

EFFICIENT BOUNDS AND ESTIMATES FOR CANONICAL
ANGLES IN RANDOMIZED SUBSPACE APPROXIMATIONS*YIJUN DONG[†], PER-GUNNAR MARTINSSON[‡], AND YUJI NAKATSUKASA[§]

Abstract. Randomized subspace approximation with “matrix sketching” is an effective approach for constructing approximate partial singular value decompositions (SVDs) of large matrices. The performance of such techniques has been extensively analyzed, and very precise estimates on the distribution of the residual errors have been derived. However, our understanding of the accuracy of the computed singular vectors (measured in terms of the canonical angles between the spaces spanned by the exact and the computed singular vectors, respectively) remains relatively limited. In this work, we present *practical bounds and estimates for canonical angles* of randomized subspace approximation that *can be computed efficiently either a priori or a posteriori*, without assuming prior knowledge of the true singular subspaces. Under moderate oversampling in the randomized SVD, our prior probabilistic bounds are asymptotically tight and can be computed efficiently, while bringing a clear insight into the balance between oversampling and power iterations given a fixed budget on the number of matrix–vector multiplications. The numerical experiments demonstrate the empirical effectiveness of these canonical angle bounds and estimates on different matrices under various algorithmic choices for the randomized SVD.

Key words. randomized subspace iteration, singular vectors, canonical angles

MSC codes. 15A18, 15A42, 65F15, 68W20

DOI. 10.1137/23M1584733

1. Introduction. In light of the ubiquity of high-dimensional data in modern computation, dimension reduction tools like low-rank matrix approximations are becoming indispensable tools for managing large data sets. In general, the goal of low-rank matrix approximation is to identify bases of proper low-dimensional subspaces that well encapsulate the dominant components in the original column and row spaces. As one of the most well-established forms of matrix decompositions, the truncated singular value decomposition (SVD) is known to achieve the optimal low-rank approximation errors for any given ranks [9]. Moreover, the corresponding left and right leading singular subspaces can be broadly leveraged for problems like principal component analysis, canonical correlation analysis, spectral clustering [3], and leverage score sampling for matrix skeleton selection [8, 17, 7].

However, for large matrices, the computational cost of classical algorithms for computing the SVD (cf. [30, Lec. 31] or [12, sec. 8.6.3]) quickly becomes prohibitive. Fortunately, a randomization framework known as “matrix sketching” [14, 37] provides a simple yet effective remedy for this challenge by projecting large matrices to

*Received by the editors July 11, 2023; accepted for publication (in revised form) June 24, 2024; published electronically October 18, 2024.

<https://doi.org/10.1137/23M1584733>

Funding: The work of the authors was supported by the Office of Naval Research grant N00014-18-1-2354, by the National Science Foundation grants DMS-1952735 and DMS-2012606, and by the Department of Energy ASCR grant DE-SC0022251.

[†]Courant Institute, New York University, New York, NY 10012 USA (yd1319@nyu.edu).

[‡]Department of Mathematics & Oden Institute, University of Texas at Austin, Austin, TX 78712 USA (pgm@oden.utexas.edu).

[§]Mathematical Institute, University of Oxford, Oxford, OX2 6GG, UK (nakatsukasa@maths.ox.ac.uk).

random low-dimensional subspaces where the classical SVD algorithms can be executed efficiently.

Concretely, for an input matrix $\mathbf{A} \in \mathbb{C}^{m \times n}$ and a target rank $k \ll \min(m, n)$, the basic version of the randomized SVD [14, Alg. 4.1] starts by drawing a Gaussian random matrix $\Omega \in \mathbb{C}^{n \times l}$ for a sample size l that is slightly larger than k so that $k < l \ll \min(m, n)$. Then through a matrix-matrix multiplication $\mathbf{X} = \mathbf{A}\Omega$ with $O(mnl)$ complexity, the n -dimensional row space of \mathbf{A} is projected to a random l -dimensional subspace. With the low-dimensional range approximation \mathbf{X} , a rank- l randomized SVD $\hat{\mathbf{A}}_l = \hat{\mathbf{U}}_l \hat{\Sigma}_l \hat{\mathbf{V}}_l^*$ can be constructed efficiently by computing the QR and SVD of small matrices in $O((m+n)l^2)$ time. When the spectral decay in \mathbf{A} is slow, a few power iterations $\mathbf{X} = (\mathbf{A}\mathbf{A}^*)^q \mathbf{A}\Omega$ (usually $q = 1, 2$) can be incorporated to enhance the accuracy; cf. [14, Algorithms 4.3 and 4.4].

Let $\mathbf{A} = \mathbf{U}\Sigma\mathbf{V}^*$ ¹ denote the (unknown) full SVD of \mathbf{A} . In this work, we explore the alignment between the true leading rank- k singular subspaces $\mathbf{U}_k, \mathbf{V}_k$ and their respective rank- l approximations $\hat{\mathbf{U}}_l, \hat{\mathbf{V}}_l$ in terms of the canonical angles $\angle(\mathbf{U}_k, \hat{\mathbf{U}}_l)$ and $\angle(\mathbf{V}_k, \hat{\mathbf{V}}_l)$. We introduce prior statistical guarantees and unbiased estimates for these angles with respect to Σ , as well as posterior deterministic bounds with additional dependence on $\hat{\mathbf{A}}_l = \hat{\mathbf{U}}_l \hat{\Sigma}_l \hat{\mathbf{V}}_l^*$. In particular, we are interested in *practical bounds and estimates that can be computed efficiently*, without assuming prior knowledge of the true singular subspaces.

1.1. Our contributions.

Prior probabilistic bounds and estimates with insight on oversampling-power iteration balance. Evaluating the randomized SVD with a fixed budget on the number of matrix-vector multiplications, the computational resource can be leveraged in two ways—oversampling (characterized by $l - k$) and power iterations (characterized by q). A natural question is *how can we distribute the computation between oversampling and power iterations for better subspace approximations?*

Answers to this question are problem-dependent: when aiming to minimize the canonical angles between the true and approximated leading singular subspaces, the prior probabilistic bounds and estimates on the canonical angles provide primary insights. To be precise, with isotropic random subspace embeddings and sufficient oversampling, the accuracy of subspace approximations depends jointly on the spectra of the target matrices, oversampling, and the number of power iterations. In this work, we present a set of prior probabilistic bounds that precisely quantify the relative benefits of oversampling versus power iterations. Specifically, the canonical angle bounds in Theorem 1

- (i) provide statistical guarantees that are asymptotically tight under sufficient oversampling (i.e., $l = \Omega(k)$),
- (ii) unveil a clear balance between oversampling and power iterations for random subspace approximations with given spectra,
- (iii) can be evaluated in $O(\text{rank}(\mathbf{A}))$ time given access to the (true/estimated) spectra while providing valuable estimations for canonical angles in practice with moderate oversampling (e.g., $l \geq 1.6k$).

Further, inspired by the derivation of the prior probabilistic bounds, we propose unbiased estimates for the canonical angles with respect to given spectra that admit efficient evaluation and concentrate well empirically.

¹Here $\mathbf{U} \in \mathbb{C}^{m \times r}$, $\mathbf{V} \in \mathbb{C}^{n \times r}$, and $\Sigma \in \mathbb{C}^{r \times r}$. Σ is a diagonal matrix with positive nonincreasing diagonal entries, and $r \leq \min(m, n)$.

Posterior residual-based guarantees. Alongside the prior probabilistic bounds, we present two sets of posterior canonical angle bounds that hold in the deterministic sense and can be approximated efficiently based on the residuals and the spectrum of \mathbf{A} .

Numerical comparisons. With numerical comparisons among different canonical angle bounds on a variety of data matrices, we aim to explore the question on *how the spectral decay and different algorithmic choices of randomized SVD affect the empirical effectiveness of different canonical angle bounds*. Our numerical experiments suggest that, for matrices with subexponential spectral decay, the prior probabilistic bounds usually provide tighter (statistical) guarantees than the (deterministic) guarantees from the posterior residual-based bounds, especially with power iterations. By contrast, for matrices with exponential spectral decay, the posterior residual-based bounds can be as tight as the prior probabilistic bounds, especially with large oversampling. The code for numerical comparisons is available at https://github.com/dyjdongyijun/Randomized_Subspace_Approximation.

1.2. Related work. The randomized SVD algorithm (with power iterations) [19, 14] has been extensively analyzed as a low-rank approximation problem where the accuracy is usually measured in terms of residual norms, as well as the discrepancy between the approximated and true spectra [14, 13, 18, 20]. For instance, [14, Thm. 10.7, Thm. 10.8] shows that for a given target rank k (usually $k \ll \min(m, n)$ for the randomized acceleration to be useful), a small constant oversampling $l \geq 1.1k$ is sufficient to guarantee that the residual norm of the resulting rank- l approximation is close to the optimal rank- k approximation (i.e., the rank- k truncated SVD) error with high probability. Alternatively, [13] investigates the accuracy of the individual approximated singular values $\hat{\sigma}_i$ and provides upper and lower bounds for each $\hat{\sigma}_i$ with respect to the true singular value σ_i .

In addition to providing accurate low-rank approximations, the randomized SVD algorithm also produces estimates of the leading left and right singular subspaces corresponding to the top singular values. When coupled with power iterations [14, Algorithms 4.3 and 4.4], such randomized subspace approximations are commonly known as randomized power (subspace) iterations. Their accuracy is explored in terms of canonical angles that measure differences between the unknown true subspaces and their random approximations [3, 27, 24]. Generally, upper bounds on the canonical angles can be categorized into two types:

- (i) probabilistic bounds that establish prior statistical guarantees by exploring the concentration of the alignment between random subspace embeddings and the unknown true subspace, and
- (ii) residual-based bounds that can be computed a posteriori from the residual of the resulting approximation.

The existing prior probabilistic bounds on canonical angles [3, 27, 21] mainly focus on the setting where the randomized SVD is evaluated without oversampling or with a small constant oversampling. Concretely, [3] derives guarantees for the canonical angles evaluated without oversampling (i.e., $l = k$) in the context of spectral clustering. Further, by taking a small constant oversampling (e.g., $l \geq k + 2$) into account, Saibaba [27] provides a comprehensive analysis for an assortment of canonical angles between the true and approximated leading singular spaces. Compared with our results (Theorem 1), in both no-oversampling and constant-oversampling regimes, the basic forms of the existing prior probabilistic bounds (e.g., [27, Theorem 1]) generally depend on the unknown singular subspace \mathbf{V} . Although such dependence is

later lifted using the isotropicity and the concentration of the randomized subspace embedding Ω (e.g., [27, Theorem 6]), the separated consideration on the spectra and the singular subspaces introduces unnecessary compromise to the upper bounds (as we will discuss in Remark 2). In contrast, by allowing a more generous choice of multiplicative oversampling $l = \Omega(k)$, we present a set of space-agnostic bounds (i.e., bounds that hold regardless of the singular vectors of \mathbf{A}) based on an integrated analysis of the spectra and the singular subspaces that appears to be tighter both from derivation and in practice.

The classical Davis–Kahan $\sin \theta$ and $\tan \theta$ theorems [5] for eigenvector perturbation can be used to compute deterministic and computable bounds for the canonical angles. These bounds have the advantage that they give strict bounds (up to the estimation of the so-called gap) rather than estimates or bounds that hold with high probability (although, as we argue below, the failure probability can be taken to be negligibly low). The Davis–Kahan theorems have been extended to perturbation of singular vectors by Wedin [36], and recent work [24] derives perturbation bounds for singular vectors computed using a subspace projection method. In this work, we establish canonical angle bounds for the singular vectors in the context of (randomized) subspace iterations. Our results indicate clearly that the accuracy of the right and left singular vectors is usually not identical (i.e., \mathbf{V} is more accurate with Algorithm 1).

As a roadmap, we formalize the problem setup in section 2, including a brief review of the randomized SVD and canonical angles. In section 3, we present the prior probabilistic space-agnostic bounds. Subsequently, in section 4, we describe a set of unbiased canonical angle estimates that is closely related to the space-agnostic bounds. Then in section 5, we introduce two sets of posterior residual-based bounds. Finally, in section 6, we instantiate the insight cast by the space-agnostic bounds on the balance between oversampling and power iterations and demonstrate the empirical effectiveness of different canonical angle bounds and estimates with numerical comparisons.

