




Contents lists available at ScienceDirect

Linear Algebra and its Applications

www.elsevier.com/locate/laa



A Schatten- q low-rank matrix perturbation analysis via perturbation projection error bound



Yuetian Luo^a, Rungang Han^a, Anru R. Zhang^{a,b,*}

^a Department of Statistics, University of Wisconsin, Madison, United States of America

^b Department of Biostatistics and Bioinformatics, Duke University, United States of America

ARTICLE INFO

Article history:

Received 19 November 2020

Accepted 9 August 2021

Available online 13 August 2021

Submitted by V. Mehrmann

MSC:

15A42

65F55

Keywords:

Perturbation theory

Schatten- q norm

Singular value decomposition

Low-rank matrix estimation

Matrix perturbation projection

Sin-theta distance

ABSTRACT

This paper studies the Schatten- q error of low-rank matrix estimation by singular value decomposition under perturbation. We specifically establish a perturbation bound on the low-rank matrix estimation via a perturbation projection error bound. Then, we establish lower bounds to justify the tightness of the upper bound on the low-rank matrix estimation error. We further develop a user-friendly $\sin\Theta$ bound for singular subspace perturbation based on the matrix perturbation projection error bound. Finally, we demonstrate the advantage of our results over the ones in the literature by simulation.

© 2021 Elsevier Inc. All rights reserved.

* Corresponding author at: Department of Biostatistics and Bioinformatics, Duke University, United States of America.

E-mail addresses: ylo86@wisc.edu (Y. Luo), rhan32@stat.wisc.edu (R. Han), anru.zhang@duke.edu (A.R. Zhang).

1. Introduction

Let $\mathbf{A}$ be an m -by- n real-valued matrix with singular value decomposition (SVD)

$$\mathbf{A} = [\mathbf{U} \ \mathbf{U}_\perp] \begin{bmatrix} \boldsymbol{\Sigma}_1 & 0 \\ 0 & \boldsymbol{\Sigma}_2 \end{bmatrix} \begin{bmatrix} \mathbf{V}^\top \\ \mathbf{V}_\perp^\top \end{bmatrix},$$

where $\mathbf{U} \in \mathbb{R}^{m \times r}$, $\mathbf{V} \in \mathbb{R}^{n \times r}$, $[\mathbf{U} \ \mathbf{U}_\perp], [\mathbf{V} \ \mathbf{V}_\perp]$ are orthogonal and $\boldsymbol{\Sigma}_1, \boldsymbol{\Sigma}_2$ are (pseudo) diagonal matrices with decreasing singular values of $\mathbf{A}$. Suppose $\mathbf{B} = \mathbf{A} + \mathbf{Z} \in \mathbb{R}^{m \times n}$, where $\mathbf{Z}$ is some perturbation matrix. We similarly write down the SVD of $\mathbf{B}$ as

$$\mathbf{B} = [\hat{\mathbf{U}} \ \hat{\mathbf{U}}_\perp] \begin{bmatrix} \hat{\boldsymbol{\Sigma}}_1 & 0 \\ 0 & \hat{\boldsymbol{\Sigma}}_2 \end{bmatrix} \begin{bmatrix} \hat{\mathbf{V}}^\top \\ \hat{\mathbf{V}}_\perp^\top \end{bmatrix}$$

such that $\hat{\mathbf{U}}$ and $\hat{\mathbf{V}}$ share the same dimensions as $\mathbf{U}$ and $\mathbf{V}$, respectively. The relationship between the singular structures of $\mathbf{A}$ and $\mathbf{B}$ is a central topic in matrix perturbation theory. Since the seminal work by Weyl [1], Davis-Kahan [2], Wedin [3], the perturbation analysis for singular values (i.e., $\boldsymbol{\Sigma}_1, \boldsymbol{\Sigma}_2$ versus $\hat{\boldsymbol{\Sigma}}_1, \hat{\boldsymbol{\Sigma}}_2$) and the leading singular vectors (i.e., $\mathbf{U}, \mathbf{V}$ versus $\hat{\mathbf{U}}, \hat{\mathbf{V}}$) have attracted enormous attentions. For example, [4–6] studied perturbation expansion for singular value decomposition; [7–10] established the relative perturbation theory for eigenvectors of Hermitian matrices and singular vectors of general matrices; [11–14] studied the numeric computation accuracy for singular values and vectors; more recently, [15–17] developed several new perturbation results under specific structural assumptions motivated by emerging applications in statistics and data science. The readers are referred to [18–20] for overviews of the historical development of matrix perturbation theory.

While most of the existing works focused on $\mathbf{U}$ and $\mathbf{V}$ or $\boldsymbol{\Sigma}_1$ and $\boldsymbol{\Sigma}_2$, there are fewer studies on the perturbation analysis of the true matrix $\mathbf{A}$ itself. In this paper, we consider the estimation of rank- r matrix $\mathbf{A}$ (i.e., $\boldsymbol{\Sigma}_2 = 0$) via rank- r truncated SVD (i.e., best rank- r approximation) of $\mathbf{B}$: $\hat{\mathbf{A}} := \hat{\mathbf{U}}\hat{\boldsymbol{\Sigma}}_1\hat{\mathbf{V}}^\top$. Such a low-rank assumption and estimation method are widely used in many applications including matrix denoising [21,22], signal processing [23,24] and multivariate statistical analysis [25], etc. We focus on the estimation error in matrix Schatten- q norm: $\|\hat{\mathbf{A}} - \mathbf{A}\|_q$. A tight upper bound on $\|\hat{\mathbf{A}} - \mathbf{A}\|_q$ can provide an important benchmark for both algorithmic and statistical analysis in the applications mentioned above; moreover, it can be used to study some other basic perturbation quantities, such as the pseudo-inverse perturbation $\|\hat{\mathbf{A}}^\dagger - \mathbf{A}^\dagger\|_q$ [26,27].

As a starting point, it is straightforward to apply the classical perturbation bounds for singular values and vectors to obtain an upper bound on $\|\hat{\mathbf{A}} - \mathbf{A}\|_q$. For example, one can immediately have the following inequality via Wedin's $\sin \Theta$ Theorem (Eq. (4.4) in [3]),

$$\|\hat{\mathbf{A}} - \mathbf{A}\|_q \leq \|\mathbf{Z}\|_q \left(3 + \|\mathbf{B} - \hat{\mathbf{A}}\|_q / \sigma_r(\mathbf{B}) \right). \quad (1)$$

Another way is utilizing the optimality of SVD (Eckart-Young-Mirsky Theorem) and some basic norm inequalities to obtain:

$$\|\hat{\mathbf{A}} - \mathbf{A}\|_q \leq \|\hat{\mathbf{A}} - \mathbf{B}\|_q + \|\mathbf{A} - \mathbf{B}\|_q \leq 2\|\mathbf{A} - \mathbf{B}\|_q = 2\|\mathbf{Z}\|_q, \quad (2)$$

$$\|\hat{\mathbf{A}} - \mathbf{A}\|_q \leq (2r)^{1/q} \|\hat{\mathbf{A}} - \mathbf{A}\| \stackrel{(2)}{\leq} 2^{1+1/q} r^{1/q} \|\mathbf{Z}\|. \quad (3)$$

In contrast, we establish the following result in this paper:

Theorem 1. Suppose $\mathbf{B} = \mathbf{A} + \mathbf{Z}$, where $\mathbf{A}$ is an unknown rank- r matrix, $\mathbf{B}$ is the observation, and $\mathbf{Z}$ is the perturbation. Let $\hat{\mathbf{A}} = \hat{\mathbf{U}}\hat{\Sigma}_1\hat{\mathbf{V}}^\top$ be the best rank- r approximation of $\mathbf{B}$. Then,

$$\|\hat{\mathbf{A}} - \mathbf{A}\|_q \leq \begin{cases} (2^q + 1)^{1/q} \|\mathbf{Z}_{\max(r)}\|_q, & 1 \leq q \leq 2; \\ \sqrt{5} \|\mathbf{Z}_{\max(r)}\|_q, & 2 \leq q < \infty; \\ 2\|\mathbf{Z}_{\max(r)}\|, & q = \infty. \end{cases} \quad (4)$$

Here $\mathbf{Z}_{\max(r)}$ is defined as the best rank- r approximation of $\mathbf{Z}$.

