishes the claim, as long as the spectral norm condition on $||\mathbf{H}||$ can be guaranteed. In view of Lemma 1, we have $||\mathbf{H}|| \ll 1/\kappa^2$ under the condition (38), thus concluding the proof.

H Proof of Corollary 6

Since $\|\boldsymbol{a}\|_2 = 1$, it follows immediately from (36) that

$$\begin{aligned} \left\| \boldsymbol{a}^{\top} \boldsymbol{U} \right\|_{2} &= \sqrt{\sum_{l=1}^{r} \left| \boldsymbol{a}^{\top} \boldsymbol{u}_{l} \right|^{2}} \\ &\lesssim \sqrt{\sum_{l=1}^{r} \left(\left| \sum_{j=1}^{r} \frac{\lambda_{j}^{\star} \boldsymbol{u}_{j}^{\star \top} \boldsymbol{u}_{l}}{\lambda_{l}} \boldsymbol{a}^{\top} \boldsymbol{u}_{j}^{\star} \right| + \max\left\{ B \log n, \sqrt{n\sigma^{2} \log n} \right\} \kappa^{2} \sqrt{\frac{\mu r}{n}} \right)^{2}}. \end{aligned}$$
(82)

Further, we have

$$(82) \stackrel{(i)}{\leq} \sqrt{\sum_{l=1}^{r} \left| \sum_{j=1}^{r} \frac{\lambda_{j}^{\star} \boldsymbol{u}_{j}^{\star^{\top}} \boldsymbol{u}_{l}}{\lambda_{l}} \boldsymbol{a}^{\top} \boldsymbol{u}_{j}^{\star} \right|^{2}} + \sqrt{\sum_{l=1}^{r} \left(\max\left\{ B \log n, \sqrt{n\sigma^{2} \log n} \right\} \kappa^{2} \sqrt{\frac{\mu r}{n}} \right)^{2}}$$

$$\stackrel{(ii)}{\leq} \sqrt{\sum_{l=1}^{r} \left(\sum_{j=1}^{r} \left| \frac{\lambda_{j}^{\star} \boldsymbol{u}_{j}^{\star^{\top}} \boldsymbol{u}_{l}}{\lambda_{l}} \right|^{2} \right) \left(\sum_{j=1}^{r} \left| \boldsymbol{a}^{\top} \boldsymbol{u}_{j}^{\star} \right|^{2} \right)} + \max\left\{ B \log n, \sqrt{n\sigma^{2} \log n} \right\} \kappa^{2} r \sqrt{\frac{\mu}{n}}$$

$$\stackrel{(iii)}{\lesssim} \sqrt{r\kappa^{2} \left(\sum_{j=1}^{r} \left| \boldsymbol{a}^{\top} \boldsymbol{u}_{j}^{\star} \right|^{2} \right)} + \max\left\{ B \log n, \sqrt{n\sigma^{2} \log n} \right\} \kappa^{2} r \sqrt{\frac{\mu}{n}}$$

$$= \kappa \sqrt{r} \left\| \boldsymbol{a}^{\top} \boldsymbol{U}^{\star} \right\|_{2} + \max\left\{ B \log n, \sqrt{n\sigma^{2} \log n} \right\} \kappa^{2} r \sqrt{\frac{\mu}{n}}, \qquad (83)$$

where (i) and (ii) make use of the Minkowski and the Cauchy-Schwarz inequalities, respectively, and (iii) result from the facts $|\lambda_j^*/\lambda_l| \lesssim \kappa$ (which follows from (39) and (38)) and $\sum_{j=1}^r |\boldsymbol{u}_j^* \top \boldsymbol{u}_l|^2 \leq 1$. This concludes the proof.