2. Problem setup. In this section, we first recapitulate the randomized SVD algorithm (with power iterations) [14] for which we analyze the accuracy of the resulting singular subspace approximations. Then, we review the notion of canonical angles [12] that quantify the distance between two subspaces of the same Euclidean space.

2.1. Notation. For any $k \in \mathbb{N}$, we denote $[k] = \{1, \dots, k\}$. We adapt the MATLAB notation for matrix slicing and stacking throughout this work. In particular, for any set of matrices $\{\mathbf{M}_i \in \mathbb{C}^{d \times d_i} \mid i \in [n]\}$, we denote $[\mathbf{M}_1, \dots, \mathbf{M}_n] \in \mathbb{C}^{d \times (\sum_{i=1}^n d_i)}$ as the horizontal concatenation.

We start by introducing notation for the SVD of a given matrix $\mathbf{A} \in \mathbb{C}^{m \times n}$ of rank r :

$$\mathbf{A} = \mathbf{U}_{m \times r} \mathbf{\Sigma}_{r \times r} \mathbf{V}^*_{r \times n} = [\mathbf{u}_1 \ \dots \ \mathbf{u}_r] \begin{bmatrix} \sigma_1 & & \\ & \ddots & \\ & & \sigma_r \end{bmatrix} \begin{bmatrix} \mathbf{v}_1^* \\ \vdots \\ \mathbf{v}_r^* \end{bmatrix}.$$

For any $1 \leq k \leq r$, we let $\mathbf{U}_k \triangleq [\mathbf{u}_1, \dots, \mathbf{u}_k]$ and $\mathbf{V}_k \triangleq [\mathbf{v}_1, \dots, \mathbf{v}_k]$ denote the orthonormal bases of the dimension- k left and right singular subspaces of \mathbf{A} corresponding to the top- k singular values, while $\mathbf{U}_{r \setminus k} \triangleq [\mathbf{u}_{k+1}, \dots, \mathbf{u}_r]$ and $\mathbf{V}_{r \setminus k} \triangleq [\mathbf{v}_{k+1}, \dots, \mathbf{v}_r]$ are orthonormal bases of the respective orthogonal complements. The diagonal submatrices consisting of the spectrum, $\mathbf{\Sigma}_k \triangleq \text{diag}(\sigma_1, \dots, \sigma_k)$ and $\mathbf{\Sigma}_{r \setminus k} \triangleq \text{diag}(\sigma_{k+1}, \dots, \sigma_r)$, follow analogously. We use $\|\cdot\|_2$ to denote the spectral norm (largest singular value) of

Algorithm 1 Randomized SVD (with power iterations) [14]

Require: $\mathbf{A} \in \mathbb{C}^{m \times n}$, power $q \in \{0, 1, 2, \dots\}$, oversampled rank $l \in \mathbb{N}$ ($l < r = \text{rank}(\mathbf{A})$)

Ensure: $\widehat{\mathbf{U}}_l \in \mathbb{C}^{m \times l}$, $\widehat{\mathbf{V}}_l \in \mathbb{C}^{n \times l}$, $\widehat{\Sigma}_l \in \mathbb{C}^{l \times l}$ such that $\widehat{\mathbf{A}}_l = \widehat{\mathbf{U}}_l \widehat{\Sigma}_l \widehat{\mathbf{V}}_l^*$

- 1: Draw $\boldsymbol{\Omega} \sim P(\mathbb{C}^{n \times l})$ with $\Omega_{ij} \sim \mathcal{N}(0, l^{-1})$ i.i.d. such that $\mathbb{E}[\boldsymbol{\Omega} \boldsymbol{\Omega}^*] = \mathbf{I}_n$
- 2: $\mathbf{X}^{(q)} = (\mathbf{A} \mathbf{A}^*)^q \mathbf{A} \boldsymbol{\Omega}$
- 3: $\mathbf{Q}_{\mathbf{X}} = \text{ortho}(\mathbf{X}^{(q)})$
- 4: $[\widetilde{\mathbf{U}}_l, \widehat{\Sigma}_l, \widehat{\mathbf{V}}_l^*] = \text{svd}(\mathbf{Q}_{\mathbf{X}}^* \mathbf{A})$ (where $\widetilde{\mathbf{U}}_l \in \mathbb{C}^{l \times l}$)
- 5: $\widehat{\mathbf{U}}_l = \mathbf{Q}_{\mathbf{X}} \widetilde{\mathbf{U}}_l$

a matrix, and $\|\cdot\|_F$ denotes the Frobenius norm. Equations and inequalities involving $\|\cdot\|$ hold for any (fixed) unitarily invariant norm.

Meanwhile, for the QR decomposition of an arbitrary matrix $\mathbf{M} \in \mathbb{C}^{d \times l}$ ($d \geq l$) with full column rank $\text{rank}(\mathbf{M}) = l$, we denote $\mathbf{M} = [\mathbf{Q}_{\mathbf{M}} \quad \mathbf{Q}_{\mathbf{M}, \perp}] [\mathbf{R}_{\mathbf{M}}]$ such that $\mathbf{Q}_{\mathbf{M}} = \text{ortho}(\mathbf{M}) \in \mathbb{C}^{d \times l}$ and $\mathbf{Q}_{\mathbf{M}, \perp} \in \mathbb{C}^{d \times (d-l)}$ consist of orthonormal bases of the subspace spanned by the columns of \mathbf{M} and its orthogonal complement. Furthermore, we denote the singular values of \mathbf{M} by $\sigma(\mathbf{M})$, a $\text{rank}(\mathbf{M}) \times \text{rank}(\mathbf{M})$ diagonal matrix with singular values $\sigma_1(\mathbf{M}) \geq \dots \geq \sigma_{\text{rank}(\mathbf{M})}(\mathbf{M}) > 0$ on the diagonal.

Throughout, it is helpful to have in mind the ordering $k \leq l \ll r \leq \min(m, n)$, where k is the dimension of the subspace of interest, l is the (oversampled) dimension of subspace used in the algorithm, and $r = \text{rank}(\mathbf{A}) \leq \min(m, n)$.

Given an arbitrary distribution $P : \mathbb{C} \rightarrow [0, 1]$, we denote $\boldsymbol{\Omega} \sim P(\mathbb{C}^{n \times l})$ as a random matrix with i.i.d. entries $\Omega_{ij} \sim P$ for all $i \in [n], j \in [l]$.

We adapt the standard asymptotic notation as follows: for any functions $f, g : \mathbb{R}_+ \rightarrow \mathbb{R}_+$, we write $f = O(g)$ if there exists some constant $C > 0$ such that $f(x) \leq Cg(x)$ for all $x \in \mathbb{R}_+$; $f = \Omega(g)$ if $g = O(f)$; $f = \Theta(g)$ if $f = O(g)$ and $f = \Omega(g)$.

2.2. Randomized SVD and power iterations. As described in Algorithm 1, the randomized SVD provides a rank- l ($l \ll \min(m, n)$) approximation of $\mathbf{A} \in \mathbb{C}^{m \times n}$ while granting provable acceleration to the truncated SVD computation— $O(mnl(2q + 1))$ when $\boldsymbol{\Omega}$ is chosen to be a Gaussian random matrix.² Such efficiency improvement is achieved by first projecting the high-dimensional row (column) space of \mathbf{A} to a low-dimensional subspace via a Johnson–Lindenstrauss transform³ (JLT) $\boldsymbol{\Omega}$ (known as “sketching”). Then, the SVD of the resulting column (row) sketch $\mathbf{X} = \mathbf{A} \boldsymbol{\Omega}$ can be evaluated efficiently in $O(ml^2)$ time, and the rank- l approximation can be constructed accordingly.

The spectral decay in \mathbf{A} has a significant impact on the accuracy of the resulting low-rank approximation from Algorithm 1 (as suggested in [14, Thm. 10.7, Thm. 10.8]). To remediate the performance of Algorithm 1 on matrices with flat spectra, power iterations (Algorithm 1, Line 2) ($q \geq 1$) are usually incorporated to enhance

²Asymptotically, there exist deterministic iterative algorithms for the truncated SVD (e.g., based on Lanczos iterations [30, Algorithm 36.1]) that run in $O(mnl)$ time. However, compared with these inherently sequential iterative algorithms, the randomized SVD can be executed much more efficiently in practice, even with power iterations (i.e., $q > 0$), since the $O(mnl(2q + 1))$ computation bottleneck in Algorithm 1 involves only matrix-matrix multiplications, which are easily parallelizable and highly optimized.

³Throughout this work, we focus on Gaussian random matrices (Algorithm 1, Line 1) for the sake of theoretical guarantees, i.e., $\boldsymbol{\Omega}$ being isotropic and rotationally invariant.

the spectral decay. However, without proper orthogonalization, plain power iterations can be numerically unstable, especially for ill-conditioned matrices. For stable power iterations, starting with $\mathbf{X} = \mathbf{A}\Omega \in \mathbb{C}^{m \times l}$ of full column rank (which holds almost surely for Gaussian random matrices), we incorporate orthogonalization in each power iteration (or one can do so selectively [29, section 6.1]) via the reduced unpivoted QR factorization (each with complexity $O(ml^2)$). Let $\text{ortho}(\mathbf{X}) = \mathbf{Q}_X \in \mathbb{C}^{m \times l}$ be an orthonormal basis of \mathbf{X} produced by the QR factorization. Then, the stable evaluation of q power iterations (Algorithm 1, Line 2) can be expressed as

$$(2.1) \quad \mathbf{X}^{(0)} \leftarrow \text{ortho}(\mathbf{A}\Omega), \quad \mathbf{X}^{(i)} \leftarrow \text{ortho} \left(\mathbf{A} \text{ortho} \left(\mathbf{A}^* \mathbf{X}^{(i-1)} \right) \right) \quad \forall i \in [q].$$

Notice that in Algorithm 1, with $\mathbf{X} = \mathbf{X}^{(q)}$, the approximated rank- l SVD of \mathbf{A} can be expressed as

$$(2.2) \quad \widehat{\mathbf{A}}_l = \widehat{\mathbf{U}}_l \widehat{\Sigma}_l \widehat{\mathbf{V}}_l^* = \mathbf{X} \mathbf{Y}^*, \quad \text{where } \mathbf{Y} = \mathbf{X}^\dagger \mathbf{A}.$$

With $\widehat{\mathbf{U}}_l$ and $\widehat{\mathbf{V}}_l$ characterizing the approximated l -dimensional left and right leading singular subspaces, let $\widehat{\mathbf{U}}_{m \setminus l} \in \mathbb{C}^{m \times (m-l)}$ and $\widehat{\mathbf{V}}_{n \setminus l} \in \mathbb{C}^{n \times (n-l)}$ denote an arbitrary pair of their respective orthogonal complements. For any $1 \leq k < l$, we further denote the partitions $\widehat{\mathbf{U}}_l = [\widehat{\mathbf{U}}_k, \widehat{\mathbf{U}}_{l \setminus k}]$ and $\widehat{\mathbf{V}}_l = [\widehat{\mathbf{V}}_k, \widehat{\mathbf{V}}_{l \setminus k}]$, where $\widehat{\mathbf{U}}_k \in \mathbb{C}^{m \times k}$ and $\widehat{\mathbf{V}}_k \in \mathbb{C}^{n \times k}$, respectively.

2.3. Canonical angles. Now, we review the notion of canonical angles [12] that measure distances between two subspaces \mathcal{U}, \mathcal{V} of an arbitrary Euclidean space \mathbb{C}^d .