The proof of Theorem 1 relies on a careful characterization of $\|P_{\hat{\mathbf{U}}_\perp} \mathbf{A}\|_q$ (where $P_{\hat{\mathbf{U}}_\perp}$ is the projection onto the subspace spanned by $\hat{\mathbf{U}}_\perp$) in Theorem 2, which we refer as the *perturbation projection error bound*. The details of Theorem 2 and the proof of Theorem 1 will be presented in Section 2.

The established bound (4) is sharper than the classic results (1), (2) and (3) since $\|\mathbf{Z}_{\max(r)}\|_q \leq \min\{\|\mathbf{Z}\|_q, r^{1/q}\|\mathbf{Z}\|\}$ for any $\mathbf{Z}$. When $m, n \gg r$ and the first r singular values of $\mathbf{Z}$ decay fast, which commonly happens in many large-scale matrix datasets [28], $\|\mathbf{Z}_{\max(r)}\|_q$ can be much smaller than $\|\mathbf{Z}\|_q, r^{1/q}\|\mathbf{Z}\|$ (see an example in Section 2) so that the upper bound of (4) can be much smaller than (1), (2) and (3).

Then, we further introduce two lower bounds to justify the tightness of the upper bound in Theorem 1. Specifically for any $\epsilon > 0$, $1 \leq q \leq \infty$, we construct a triplet of matrices $(\mathbf{A}, \mathbf{Z}, \mathbf{B})$ such that

$$\|\hat{\mathbf{A}} - \mathbf{A}\|_q \geq ((2^q + 1)^{1/q} - \epsilon) \|\mathbf{Z}_{\max(r)}\|_q > 0, \quad (5)$$

which suggests that the constant in (4) cannot be further improved for $q \in [1, 2] \cup \{\infty\}$. In addition, we introduce an estimation error lower bound to show that the rank- r truncated SVD estimator (i.e., $\hat{\mathbf{A}}$) is minimax rate-optimal over the class of all rank- r matrices.

As a byproduct of the theory in this paper, we derive a subspace (singular vectors) $\sin\Theta$ perturbation bound (definition of Schatten- q $\sin\Theta$ distance is in Section 1.1) under the same condition as Theorem 1:

$$\max \left\{ \|\sin \Theta(\widehat{\mathbf{U}}, \mathbf{U})\|_q, \|\sin \Theta(\widehat{\mathbf{V}}, \mathbf{V})\|_q \right\} \leq \frac{2\|\mathbf{Z}_{\max(r)}\|_q}{\sigma_r(\mathbf{A})}.$$

This bound is “user-friendly” as it is free of $\mathbf{B}$, $\widehat{\mathbf{U}}$, and $\widehat{\mathbf{V}}$, which are often perturbed and uncontrolled quantities in practice (see more discussions in Section 4).

The rest of this paper is organized as follows. After a brief introduction on notation and preliminaries in Section 1.1, we present the proof of Theorem 1 in Section 2 and develop the corresponding lower bounds in Section 3. The new $\sin \Theta$ perturbation analysis is done in Section 4. We provide numerical studies to corroborate our theoretical findings in Section 5. Conclusion and discussions are made in Section 6.

1.1. Notation and preliminaries

The following notation will be used throughout this paper. The lowercase letters (e.g., a, b), lowercase boldface letters (e.g., $\mathbf{u}, \mathbf{v}$), uppercase boldface letters (e.g., $\mathbf{U}, \mathbf{V}$) are used to denote scalars, vectors, matrices, respectively. For any two numbers a, b , let $a \wedge b = \min\{a, b\}$, $a \vee b = \max\{a, b\}$. For any matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$ with singular value decomposition $\sum_{i=1}^{m \wedge n} \sigma_i(\mathbf{A}) \mathbf{u}_i \mathbf{v}_i^\top$, let $\mathbf{A}_{\max(r)} = \sum_{i=1}^r \sigma_i(\mathbf{A}) \mathbf{u}_i \mathbf{v}_i^\top$ be the best rank- r approximation of $\mathbf{A}$, and $\mathbf{A}_{-\max(r)} = \sum_{i=r+1}^{m \wedge n} \sigma_i(\mathbf{A}) \mathbf{u}_i \mathbf{v}_i^\top$ be the remainder. For $q \in [1, \infty]$, the Schatten- q norm of matrix $\mathbf{A}$ is defined as $\|\mathbf{A}\|_q := (\sum_{i=1}^{m \wedge n} \sigma_i^q(\mathbf{A}))^{1/q}$. Especially, Frobenius norm $\|\cdot\|_F$ and spectral norm $\|\cdot\|$ are Schatten-2 norm and Schatten- ∞ norm, respectively. In addition, let $\mathbf{I}_r$ be the r -by- r identity matrix. Let $\mathbb{O}_r$ be the set of r -by- r orthogonal matrices, $\mathbb{O}_{p,r} = \{\mathbf{U} \in \mathbb{R}^{p \times r} : \mathbf{U}^\top \mathbf{U} = \mathbf{I}_r\}$ be the set of all p -by- r matrices with orthonormal columns. For any $\mathbf{U} \in \mathbb{O}_{p,r}$, $P_{\mathbf{U}} = \mathbf{U} \mathbf{U}^\top$ is the projection matrix onto the column span of $\mathbf{U}$. We also use $\mathbf{U}_\perp \in \mathbb{O}_{p,p-r}$ to represent the orthonormal complement of $\mathbf{U}$. We use bracket subscripts to denote sub-matrices. For example, $\mathbf{A}_{[i_1, i_2]}$ is the entry of $\mathbf{A}$ on the i_1 -th row and i_2 -th column; $\mathbf{A}_{[(r+1):m, :]}$ contains the $(r+1)$ -th to the m -th rows of $\mathbf{A}$.

We use the $\sin \Theta$ norm to quantify the distance between singular subspaces. Suppose $\mathbf{U}_1$ and $\mathbf{U}_2$ are two p -by- r matrices with orthonormal columns. Let the singular values of $\mathbf{U}_1^\top \mathbf{U}_2$ be $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_r \geq 0$. Then $\Theta(\mathbf{U}_1, \mathbf{U}_2)$ is defined as a diagonal matrix with principal angles between $\mathbf{U}_1$ and $\mathbf{U}_2$:

$$\Theta(\mathbf{U}_1, \mathbf{U}_2) = \text{diag}(\cos^{-1}(\sigma_1), \dots, \cos^{-1}(\sigma_r)).$$

Then the Schatten- q $\sin \Theta$ distance is defined as

$$\|\sin \Theta(\mathbf{U}_1, \mathbf{U}_2)\|_q = \|\text{diag}(\sin \cos^{-1}(\sigma_1), \dots, \sin \cos^{-1}(\sigma_r))\|_q = \left(\sum_{i=1}^r (1 - \sigma_i^2)^{q/2} \right)^{1/q}. \quad (6)$$

Importantly, $\|\mathbf{U}_{1\perp}^\top \mathbf{U}_2\|_q = \|\sin \Theta(\mathbf{U}_1, \mathbf{U}_2)\|_q$ for any $q \in [1, \infty]$ [8, Lemma 2.1].

Finally, a function $\Phi: \mathbb{R}^n \rightarrow \mathbb{R}$ is called a symmetric gauge function if (1) $\mathbf{x} \neq 0 \implies \Phi(\mathbf{x}) > 0$, (2) $\Phi(\rho\mathbf{x}) = |\rho|\Phi(\mathbf{x})$ for $\rho \in \mathbb{R}$, (3) $\Phi(\mathbf{x} + \mathbf{y}) \leq \Phi(\mathbf{x}) + \Phi(\mathbf{y})$ for any $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$, and (4) for any permutation matrix $\mathbf{P}$, we have $\Phi(\mathbf{P}\mathbf{x}) = \Phi(\mathbf{x})$ [18, Definition II.3.3].

2. Proof of Theorem 1

The roadmap of the proof of Theorem 1 is the following. We first introduce Theorem 2, which quantifies the projection error, $\|P_{\hat{\mathbf{U}}_\perp} \mathbf{A}\|_q$ and $\|\mathbf{A}P_{\hat{\mathbf{V}}_\perp}\|_q$, under the perturbation model. This result plays a crucial role in the proof of Theorem 1 and may also be of independent interest. Next, we present Lemma 1 with proof and then give the proof for Theorem 1. Since the proof of Theorem 2 is relatively long, we present its full proof and discussions in Section 2.1.