I Proof of Theorem 5

To simplify presentation, we introduce the following notation throughout this section:

$$\widetilde{u}_1^{\star} = rac{1}{\sqrt{2}} \left(egin{array}{c} m{u}^{\star} \\ m{v}^{\star} \end{array}
ight), \qquad \widetilde{u}_2^{\star} = rac{1}{\sqrt{2}} \left(egin{array}{c} m{u}^{\star} \\ -m{v}^{\star} \end{array}
ight), \qquad \widetilde{u}_1 = m{u}_1^{ ext{dilation}}, \qquad \widetilde{u}_2 = m{u}_2^{ ext{dilation}};$$

$$\widetilde{\lambda}_1^{\star} = \lambda_1(\boldsymbol{M}_{\text{dilation}}^{\star}) = 1, \qquad \widetilde{\lambda}_2^{\star} = \lambda_2(\boldsymbol{M}_{\text{dilation}}^{\star}) = -1, \qquad \widetilde{\lambda}_1 = \lambda_1^{\text{dilation}}, \qquad \widetilde{\lambda}_2 = \lambda_2^{\text{dilation}}.$$

In addition, we denote $\widetilde{\boldsymbol{u}}_{1,1} = \boldsymbol{u}_{1,1}^{\text{dilation}}$ and $\widetilde{\boldsymbol{u}}_{1,2} = \boldsymbol{u}_{1,2}^{\text{dilation}}$. Recall that we assume $\widetilde{\lambda}_1 \geq \widetilde{\lambda}_2$. We also denote $\min\{\|\boldsymbol{x} \pm \boldsymbol{y}\|_2\} = \min\{\|\boldsymbol{x} - \boldsymbol{y}\|_2, \|\boldsymbol{x} + \boldsymbol{y}\|_2\}.$

I.1 ℓ_2 eigenvector perturbation bounds

The ℓ_2 perturbation bound results from near orthogonality between \tilde{u}_1 and \tilde{u}_2^* . By symmetry, it suffices to establish the result for u.

To begin with, applying Theorem 4 on $M_{\sf dilation}$ and taking $a = \widetilde{u}_2^{\star}$, we derive

$$\left| \widetilde{\boldsymbol{u}}_{2}^{\star \top} \widetilde{\boldsymbol{u}}_{1} \left(1 - \frac{\widetilde{\lambda}_{2}^{\star}}{\widetilde{\lambda}_{1}} \right) \right| = \left| \widetilde{\boldsymbol{u}}_{2}^{\star \top} \left(\widetilde{\boldsymbol{u}}_{1} - \frac{\widetilde{\lambda}_{1}^{\star} \widetilde{\boldsymbol{u}}_{1}^{\star \top} \widetilde{\boldsymbol{u}}_{1}}{\widetilde{\lambda}_{1}} \widetilde{\boldsymbol{u}}_{1}^{\star} - \frac{\widetilde{\lambda}_{2}^{\star} \widetilde{\boldsymbol{u}}_{2}^{\star \top} \widetilde{\boldsymbol{u}}_{1}}{\widetilde{\lambda}_{1}} \widetilde{\boldsymbol{u}}_{2}^{\star} \right) \\ \lesssim \max \left\{ B \log n, \sqrt{n\sigma^{2} \log n} \right\} \sqrt{\frac{\mu}{n}},$$

where the identity arises since $\tilde{u}_2^{\star \top} \tilde{u}_1^{\star} = 0$. Given that $\tilde{\lambda}_2^{\star} = -1$ and $\tilde{\lambda}_1 > 0$ (see Corollary 8), we have $1 - \tilde{\lambda}_2^{\star} / \tilde{\lambda}_1 > 1$. The near orthogonality property can then be described as follows

$$\left|\widetilde{\boldsymbol{u}}_{2}^{\star\top}\widetilde{\boldsymbol{u}}_{1}\right| \leq \left|\widetilde{\boldsymbol{u}}_{2}^{\star\top}\widetilde{\boldsymbol{u}}_{1}\left(1 - \widetilde{\lambda}_{2}^{\star}/\widetilde{\lambda}_{1}\right)\right| \lesssim \max\left\{B\log n, \sqrt{n\sigma^{2}\log n}\right\}\sqrt{\frac{\mu}{n}}.$$
(84)