DEFINITION 1 (canonical angles [12]). *Given two subspaces $\mathcal{U}, \mathcal{V} \subseteq \mathbb{C}^d$ with dimensions $\dim(\mathcal{U}) = l$ and $\dim(\mathcal{V}) = k$ (assuming $l \geq k$ without loss of generality), the canonical angles, denoted by $\angle(\mathcal{U}, \mathcal{V}) = \text{diag}(\theta_1, \dots, \theta_k)$, consist of k angles that measure the alignment between \mathcal{U} and \mathcal{V} , defined recursively such that*

$$\begin{aligned} \mathbf{u}_i, \mathbf{v}_i &\triangleq \text{argmax } \mathbf{u}_i^* \mathbf{v}_i \\ \text{s.t. } \mathbf{u}_i &\in \left(\mathcal{U} \setminus \text{span} \{ \mathbf{u}_\ell \}_{\ell=1}^{i-1} \right) \cap \mathbb{S}^{d-1}, \\ \mathbf{v}_i &\in \left(\mathcal{V} \setminus \text{span} \{ \mathbf{v}_\ell \}_{\ell=1}^{i-1} \right) \cap \mathbb{S}^{d-1}, \\ \cos \theta_i &= \mathbf{u}_i^* \mathbf{v}_i \quad \forall i = 1, \dots, k, \quad 0 \leq \theta_1 \leq \dots \leq \theta_k \leq \pi/2. \end{aligned}$$

For arbitrary full-rank matrices $\mathbf{M} \in \mathbb{C}^{d \times l}$ and $\mathbf{N} \in \mathbb{C}^{d \times k}$ (assuming $k \leq l \leq d$ without loss of generality), let $\angle(\mathbf{M}, \mathbf{N}) \triangleq \angle(\text{span}(\mathbf{M}), \text{span}(\mathbf{N}))$ denote the canonical angles between the corresponding spanning subspaces in \mathbb{C}^d . For each $i \in [k]$, let $\angle_i(\mathbf{M}, \mathbf{N})$ be the i th (smallest) canonical angle defined as $\cos \angle_i(\mathbf{M}, \mathbf{N}) = \sigma_i(\mathbf{Q}_M^* \mathbf{Q}_N)$ such that $\sin \angle_i(\mathbf{M}, \mathbf{N}) = \sigma_{k-i+1}((\mathbf{I} - \mathbf{Q}_M \mathbf{Q}_M^*) \mathbf{Q}_N)$ (cf. [1, section 3]).

With the unknown true rank- k truncated SVD $\mathbf{A}_k = \mathbf{U}_k \Sigma_k \mathbf{V}_k^*$ and an approximated rank- l SVD $\widehat{\mathbf{A}}_l = \widehat{\mathbf{U}}_l \widehat{\Sigma}_l \widehat{\mathbf{V}}_l^*$ from Algorithm 1, in this work, we mainly focus on prior and posterior guarantees for the canonical angles $\angle(\mathbf{U}_k, \widehat{\mathbf{U}}_l)$ and $\angle(\mathbf{V}_k, \widehat{\mathbf{V}}_l)$. Meanwhile, in Theorem 4, we present a set of posterior residual-based upper bounds for the canonical angles $\angle(\mathbf{U}_k, \widehat{\mathbf{U}}_k)$ and $\angle(\mathbf{V}_k, \widehat{\mathbf{V}}_k)$ as corollaries.

3. Space-agnostic bounds under sufficient oversampling. We start by pointing out the intuition that, under sufficient oversampling, with Gaussian random matrices whose distribution is orthogonally invariant, the alignment between the approximated and true subspaces is independent of the unknown true subspaces; i.e., the canonical angles are space-agnostic, as reflected in the following theorem.

THEOREM 1. *For a rank- l randomized SVD (Algorithm 1) with a Gaussian embedding Ω and $q \geq 0$ power iterations, when the oversampled rank l satisfies $l = \Omega(k)$ (where k is the target rank, $k < l < r = \text{rank}(\mathbf{A})$) and q is reasonably small such that $\eta \triangleq \frac{(\sum_{j=k+1}^r \sigma_j^{4q+2})^2}{\sum_{j=k+1}^r \sigma_j^{2(4q+2)}}^4$ satisfies $\eta = \Omega(l)$, with high probability (at least $1 - e^{-\Theta(k)} - e^{-\Theta(l)}$), there exist distortion factors $0 < \epsilon_1, \epsilon_2 < 1$ such that*

$$(3.1) \quad \sin \angle_i (\mathbf{U}_k, \widehat{\mathbf{U}}_l) \leq \left(1 + \frac{1 - \epsilon_1}{1 + \epsilon_2} \cdot \frac{l}{\sum_{j=k+1}^r \sigma_j^{4q+2} \cdot \sigma_i^{4q+2}} \right)^{-\frac{1}{2}},$$

$$(3.2) \quad \sin \angle_i (\mathbf{V}_k, \widehat{\mathbf{V}}_l) \leq \left(1 + \frac{1 - \epsilon_1}{1 + \epsilon_2} \cdot \frac{l}{\sum_{j=k+1}^r \sigma_j^{4q+4} \cdot \sigma_i^{4q+4}} \right)^{-\frac{1}{2}}$$

for all $i \in [k]$, where $\epsilon_1 = \Theta(\sqrt{\frac{k}{l}})$ and $\epsilon_2 = \Theta(\sqrt{\frac{l}{\eta}})$. Furthermore, both bounds are asymptotically tight:

$$(3.3) \quad \sin \angle_i (\mathbf{U}_k, \widehat{\mathbf{U}}_l) \geq \left(1 + O \left(\frac{l \cdot \sigma_i^{4q+2}}{\sum_{j=k+1}^r \sigma_j^{4q+2}} \right) \right)^{-\frac{1}{2}},$$

$$(3.4) \quad \sin \angle_i (\mathbf{V}_k, \widehat{\mathbf{V}}_l) \geq \left(1 + O \left(\frac{l \cdot \sigma_i^{4q+4}}{\sum_{j=k+1}^r \sigma_j^{4q+4}} \right) \right)^{-\frac{1}{2}},$$

where $O(\cdot)$ suppresses the distortion factors $\frac{1+\epsilon_1}{1-\epsilon_2}$.⁴

The main insights provided by Theorem 1 include

- (i) improved statistical guarantees for canonical angles under sufficient oversampling (i.e., $l = \Omega(k)$), as discussed later in Remark 2,
- (ii) a clear view of the balance between oversampling and power iterations for random subspace approximations with given spectra, as instantiated in subsection 6.3, and
- (iii) affordable upper bounds that can be evaluated in $O(\text{rank}(\mathbf{A}))$ time with access to the (true/estimated) spectra and hold in practice with only moderate oversampling (e.g., $l \geq 1.6k$), as shown in subsection 6.2.

We also note that when the true singular values in Theorem 1 are unknown, they may be replaced by the approximated singular values $\widehat{\Sigma}_l$ from Algorithm 1 in practice; cf. subsection 6.2.

⁴Notice that $1 < \eta \leq r - k$. To the extremes, $\eta = r - k$ when the tail is flat and $\sigma_{k+1} = \dots = \sigma_r$, while $\eta \rightarrow 1$ when $\sigma_{k+1} \gg \sigma_j$ for all $j = k + 2, \dots, r$. In particular, with a relatively flat tail Σ_{r-k} and a reasonably small q (recall that $q = 1, 2$ is usually sufficient in practice), we have $\eta = \Theta(r - k)$, and the assumption can be simplified as $r - k = \Omega(l)$. Although exponential tail decay can lead to small η and may render the assumption infeasible in theory, in practice, simply taking $r - k = \Omega(l)$, $l = \Omega(k)$, $\epsilon_1 = \sqrt{\frac{k}{l}}$, and $\epsilon_2 = \sqrt{\frac{l}{r-k}}$ is sufficient to ensure the validity of upper bounds when $q \leq 10$ even for matrices with rapid tail decay, as shown in subsection 6.2.

⁵Despite the asymptotic tightness of Theorem 1 theoretically, in practice, we observe that the empirical validity of lower bounds is more restrictive on oversampling than that of upper bounds. In specific, the numerical observations in subsection 6.2 suggest that $l \geq 1.6k$ is usually sufficient for the upper bounds to hold, whereas the empirical validity of lower bounds generally requires more aggressive oversampling of at least $l \geq 4k$, also with slightly larger constants associated with ϵ_1 and ϵ_2 , as demonstrated in Appendix B.

Proof of Theorem 1. We show the derivation of (3.1) for left canonical angles $\sin \angle_i(\mathbf{U}_k, \widehat{\mathbf{U}}_l)$. The derivation for $\sin \angle_i(\mathbf{V}_k, \widehat{\mathbf{V}}_l)$ in (3.2) follows directly by replacing the exponent $4q + 2$ in (3.1) with $4q + 4$ in (3.2). This slightly larger exponent comes from the additional half power iteration associated with $\widehat{\mathbf{V}}_l$ in Algorithm 1 (as discussed in Remark 3), an observation made also in [27].

For rank- l randomized SVD with Gaussian embedding $\boldsymbol{\Omega} \in \mathbb{C}^{n \times l}$ and q power iterations, we denote $\boldsymbol{\Omega}_1 \triangleq \mathbf{V}_k^* \boldsymbol{\Omega}$ and $\boldsymbol{\Omega}_2 \triangleq \mathbf{V}_{r \setminus k}^* \boldsymbol{\Omega}$, as well as their weighted correspondences $\tilde{\boldsymbol{\Omega}}_1 \triangleq \boldsymbol{\Sigma}_k^{2q+1} \boldsymbol{\Omega}_1$ and $\tilde{\boldsymbol{\Omega}}_2 \triangleq \boldsymbol{\Sigma}_{r \setminus k}^{2q+1} \boldsymbol{\Omega}_2$, such that

$$\tilde{\mathbf{X}} \triangleq \mathbf{U}^* \mathbf{X} = \mathbf{U}^* (\mathbf{A} \mathbf{A}^*)^q \mathbf{A} \boldsymbol{\Omega} = \begin{bmatrix} \boldsymbol{\Sigma}_k^{2q+1} \boldsymbol{\Omega}_1 \\ \boldsymbol{\Sigma}_{r \setminus k}^{2q+1} \boldsymbol{\Omega}_2 \end{bmatrix} = \begin{bmatrix} \tilde{\boldsymbol{\Omega}}_1 \\ \tilde{\boldsymbol{\Omega}}_2 \end{bmatrix}.$$

Since $\text{span}(\mathbf{U}_k) \subseteq \text{span}(\mathbf{U})$ and $\text{span}(\mathbf{X}) \subseteq \text{span}(\mathbf{U})$, we have $\mathbf{U} \mathbf{U}^* \mathbf{U}_k = \mathbf{U}_k$ and $\mathbf{U} \mathbf{U}^* \mathbf{X} = \mathbf{X}$, respectively. Then for all $i \in [k]$,

$$\begin{aligned} \sin \angle_{k-i+1}(\mathbf{U}_k, \widehat{\mathbf{U}}_l) &= \sigma_i \left((\mathbf{I}_m - \widehat{\mathbf{U}}_l \widehat{\mathbf{U}}_l^*) \mathbf{U}_k \right), \\ (\mathbf{U} \mathbf{U}^* \mathbf{U}_k = \mathbf{U}_k) &= \sigma_i \left((\mathbf{U} \mathbf{U}^* - \widehat{\mathbf{U}}_l \widehat{\mathbf{U}}_l^*) \mathbf{U}_k \right), \\ (\text{span}(\mathbf{U}_l) = \text{span}(\mathbf{X}) \text{ by (2.2)}) &= \sigma_i \left((\mathbf{U} \mathbf{U}^* - \mathbf{X} \mathbf{X}^\dagger) \mathbf{U}_k \right), \\ (\mathbf{U} \mathbf{U}^* \mathbf{X} = \mathbf{X}) &= \sigma_i \left(\mathbf{U} \mathbf{U}^* (\mathbf{I}_m - \mathbf{X} \mathbf{X}^\dagger) \mathbf{U} \mathbf{U}^* \mathbf{U}_k \right), \\ (\mathbf{U} \text{ consists of orthonormal columns}) &= \sigma_i \left(\mathbf{U}^* (\mathbf{I}_m - \mathbf{X} \mathbf{X}^\dagger) \mathbf{U} \mathbf{U}^* \mathbf{U}_k \right), \\ \left(\mathbf{U}^* \mathbf{U} = \mathbf{I}_r, \mathbf{X}^\dagger \mathbf{U} = (\mathbf{U}^* \mathbf{X})^\dagger = \tilde{\mathbf{X}}^\dagger \right) &= \sigma_i \left((\mathbf{I}_r - \tilde{\mathbf{X}} \tilde{\mathbf{X}}^\dagger) \begin{bmatrix} \mathbf{I}_k \\ \mathbf{0} \end{bmatrix} \right). \end{aligned}$$