Theorem 2 (A perturbation projection error bound). Suppose $\mathbf{B} = \mathbf{A} + \mathbf{Z}$ for some rank- r matrix $\mathbf{A}$ and perturbation matrix $\mathbf{Z}$. Then for any $q \in [1, \infty]$,

$$\max \left\{ \|P_{\hat{\mathbf{U}}_\perp} \mathbf{A}\|_q, \|\mathbf{A}P_{\hat{\mathbf{V}}_\perp}\|_q \right\} \leq 2\|\mathbf{Z}_{\max(r)}\|_q. \quad (7)$$

Next, the following Lemma 1 characterizes the Schatten- q norm of matrix orthogonal projections.

Lemma 1. Suppose $\mathbf{A}, \mathbf{B} \in \mathbb{R}^{m \times n}$, $\mathbf{U} \in \mathbb{O}_{m,r}$, $q \geq 1$. Then,

$$\|P_{\mathbf{U}} \mathbf{A} + P_{\mathbf{U}_\perp} \mathbf{B}\|_q \leq \begin{cases} (\|P_{\mathbf{U}} \mathbf{A}\|_q^2 + \|P_{\mathbf{U}_\perp} \mathbf{B}\|_q^2)^{1/2}, & 2 \leq q \leq \infty; \\ (\|P_{\mathbf{U}} \mathbf{A}\|_q^q + \|P_{\mathbf{U}_\perp} \mathbf{B}\|_q^q)^{1/q}, & 1 \leq q \leq 2. \end{cases}$$

Proof. Let $\mathbf{T} = P_{\mathbf{U}} \mathbf{A} + P_{\mathbf{U}_\perp} \mathbf{B}$. We construct $\mathbf{T}_1 = P_{\mathbf{U}} \mathbf{T} = P_{\mathbf{U}} \mathbf{A}$, $\mathbf{T}_2 = P_{\mathbf{U}_\perp} \mathbf{T} = P_{\mathbf{U}_\perp} \mathbf{B}$. First we have $\mathbf{T}^\top \mathbf{T} = \mathbf{T}_1^\top \mathbf{T}_1 + \mathbf{T}_2^\top \mathbf{T}_2$. So for $p \geq 1$,

$$\|\mathbf{T}\|_{2p}^2 = \|\mathbf{T}^\top \mathbf{T}\|_p = \|\mathbf{T}_1^\top \mathbf{T}_1 + \mathbf{T}_2^\top \mathbf{T}_2\|_p \leq \|\mathbf{T}_1^\top \mathbf{T}_1\|_p + \|\mathbf{T}_2^\top \mathbf{T}_2\|_p = \|\mathbf{T}_1\|_{2p}^2 + \|\mathbf{T}_2\|_{2p}^2,$$

which has proved the first part.

For the second part, note that when $q = 1$, the inequality holds by triangle inequality. Next we show the inequality holds when $1 < q \leq 2$. Let $\mathbf{X} = [(\mathbf{T}_1^\top \mathbf{T}_1)^{1/2} \quad (\mathbf{T}_2^\top \mathbf{T}_2)^{1/2}]$. For any $p \geq 1$ we have

$$\|\mathbf{T}_1^\top \mathbf{T}_1 + \mathbf{T}_2^\top \mathbf{T}_2\|_p^p = \|\mathbf{X}\mathbf{X}^\top\|_p^p = \|\mathbf{X}^\top \mathbf{X}\|_p^p \stackrel{(a)}{\geq} \|\mathbf{T}_1^\top \mathbf{T}_1\|_p^p + \|\mathbf{T}_2^\top \mathbf{T}_2\|_p^p, \quad (8)$$

where (a) is because the norm of the diagonal part of a matrix is no greater than the norm of the whole matrix [29]. So we have

$$\|\mathbf{T}\|_{2p}^{2p} = \|\mathbf{T}^\top \mathbf{T}\|_p^p = \|\mathbf{T}_1^\top \mathbf{T}_1 + \mathbf{T}_2^\top \mathbf{T}_2\|_p^p \stackrel{(8)}{\geq} \|\mathbf{T}_1^\top \mathbf{T}_1\|_p^p + \|\mathbf{T}_2^\top \mathbf{T}_2\|_p^p = \|\mathbf{T}_1\|_{2p}^{2p} + \|\mathbf{T}_2\|_{2p}^{2p}. \quad (9)$$

Define subspaces

$$\mathcal{T} = \{\bar{\mathbf{T}} : \bar{\mathbf{T}} = \bar{\mathbf{T}}_1 + \bar{\mathbf{T}}_2, \bar{\mathbf{T}}_1 = P_{\mathbf{U}}\bar{\mathbf{A}}, \bar{\mathbf{T}}_2 = P_{\mathbf{U}^\perp}\bar{\mathbf{B}} \text{ and } \bar{\mathbf{A}}, \bar{\mathbf{B}} \in \mathbb{R}^{m \times n}\},$$

$$\mathcal{T}' = \left\{ \bar{\mathbf{T}}' : \bar{\mathbf{T}}' = \begin{bmatrix} \bar{\mathbf{T}}'_1 & \mathbf{0} \\ \mathbf{0} & \bar{\mathbf{T}}'_2 \end{bmatrix}, \bar{\mathbf{T}}'_1 = P_{\mathbf{U}}\bar{\mathbf{A}}, \bar{\mathbf{T}}'_2 = P_{\mathbf{U}^\perp}\bar{\mathbf{B}} \text{ and } \bar{\mathbf{A}}, \bar{\mathbf{B}} \in \mathbb{R}^{m \times n} \right\}.$$

Consider the linear map $\mathcal{A} : \bar{\mathbf{T}} \in \mathcal{T} \rightarrow \begin{bmatrix} P_{\mathbf{U}}\bar{\mathbf{T}} & \mathbf{0} \\ \mathbf{0} & P_{\mathbf{U}^\perp}\bar{\mathbf{T}} \end{bmatrix} \in \mathcal{T}'$. We can verify the adjoint map of $\mathcal{A}$, $\mathcal{A}^*$, satisfies $\mathcal{A}^* \left(\begin{bmatrix} \bar{\mathbf{T}}'_1 & \mathbf{0} \\ \mathbf{0} & \bar{\mathbf{T}}'_2 \end{bmatrix} \right) = \bar{\mathbf{T}}_1 + \bar{\mathbf{T}}_2$ as for any $\bar{\mathbf{T}} \in \mathcal{T}$ and $\bar{\mathbf{T}}' = \begin{bmatrix} \bar{\mathbf{T}}'_1 & \mathbf{0} \\ \mathbf{0} & \bar{\mathbf{T}}'_2 \end{bmatrix} \in \mathcal{T}'$, we have $\langle \mathcal{A}(\bar{\mathbf{T}}), \bar{\mathbf{T}}' \rangle = \langle P_{\mathbf{U}}\bar{\mathbf{T}}, \bar{\mathbf{T}}'_1 \rangle + \langle P_{\mathbf{U}^\perp}\bar{\mathbf{T}}, \bar{\mathbf{T}}'_2 \rangle = \langle \bar{\mathbf{T}}, \mathcal{A}^*(\bar{\mathbf{T}}') \rangle$. From (9) we have shown $\|\mathcal{A}(\mathbf{T})\|_p \leq \|\mathbf{T}\|_p$, i.e., $\mathcal{A}$ is contractive with respect to $\|\cdot\|_p$ for $p \geq 2$. Set $1 \leq p < \infty$ such that $1/q + 1/p = 1$. Since $\|\cdot\|_p$ and $\|\cdot\|_q$ are dual norms, we have for $1 < q \leq 2$:

$$\begin{aligned} \|\mathbf{T}\|_q &= \left\| \mathcal{A}^* \left(\begin{bmatrix} \mathbf{T}_1 & \mathbf{0} \\ \mathbf{0} & \mathbf{T}_2 \end{bmatrix} \right) \right\|_q = \sup_{\mathbf{X} : \|\mathbf{X}\|_p \leq 1} \left\langle \mathcal{A}^* \left(\begin{bmatrix} \mathbf{T}_1 & \mathbf{0} \\ \mathbf{0} & \mathbf{T}_2 \end{bmatrix} \right), \mathbf{X} \right\rangle \\ &= \sup_{\mathbf{X} : \|\mathbf{X}\|_p \leq 1} \left\langle \begin{bmatrix} \mathbf{T}_1 & \mathbf{0} \\ \mathbf{0} & \mathbf{T}_2 \end{bmatrix}, \mathcal{A}(\mathbf{X}) \right\rangle \\ &\stackrel{(a)}{\leq} \sup_{\mathbf{X} : \|\mathbf{X}\|_p \leq 1} \left\| \begin{bmatrix} \mathbf{T}_1 & \mathbf{0} \\ \mathbf{0} & \mathbf{T}_2 \end{bmatrix} \right\|_q \|\mathcal{A}(\mathbf{X})\|_p \\ &\leq (\|\mathbf{T}_1\|_q^q + \|\mathbf{T}_2\|_q^q)^{1/q}. \end{aligned}$$

Here (a) is an application of [18, Lemma II.3.4] and Hölder's inequality. This shows $\|\mathbf{T}\|_q^q \leq \|\mathbf{T}_1\|_q^q + \|\mathbf{T}_2\|_q^q$ and finishes the proof. $\square$

Next, we prove Theorem 1 based on Theorem 2 and Lemma 1.