This combined with Lemma 3 yields

$$\min\{\|\widetilde{\boldsymbol{u}}_{1} \pm \widetilde{\boldsymbol{u}}_{1}^{\star}\|_{2}\} = \sqrt{\|\widetilde{\boldsymbol{u}}_{1}^{\star}\|_{2}^{2} + \|\widetilde{\boldsymbol{u}}_{1}\|_{2}^{2} - 2\left|\widetilde{\boldsymbol{u}}_{1}^{\star\top}\widetilde{\boldsymbol{u}}_{1}\right|}$$
$$= \sqrt{2\left(1 - \left|\widetilde{\boldsymbol{u}}_{1}^{\star\top}\widetilde{\boldsymbol{u}}_{1}\right|\right)} \leq \sqrt{2\left(1 - \left|\widetilde{\boldsymbol{u}}_{1}^{\star\top}\widetilde{\boldsymbol{u}}_{1}\right|^{2}\right)}$$
$$\stackrel{(i)}{\leq} \sqrt{\frac{128}{9} \|\boldsymbol{H}_{\text{dilation}}\|^{2} + 2\left|\widetilde{\boldsymbol{u}}_{2}^{\star\top}\widetilde{\boldsymbol{u}}_{1}\right|^{2}}$$
$$\stackrel{(ii)}{\lesssim} \|\boldsymbol{H}_{\text{dilation}}\| + \left|\widetilde{\boldsymbol{u}}_{2}^{\star\top}\widetilde{\boldsymbol{u}}_{1}\right|$$
$$\stackrel{(iii)}{\lesssim} \max\left\{B\log n, \sqrt{n\sigma^{2}\log n}\right\},$$
(85)

where (i) results from Lemma 3 and the fact that the condition number of $M_{\text{dilation}}^{\star}$ is 1, (ii) follows since $\sqrt{x^2 + y^2} \leq x + y$ for all $x, y \geq 0$, and (iii) arises from (84), Lemma 1, the identity $\|H_{\text{dilation}}\| = \max\{\|H_1\|, \|H_2\|\}$, and the fact $\sqrt{\mu/n} \leq 1$.

It then boils down to showing that $\min\{\|u \pm u^*\|_2\} \lesssim \min\{\|\widetilde{u}_1 \pm \widetilde{u}_1^*\|_2\}$. To this end, we see that the estimate (46) satisfies

$$\|\boldsymbol{u} - \boldsymbol{u}^{\star}\|_{2} = \left\| \left(\frac{\widetilde{\boldsymbol{u}}_{1,1}}{\|\widetilde{\boldsymbol{u}}_{1,1}\|_{2}} - \sqrt{2}\widetilde{\boldsymbol{u}}_{1,1} \right) + \left(\sqrt{2}\widetilde{\boldsymbol{u}}_{1,1} - \boldsymbol{u}^{\star} \right) \right\|_{2}$$

$$\leq \left| \frac{1}{\|\widetilde{\boldsymbol{u}}_{1,1}\|_{2}} - \sqrt{2} \right| + \sqrt{2} \left\| \widetilde{\boldsymbol{u}}_{1,1} - \frac{\boldsymbol{u}^{\star}}{\sqrt{2}} \right\|_{2}$$

$$= \frac{\sqrt{2}}{\|\widetilde{\boldsymbol{u}}_{1,1}\|_{2}} \left\| \|\widetilde{\boldsymbol{u}}_{1,1}\|_{2} - \left\| \frac{\boldsymbol{u}^{\star}}{\sqrt{2}} \right\|_{2} \right| + \sqrt{2} \left\| \widetilde{\boldsymbol{u}}_{1,1} - \frac{\boldsymbol{u}^{\star}}{\sqrt{2}} \right\|_{2}$$

$$\leq \sqrt{2} \left(\frac{1}{\|\widetilde{\boldsymbol{u}}_{1,1}\|_{2}} + 1 \right) \left\| \widetilde{\boldsymbol{u}}_{1,1} - \frac{\boldsymbol{u}^{\star}}{\sqrt{2}} \right\|_{2}$$
(86)

$$\leq \sqrt{2} \left(1 + \frac{\sqrt{2}}{1 - \sqrt{2} \|\widetilde{\boldsymbol{u}}_1 - \widetilde{\boldsymbol{u}}_1^\star\|_2} \right) \|\widetilde{\boldsymbol{u}}_1 - \widetilde{\boldsymbol{u}}_1^\star\|_2.$$
(87)