Since \mathbf{X} is assumed to have full column rank, we have $\tilde{\mathbf{X}} \tilde{\mathbf{X}}^\dagger = \tilde{\mathbf{X}} (\tilde{\mathbf{X}}^* \tilde{\mathbf{X}})^{-1} \tilde{\mathbf{X}}^*$ (which is an orthogonal projection), and

$$\begin{aligned} &[\mathbf{I}_k \quad \mathbf{0}] (\mathbf{I}_r - \tilde{\mathbf{X}} \tilde{\mathbf{X}}^\dagger) \begin{bmatrix} \mathbf{I}_k \\ \mathbf{0} \end{bmatrix} \\ &= [\mathbf{I}_k \quad \mathbf{0}] \left(\mathbf{I}_r - \begin{bmatrix} \tilde{\boldsymbol{\Omega}}_1 \\ \tilde{\boldsymbol{\Omega}}_2 \end{bmatrix} (\tilde{\boldsymbol{\Omega}}_1^* \tilde{\boldsymbol{\Omega}}_1 + \tilde{\boldsymbol{\Omega}}_2^* \tilde{\boldsymbol{\Omega}}_2)^{-1} \begin{bmatrix} \tilde{\boldsymbol{\Omega}}_1^* & \tilde{\boldsymbol{\Omega}}_2^* \end{bmatrix} \right) \begin{bmatrix} \mathbf{I}_k \\ \mathbf{0} \end{bmatrix} \\ &= \mathbf{I}_k - \tilde{\boldsymbol{\Omega}}_1 (\tilde{\boldsymbol{\Omega}}_1^* \tilde{\boldsymbol{\Omega}}_1 + \tilde{\boldsymbol{\Omega}}_2^* \tilde{\boldsymbol{\Omega}}_2)^{-1} \tilde{\boldsymbol{\Omega}}_1^* \\ &\text{(Woodbury identity)} = \left(\mathbf{I}_k + \tilde{\boldsymbol{\Omega}}_1 (\tilde{\boldsymbol{\Omega}}_2^* \tilde{\boldsymbol{\Omega}}_2)^{-1} \tilde{\boldsymbol{\Omega}}_1^* \right)^{-1}. \end{aligned}$$

Therefore, for all $i \in [k]$,

$$\begin{aligned} \sin^2 \angle_{k-i+1}(\mathbf{U}_k, \widehat{\mathbf{U}}_l) &= \sigma_i \left([\mathbf{I}_k \quad \mathbf{0}] (\mathbf{I}_r - \tilde{\mathbf{X}} \tilde{\mathbf{X}}^\dagger) \begin{bmatrix} \mathbf{I}_k \\ \mathbf{0} \end{bmatrix} \right) \\ (3.5) \quad &= \sigma_i \left(\left(\mathbf{I}_k + \tilde{\boldsymbol{\Omega}}_1 (\tilde{\boldsymbol{\Omega}}_2^* \tilde{\boldsymbol{\Omega}}_2)^{-1} \tilde{\boldsymbol{\Omega}}_1^* \right)^{-1} \right). \end{aligned}$$

By the orthogonal invariance of the distribution of Gaussian embeddings $\boldsymbol{\Omega}$ together with the orthonormality of $\mathbf{V}_k \perp \mathbf{V}_{r \setminus k}$, we see that $\boldsymbol{\Omega}_1 \sim P(\mathbb{C}^{k \times l})$ and $\boldsymbol{\Omega}_2 \sim P(\mathbb{C}^{(r-k) \times l})$ are independent Gaussian random matrices with the same entrywise i.i.d. distribution $\mathcal{N}(0, l^{-1})$ as $\boldsymbol{\Omega}$. Therefore, by Lemma 1, when $l = \Omega(k)$, with high probability (at least $1 - e^{-\Theta(k)}$),

$$(1 - \epsilon_1) \boldsymbol{\Sigma}_k^{4q+2} \preceq \tilde{\boldsymbol{\Omega}}_1 \tilde{\boldsymbol{\Omega}}_1^* \preceq (1 + \epsilon_1) \boldsymbol{\Sigma}_k^{4q+2}$$

for some $\epsilon_1 = \Theta(\sqrt{\frac{k}{l}})$.

Analogously, when $r - k \geq \eta = \frac{\text{tr}(\Sigma_{r \setminus k}^{4q+2})^2}{\text{tr}(\Sigma_{r \setminus k}^{2(4q+2)})} = \Omega(l)$, with high probability (at least $1 - e^{-\Theta(l)}$),

$$\frac{1 - \epsilon_2}{l} \text{tr}(\Sigma_{r \setminus k}^{4q+2}) \mathbf{I}_l \preccurlyeq \tilde{\Omega}_2^* \tilde{\Omega}_2 \preccurlyeq \frac{1 + \epsilon_2}{l} \text{tr}(\Sigma_{r \setminus k}^{4q+2}) \mathbf{I}_l$$

for some $\epsilon_2 = \Theta(\sqrt{\frac{l}{\eta}})$.

Therefore by the union bound, we have with high probability (at least $1 - e^{-\Theta(k)} - e^{-\Theta(l)}$) that

$$\left(\mathbf{I}_k + \tilde{\Omega}_1 \left(\tilde{\Omega}_2^* \tilde{\Omega}_2 \right)^{-1} \tilde{\Omega}_1^* \right)^{-1} \succcurlyeq \left(\mathbf{I}_k + \frac{1 - \epsilon_1}{1 + \epsilon_2} \cdot \frac{l}{\text{tr}(\Sigma_{r \setminus k}^{4q+2})} \cdot \Sigma_k^{4q+2} \right)^{-1},$$

which leads to (3.1), while the tightness is implied by

$$\left(\mathbf{I}_k + \tilde{\Omega}_1 \left(\tilde{\Omega}_2^* \tilde{\Omega}_2 \right)^{-1} \tilde{\Omega}_1^* \right)^{-1} \succcurlyeq \left(\mathbf{I}_k + \frac{1 + \epsilon_1}{1 - \epsilon_2} \cdot \frac{l}{\text{tr}(\Sigma_{r \setminus k}^{4q+2})} \cdot \Sigma_k^{4q+2} \right)^{-1}.$$

The proof of (3.2) follows analogously by replacing the exponents $2q+1$ and $4q+2$ with $2q+2$ and $4q+4$, respectively. \square

Remark 2 (comparison with existing probabilistic bounds). With access to the unknown right singular subspace \mathbf{V} , let $\Omega_1 \triangleq \mathbf{V}_k^* \Omega$ and $\Omega_2 \triangleq \mathbf{V}_{r \setminus k}^* \Omega$. Then, Saibaba [27, Thm. 1] indicates that, for all $i \in [k]$,

$$(3.6) \quad \sin \angle_i(\mathbf{U}_k, \widehat{\mathbf{U}}_l) \leq \left(1 + \frac{\sigma_i^{4q+2}}{\sigma_{k+1}^{4q+2} \|\Omega_2 \Omega_1^\dagger\|_2^2} \right)^{-\frac{1}{2}},$$

$$(3.7) \quad \sin \angle_i(\mathbf{V}_k, \widehat{\mathbf{V}}_l) \leq \left(1 + \frac{\sigma_i^{4q+4}}{\sigma_{k+1}^{4q+4} \|\Omega_2 \Omega_1^\dagger\|_2^2} \right)^{-\frac{1}{2}}.$$

Further, leveraging existing results on concentration properties of the independent and isotropic Gaussian random matrices Ω_1 and Ω_2 (e.g., from the proof of [14, Theorem 10.8]), [27] shows that, when $l \geq k+2$, for any $\delta \in (0, 1)$, with probability at least $1 - \delta$,

$$\|\Omega_2 \Omega_1^\dagger\|_2 \leq \frac{e\sqrt{l}}{l - k + 1} \left(\frac{2}{\delta} \right)^{\frac{1}{l-k+1}} \left(\sqrt{n-k} + \sqrt{l} + \sqrt{2 \log \frac{2}{\delta}} \right).$$

Without loss of generality, we consider the bounds on $\sin \angle(\mathbf{U}_k, \widehat{\mathbf{U}}_l)$. Comparing to the existing bound in (3.6), under multiplicative oversampling ($l = \Omega(k)$, $r = \Omega(l)$), (3.1) in Theorem 1 captures the spectral decay on the tail by replacing the denominator term

$$\sigma_{k+1}^{4q+2} \|\Omega_2 \Omega_1^\dagger\|_2^2 \quad \text{with} \quad \frac{1 + \epsilon_2}{1 - \epsilon_1} \cdot \frac{1}{l} \sum_{j=k+1}^r \sigma_j^{4q+2} = \Theta \left(\frac{1}{l} \sum_{j=k+1}^r \sigma_j^{4q+2} \right).$$

We observe that $\frac{1}{l} \sum_{j=k+1}^r \sigma_j^{4q+2} \leq \frac{r-k}{l} \sigma_{k+1}^{4q+2}$, while Lemma 1 implies that, for independent Gaussian random matrices $\Omega_1 \sim P(\mathbb{C}^{k \times l})$ and $\Omega_2 \sim P(\mathbb{C}^{(r-k) \times l})$ with i.i.d. entries from $\mathcal{N}(0, l^{-1})$,

$$\mathbb{E} \left[\left\| \Omega_2 \Omega_1^\dagger \right\|_2^2 \right] = \mathbb{E}_{\Omega_1} \left[\left\| \left(\Omega_1^\dagger \right)^* \mathbb{E}_{\Omega_2} [\Omega_2^* \Omega_2] \Omega_1^\dagger \right\|_2 \right] = \frac{r-k}{l} \mathbb{E}_{\Omega_1} \left[\left\| (\Omega_1 \Omega_1^*)^\dagger \right\|_2 \right].$$

With nonnegligible spectral decay on the tail such that $\frac{1}{l} \sum_{j=k+1}^r \sigma_j^{4q+2} \ll \frac{r-k}{l} \sigma_{k+1}^{4q+2}$, when $\frac{1+\epsilon_2}{1-\epsilon_1} \cdot \frac{1}{l} \sum_{j=k+1}^r \sigma_j^{4q+2} \ll \sigma_{k+1}^{4q+2} \|\Omega_2 \Omega_1^\dagger\|_2^2$, (3.1) provides a tighter statistical guarantee than (3.6), which is also confirmed by numerical observations in subsection 6.2.

From the derivation perspective, such improvement is achieved by taking an integrated view on the concentration of $\Sigma_{r \setminus k}^{2q+1} \Omega_2$ (commonly used in analyzing randomized low-rank approximation error [14, 11, 10, 32]), instead of considering the spectrum and the unknown singular subspace separately.

4. Unbiased space-agnostic estimates. A natural corollary from the proof of Theorem 4 is unbiased estimates for the canonical angles that hold for arbitrary oversampling (i.e., for all $l \geq k$). Further, we will subsequently see in section 6 that such unbiased estimates also enjoy good empirical concentration.

PROPOSITION 1. *For a rank- l randomized SVD (Algorithm 1) with the Gaussian embedding $\Omega \sim P(\mathbb{C}^{n \times l})$ such that $\Omega_{ij} \sim \mathcal{N}(0, l^{-1})$ i.i.d. and $q \geq 0$ power iterations, for all $i \in [k]$,*

(4.1)

$$\mathbb{E}_\Omega \left[\sin \angle_i \left(\mathbf{U}_k, \widehat{\mathbf{U}}_l \right) \right] = \mathbb{E}_{\Omega'_1, \Omega'_2} \left[\sigma_i^{-\frac{1}{2}} \left(\mathbf{I}_k + \Sigma_k^{2q+1} \Omega'_1 \left(\Omega'_2 \Sigma_{r \setminus k}^{4q+2} \Omega'_2 \right)^{-1} \Omega'_1 \Sigma_k^{2q+1} \right) \right],$$

and analogously,

(4.2)

$$\mathbb{E}_\Omega \left[\sin \angle_i \left(\mathbf{V}_k, \widehat{\mathbf{V}}_l \right) \right] = \mathbb{E}_{\Omega'_1, \Omega'_2} \left[\sigma_i^{-\frac{1}{2}} \left(\mathbf{I}_k + \Sigma_k^{2q+2} \Omega'_1 \left(\Omega'_2 \Sigma_{r \setminus k}^{4q+4} \Omega'_2 \right)^{-1} \Omega'_1 \Sigma_k^{2q+2} \right) \right],$$

where $\Omega'_1 \sim P(\mathbb{C}^{k \times l})$ and $\Omega'_2 \sim P(\mathbb{C}^{(r-k) \times l})$ are independent Gaussian random matrices with i.i.d. entries drawn from $\mathcal{N}(0, l^{-1})$.