Proof of Theorem 1. For $1 \leq q < \infty$, since $\hat{\mathbf{A}} = \mathbf{B}_{\max(r)}$ and $\hat{\mathbf{U}}$ is composed of the first r left singular vectors of $\mathbf{B}$, we have $\hat{\mathbf{A}} = P_{\hat{\mathbf{U}}}\mathbf{B}$ and

$$\begin{aligned} \|\hat{\mathbf{A}} - \mathbf{A}\|_q &= \|P_{\hat{\mathbf{U}}}\mathbf{B} - P_{\hat{\mathbf{U}}}\mathbf{A} - P_{\hat{\mathbf{U}}^\perp}\mathbf{A}\|_q = \|P_{\hat{\mathbf{U}}}\mathbf{Z} - P_{\hat{\mathbf{U}}^\perp}\mathbf{A}\|_q \\ &\stackrel{(a)}{\leq} \begin{cases} \left(\|P_{\hat{\mathbf{U}}}\mathbf{Z}\|_q^q + \|P_{\hat{\mathbf{U}}^\perp}\mathbf{A}\|_q^q \right)^{1/q}, & 1 \leq q \leq 2; \\ \left(\|P_{\hat{\mathbf{U}}}\mathbf{Z}\|_q^2 + \|P_{\hat{\mathbf{U}}^\perp}\mathbf{A}\|_q^2 \right)^{1/2}, & 2 \leq q < \infty \end{cases} \\ &\stackrel{(b)}{\leq} \begin{cases} (2^q + 1)^{1/q} \|\mathbf{Z}_{\max(r)}\|_q, & 1 \leq q \leq 2; \\ \sqrt{5} \|\mathbf{Z}_{\max(r)}\|_q, & 2 \leq q < \infty. \end{cases} \end{aligned}$$

Here, (a) is due to Lemma 1 and (b) is due to Theorem 2. For $q = \infty$,

$$\|\hat{\mathbf{A}} - \mathbf{A}\| \leq \|\hat{\mathbf{A}} - \mathbf{B}\| + \|\mathbf{A} - \mathbf{B}\| \stackrel{(a)}{\leq} 2\|\mathbf{A} - \mathbf{B}\| \leq 2\|\mathbf{Z}\| = 2\|\mathbf{Z}_{\max(r)}\|.$$

Here (a) comes from the fact that $\hat{\mathbf{A}}$ is the best rank- r approximation of $\mathbf{B}$. $\square$

As discussed in Section 1, one can derive the matrix estimation error bounds relying on $\|\mathbf{Z}\|_q$ or $r^{1/q}\|\mathbf{Z}\|$ via the existing perturbation theory in the literature. The following example illustrates that our result can be much sharper when the singular values of $\mathbf{Z}$ have some polynomial decay.

Example 1. Suppose $\mathbf{Z}$ satisfies that $\sigma_k(\mathbf{Z}) = k^{-1/q}$ for $q > 1$. Then

$$\|\mathbf{Z}_{\max(r)}\|_q = \left(\sum_{k=1}^r k^{-1} \right)^{1/q} \approx (1 + \log r)^{1/q},$$

which can be much smaller than

$$\|\mathbf{Z}\|_q = \left(\sum_{k=1}^{m \wedge n} k^{-1} \right)^{1/q} \approx (1 + \log(m \wedge n))^{1/q}, \quad r^{1/q}\|\mathbf{Z}\| = r^{1/q}.$$

2.1. Proof of Theorem 2

In this section, we focus on the proof of Theorem 2. We first introduce several additional lemmas on the properties of matrix singular values and norms, then present the proof of Theorem 2 and discussions. Recall that a matrix norm $\|\cdot\|$ is unitarily invariant if $\|\mathbf{A}\| = \|\mathbf{UAV}\|$ for any matrix $\mathbf{A}$ and orthogonal matrices $\mathbf{U}, \mathbf{V}$. Define $\Phi_q(\mathbf{x}) := \max_{1 \leq i_1 < \dots < i_r \leq n} \left(\sum_{j=1}^r |x_{i_j}|^q \right)^{1/q}$ with $q \geq 1$ for any $\mathbf{x} \in \mathbb{R}^n$. We have the following Lemmas for $\|(\cdot)_{\max(r)}\|_q$.

Lemma 2. Suppose $q \geq 1$. Then $\Phi_q(\cdot)$ is a symmetric gauge function and $\|(\cdot)_{\max(r)}\|_q$ is a unitarily invariant matrix norm.

Proof. First, $\|(\cdot)_{\max(r)}\|_q$ is a unitarily invariant matrix norm follows by von Neumann's Theory [18, Theorem II.3.6] if we can show $\Phi_q(\cdot)$ is a symmetric gauge function. Recall the definition of symmetric gauge function from Section 1.1, it is easy to see we just need to show $\Phi_q(\mathbf{x} + \mathbf{y}) \leq \Phi_q(\mathbf{x}) + \Phi_q(\mathbf{y})$ for any $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$. To show this, for any $\mathbf{x} \in \mathbb{R}^n$ and $\mathbf{z} \in \mathbb{R}^n$ with $\|\mathbf{z}\|_0 \leq r$ (here $\|\mathbf{z}\|_0$ denotes the number of non-zero entries in $\mathbf{z}$), by Hölder's inequality and the definition of $\Phi_q(\cdot)$, we have $\langle \mathbf{x}, \mathbf{z} \rangle \leq \|\mathbf{z}\|_p \Phi_q(\mathbf{x})$ for p such that $1/p + 1/q = 1$. Moreover, the equality is achieved when the equality condition in Hölder's inequality is satisfied. So

$$\Phi_q(\mathbf{x}) = \sup_{\mathbf{z}: \|\mathbf{z}\|_p \leq 1, \|\mathbf{z}\|_0 \leq r} \langle \mathbf{x}, \mathbf{z} \rangle. \quad (10)$$

Thus, we have

$$\begin{aligned}\Phi_q(\mathbf{x} + \mathbf{y}) &\stackrel{(10)}{=} \sup_{\mathbf{z}: \|\mathbf{z}\|_p \leq 1, \|\mathbf{z}\|_0 \leq r} \langle \mathbf{x} + \mathbf{y}, \mathbf{z} \rangle \leq \sup_{\mathbf{z}: \|\mathbf{z}\|_p \leq 1, \|\mathbf{z}\|_0 \leq r} \langle \mathbf{x}, \mathbf{z} \rangle + \sup_{\mathbf{z}: \|\mathbf{z}\|_p \leq 1, \|\mathbf{z}\|_0 \leq r} \langle \mathbf{y}, \mathbf{z} \rangle \\ &\stackrel{(10)}{=} \Phi_q(\mathbf{x}) + \Phi_q(\mathbf{y}).\end{aligned}$$

This shows $\Phi_q(\cdot)$ is a symmetric gauge function and finishes the proof of this lemma. $\square$

Next, the following lemma introduces a dual characterization of the truncated matrix Schatten- q norm.