Here, (86) makes use of the triangle inequality, and (87) follows since

$$\frac{1}{\|\tilde{\boldsymbol{u}}_{1,1}\|_{2}} \leq \frac{1}{\|\boldsymbol{u}^{\star}\|_{2}/\sqrt{2} - \|\tilde{\boldsymbol{u}}_{1,1} - \boldsymbol{u}^{\star}/\sqrt{2}\|_{2}} = \frac{\sqrt{2}}{1 - \|\sqrt{2}\tilde{\boldsymbol{u}}_{1,1} - \boldsymbol{u}^{\star}\|_{2}} \leq \frac{\sqrt{2}}{1 - \sqrt{2}\|\tilde{\boldsymbol{u}}_{1} - \tilde{\boldsymbol{u}}_{1}^{\star}\|_{2}},$$

where the last inequality relies on the fact that $\|\tilde{u}_{1,1} - u^*/\sqrt{2}\|_2 \leq \|\tilde{u}_1 - \tilde{u}_1^*\|_2$ (the first n_1 coordinates). Similarly, one can derive the above bounds for $\|u + u^*\|_2$ as well. Therefore, we are left with

$$\min\left\{\left\|\boldsymbol{u} \pm \boldsymbol{u}^{\star}\right\|_{2}\right\} \leq \sqrt{2} \left(1 + \frac{\sqrt{2}}{1 - \sqrt{2}\min\left\{\left\|\widetilde{\boldsymbol{u}}_{1} \pm \widetilde{\boldsymbol{u}}_{1}^{\star}\right\|_{2}\right\}}\right) \min\left\{\left\|\widetilde{\boldsymbol{u}}_{1} \pm \widetilde{\boldsymbol{u}}_{1}^{\star}\right\|_{2}\right\}$$
(88)

$$\lesssim \min\{\|\widetilde{\boldsymbol{u}}_1 \pm \widetilde{\boldsymbol{u}}_1^*\|_2\} \tag{89}$$

$$\lesssim \max\left\{B\log n, \sqrt{n\sigma^2\log n}\right\}.$$
(90)

This concludes the proof.

I.2 Perturbation bounds for linear forms of eigenvectors

Fix any unit vector $a \in \mathbb{R}^{n_1}$. Considering the linear transformation by the vector $\begin{pmatrix} a \\ 0 \end{pmatrix}$, we can apply Theorem 4 to M_{dilation} and get

$$\left| \begin{pmatrix} \boldsymbol{a} \\ \boldsymbol{0} \end{pmatrix}^{\top} \left(\widetilde{\boldsymbol{u}}_1 - \frac{\widetilde{\lambda}_1^{\star} \widetilde{\boldsymbol{u}}_1^{\star \top} \widetilde{\boldsymbol{u}}_1}{\widetilde{\lambda}_1} \widetilde{\boldsymbol{u}}_1^{\star} - \frac{\widetilde{\lambda}_2^{\star} \widetilde{\boldsymbol{u}}_2^{\star \top} \widetilde{\boldsymbol{u}}_1}{\widetilde{\lambda}_1} \widetilde{\boldsymbol{u}}_2^{\star} \right) \right| \lesssim \max\left\{ B \log n, \sqrt{n\sigma^2 \log n} \right\} \sqrt{\frac{\mu}{n}}.$$
(91)

Given that the first n_1 coordinates of \tilde{u}_1^{\star} and \tilde{u}_2^{\star} are both $u^{\star}/\sqrt{2}$, we can invoke Corollary 8 (i.e. $\tilde{\lambda}_1 \approx 1$) to obtain

$$\left| \boldsymbol{a}^{\top} \left(\widetilde{\boldsymbol{u}}_{1,1} - \frac{\widetilde{\boldsymbol{u}}_{1}^{\star \top} \widetilde{\boldsymbol{u}}_{1}}{\sqrt{2} \widetilde{\lambda}_{1}} \boldsymbol{u}^{\star} \right) \right| \lesssim \left| \frac{\widetilde{\boldsymbol{u}}_{2}^{\star \top} \widetilde{\boldsymbol{u}}_{1}}{\sqrt{2}} \right| \left| \boldsymbol{a}^{\top} \boldsymbol{u}^{\star} \right| + \max \left\{ B \log n, \sqrt{n\sigma^{2} \log n} \right\} \sqrt{\frac{\mu}{n}}.$$
(92)