To calculate the unbiased estimate, for a modest integer N we draw a set of independent Gaussian random matrices $\{\Omega_1^{(j)} \sim P(\mathbb{C}^{k \times l}), \Omega_2^{(j)} \sim P(\mathbb{C}^{(r-k) \times l}) \mid j \in [N]\}$ and evaluate

$$\begin{aligned} \sin \angle_i \left(\mathbf{U}_k, \widehat{\mathbf{U}}_l \right) &\approx \alpha_i = \frac{1}{N} \sum_{j=1}^N \left(1 + \sigma_i^2 \left(\Sigma_k^{2q+1} \Omega_1^{(j)} \left(\Sigma_{r \setminus k}^{2q+1} \Omega_2^{(j)} \right)^\dagger \right) \right)^{-\frac{1}{2}}, \\ \sin \angle_i \left(\mathbf{V}_k, \widehat{\mathbf{V}}_l \right) &\approx \beta_i = \frac{1}{N} \sum_{j=1}^N \left(1 + \sigma_i^2 \left(\Sigma_k^{2q+2} \Omega_1^{(j)} \left(\Sigma_{r \setminus k}^{2q+2} \Omega_2^{(j)} \right)^\dagger \right) \right)^{-\frac{1}{2}}, \end{aligned}$$

for all $i \in [k]$, which can be conducted efficiently in $O(Nrl^2)$ time. Algorithm 2 demonstrates the construction of unbiased estimates for $\mathbb{E}[\sin \angle_i(\mathbf{U}_k, \widehat{\mathbf{U}}_l)] = \alpha_i$, while the unbiased estimates for $\mathbb{E}[\sin \angle_i(\mathbf{V}_k, \widehat{\mathbf{V}}_l)] = \beta_i$ can be evaluated analogously by replacing Line 4 with $\widetilde{\Omega}_1^{(j)} = \Sigma_k^{2q+2} \Omega_1^{(j)}(1:k,:)$, $\widetilde{\Omega}_2^{(j)} = \Sigma_{r \setminus k}^{2q+2} \Omega_2^{(j)}(k+1:r,:)$.

Algorithm 2 Unbiased canonical angle estimates

Require: (Exact or estimated) singular values Σ , rank k , sample size $l \geq k$, number of power iterations q , number of trials N

Ensure: Unbiased estimates $\mathbb{E}[\sin \angle_i(\mathbf{U}_k, \widehat{\mathbf{U}}_l)] = \alpha_i$ for all $i \in [k]$

- 1: Partition Σ into $\Sigma_k = \Sigma(1:k, 1:k)$ and $\Sigma_{r \setminus k} = \Sigma(k+1:r, k+1:r)$
- 2: **for** $j = 1, \dots, N$ **do**
- 3: Draw $\Omega^{(j)} \sim P(\mathbb{C}^{r \times l})$ such that $\Omega_{ij}^{(j)} \sim \mathcal{N}(0, l^{-1})$ i.i.d.
- 4: $\tilde{\Omega}_1^{(j)} = \Sigma_k^{2q+1} \Omega^{(j)}(1:k, :)$, $\tilde{\Omega}_2^{(j)} = \Sigma_{r \setminus k}^{2q+1} \Omega^{(j)}(k+1:r, :)$
- 5: $[\mathbf{U}_{\tilde{\Omega}_2^{(j)}}, \Sigma_{\tilde{\Omega}_2^{(j)}}, \mathbf{V}_{\tilde{\Omega}_2^{(j)}}, \mathbf{U}_{\tilde{\Omega}_2^{(j)}}^*] = \text{svd}(\tilde{\Omega}_2^{(j)}, \text{“econ”})$
- 6: $\nu^{(j)} = \text{svd}(\tilde{\Omega}_1^{(j)} \mathbf{V}_{\tilde{\Omega}_2^{(j)}} \Sigma_{\tilde{\Omega}_2^{(j)}}^{-1} \mathbf{U}_{\tilde{\Omega}_2^{(j)}}^*)$
- 7: $\theta_i^{(j)} = 1/\sqrt{1 + (\nu_i^{(j)})^2}$ for all $i \in [k]$
- 8: **end for**
- 9: $\alpha_i = \frac{1}{N} \sum_{j=1}^N \theta_i^{(j)}$ for all $i \in [k]$

Experiments in section 6 show that the unbiased estimates concentrate well in practice: a sample size as small as $N = 3$ is seen to be sufficient to provide good estimates. Further, with independent Gaussian random matrices, the unbiased estimates in Proposition 1 are space-agnostic; i.e., (4.1) and (4.2) only depend on the spectrum Σ but not on the unknown true singular subspaces \mathbf{U} and \mathbf{V} . As in Theorem 1, the true singular values may be replaced by their approximations from Algorithm 1 (cf. subsection 6.2).

Proof of Proposition 1. To show (4.1), we recall from the proof of Theorem 1 that, for the rank- l randomized SVD with a Gaussian embedding $\Omega \sim P(\mathbb{C}^{n \times l})$ and q power iterations, $\Omega_1 \triangleq \mathbf{V}_k^* \Omega$ and $\Omega_2 \triangleq \mathbf{V}_{r \setminus k}^* \Omega$ are independent Gaussian random matrices with the same entrywise distribution as Ω . Therefore, with $\text{rank}(\mathbf{A}) = r$, for all $i \in [k]$,

$$\begin{aligned} & \mathbb{E}_\Omega [\sin \angle_i(\mathbf{U}_k, \widehat{\mathbf{U}}_l)] \quad (\text{Recall (3.5)}) \\ &= \mathbb{E}_\Omega \left[\sigma_{k-i+1}^{\frac{1}{2}} \left(\left(\mathbf{I}_k + \Sigma_k^{2q+1} \Omega_1 \left(\Omega_2^* \Sigma_{r \setminus k}^{4q+2} \Omega_2 \right)^{-1} \Omega_1^* \Sigma_k^{2q+1} \right)^{-1} \right) \right] \\ &= \mathbb{E}_\Omega \left[\sigma_i^{-\frac{1}{2}} \left(\mathbf{I}_k + \Sigma_k^{2q+1} \Omega_1 \left(\Omega_2^* \Sigma_{r \setminus k}^{4q+2} \Omega_2 \right)^{-1} \Omega_1^* \Sigma_k^{2q+1} \right) \right] \\ &= \mathbb{E}_{\Omega'_1, \Omega'_2} \left[\sigma_i^{-\frac{1}{2}} \left(\mathbf{I}_k + \Sigma_k^{2q+1} \Omega'_1 \left(\Omega'_2^* \Sigma_{r \setminus k}^{4q+2} \Omega'_2 \right)^{-1} \Omega'_1^* \Sigma_k^{2q+1} \right) \right]. \end{aligned}$$

The unbiased estimate in (4.2) follows analogously. \square

As a side note, we point out that, compared with the probabilistic upper bounds (3.1) and (3.2), the estimates (4.1) and (4.2) circumvent overestimation from the operator-convexity of inversion $\sigma \rightarrow \sigma^{-1}$,

$$\mathbb{E}_{\Omega'_2} \left[\left(\Omega'_2^* \Sigma_{r \setminus k}^{4q+2} \Omega'_2 \right)^{-1} \right] \succcurlyeq \left(\mathbb{E}_{\Omega'_2} [\Omega'_2^* \Sigma_{r \setminus k}^{4q+2} \Omega'_2] \right)^{-1},$$

which implies that

$$\begin{aligned} & \mathbb{E}_{\Omega'_1, \Omega'_2} \left[\left(\mathbf{I}_k + \Sigma_k^{2q+1} \Omega'_1 \left(\Omega_2'^* \Sigma_{r \setminus k}^{4q+2} \Omega'_2 \right)^{-1} \Omega_1'^* \Sigma_k^{2q+1} \right)^{-1} \right] \\ & \geq \mathbb{E}_{\Omega'_1} \left[\left(\mathbf{I}_k + \Sigma_k^{2q+1} \Omega'_1 \left(\mathbb{E}_{\Omega'_2} \left[\Omega_2'^* \Sigma_{r \setminus k}^{4q+2} \Omega'_2 \right] \right)^{-1} \Omega_1'^* \Sigma_k^{2q+1} \right)^{-1} \right]. \end{aligned}$$

5. Posterior residual-based bounds. In addition to the prior probabilistic bounds and unbiased estimates, in this section, we introduce two sets of posterior guarantees for the canonical angles that hold deterministically and can be evaluated/approximated efficiently based on the residual of the resulting low-rank approximation $\hat{\mathbf{A}}_l = \hat{\mathbf{U}}_l \hat{\Sigma}_l \hat{\mathbf{V}}_l^*$ from Algorithm 1.

Remark 1 (generality of residual-based bounds). It is worth highlighting that both the statements and the proofs of the posterior guarantees Theorems 2 and 4 to be presented are algorithm-independent. In contrast to Theorem 1 and Proposition 1 whose derivation depends explicitly on the algorithm (e.g., assuming Ω being Gaussian in Algorithm 1), the residual-based bounds in Theorems 2 and 4 hold for general approximated low-rank SVDs, $\mathbf{A} \approx \hat{\mathbf{U}}_l \hat{\Sigma}_l \hat{\mathbf{V}}_l^T$.

We start with the following proposition, which establishes relations between the canonical angles, the residuals, and the true spectrum $\sigma(\mathbf{A})$.

THEOREM 2. *Given any $\hat{\mathbf{U}}_l \in \mathbb{C}^{m \times l}$ and $\hat{\mathbf{V}}_l \in \mathbb{C}^{n \times l}$ with orthonormal columns such that $\text{Range}(\hat{\mathbf{U}}_l) \subseteq \text{Col}(\mathbf{A})$ and $\text{Range}(\hat{\mathbf{V}}_l) \subseteq \text{Row}(\mathbf{A})$, we have, for each $i = 1, \dots, k$ ($k \leq l$),*

$$(5.1) \quad \sin \angle_i (\mathbf{U}_k, \hat{\mathbf{U}}_l) \leq \min \left\{ \frac{\sigma_{k-i+1} \left((\mathbf{I}_m - \hat{\mathbf{U}}_l \hat{\mathbf{U}}_l^*) \mathbf{A} \right)}{\sigma_k}, \frac{\sigma_1 \left((\mathbf{I}_m - \hat{\mathbf{U}}_l \hat{\mathbf{U}}_l^*) \mathbf{A} \right)}{\sigma_i} \right\},$$

while

$$(5.2) \quad \sin \angle_i (\mathbf{V}_k, \hat{\mathbf{V}}_l) \leq \min \left\{ \frac{\sigma_{k-i+1} \left(\mathbf{A} \left(\mathbf{I}_n - \hat{\mathbf{V}}_l \hat{\mathbf{V}}_l^* \right) \right)}{\sigma_k}, \frac{\sigma_1 \left(\mathbf{A} \left(\mathbf{I}_n - \hat{\mathbf{V}}_l \hat{\mathbf{V}}_l^* \right) \right)}{\sigma_i} \right\}.$$

It is worth highlighting that (5.1) and (5.2) are independent of the unknown true singular subspaces \mathbf{U}_k and \mathbf{V}_k but only depend on the singular values Σ_k (which may be replaced with the approximations $\hat{\Sigma}_l$ from Algorithm 1; cf. subsection 6.2) and approximated singular subspaces $\hat{\mathbf{U}}_l$ and $\hat{\mathbf{V}}_l$. Therefore, these upper bounds can be evaluated/approximated efficiently based on the residual of the resulting low-rank approximation $\hat{\mathbf{A}}_l = \hat{\mathbf{U}}_l \hat{\Sigma}_l \hat{\mathbf{V}}_l^*$.

Remark 3 (left versus right singular subspaces). When $\hat{\mathbf{U}}_l$ and $\hat{\mathbf{V}}_l$ consist of approximated left and right singular vectors from Algorithm 1, upper bounds on $\sin \angle_i (\mathbf{V}_k, \hat{\mathbf{V}}_l)$ tend to be smaller than those on $\sin \angle_i (\mathbf{U}_k, \hat{\mathbf{U}}_l)$. This is induced by the algorithmic fact that, in Algorithm 1, $\hat{\mathbf{V}}_l$ is an orthonormal basis of $\mathbf{A}^* \mathbf{Q}_X$, while $\hat{\mathbf{U}}_l$ and \mathbf{Q}_X are orthonormal bases of $\mathbf{X} = \mathbf{A} \Omega$. That is, the evaluation of $\hat{\mathbf{V}}_l$ is enhanced by an additional half power iteration compared with that of $\hat{\mathbf{U}}_l$, which is also reflected by the differences in exponents on Σ (i.e., $2q+1$ versus $2q+2$) in Theorem 1 and Proposition 1. This difference can be important especially when q is small (e.g., $q=0$). When higher accuracy in the left singular subspace is desirable, one can work with \mathbf{A}^* .