Lemma 3 (Dual representation of Truncated Schatten- q norm). *Let $\mathbf{X}$ be a m -by- n real-valued matrix. For any non-negative integer $r \leq m \wedge n$, $q \geq 1$ and $1/p + 1/q = 1$, we have*

$$\|\mathbf{X}_{\max(r)}\|_p = \sup_{\|\mathbf{B}\|_q \leq 1, \text{rank}(\mathbf{B}) \leq r} \langle \mathbf{B}, \mathbf{X} \rangle. \quad (11)$$

If $\text{rank}(\mathbf{X}) \leq r$, then

$$\|\mathbf{X}\|_p = \sup_{\|\mathbf{B}\|_q \leq 1, \text{rank}(\mathbf{B}) \leq r} \langle \mathbf{B}, \mathbf{X} \rangle. \quad (12)$$

Proof. We first prove (11). Since $\Phi_q(\mathbf{x})$ is a symmetric gauge function as we have shown in Lemma 2 and its dual is $\Phi_p(\cdot)$ with $1/p + 1/q = 1$, (11) follows from [18, Lemma II.3.5]. Finally, (12) is a special case of (11). $\square$

Lemma 4. *Given matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$ and any non-negative integer $k \leq m \wedge n$, for any matrix $\mathbf{M}$ with $\text{rank}(\mathbf{M}) \leq r$, we have*

$$\left\| (\mathbf{A}_{-\max(r)})_{\max(k)} \right\|_q \leq \|(\mathbf{A} - \mathbf{M})_{\max(k)}\|_q.$$

The equality is achieved when $\mathbf{M} = \mathbf{A}_{\max(r)}$.

Proof. By the well-known Eckart-Young-Mirsky Theorem [30–32], the truncated SVD achieves the best low-rank matrix approximation in any unitarily invariant norm. This lemma follows from the Eckart-Young-Mirsky Theorem and the fact that $\|(\cdot)_{\max(k)}\|_q$ is a unitarily invariant matrix norm (Lemma 2). $\square$

Now we are in position to prove Theorem 2.

Proof of Theorem 2. We only study $\|P_{\hat{\mathbf{U}}_\perp} \mathbf{A}\|_q$ since the proof of the upper bound of $\|\mathbf{A}P_{\hat{\mathbf{V}}_\perp}\|_q$ follows by symmetry. Denote $\sum_{k=1}^r \sigma_k(\mathbf{A}) \mathbf{u}_k \mathbf{v}_k^\top$ as a singular value decomposition of $\mathbf{A}$. Since $\text{rank}(P_{\hat{\mathbf{U}}_\perp} \mathbf{A}) \leq \text{rank}(\mathbf{A}) = r$, for $p \geq 1$ satisfying $1/p + 1/q = 1$, we have

$$\begin{aligned}
\|P_{\hat{\mathbf{U}}^\perp} \mathbf{A}\|_q &\stackrel{(a)}{=} \sup_{\|\mathbf{X}\|_p \leq 1, \text{rank}(\mathbf{X}) \leq r} \langle P_{\hat{\mathbf{U}}^\perp} \mathbf{A}, \mathbf{X} \rangle \\
&= \sup_{\|\mathbf{X}\|_p \leq 1, \text{rank}(\mathbf{X}) \leq r} \langle P_{\hat{\mathbf{U}}^\perp} (\mathbf{A} + \mathbf{Z}) - P_{\hat{\mathbf{U}}^\perp} \mathbf{Z}, \mathbf{X} \rangle \\
&\leq \sup_{\|\mathbf{X}\|_p \leq 1, \text{rank}(\mathbf{X}) \leq r} \langle P_{\hat{\mathbf{U}}^\perp} (\mathbf{A} + \mathbf{Z}), \mathbf{X} \rangle + \sup_{\|\mathbf{X}\|_p \leq 1, \text{rank}(\mathbf{X}) \leq r} \langle P_{\hat{\mathbf{U}}^\perp} \mathbf{Z}, \mathbf{X} \rangle \\
&\stackrel{(b)}{\leq} \sup_{\|\mathbf{X}\|_p \leq 1, \text{rank}(\mathbf{X}) \leq r} \|\mathbf{X}\|_p \left\| \left(P_{\hat{\mathbf{U}}^\perp} (\mathbf{A} + \mathbf{Z}) \right)_{\max(r)} \right\|_q \\
&\quad + \sup_{\|\mathbf{X}\|_p \leq 1, \text{rank}(\mathbf{X}) \leq r} \|\mathbf{X}\|_p \left\| \left(P_{\hat{\mathbf{U}}^\perp} \mathbf{Z} \right)_{\max(r)} \right\|_q \\
&\stackrel{(c)}{\leq} \min_{\text{rank}(\mathbf{M}) \leq r} \left\| (\mathbf{A} + \mathbf{Z} - \mathbf{M})_{\max(r)} \right\|_q + \left\| \left(P_{\hat{\mathbf{U}}^\perp} \mathbf{Z} \right)_{\max(r)} \right\|_q \\
&\leq \left\| (\mathbf{A} + \mathbf{Z} - P_{\mathbf{U}} (\mathbf{A} + \mathbf{Z}))_{\max(r)} \right\|_q + \left\| \left(P_{\hat{\mathbf{U}}^\perp} \mathbf{Z} \right)_{\max(r)} \right\|_q \\
&\leq \left\| (P_{\mathbf{U}^\perp} \mathbf{Z})_{\max(r)} \right\|_q + \left\| (P_{\hat{\mathbf{U}}^\perp} \mathbf{Z})_{\max(r)} \right\|_q \\
&\leq \|\mathbf{Z}_{\max(r)}\|_q + \|\mathbf{Z}_{\max(r)}\|_q \leq 2\|\mathbf{Z}_{\max(r)}\|_q.
\end{aligned} \tag{13}$$

Here (a), (b) are due to Lemma 3 and (c) is due to Lemma 4. $\square$

We make several remarks on Theorem 2.

First, Theorem 2 may not be simply implied by the classic results. For example, the classic Wedin's $\sin \Theta$ Theorem [3],

$$\max \left\{ \|\sin \Theta(\mathbf{U}, \hat{\mathbf{U}})\|_q, \|\sin \Theta(\mathbf{V}, \hat{\mathbf{V}})\|_q \right\} \leq \frac{\max\{\|\mathbf{Z}\hat{\mathbf{V}}\|_q, \|\hat{\mathbf{U}}^\top \mathbf{Z}\|_q\}}{\sigma_r(\mathbf{B})}, \tag{14}$$

yields

$$\begin{aligned}
\|P_{\hat{\mathbf{U}}^\perp} \mathbf{A}\|_q &= \|\hat{\mathbf{U}}^\top \mathbf{U} \Sigma_1 \mathbf{V}^\top\|_q \stackrel{(a)}{\leq} \|\hat{\mathbf{U}}^\top \mathbf{U}\|_q \sigma_1(\mathbf{A}) = \|\sin \Theta(\mathbf{U}, \hat{\mathbf{U}})\|_q \sigma_1(\mathbf{A}) \\
&\leq \max \left\{ \|\mathbf{Z}\hat{\mathbf{V}}\|_q, \|\hat{\mathbf{U}}^\top \mathbf{Z}\|_q \right\} \frac{\sigma_1(\mathbf{A})}{\sigma_r(\mathbf{B})},
\end{aligned} \tag{15}$$

here (a) is by [18, Theorem II.3.9].

This bound (15) can be less sharp or practical for its dependency on $\sigma_1(\mathbf{A})/\sigma_r(\mathbf{B})$. As pointed out by [28], the spectrum of large matrix datasets arising from applications often decays fast. If the singular values of $\mathbf{A}, \mathbf{B}$ decay fast, $\sigma_1(\mathbf{A})/\sigma_r(\mathbf{B}) \gg 1$ and (15) can be loose. In contrast, our bound (7) in Theorem 2 is free of any ratio of singular values, which can be a significant advantage in practice. We will further illustrate the difference between (7) and (15) by simulation in Section 5.2.