In view of (84) and the fact $|\boldsymbol{a}^{\top}\boldsymbol{u}^{\star}| \leq 1$, one has

$$\left| \boldsymbol{a}^{\top} \left(\widetilde{\boldsymbol{u}}_{1,1} - \frac{\widetilde{\boldsymbol{u}}_{1}^{\star \top} \widetilde{\boldsymbol{u}}_{1}}{\sqrt{2} \,\widetilde{\lambda}_{1}} \boldsymbol{u}^{\star} \right) \right| \lesssim \max \left\{ B \log n, \sqrt{n\sigma^{2} \log n} \right\} \sqrt{\frac{\mu}{n}}.$$
(93)

Recall that $\widetilde{\lambda}_1 > 0$ (cf. Corollary 8). If we further have

$$\left| \frac{\left| \widetilde{\boldsymbol{u}}_{1}^{\star \top} \widetilde{\boldsymbol{u}}_{1} \right|}{\sqrt{2} \,\widetilde{\lambda}_{1} \left\| \widetilde{\boldsymbol{u}}_{1,1} \right\|_{2}} - 1 \right| \lesssim \max \left\{ B \log n, \sqrt{n\sigma^{2} \log n} \right\},\tag{94}$$

then we can use the triangle inequality to reach

$$\min\left\{\left|\boldsymbol{a}^{\top}(\boldsymbol{u}\pm\boldsymbol{u}^{\star})\right|\right\} \leq \frac{1}{\|\widetilde{\boldsymbol{u}}_{1,1}\|_{2}} \left|\boldsymbol{a}^{\top}\left(\widetilde{\boldsymbol{u}}_{1,1}-\frac{\widetilde{\boldsymbol{u}}_{1}^{\top}\widetilde{\boldsymbol{u}}_{1}}{\sqrt{2}\,\widetilde{\lambda}_{1}}\,\boldsymbol{u}^{\star}\right)\right| + \left|\frac{\|\widetilde{\boldsymbol{u}}_{1}^{\top}\widetilde{\boldsymbol{u}}_{1}\|}{\sqrt{2}\,\widetilde{\lambda}_{1}}\,\|\widetilde{\boldsymbol{u}}_{1,1}\|_{2} - 1\right| \left|\boldsymbol{a}^{\top}\boldsymbol{u}^{\star}\right| \\ \lesssim \left(\sqrt{\frac{\mu}{n}} + |\boldsymbol{a}^{\top}\boldsymbol{u}^{\star}|\right) \max\left\{B\log n, \sqrt{n\sigma^{2}\log n}\right\}$$
(95)

as claimed. It then remains to justify (94). Towards this, it suffices to combine a series of consequences from (85), Corollary 8, and (90), namely,

$$1 - \left|\widetilde{\boldsymbol{u}}_{1}^{\star \top}\widetilde{\boldsymbol{u}}_{1}\right| = \frac{1}{2}\min\{\|\widetilde{\boldsymbol{u}}_{1} \pm \widetilde{\boldsymbol{u}}_{1}^{\star}\|_{2}^{2}\} \lesssim \max\left\{B\log n, \sqrt{n\sigma^{2}\log n}\right\},\tag{96}$$

$$\left|\widetilde{\lambda}_{1} - \widetilde{\lambda}_{1}^{\star}\right| = \left|\widetilde{\lambda}_{1} - 1\right| \lesssim \max\left\{B\log n, \sqrt{n\sigma^{2}\log n}\right\} \sqrt{\frac{\mu}{n}},\tag{97}$$

$$\left\| \widetilde{\boldsymbol{u}}_{1,1} \|_2 - \frac{1}{\sqrt{2}} \right\| \le \min\left\{ \left\| \widetilde{\boldsymbol{u}}_{1,1} \pm \frac{\boldsymbol{u}^*}{\sqrt{2}} \right\|_2 \right\} \le \min\left\{ \left\| \widetilde{\boldsymbol{u}}_1 \pm \widetilde{\boldsymbol{u}}_1^* \right\|_2 \right\} \le \max\left\{ B \log n, \sqrt{n\sigma^2 \log n} \right\}.$$
(98)

The proof for the bounds on v is similar and is thus omitted.