Proof of Theorem 2. Starting with the leading left singular subspace, by definition, for each $i = 1, \dots, k$, we have

$$\begin{aligned}\sin \angle_i (\mathbf{U}_k, \widehat{\mathbf{U}}_l) &= \sigma_{k-i+1} \left(\left(\mathbf{I}_m - \widehat{\mathbf{U}}_l \widehat{\mathbf{U}}_l^* \right) \mathbf{U}_k \right) \\ &= \sigma_{k-i+1} \left(\left(\mathbf{I}_m - \widehat{\mathbf{U}}_l \widehat{\mathbf{U}}_l^* \right) \mathbf{A} (\mathbf{V} \Sigma^{-1} \mathbf{U}^* \mathbf{U}_k) \right) \\ &= \sigma_{k-i+1} \left(\left(\left(\mathbf{I}_m - \widehat{\mathbf{U}}_l \widehat{\mathbf{U}}_l^* \right) \mathbf{A} \mathbf{V}_k \right) \Sigma_k^{-1} \right).\end{aligned}$$

Then, we observe that the following hold simultaneously:

$$\begin{aligned}\sigma_{k-i+1} \left(\left(\left(\mathbf{I}_m - \widehat{\mathbf{U}}_l \widehat{\mathbf{U}}_l^* \right) \mathbf{A} \mathbf{V}_k \right) \Sigma_k^{-1} \right) &\leq \sigma_1 \left(\left(\mathbf{I}_m - \widehat{\mathbf{U}}_l \widehat{\mathbf{U}}_l^* \right) \mathbf{A} \mathbf{V}_k \right) \cdot \sigma_{k-i+1} (\Sigma_k^{-1}), \\ \sigma_{k-i+1} \left(\left(\left(\mathbf{I}_m - \widehat{\mathbf{U}}_l \widehat{\mathbf{U}}_l^* \right) \mathbf{A} \mathbf{V}_k \right) \Sigma_k^{-1} \right) &\leq \sigma_{k-i+1} \left(\left(\mathbf{I}_m - \widehat{\mathbf{U}}_l \widehat{\mathbf{U}}_l^* \right) \mathbf{A} \mathbf{V}_k \right) \cdot \sigma_1 (\Sigma_k^{-1}),\end{aligned}$$

where $\sigma_{k-i+1} (\Sigma_k^{-1}) = 1/\sigma_i$ and $\sigma_1 (\Sigma_k^{-1}) = 1/\sigma_k$. Finally by Lemma 2, we have

$$\sigma_{k-i+1} \left(\left(\mathbf{I}_m - \widehat{\mathbf{U}}_l \widehat{\mathbf{U}}_l^* \right) \mathbf{A} \mathbf{V}_k \right) \leq \sigma_{k-i+1} \left(\left(\mathbf{I}_m - \widehat{\mathbf{U}}_l \widehat{\mathbf{U}}_l^* \right) \mathbf{A} \right).$$

Meanwhile, the upper bound for the leading right singular subspace can be derived analogously by observing that

$$\sin \angle_i (\mathbf{V}_k, \widehat{\mathbf{V}}_l) = \sigma_{k-i+1} \left(\mathbf{V}_k^* \left(\mathbf{I}_n - \widehat{\mathbf{V}}_l \widehat{\mathbf{V}}_l^* \right) \right) = \sigma_{k-i+1} \left(\Sigma_k^{-1} \mathbf{U}_k^* \mathbf{A} \left(\mathbf{I}_n - \widehat{\mathbf{V}}_l \widehat{\mathbf{V}}_l^* \right) \right). \square$$

As a potential drawback, although the residuals $(\mathbf{I}_m - \widehat{\mathbf{U}}_l \widehat{\mathbf{U}}_l^*) \mathbf{A}$ and $\mathbf{A} (\mathbf{I}_n - \widehat{\mathbf{V}}_l \widehat{\mathbf{V}}_l^*)$ in Theorem 2 can be evaluated efficiently in $O(mn)$ and $O(mnl)$ time,⁶ respectively, the exact evaluation of their full spectra can be unaffordable. A straightforward remedy for this problem is to use only the second terms in the right-hand sides of (5.1) and (5.2) while estimating $\|(\mathbf{I}_m - \widehat{\mathbf{U}}_l \widehat{\mathbf{U}}_l^*) \mathbf{A}\|_2$ and $\|\mathbf{A} (\mathbf{I}_n - \widehat{\mathbf{V}}_l \widehat{\mathbf{V}}_l^*)\|_2$ with the randomized power method (cf. [16], [20, Algorithm 4]). Alternatively, we leverage the analysis from [24, Theorem 6.1] and present the following posterior bounds based only on norms of the residuals which can be estimated efficiently via sampling.

THEOREM 4. *For any approximated rank- l SVD $\mathbf{A} \approx \widehat{\mathbf{U}}_l \widehat{\Sigma}_l \widehat{\mathbf{V}}_l^*$ (not necessarily obtained by Algorithm 1), recall the notation that $\widehat{\mathbf{U}}_l = [\widehat{\mathbf{U}}_k, \widehat{\mathbf{U}}_{l \setminus k}]$, $\widehat{\mathbf{V}}_l = [\widehat{\mathbf{V}}_k, \widehat{\mathbf{V}}_{l \setminus k}]$, while $\widehat{\mathbf{U}}_{m \setminus l}$, $\widehat{\mathbf{U}}_{n \setminus l}$ are the orthogonal complements of $\widehat{\mathbf{U}}_l$, $\widehat{\mathbf{V}}_l$, respectively. Then, with $\mathbf{E}_{31} \triangleq \widehat{\mathbf{U}}_{m \setminus l}^* \mathbf{A} \widehat{\mathbf{V}}_k$, $\mathbf{E}_{32} \triangleq \widehat{\mathbf{U}}_{m \setminus l}^* \mathbf{A} \widehat{\mathbf{V}}_{l \setminus k}$, and $\mathbf{E}_{33} \triangleq \widehat{\mathbf{U}}_{m \setminus l}^* \mathbf{A} \widehat{\mathbf{V}}_{n \setminus l}$, assuming $\sigma_k > \widehat{\sigma}_{k+1}$ and $\sigma_k > \|\mathbf{E}_{33}\|_2$, we define the spectral gaps*

$$\gamma_1 \triangleq \frac{\sigma_k^2 - \widehat{\sigma}_{k+1}^2}{\sigma_k}, \quad \gamma_2 \triangleq \frac{\sigma_k^2 - \widehat{\sigma}_{k+1}^2}{\widehat{\sigma}_{k+1}}, \quad \Gamma_1 = \frac{\sigma_k^2 - \|\mathbf{E}_{33}\|_2^2}{\sigma_k}, \quad \Gamma_2 = \frac{\sigma_k^2 - \|\mathbf{E}_{33}\|_2^2}{\|\mathbf{E}_{33}\|_2}.$$

Then for an arbitrary unitarily invariant norm $\|\cdot\|$,

$$(5.3) \quad \left\| \sin \angle_i (\mathbf{U}_k, \widehat{\mathbf{U}}_l) \right\| \leq \frac{\|[\mathbf{E}_{31}, \mathbf{E}_{32}]\|}{\Gamma_1},$$

⁶For the $O(mn)$ complexity of computing $(\mathbf{I}_m - \widehat{\mathbf{U}}_l \widehat{\mathbf{U}}_l^*) \mathbf{A}$, we assume that $\widehat{\mathbf{A}}_l = \widehat{\mathbf{U}}_l \widehat{\mathbf{U}}_l^* \mathbf{A} = \widehat{\mathbf{U}}_l \widehat{\Sigma}_l \widehat{\mathbf{V}}_l^*$ is readily available from Algorithm 1. Otherwise (e.g., when Algorithm 1 returns $(\widehat{\mathbf{U}}_l, \widehat{\Sigma}_l, \widehat{\mathbf{V}}_l)$ but not $\widehat{\mathbf{A}}_l$), the evaluation of $(\mathbf{I}_m - \widehat{\mathbf{U}}_l \widehat{\mathbf{U}}_l^*) \mathbf{A}$ will inevitably take $O(mnl)$ as that of $\mathbf{A} (\mathbf{I}_n - \widehat{\mathbf{V}}_l \widehat{\mathbf{V}}_l^*)$.

$$(5.4) \quad \left\| \sin \angle (\mathbf{V}_k, \widehat{\mathbf{V}}_l) \right\| \leq \frac{\|[\mathbf{E}_{31}, \mathbf{E}_{32}]\|}{\Gamma_2},$$

and specifically for the spectral or Frobenius norm $\|\cdot\|_\xi$ ($\xi = 2, F$),

$$(5.5) \quad \left\| \sin \angle (\mathbf{U}_k, \widehat{\mathbf{U}}_l) \right\|_\xi \leq \frac{\|[\mathbf{E}_{31}, \mathbf{E}_{32}]\|_\xi}{\Gamma_1} \sqrt{1 + \frac{\|\mathbf{E}_{32}\|_2^2}{\gamma_2^2}},$$

$$(5.6) \quad \left\| \sin \angle (\mathbf{V}_k, \widehat{\mathbf{V}}_l) \right\|_\xi \leq \frac{\|[\mathbf{E}_{31}, \mathbf{E}_{32}]\|_\xi}{\Gamma_1} \sqrt{\frac{\|\mathbf{E}_{32}\|_2^2}{\gamma_1^2} + \frac{\|\mathbf{E}_{33}\|_2^2}{\sigma_k^2}}.$$

Furthermore, for all $i \in [k]$,

$$(5.7) \quad \sin \angle_i (\mathbf{U}_k, \widehat{\mathbf{U}}_l) \leq \frac{\sigma_k}{\sigma_{k-i+1}} \cdot \frac{\|[\mathbf{E}_{31}, \mathbf{E}_{32}]\|_2}{\Gamma_1},$$

$$(5.8) \quad \sin \angle_i (\mathbf{V}_k, \widehat{\mathbf{V}}_l) \leq \frac{\sigma_k}{\sigma_{k-i+1}} \cdot \frac{\|[\mathbf{E}_{31}, \mathbf{E}_{32}]\|_2}{\Gamma_2},$$

$$(5.9) \quad \sin \angle_i (\mathbf{U}_k, \widehat{\mathbf{U}}_k) \leq \frac{\|[\mathbf{E}_{31}, \mathbf{E}_{32}]\|_2}{\Gamma_1} \sqrt{1 + \left(\frac{\sigma_k}{\sigma_{k-i+1}} \cdot \frac{\|\mathbf{E}_{32}\|_2}{\gamma_2} \right)^2},$$

$$(5.10) \quad \sin \angle_i (\mathbf{V}_k, \widehat{\mathbf{V}}_k) \leq \frac{\|[\mathbf{E}_{31}, \mathbf{E}_{32}]\|_2}{\Gamma_1} \sqrt{\left(\frac{\sigma_k}{\sigma_{k-i+1}} \cdot \frac{\|\mathbf{E}_{32}\|_2}{\gamma_1} \right)^2 + \left(\frac{\|\mathbf{E}_{33}\|_2}{\sigma_k} \right)^2}.$$

In practice, norms of the residuals can be computed as

$$\|[\mathbf{E}_{31}, \mathbf{E}_{32}]\| = \left\| \widehat{\mathbf{U}}_{m \setminus l}^* \mathbf{A} \widehat{\mathbf{V}}_l \right\| = \left\| (\mathbf{I}_m - \widehat{\mathbf{U}}_l \widehat{\mathbf{U}}_l^*) \mathbf{A} \widehat{\mathbf{V}}_l \right\| = \left\| (\mathbf{A} - \widehat{\mathbf{A}}_l) \widehat{\mathbf{V}}_l \right\|,$$

$$\|\mathbf{E}_{32}\|_2 = \left\| \widehat{\mathbf{U}}_{m \setminus l}^* \mathbf{A} \widehat{\mathbf{V}}_{l \setminus k} \right\|_2 = \left\| (\mathbf{A} - \widehat{\mathbf{A}}_l) \widehat{\mathbf{V}}_{l \setminus k} \right\|_2,$$

$$\|\mathbf{E}_{33}\|_2 = \left\| \widehat{\mathbf{U}}_{m \setminus l}^* \mathbf{A} \widehat{\mathbf{V}}_{n \setminus l} \right\|_2 = \left\| (\mathbf{A} - \widehat{\mathbf{A}}_l) (\mathbf{I}_n - \widehat{\mathbf{V}}_l \widehat{\mathbf{V}}_l^*) \right\|_2 = \left\| \mathbf{A} - \mathbf{A} \widehat{\mathbf{V}}_l \widehat{\mathbf{V}}_l^* \right\|_2,$$

where the construction of $(\mathbf{A} - \widehat{\mathbf{A}}_l) \widehat{\mathbf{V}}_l$, $(\mathbf{A} - \widehat{\mathbf{A}}_l) \widehat{\mathbf{V}}_{l \setminus k}$, and $\mathbf{A} - \mathbf{A} \widehat{\mathbf{V}}_l \widehat{\mathbf{V}}_l^*$ takes $O(mnl)$ time, while the respective norms can be estimated efficiently via sampling (cf. [20, Algorithms 1–4], [23], etc.). With an unknown true spectrum in practice, replacing the true singular values in Theorem 4 with their corresponding approximations from Algorithm 1 usually yields similar upper bounds (cf. subsection 6.2).