Second, it is noteworthy by (13) in the proof of Theorem 2, we have actually proved

$$\begin{aligned}\|P_{\hat{\mathbf{U}}^\perp} \mathbf{A}\|_q &\leq \left\| \left(P_{\hat{\mathbf{U}}^\perp} \mathbf{Z} \right)_{\max(r)} \right\|_q + \left\| (P_{\mathbf{U}^\perp} \mathbf{Z})_{\max(r)} \right\|_q \\ \| \mathbf{A} P_{\hat{\mathbf{V}}^\perp} \|_q &\leq \left\| \left(\mathbf{Z} P_{\hat{\mathbf{V}}^\perp} \right)_{\max(r)} \right\|_q + \left\| (\mathbf{Z} P_{\mathbf{V}^\perp})_{\max(r)} \right\|_q\end{aligned}\quad (16)$$

under the setting of Theorem 2. The bound (16) can be better than the one in Theorem 2 in some scenarios. For example, when $\mathbf{Z}$ is (or is close to) $\mathbf{U} \Sigma_{\mathbf{Z}} \mathbf{V}^\top$ for some r -by- r matrix $\Sigma_{\mathbf{Z}}$, the bound in (16) is smaller than $\|\mathbf{Z}_{\max(r)}\|_q$. On the other hand, the proposed bound in Theorem 2 is strong enough for proving Theorem 1, does not involve $P_{\hat{\mathbf{U}}^\perp}$ or $P_{\hat{\mathbf{V}}^\perp}$, and can be more convenient to use.

3. Lower bounds

The following Theorem 3 shows that the error upper bound for the rank- r truncated SVD estimator $\hat{\mathbf{A}}$ in Theorem 1 is sharp.

Theorem 3. *For any $\varepsilon > 0$ and $q \geq 1$, there exist $\mathbf{A}$, $\mathbf{B}$, and $\mathbf{Z} \neq 0$ such that $\text{rank}(\mathbf{A}) = r$, $\mathbf{B} = \mathbf{A} + \mathbf{Z}$, and*

$$\|\hat{\mathbf{A}} - \mathbf{A}\|_q > ((2^q + 1)^{1/q} - \varepsilon) \|\mathbf{Z}_{\max(r)}\|_q.$$

Proof. Without loss of generality we assume $0 < \varepsilon < 1$. We choose a value $\eta \in (0, \frac{(2^q+1)^{1/q}}{(2^q+1)^{1/q}-\varepsilon} - 1)$. Define

$$\mathbf{A} = \begin{bmatrix} 2\mathbf{I}_r & \mathbf{0}_{r \times r} & \mathbf{0} \\ \mathbf{0}_{r \times r} & \mathbf{0}_{r \times r} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} \end{bmatrix}, \quad \mathbf{Z} = \begin{bmatrix} -(1+\eta)\mathbf{I}_r & \mathbf{0}_{r \times r} & \mathbf{0} \\ \mathbf{0}_{r \times r} & \mathbf{I}_r & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} \end{bmatrix},$$

and

$$\mathbf{B} = \begin{bmatrix} (1-\eta)\mathbf{I}_r & \mathbf{0}_{r \times r} & \mathbf{0} \\ \mathbf{0}_{r \times r} & \mathbf{I}_r & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} \end{bmatrix}.$$

Then,

$$\|\hat{\mathbf{A}} - \mathbf{A}\|_q = \left\| \begin{bmatrix} -2\mathbf{I}_r & \mathbf{0}_{r \times r} \\ \mathbf{0}_{r \times r} & \mathbf{I}_r \end{bmatrix} \right\|_q = (2^q r + r)^{1/q}, \quad \|\mathbf{Z}_{\max(r)}\|_q = (1+\eta)r^{1/q}.$$

We thus have

$$\|\hat{\mathbf{A}} - \mathbf{A}\|_q > ((2^q + 1)^{1/q} - \varepsilon) \|\mathbf{Z}_{\max(r)}\|_q. \quad \square$$

Theorems 1 and 3 together imply that the constants in (4) are not improvable when $1 \leq q \leq 2$ and $q = \infty$. For $2 < p < \infty$, it would be an interesting future work to close the gap between the upper bound $(\sqrt{5}\|\mathbf{Z}_{\max(r)}\|_q)$ and the lower bound $((2^q + 1)^{1/q}\|\mathbf{Z}_{\max(r)}\|_q)$.

Apart from checking the sharpness of the upper bound (4), another natural question is, whether the rank- r truncated SVD estimator is an optimal estimator in estimating $\mathbf{A}$. To answer this question, we consider the minimax estimation error lower bound among all possible data-dependent procedures $\check{\mathbf{A}} = \check{\mathbf{A}}(\mathbf{B})$ (i.e., $\check{\mathbf{A}}$ is a deterministic or random function of matrix $\mathbf{B}$). We specifically focus on the following class of $(\tilde{\mathbf{A}}, \tilde{\mathbf{Z}}, \tilde{\mathbf{B}})$ triplets:

$$\mathcal{F}_r(\xi) = \left\{ (\tilde{\mathbf{A}}, \tilde{\mathbf{Z}}, \tilde{\mathbf{B}}) : \tilde{\mathbf{B}} = \tilde{\mathbf{A}} + \tilde{\mathbf{Z}}, \text{rank}(\tilde{\mathbf{A}}) = r, \|\tilde{\mathbf{Z}}_{\max(r)}\|_q \leq \xi \right\}.$$

Here, ξ corresponds to $\|\mathbf{Z}_{\max(r)}\|_q$ in the context of Theorem 1.

Theorem 4 (*Schatten- q minimax lower bound*). *For the low-rank perturbation model, if $m \wedge n \geq 2r$, then, for any $q \geq 1$, we have*

$$\inf_{\check{\mathbf{A}}} \sup_{(\tilde{\mathbf{A}}, \tilde{\mathbf{Z}}, \tilde{\mathbf{B}}) \in \mathcal{F}_r(\xi)} \|\check{\mathbf{A}} - \tilde{\mathbf{A}}\|_q \geq 2^{1/q-1}\xi.$$

Here the infimum is taken over all the estimation procedures.

Proof. The proof is done by construction. We construct

$$\mathbf{Z}_1 = \begin{pmatrix} \mathbf{0}_{r \times r} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \frac{\xi}{r^{1/q}}\mathbf{I}_r & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} \end{pmatrix}, \quad \bar{\mathbf{A}}_1 = \begin{pmatrix} \frac{\xi}{r^{1/q}}\mathbf{I}_r & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0}_{r \times r} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} \end{pmatrix},$$

and

$$\mathbf{Z}_2 = \begin{pmatrix} \frac{\xi}{r^{1/q}}\mathbf{I}_r & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0}_{r \times r} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} \end{pmatrix}, \quad \bar{\mathbf{A}}_2 = \begin{pmatrix} \mathbf{0}_{r \times r} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \frac{\xi}{r^{1/q}}\mathbf{I}_r & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} \end{pmatrix}.$$

By the construction above, we have $\|(\mathbf{Z}_1)_{\max(r)}\|_q = \xi$, $\|(\mathbf{Z}_2)_{\max(r)}\|_q = \xi$, and $\bar{\mathbf{A}}_1 + \mathbf{Z}_1 = \bar{\mathbf{A}}_2 + \mathbf{Z}_2$. So

$$\begin{aligned} & \inf_{\check{\mathbf{A}}} \sup_{(\tilde{\mathbf{A}}, \tilde{\mathbf{Z}}, \tilde{\mathbf{B}}) \in \mathcal{F}_r(\xi)} \|\check{\mathbf{A}} - \tilde{\mathbf{A}}\|_q \geq \inf_{\check{\mathbf{A}}} \left(\max \left\{ \|\check{\mathbf{A}} - \bar{\mathbf{A}}_1\|_q, \|\check{\mathbf{A}} - \bar{\mathbf{A}}_2\|_q \right\} \right) \\ & \geq \frac{1}{2} \inf_{\check{\mathbf{A}}} \left(\|\check{\mathbf{A}} - \bar{\mathbf{A}}_1\|_q + \|\check{\mathbf{A}} - \bar{\mathbf{A}}_2\|_q \right) \geq \frac{1}{2} \|\bar{\mathbf{A}}_1 - \bar{\mathbf{A}}_2\|_q = 2^{1/q-1}\xi. \quad \square \end{aligned}$$

Combining Theorems 1 and 4, we conclude that the truncated SVD $\hat{\mathbf{A}}$ achieves the optimal rate of low-rank matrix estimation error among all possible procedures $\check{\mathbf{A}}$ in the class of $\mathcal{F}_r(\xi)$.