I.3 Entrywise eigenvector perturbation bounds

Recognizing that $\|\boldsymbol{u} - \boldsymbol{u}^{\star}\|_{\infty} = \max_i |\boldsymbol{e}_i^{\top}\boldsymbol{u} - \boldsymbol{e}_i^{\top}\boldsymbol{u}^{\star}|$ and using the incoherence $|\boldsymbol{e}_i^{\top}\boldsymbol{u}| \leq \sqrt{\mu/n}$, we can prove this claim directly by invoking the results established in Appendix I.2 and taking the union bound.

J Asymmetrization of data samples: two examples

As mentioned earlier, an independent and asymmetric noise matrix arises when we collect two samples for each entry of the matrix of interest and arrange the samples in an asymmetric manner (i.e. placing 1 sample on the upper triangular part and the other on the lower triangular part). Interestingly, our results might be applicable for some cases where we only have 1 sample for each entry. In what follows, we describe two examples of this kind similar to the ones discussed in Section 4.2, but with a symmetric noise matrix. Once again, it is assumed that M^* is a rank-1 matrix with leading eigenvalue 1 and incoherence parameter μ .

Low-rank matrix estimation from Gaussian noise.³ Suppose that the noise matrix H is a symmetric matrix, and that the upper triangular part of H is composed of i.i.d. Gaussian random variables $\mathcal{N}(0, \sigma^2)$.

When σ is known, one can decouple the upper and low triangular parts of H by adding a *skew-symmetric* Gaussian matrix Δ . Specifically, our strategy is:

- (1) For each $1 \leq i \leq j \leq n$, generate $\Delta_{ij} \sim \mathcal{N}(0, \sigma^2)$ independently, and set $\Delta_{ji} = -\Delta_{ij}$;
- (2) Compute the leading eigenvalue and eigenvector of $M + \Delta$.

This is motivated by a simple observation from Gaussianality: $H + \Delta$ is now an asymmetric matrix whose offdiagonal entries are i.i.d. $\mathcal{N}(0, 2\sigma^2)$; in fact, it is easy to verify that $H_{ij} + \Delta_{ij}$ and $H_{ji} - \Delta_{ji}$ are independent Gaussian random variables. As a result, the orderwise bounds (26) continue to hold if λ and u are taken to be the leading eigenvalue and eigenvector of $M + \Delta$, respectively.

While this asymmetrization procedure comes with the price of doubling the noise variance, the eigenvalue perturbation bound may still be significantly smaller — up to a factor of $O(\sqrt{n})$ — than the bound for the SVD approach. This is also confirmed in the numerical simulations reported in Fig. 8.



Figure 8: Numerical simulation for rank-1 matrix estimation when \boldsymbol{H} is symmetric and its upper triangular part consists of i.i.d. Gaussian entries $\mathcal{N}(0, \sigma^2)$. The truth \boldsymbol{M}^* is rank-1 and is randomly generated with leading eigenvalue $\lambda^* = 1$. (a) $|\lambda - \lambda^*|$ vs. n with $\sigma = 1/\sqrt{n \log n}$; (b) $|\lambda - \lambda^*|$ vs. σ with n = 1000; (c) ℓ_{∞} eigenvector perturbation vs. n with $\sigma = 1/\sqrt{n \log n}$. We compare eigen-decomposition applied to symmetric matrix samples (the red lines) and asymmetrized data (the blue lines). In the green line, the estimate \boldsymbol{u} is obtained by generating 10 independent copies of $\boldsymbol{\Delta}$ and aggregating the leading eigenvectors of these copies of $\boldsymbol{M} + \boldsymbol{\Delta}$ (see Remark 9).