The proof of Theorem 4 is similar to that of [24, Thm. 6.1].

Proof of Theorem 4. Let $\tilde{\mathbf{U}}_{11} \triangleq \widehat{\mathbf{U}}_k^* \mathbf{U}_k$, $\tilde{\mathbf{U}}_{21} \triangleq \widehat{\mathbf{U}}_{l \setminus k}^* \mathbf{U}_k$, $\tilde{\mathbf{U}}_{31} \triangleq \widehat{\mathbf{U}}_{m \setminus l}^* \mathbf{U}_k$, and $\tilde{\mathbf{V}}_{11} \triangleq \widehat{\mathbf{V}}_k^* \mathbf{V}_k$, $\tilde{\mathbf{V}}_{21} \triangleq \widehat{\mathbf{V}}_{l \setminus k}^* \mathbf{V}_k$, and $\tilde{\mathbf{V}}_{31} \triangleq \widehat{\mathbf{V}}_{n \setminus l}^* \mathbf{V}_k$. We start by expressing the canonical angles in terms of $\tilde{\mathbf{U}}_{31}$ and $\tilde{\mathbf{U}}_{21}$:

$$\sin \angle (\mathbf{U}_k, \widehat{\mathbf{U}}_l) = \sigma (\widehat{\mathbf{U}}_{m \setminus l}^* \mathbf{U}_k) = \sigma (\tilde{\mathbf{U}}_{31}),$$

$$\sin \angle (\mathbf{V}_k, \widehat{\mathbf{V}}_l) = \sigma (\widehat{\mathbf{V}}_{n \setminus l}^* \mathbf{V}_k) = \sigma (\tilde{\mathbf{V}}_{31}),$$

$$\sin \angle (\mathbf{U}_k, \widehat{\mathbf{U}}_k) = \sigma \left(\begin{bmatrix} \widehat{\mathbf{U}}_{l \setminus k}^* \\ \widehat{\mathbf{U}}_{m \setminus l}^* \end{bmatrix} \mathbf{U}_k \right) = \sigma \left(\begin{bmatrix} \tilde{\mathbf{U}}_{21} \\ \tilde{\mathbf{U}}_{31} \end{bmatrix} \right),$$

$$\sin \angle (\mathbf{V}_k, \widehat{\mathbf{V}}_k) = \sigma \left(\begin{bmatrix} \widehat{\mathbf{V}}_{l \setminus k}^* \\ \widehat{\mathbf{V}}_{n \setminus l}^* \end{bmatrix} \mathbf{V}_k \right) = \sigma \left(\begin{bmatrix} \tilde{\mathbf{V}}_{21} \\ \tilde{\mathbf{V}}_{31} \end{bmatrix} \right).$$

By observing that for any rank- l approximation in the SVD form $\mathbf{A} \approx \widehat{\mathbf{U}}_l \widehat{\Sigma}_l \widehat{\mathbf{V}}_l^*$,

$$\mathbf{A} = \widehat{\mathbf{U}}_l \widehat{\Sigma}_l \widehat{\mathbf{V}}_l^* + \widehat{\mathbf{U}}_{m \setminus l} \widehat{\mathbf{U}}_{m \setminus l}^* \mathbf{A} = \begin{bmatrix} \widehat{\mathbf{U}}_k & \widehat{\mathbf{U}}_{l \setminus k} & \widehat{\mathbf{U}}_{m \setminus l} \end{bmatrix} \begin{bmatrix} \widehat{\Sigma}_k & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \widehat{\Sigma}_{l \setminus k} & \mathbf{0} \\ \mathbf{E}_{31} & \mathbf{E}_{32} & \mathbf{E}_{33} \end{bmatrix} \begin{bmatrix} \widehat{\mathbf{V}}_k^* \\ \widehat{\mathbf{V}}_{l \setminus k}^* \\ \widehat{\mathbf{V}}_{n \setminus l}^* \end{bmatrix},$$

we left multiply by $\widehat{\mathbf{U}}^* = [\widehat{\mathbf{U}}_k, \widehat{\mathbf{U}}_{l \setminus k}, \widehat{\mathbf{U}}_{m \setminus l}]^* \in \mathbb{C}^{m \times m}$ and right multiply by \mathbf{V}_k on both sides and get

$$(5.11) \quad \begin{bmatrix} \widetilde{\mathbf{U}}_{11} \\ \widetilde{\mathbf{U}}_{21} \\ \widetilde{\mathbf{U}}_{31} \end{bmatrix} \Sigma_k = \begin{bmatrix} \widehat{\Sigma}_k & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \widehat{\Sigma}_{l \setminus k} & \mathbf{0} \\ \mathbf{E}_{31} & \mathbf{E}_{32} & \mathbf{E}_{33} \end{bmatrix} \begin{bmatrix} \widetilde{\mathbf{V}}_{11} \\ \widetilde{\mathbf{V}}_{21} \\ \widetilde{\mathbf{V}}_{31} \end{bmatrix},$$

while left multiplying \mathbf{U}_k^* and right multiplying $\widehat{\mathbf{V}} = [\widehat{\mathbf{V}}_k, \widehat{\mathbf{V}}_{l \setminus k}, \widehat{\mathbf{V}}_{n \setminus l}]$ yield

$$(5.12) \quad \Sigma_k \begin{bmatrix} \widetilde{\mathbf{V}}_{11}^* & \widetilde{\mathbf{V}}_{21}^* & \widetilde{\mathbf{V}}_{31}^* \end{bmatrix} = \begin{bmatrix} \widetilde{\mathbf{U}}_{11}^* & \widetilde{\mathbf{U}}_{21}^* & \widetilde{\mathbf{U}}_{31}^* \end{bmatrix} \begin{bmatrix} \widehat{\Sigma}_k & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \widehat{\Sigma}_{l \setminus k} & \mathbf{0} \\ \mathbf{E}_{31} & \mathbf{E}_{32} & \mathbf{E}_{33} \end{bmatrix}.$$

Bounding $\sigma(\widetilde{\mathbf{U}}_{31})$ and $\sigma(\widetilde{\mathbf{V}}_{31})$. To bound $\sigma(\widetilde{\mathbf{U}}_{31})$, we observe the following from the third row of (5.11) and the third column of (5.12),

$$\widetilde{\mathbf{U}}_{31} \Sigma_k = \mathbf{E}_{31} \widetilde{\mathbf{V}}_{11} + \mathbf{E}_{32} \widetilde{\mathbf{V}}_{21} + \mathbf{E}_{33} \widetilde{\mathbf{V}}_{31}, \quad \widetilde{\mathbf{U}}_{31}^* \mathbf{E}_{33} = \Sigma_k \widetilde{\mathbf{V}}_{31}^*.$$

Noticing that $[\begin{smallmatrix} \widetilde{\mathbf{V}}_{11} \\ \widetilde{\mathbf{V}}_{21} \end{smallmatrix}] = \widehat{\mathbf{V}}_l^* \mathbf{V}_k$ and $\|\widehat{\mathbf{V}}_l^* \mathbf{V}_k\|_2 \leq 1$, we have

$$\begin{aligned} \|\widetilde{\mathbf{U}}_{31} \Sigma_k\| &\leq \|[\mathbf{E}_{31}, \mathbf{E}_{32}]\| \|\widehat{\mathbf{V}}_l^* \mathbf{V}_k\|_2 + \|[\mathbf{E}_{33} \mathbf{E}_{33}^* \widetilde{\mathbf{U}}_{31} \Sigma_k^{-1}]\| \\ &\leq \|[\mathbf{E}_{31}, \mathbf{E}_{32}]\| + \frac{\|\mathbf{E}_{33}\|_2^2}{\sigma_k^2} \|\widetilde{\mathbf{U}}_{31} \Sigma_k\| \end{aligned}$$

for all $i \in [\min(k, m - l)]$, which implies that

$$\|\widetilde{\mathbf{U}}_{31} \Sigma_k\| \leq \left(1 - \frac{\|\mathbf{E}_{33}\|_2^2}{\sigma_k^2}\right)^{-1} \|[\mathbf{E}_{31}, \mathbf{E}_{32}]\| = \sigma_k \cdot \frac{\|[\mathbf{E}_{31}, \mathbf{E}_{32}]\|}{\Gamma_1},$$

and leads to

$$\begin{aligned} \|\widetilde{\mathbf{U}}_{31}\| &\leq \frac{1}{\sigma_k} \|\widetilde{\mathbf{U}}_{31} \Sigma_k\| \leq \frac{\|[\mathbf{E}_{31}, \mathbf{E}_{32}]\|}{\Gamma_1}, \\ \sigma_i(\widetilde{\mathbf{U}}_{31}) &\leq \frac{1}{\sigma_{k-i+1}} \|\widetilde{\mathbf{U}}_{31} \Sigma_k\|_2 \leq \frac{\sigma_k}{\sigma_{k-i+1}} \cdot \frac{\|[\mathbf{E}_{31}, \mathbf{E}_{32}]\|_2}{\Gamma_1} \quad \forall i \in [k], \end{aligned}$$

where the second line follows from Lemma 3. These lead to (5.3) and (5.7).

To bound $\sigma(\widetilde{\mathbf{V}}_{31})$, we use the relation $\widetilde{\mathbf{U}}_{31}^* \mathbf{E}_{33} = \Sigma_k \widetilde{\mathbf{V}}_{31}^*$,

$$\begin{aligned} \|\widetilde{\mathbf{V}}_{31}\| &\leq \frac{\|\mathbf{E}_{33}\|_2}{\sigma_k} \|\widetilde{\mathbf{U}}_{31}\| \leq \frac{\|[\mathbf{E}_{31}, \mathbf{E}_{32}]\|}{\Gamma_2}, \\ \sigma_i(\widetilde{\mathbf{V}}_{31}) &\leq \frac{1}{\sigma_k} \sigma_i(\mathbf{E}_{33}^* \widetilde{\mathbf{U}}_{31}) \leq \frac{\|\mathbf{E}_{33}\|_2}{\sigma_k} \sigma_i(\widetilde{\mathbf{U}}_{31}) \leq \frac{\sigma_k}{\sigma_{k-i+1}} \cdot \frac{\|[\mathbf{E}_{31}, \mathbf{E}_{32}]\|_2}{\Gamma_2} \quad \forall i \in [k]. \end{aligned}$$

We therefore have upper bounds (5.4) and (5.8).