4. Subspace perturbation bounds

In this section, we apply the perturbation projection error bound established in Theorem 2 to derive a user-friendly subspace (singular vectors) perturbation bound.

Theorem 5. *Consider the same perturbation setting as in Theorem 1. For any $q \geq 1$, we have*

$$\max \left\{ \|\sin \Theta(\hat{\mathbf{U}}, \mathbf{U})\|_q, \|\sin \Theta(\hat{\mathbf{V}}, \mathbf{V})\|_q \right\} \leq \frac{2\|\mathbf{Z}_{\max(r)}\|_q}{\sigma_r(\mathbf{A})}.$$

Proof. By Theorem 2, we have

$$\|P_{\hat{\mathbf{U}}^\perp} \mathbf{A}\|_q \leq 2\|\mathbf{Z}_{\max(r)}\|_q.$$

Since the left singular subspace of $\mathbf{A}$ is $\mathbf{U}$, we have $\mathbf{U}\mathbf{U}^\top \mathbf{A} = P_{\mathbf{U}} \mathbf{A} = \mathbf{A}$. Then

$$\|\sin \Theta(\hat{\mathbf{U}}, \mathbf{U})\|_q = \|\hat{\mathbf{U}}^\top \mathbf{U}\|_q \stackrel{(a)}{\leq} \frac{\|\hat{\mathbf{U}}^\top \mathbf{U}\mathbf{U}^\top \mathbf{A}\|_q}{\sigma_r(\mathbf{U}^\top \mathbf{A})} = \frac{\|P_{\hat{\mathbf{U}}^\perp} \mathbf{A}\|_q}{\sigma_r(\mathbf{A})} \leq \frac{2\|\mathbf{Z}_{\max(r)}\|_q}{\sigma_r(\mathbf{A})},$$

here (a) is by [18, Theorem II.3.9]. $\square$

We note that several similar bounds are developed towards the applications in statistics and machine learning in the past few years, for example, [33, Corollary 4.1], [15, Theorem 2], and [34, Lemma 5.1]. When the matrix is positive semidefinite, these results yield

$$\left\| \sin \Theta(\hat{\mathbf{U}}, \mathbf{U}) \right\|_F \leq \frac{\sqrt{2}\|\mathbf{Z}\|_F}{\sigma_r(\mathbf{A})}, \quad ([33, \text{Corollary 4.1}], \quad (17)$$

$$\left\| \sin \Theta(\hat{\mathbf{U}}, \mathbf{U}) \right\|_F \leq \frac{2 \min\{r^{1/2}\|\mathbf{Z}\|, \|\mathbf{Z}\|_F\}}{\sigma_r(\mathbf{A})} \quad [15, \text{Theorem 2}], [34, \text{Lemma 5.1}]. \quad (18)$$

When $\mathbf{A}, \mathbf{Z}, \mathbf{B}$ are asymmetric, [15] also proved

$$\left\| \sin \Theta(\hat{\mathbf{U}}, \mathbf{U}) \right\|_F \leq \frac{2(2\|\mathbf{A}\| + \|\mathbf{Z}\|) \min\{r^{1/2}\|\mathbf{Z}\|, \|\mathbf{Z}\|_F\}}{\sigma_r^2(\mathbf{A})} \quad [15, \text{Theorem 3}]. \quad (19)$$

The perturbation bounds (17), (18), (19), along with Theorem 5 in this paper, are “user friendly” as they do not involve $\hat{\mathbf{U}}, \hat{\mathbf{V}}$ or $\mathbf{B}$ in contrast to the classical Wedin’s $\sin \Theta$ bound (14). This advantage facilitates the application of these perturbations to many

settings when $\mathbf{A}$ and $\mathbf{Z}$ are the given arguments: one no longer needs to further bound $\|\mathbf{Z}\hat{\mathbf{V}}\|_q$, $\|\hat{\mathbf{U}}^\top \mathbf{Z}\|_q$. The “user friendly” advantage is also important in many settings as the denominator of (14), $\sigma_r(\mathbf{B})$, depends highly on the perturbation $\mathbf{Z}$ and can be rather small due to perturbation [15]. In addition, our new result in Theorem 5 has a better dependence on both $\mathbf{Z}$ and $\sigma_r(\mathbf{A})$ than (17), (18), (19) because

$$\|\mathbf{Z}_{\max(r)}\|_F \leq \min \left\{ r^{1/2} \|\mathbf{Z}\|, \|\mathbf{Z}\|_F \right\},$$

while the opposite side of this inequality does not hold. Moreover, Theorem 5 covers the more general asymmetric matrices in Schatten- q $\sin\Theta$ norms for any $q \in [1, \infty]$.

5. Simulations

In this section, we provide numerical studies to support our theoretical results. We specifically compare the low-rank matrix estimation error bound (Theorem 1) and the matrix perturbation projection error bound (Theorem 2) in Section 2 with the results in previous literature. In each setting, we randomly generate a perturbation $\mathbf{Z} = \mathbf{u}\mathbf{v}^\top + \tilde{\mathbf{Z}}$, draw $\mathbf{A}$ by a to-be-specified scheme, and construct $\mathbf{B} = \mathbf{A} + \mathbf{Z}$. Here $\mathbf{u}, \mathbf{v}$ are randomly generated unit vectors and $\tilde{\mathbf{Z}}$ has i.i.d. $N(0, \sigma^2)$ entries. Throughout the simulation studies, we consider the Schatten-2 norm (i.e., Frobenius norm) as the error metric. Each simulation setting is repeated for 100 times and the average values are reported.

5.1. Numerical comparison of low-rank matrix estimation error bounds

We first compare the low-rank matrix estimation error bound $\|\hat{\mathbf{A}} - \mathbf{A}\|_q$ in Theorem 1 and the bounds in (2) and (3). We set $n \in \{100, 300\}$, $r \in \{4, 6, \dots, 16\}$, $\sigma = 0.02$, and generate $\mathbf{A} = \mathbf{U}\mathbf{\Sigma}_1\mathbf{V}^\top$, where $\mathbf{U} \in \mathbb{R}^{n \times r}$, $\mathbf{V} \in \mathbb{R}^{n \times r}$ are independently drawn from $\mathcal{O}_{n,r}$ uniformly at random; $\mathbf{\Sigma}_1$ is a diagonal matrix with singular values decaying polynomially as: $(\mathbf{\Sigma}_1)_{[i,i]} = \frac{10}{i}$, $1 \leq i \leq r$.

The evaluations of the upper bounds in Theorem 1, (2), (3), and the true value of $\|\hat{\mathbf{A}} - \mathbf{A}\|_F$ are given in Fig. 1. It shows that the upper bound in Theorem 1 is tighter than the upper bounds in (2), (3) in all settings. In addition, when n increases from 100 to 300, the upper bound of (2) significantly increases while the upper bound of Theorem 1 remains steady. This is because the upper bounds of (2) and Theorem 1 rely on $\|\mathbf{Z}\|_F$ and $\|\mathbf{Z}_{\max(r)}\|_F$, respectively.

5.2. Numerical comparison of matrix perturbation projection error bounds

Next, we compare the matrix perturbation projection error bound in Theorem 2 with the upper bound (15) derived from Wedin’s $\sin\Theta$ Theorem. We generate $\mathbf{B}, \mathbf{Z}$ in the same way as the previous simulation setting. When generating $\mathbf{\Sigma}_1$ in $\mathbf{A}$, apart from the




Fig. 1. Low-rank matrix estimation error bound (Theorem 1), upper bounds (2), (3) and the true value of $\|\hat{\mathbf{A}} - \mathbf{A}\|_F$.




Fig. 2. Matrix perturbation projection error upper bound (Theorem 2), upper bound via Wedin's $\sin\Theta$ Theorem (15), and the true value of $\|P_{\hat{\mathbf{U}}^\perp} \mathbf{A}\|_F$.

polynomial singular value decaying pattern considered in the last setting, we also consider the following exponential singular value decaying pattern: $(\Sigma_1)_{[i,i]} = 2^{5-i}$, $1 \leq i \leq r$.

The values of the upper bounds in Theorem 2 and (15), along with the true value of $\|P_{\hat{\mathbf{U}}^\perp} \mathbf{A}\|_q$, are presented in Fig. 2. We find the bound of Theorem 2 is much tighter than the bound in (15). As r increases or singular value decaying pattern becomes exponential, i.e., $\mathbf{A}$ becomes ill-conditioned, (15) becomes loose while Theorem 2 can still be sharp.