Remark 9. While this asymmetrization procedure achieves enhanced eigenvalue estimation accuracy compared to the SVD approach (see Fig. 8(a)(b)), it results in higher eigenvector estimation errors (see Fig. 8(c)). This is perhaps not surprising as we have added extra noise to the observed matrix. One way to mitigate

 $^{^{3}}$ We thank Prof. Zhou Fan for telling us this example.

this issues is to (1) generate K independent copies of Δ ; (2) compute the leading eigenvector of each copy of $M + \Delta$, denoted by $\{u^{(l)} \mid 1 \leq l \leq K\}$; and (3) aggregate these eigenvectors, namely, compute the leading eigenvector of $\frac{1}{K} \sum_{l=1}^{K} u^{(l)} u^{(l)\top}$. As can be seen in the green line of Fig. 8(c), this allows us to mitigate the effect of the extra noise component.

Low-rank matrix completion. Suppose that each entry M_{ij}^{\star} $(i \geq j)$ is observed independently with probability p. This is different from the settings in Section 4.2, as we do not have additional samples for M_{ji}^{\star} $(i \geq j)$. In order to arrange the data samples in an asymmetric and independent fashion, we employ a simple resampling technique to decouple the statistical dependency between M_{ij} and M_{ji} :

(1) Define M^{sym} such that

$$M_{ij}^{\mathsf{sym}} = \begin{cases} \frac{1}{p} M_{ij}^{\star}, & \text{if } M_{ij}^{\star} \text{ is observed}; \\ 0, & \text{else.} \end{cases}$$

(2) Define p^{asym} so that $p = 1 - (1 - p^{\mathsf{asym}})^2$ (i.e. $p^{\mathsf{asym}} = \frac{p}{1 + \sqrt{1-p}}$). For any pair i > j, set

$$\left(M_{ij}^{\mathsf{asym}}, M_{ji}^{\mathsf{asym}}\right) = \begin{cases} \frac{p}{p^{\mathsf{asym}}} \left(M_{ij}^{\mathsf{sym}}, 0\right) & \text{with probability } \frac{1-p^{\mathsf{asym}}}{2-p^{\mathsf{asym}}}, \\ \frac{p}{p^{\mathsf{asym}}} \left(0, M_{ij}^{\mathsf{sym}}\right) & \text{with probability } \frac{1-p^{\mathsf{asym}}}{2-p^{\mathsf{asym}}}, \\ \frac{p}{p^{\mathsf{asym}}} \left(M_{ij}^{\mathsf{sym}}, M_{ij}^{\mathsf{sym}}\right) & \text{with probability } \frac{p^{\mathsf{asym}}}{2-p^{\mathsf{asym}}}. \end{cases} \tag{99}$$

For any i, set

$$M_{ii}^{\mathsf{asym}} = \begin{cases} \frac{p}{p^{\mathsf{asym}}} M_{ii}^{\mathsf{sym}} & \text{with probability } \frac{p^{\mathsf{asym}}}{p}, \\ 0 & \text{else.} \end{cases}$$

(3) Compute the leading eigenvalue and eigenvector of M^{asym} .

As can be easily verified, this scheme is equivalent to saying that (i) with probability p, either M_{ij}^{asym} or M_{ji}^{asym} is taken to be a rescaled version of M_{ij}^{\star} ; (ii) for any $i \neq j$, the entries $\{M_{ij}^{asym}\}$ are independently drawn. Since $p^{asym} \approx p$, our results (28) in Section 4.2 remain valid, as long as λ and u are set to be the leading eigenvalue and eigenvector of M^{asym} , respectively. Numerical simulations have been carried out in Fig. 9 to verify the effectiveness of this scheme.



(a) eigenvalue perturbation vs. n

(b) eigenvalue perturbation vs. p

(c) ℓ_{∞} eigenvector perturbation

Figure 9: Numerical experiments for rank-1 matrix completion, where the rank-1 truth M^* is randomly generated with leading eigenvalue 1. (a) $|\lambda - \lambda^*|$ vs. n with $p = 3 \log n / n$; (b) $|\lambda - \lambda^*|$ vs. p with n = 1000; (c) ℓ_{∞} perturbation error vs. n with $p = 3 \log n / n$. We compare eigen-decomposition applied to symmetric matrix samples and asymmetrized data. The blue (resp. red) lines represent the average errors over 100 independent trials using the eigen-decomposition (resp. SVD) approach.