6. Discussions

In this paper, we prove a sharp upper bound for estimation error of rank- r truncated SVD ($\|\hat{\mathbf{A}} - \mathbf{A}\|_q$) under perturbation, and show its optimality in low-rank matrix estimation. The key technical tool we use is a novel matrix perturbation projection error bound for $\|P_{\hat{\mathbf{U}}^\perp} \mathbf{A}\|_q$. As a byproduct, we also provide a sharper user-friendly $\sin\Theta$ perturbation bound. The numerical studies demonstrate the advantages of these new results over the ones in the literature.

The main result of this paper is the upper bound in (4), which is sharper than ones directly derived from the literature (1), (2), (3). We also comment that (1) can be conveniently extended to the general case that $\mathbf{A}$ is approximately rank r [3, Eq.(4.4)]. It is interesting future work to study if a similar bound to (4) can be obtained for the general approximately low-rank $\mathbf{A}$.

Throughout the paper, we study the additive perturbations and it is a future work to extend the results to multiplicative perturbations [7,8]. Also for convenience of presentation, we focus on the real number field in this paper. It is interesting to extend the developed results to the field of complex numbers. The main technical work for such an extension includes complex versions of Lemma 2 and 3.

Apart from the widely studied perturbation theory on singular value decomposition, the perturbation theory for other problems, such as pseudo-inverses [26,27], least squares problems [27], orthogonal projection [27,35–37], rank-one perturbation [38], are also important topics. It would be interesting to explore whether the tools developed in this paper are useful in studying the perturbation theory for these problems.

Declaration of competing interest

No competing interest.

References

- [1] H. Weyl, Das asymptotische Verteilungsgesetz der Eigenwerte linearer partieller Differentialgleichungen (mit einer Anwendung auf die Theorie der Hohlraumstrahlung), *Math. Ann.* 71 (4) (1912) 441–479.
- [2] C. Davis, W.M. Kahan, The rotation of eigenvectors by a perturbation. III, *SIAM J. Numer. Anal.* 7 (1) (1970) 1–46.
- [3] P.-A. Wedin, Perturbation bounds in connection with singular value decomposition, *BIT Numer. Math.* 12 (1) (1972) 99–111.
- [4] R.J. Vaccaro, A second-order perturbation expansion for the SVD, *SIAM J. Matrix Anal. Appl.* 15 (2) (1994) 661–671.
- [5] Z. Xu, Perturbation analysis for subspace decomposition with applications in subspace-based algorithms, *IEEE Trans. Signal Process.* 50 (11) (2002) 2820–2830.
- [6] J. Liu, X. Liu, X. Ma, First-order perturbation analysis of singular vectors in singular value decomposition, *IEEE Trans. Signal Process.* 56 (7) (2008) 3044–3049.
- [7] R.-C. Li, Relative perturbation theory: I. Eigenvalue and singular value variations, *SIAM J. Matrix Anal. Appl.* 19 (4) (1998) 956–982.
- [8] R.-C. Li, Relative perturbation theory: II. Eigenspace and singular subspace variations, *SIAM J. Matrix Anal. Appl.* 20 (2) (1998) 471–492.
- [9] T. Londré, N.H. Rhee, A note on relative perturbation bounds, *SIAM J. Matrix Anal. Appl.* 21 (2) (2000) 357–361.
- [10] M. Stewart, Perturbation of the SVD in the presence of small singular values, *Linear Algebra Appl.* 419 (1) (2006) 53–77.
- [11] J. Demmel, W. Kahan, Accurate singular values of bidiagonal matrices, *SIAM J. Sci. Stat. Comput.* 11 (5) (1990) 873–912.
- [12] J. Barlow, J. Demmel, Computing accurate eigensystems of scaled diagonally dominant matrices, *SIAM J. Numer. Anal.* 27 (3) (1990) 762–791.
- [13] J. Demmel, K. Veselić, Jacobi’s method is more accurate than QR, *SIAM J. Matrix Anal. Appl.* 13 (4) (1992) 1204–1245.
- [14] Z. Drmac, K. Veselić, New fast and accurate Jacobi SVD algorithm. I, *SIAM J. Matrix Anal. Appl.* 29 (4) (2008) 1322–1342.

- [15] Y. Yu, T. Wang, R.J. Samworth, A useful variant of the Davis–Kahan theorem for statisticians, *Biometrika* 102 (2) (2015) 315–323.
- [16] T.T. Cai, A. Zhang, Rate-optimal perturbation bounds for singular subspaces with applications to high-dimensional statistics, *Ann. Appl. Stat.* 46 (1) (2018) 60–89.
- [17] J. Cape, M. Tang, C.E. Priebe, The two-to-infinity norm and singular subspace geometry with applications to high-dimensional statistics, *Ann. Appl. Stat.* 47 (5) (2019) 2405–2439.
- [18] G.W. Stewart, J.-G. Sun, *Matrix Perturbation Theory*, 1990.
- [19] I.C. Ipsen, An overview of relative $\sin\Theta$ theorems for invariant subspaces of complex matrices, *J. Comput. Appl. Math.* 123 (1–2) (2000) 131–153.
- [20] R. Bhatia, *Matrix Analysis*, vol. 169, Springer Science & Business Media, 2013.
- [21] M. Gavish, D.L. Donoho, The optimal hard threshold for singular values is $4/\sqrt{3}$, *IEEE Trans. Inf. Theory* 60 (8) (2014) 5040–5053.
- [22] D. Donoho, M. Gavish, Minimax risk of matrix denoising by singular value thresholding, *Ann. Stat.* 42 (6) (2014) 2413–2440.
- [23] D.W. Tufts, A.A. Shah, Estimation of a signal waveform from noisy data using low-rank approximation to a data matrix, *IEEE Trans. Signal Process.* 41 (4) (1993) 1716–1721.
- [24] I. Jolliffe, *Principal Component Analysis*, 2nd edition, Springer, New York, 2002.
- [25] D.F. Morrison, L.C. Marshall, H.L. Sahlin, *Multivariate Statistical Methods*, McGraw-Hill, New York, 1976.
- [26] P.-A. Wedin, Perturbation theory for pseudo-inverses, *BIT Numer. Math.* 13 (2) (1973) 217–232.
- [27] G.W. Stewart, On the perturbation of pseudo-inverses, projections and linear least squares problems, *SIAM Rev.* 19 (4) (1977) 634–662.
- [28] M. Udell, A. Townsend, Why are big data matrices approximately low rank?, *SIAM J. Math. Data Sci.* 1 (1) (2019) 144–160.
- [29] R. Bhatia, J.A. Holbrook, On the Clarkson-McCarthy inequalities, *Math. Ann.* 281 (1988) 7–12.
- [30] C. Eckart, G. Young, The approximation of one matrix by another of lower rank, *Psychometrika* 1 (3) (1936) 211–218.
- [31] L. Mirsky, Symmetric gauge functions and unitarily invariant norms, *Chin. Q. J. Math.* 11 (1) (1960) 50–59.
- [32] G.H. Golub, A. Hoffman, G.W. Stewart, A generalization of the Eckart-Young-Mirsky matrix approximation theorem, *Linear Algebra Appl.* 88 (1987) 317–327.
- [33] V.Q. Vu, J. Lei, Minimax sparse principal subspace estimation in high dimensions, *Ann. Stat.* 41 (6) (2013) 2905–2947.
- [34] J. Lei, A. Rinaldo, Consistency of spectral clustering in stochastic block models, *Ann. Stat.* 43 (1) (2015) 215–237.
- [35] X. Xu, On the perturbation of an l_2 -orthogonal projection, *J. Comput. Appl. Math.* 368 (2020) 112327.
- [36] R.D. Fierro, J.R. Bunch, Perturbation theory for orthogonal projection methods with applications to least squares and total least squares, *Linear Algebra Appl.* 234 (1996) 71–96.
- [37] Y.M. Chen, X.S. Chen, W. Li, On perturbation bounds for orthogonal projections, *Numer. Algorithms* 73 (2) (2016) 433–444.
- [38] L. Zhu, X. Peng, H. Liu, Rank-one perturbation bounds for singular values of arbitrary matrices, *Arch. Inequal. Appl.* 2019 (1) (2019) 138.