Bounding $\sigma(\tilde{\mathbf{U}}_{21})$ and $\sigma(\tilde{\mathbf{V}}_{21})$. To bound $\sigma(\tilde{\mathbf{U}}_{21})$, we leverage the second row of (5.11) and the second column of (5.12),

$$\tilde{\mathbf{U}}_{21}\Sigma_k = \hat{\Sigma}_{l \setminus k} \tilde{\mathbf{V}}_{21}, \quad \Sigma_k \tilde{\mathbf{V}}_{21}^* = \tilde{\mathbf{U}}_{21}^* \hat{\Sigma}_{l \setminus k} + \tilde{\mathbf{U}}_{31}^* \mathbf{E}_{32}.$$

Up to rearrangement, we observe that

$$\begin{aligned} \|\tilde{\mathbf{U}}_{21}\Sigma_k\| &= \left\| \hat{\Sigma}_{l \setminus k} \left(\hat{\Sigma}_{l \setminus k} \tilde{\mathbf{U}}_{21} + \mathbf{E}_{32}^* \tilde{\mathbf{U}}_{31} \right) \Sigma_k^{-1} \right\| \\ &\leq \frac{\|\hat{\Sigma}_{l \setminus k}\|_2^2}{\sigma_k^2} \|\tilde{\mathbf{U}}_{21}\Sigma_k\| + \frac{\|\hat{\Sigma}_{l \setminus k}\|_2}{\sigma_k} \|\mathbf{E}_{32}^* \tilde{\mathbf{U}}_{31}\|, \end{aligned}$$

which implies that

$$\|\tilde{\mathbf{U}}_{21}\Sigma_k\| \leq \left(1 - \frac{\hat{\sigma}_{k+1}^2}{\sigma_k^2} \right)^{-1} \frac{\hat{\sigma}_{k+1}}{\sigma_k} \|\mathbf{E}_{32}^* \tilde{\mathbf{U}}_{31}\| \leq \sigma_k \cdot \frac{\|\mathbf{E}_{32}\|_2}{\gamma_2} \|\tilde{\mathbf{U}}_{31}\|,$$

and therefore, with Lemma 3, for all $i \in [k]$,

$$\begin{aligned} \|\tilde{\mathbf{U}}_{21}\| &\leq \frac{1}{\sigma_k} \|\tilde{\mathbf{U}}_{21}\Sigma_k\| \leq \frac{\|\mathbf{E}_{32}\|_2}{\gamma_2} \|\tilde{\mathbf{U}}_{31}\|, \\ \sigma_i(\tilde{\mathbf{U}}_{21}) &\leq \frac{1}{\sigma_{k-i+1}} \|\tilde{\mathbf{U}}_{21}\Sigma_k\|_2 \leq \frac{\sigma_k}{\sigma_{k-i+1}} \cdot \frac{\|\mathbf{E}_{32}\|_2}{\gamma_2} \|\tilde{\mathbf{U}}_{31}\|_2. \end{aligned}$$

Then, with the stronger inequality for the spectral or Frobenius norm $\|\cdot\|_\xi$ ($\xi = 2, F$),

$$\begin{aligned} \left\| \begin{bmatrix} \tilde{\mathbf{U}}_{21} \\ \tilde{\mathbf{U}}_{31} \end{bmatrix} \right\|_\xi &\leq \sqrt{\|\tilde{\mathbf{U}}_{31}\|_\xi^2 + \|\tilde{\mathbf{U}}_{21}\|_\xi^2} \leq \|\tilde{\mathbf{U}}_{31}\|_\xi \sqrt{1 + \frac{\|\mathbf{E}_{32}\|_2^2}{\gamma_2^2}} \\ &\leq \frac{\|[\mathbf{E}_{31}, \mathbf{E}_{32}]\|_\xi}{\Gamma_1} \sqrt{1 + \frac{\|\mathbf{E}_{32}\|_2^2}{\gamma_2^2}} \end{aligned}$$

leads to (5.5). Meanwhile for (5.9), the individual canonical angles are upper bounded by

$$\begin{aligned} \sigma_i \left(\begin{bmatrix} \tilde{\mathbf{U}}_{21} \\ \tilde{\mathbf{U}}_{31} \end{bmatrix} \right) &= \sqrt{\sigma_i \left(\tilde{\mathbf{U}}_{21}^* \tilde{\mathbf{U}}_{21} + \tilde{\mathbf{U}}_{31}^* \tilde{\mathbf{U}}_{31} \right)} \\ (\text{Lemma 4}) \quad &\leq \sqrt{\|\tilde{\mathbf{U}}_{31}\|_2^2 + \sigma_i(\tilde{\mathbf{U}}_{21})^2} \\ &\leq \|\tilde{\mathbf{U}}_{31}\|_2 \sqrt{1 + \left(\frac{\sigma_k}{\sigma_{k-i+1}} \cdot \frac{\|\mathbf{E}_{32}\|_2}{\gamma_2} \right)^2} \\ &\leq \frac{\|[\mathbf{E}_{31}, \mathbf{E}_{32}]\|_2}{\Gamma_1} \sqrt{1 + \left(\frac{\sigma_k}{\sigma_{k-i+1}} \cdot \frac{\|\mathbf{E}_{32}\|_2}{\gamma_2} \right)^2}. \end{aligned}$$

Analogously, by observing that

$$\|\tilde{\mathbf{V}}_{21}\Sigma_k\| = \left\| \hat{\Sigma}_{l \setminus k}^2 \tilde{\mathbf{V}}_{21}\Sigma_k^{-1} + \mathbf{E}_{32}^* \tilde{\mathbf{U}}_{31} \right\| \leq \frac{\|\hat{\Sigma}_{l \setminus k}\|_2^2}{\sigma_k^2} \|\tilde{\mathbf{V}}_{21}\Sigma_k\| + \|\mathbf{E}_{32}^* \tilde{\mathbf{U}}_{31}\|,$$

we have that, by Lemma 3, for all $i \in [k]$,

$$\begin{aligned}\|\tilde{\mathbf{V}}_{21}\Sigma_k\| &\leq \left(1 - \frac{\hat{\sigma}_{k+1}^2}{\sigma_k^2}\right)^{-1} \|\mathbf{E}_{32}^* \tilde{\mathbf{U}}_{31}\| \leq \sigma_k \cdot \frac{\|\mathbf{E}_{32}\|_2}{\gamma_1} \|\tilde{\mathbf{U}}_{31}\|, \\ \|\tilde{\mathbf{V}}_{21}\| &\leq \frac{1}{\sigma_k} \|\tilde{\mathbf{V}}_{21}\Sigma_k\| \leq \frac{\|\mathbf{E}_{32}\|_2}{\gamma_1} \|\tilde{\mathbf{U}}_{31}\|, \\ \sigma_i(\tilde{\mathbf{V}}_{21}) &\leq \frac{1}{\sigma_{k-i+1}} \|\tilde{\mathbf{V}}_{21}\Sigma_k\|_2 \leq \frac{\sigma_k}{\sigma_{k-i+1}} \cdot \frac{\|\mathbf{E}_{32}\|_2}{\gamma_1} \|\tilde{\mathbf{U}}_{31}\|_2,\end{aligned}$$

and therefore for the spectral or Frobenius norm $\|\cdot\|_\xi$ ($\xi = 2, F$),

$$\begin{aligned}\left\| \begin{bmatrix} \tilde{\mathbf{V}}_{21} \\ \tilde{\mathbf{V}}_{31} \end{bmatrix} \right\|_\xi &\leq \sqrt{\|\tilde{\mathbf{V}}_{21}\|_\xi^2 + \|\tilde{\mathbf{V}}_{31}\|_\xi^2} \leq \|\tilde{\mathbf{U}}_{31}\|_\xi \sqrt{\frac{\|\mathbf{E}_{32}\|_2^2}{\gamma_1^2} + \frac{\|\mathbf{E}_{33}\|_2^2}{\sigma_k^2}} \\ &\leq \frac{\|[\mathbf{E}_{31}, \mathbf{E}_{32}]\|_\xi}{\Gamma_1} \sqrt{\frac{\|\mathbf{E}_{32}\|_2^2}{\gamma_1^2} + \frac{\|\mathbf{E}_{33}\|_2^2}{\sigma_k^2}},\end{aligned}$$

which leads to (5.6). Additionally for individual canonical angles $i \in [k]$,

$$\begin{aligned}\sigma_i\left(\begin{bmatrix} \tilde{\mathbf{V}}_{21} \\ \tilde{\mathbf{V}}_{31} \end{bmatrix}\right) &\leq \sqrt{\sigma_i(\tilde{\mathbf{V}}_{21})^2 + \|\tilde{\mathbf{V}}_{31}\|_2^2} \quad (\text{Lemma 4}) \\ &\leq \|\tilde{\mathbf{U}}_{31}\|_2 \sqrt{\left(\frac{\sigma_k}{\sigma_{k-i+1}} \cdot \frac{\|\mathbf{E}_{32}\|_2}{\gamma_1}\right)^2 + \left(\frac{\|\mathbf{E}_{33}\|_2}{\sigma_k}\right)^2} \\ &\leq \frac{\|[\mathbf{E}_{31}, \mathbf{E}_{32}]\|_2}{\Gamma_1} \sqrt{\left(\frac{\sigma_k}{\sigma_{k-i+1}} \cdot \frac{\|\mathbf{E}_{32}\|_2}{\gamma_1}\right)^2 + \left(\frac{\|\mathbf{E}_{33}\|_2}{\sigma_k}\right)^2}.\end{aligned}$$

This yields (5.10) and completes the proof. \square

6. Numerical experiments. First, we present numerical comparisons among different canonical angle upper bounds and the unbiased estimates on the left and right leading singular subspaces of various synthetic and real data matrices. We start by describing the target matrices in subsection 6.1. In subsection 6.2, we discuss the performance of the unbiased estimates, as well as the relative tightness of the canonical angle bounds, for different algorithmic choices based on the numerical observations. Second, in subsection 6.3, we present an illustrative example that provides insight into the balance between oversampling and power iterations brought by the space-agnostic bounds.

6.1. Target matrices. We consider several different classes of target matrices, including some synthetic random matrices with different spectral patterns, as well as an empirical dataset, as summarized below:

1. A random sparse nonnegative (SNN) matrix [28] \mathbf{A} of size $m \times n$ takes the form

$$(6.1) \quad \mathbf{A} = \text{SNN}(a, r_1) := \sum_{i=1}^{r_1} \frac{a}{i} \mathbf{x}_i \mathbf{y}_i^T + \sum_{i=r_1+1}^{\min(m, n)} \frac{1}{i} \mathbf{x}_i \mathbf{y}_i^T,$$

where $a > 1$ and $r_1 < \min(m, n)$ control the spectral decay, and $\mathbf{x}_i \in \mathbb{C}^m$, $\mathbf{y}_i \in \mathbb{C}^n$ are random sparse vectors with nonnegative entries. In the experiments, we test on two random SNN matrices of size 500×500 with $r_1 = 20$ and $a = 1, 100$, respectively.

2. Gaussian dense matrices with controlled spectral decay are randomly generated via a similar construction to the SNN matrix, with $\mathbf{x}_j \in \mathbb{S}^{m-1}$ and $\mathbf{y}_j \in \mathbb{S}^{n-1}$ in (6.1) replaced by uniformly random dense orthonormal vectors. The generating procedures for $\mathbf{A} \in \mathbb{C}^{m \times n}$ with rank $r \leq \min(m, n)$ can be summarized as follows:
 - (i) Draw Gaussian random matrices, $\mathbf{G}_m \in \mathbb{C}^{m \times r}$ and $\mathbf{G}_n \in \mathbb{C}^{n \times r}$.
 - (ii) Compute $\mathbf{U} = \text{ortho}(\mathbf{G}_m) \in \mathbb{C}^{m \times r}$, $\mathbf{V} = \text{ortho}(\mathbf{G}_n) \in \mathbb{C}^{n \times r}$ as orthonormal bases.
 - (iii) Given the spectrum $\Sigma = \text{diag}(\sigma_1, \dots, \sigma_r)$, we construct $\mathbf{A} = \mathbf{U}\Sigma\mathbf{V}^*$.
- In the experiments, we consider two types of spectral decay:
 - (i) slower decay with $r_1 = 20$, $\sigma_i = 1$ for all $i \in [r_1]$, $\sigma_i = 1/\sqrt{i-r_1+1}$ for all $i = r_1+1, \dots, r$, and
 - (ii) faster decay with $r_1 = 20$, $\sigma_i = 1$ for all $i \in [r_1]$, $\sigma_i = \max(0.99^{i-r_1}, 10^{-3})$ for all $i = r_1+1, \dots, r$.
3. MNIST training set consists of 60,000 images of hand-written digits from 0 to 9. Each image is of size 28×28 . We form the target matrices by uniformly sampling $N = 800$ images from the MNIST training set. The images are flattened and normalized to form a full-rank matrix of size $N \times d$ where $d = 784$ is the size of the flattened images, with entries bounded in $[0, 1]$. The nonzero entries take approximately 20% of the matrix for both the training and the testing sets.

6.2. Canonical angle bounds and estimates. Now we present numerical comparisons of the performance of the canonical angle bounds and the unbiased estimates under different algorithmic choices. Considering the scenario where the true matrix spectra may not be available in practice, we calculate two sets of upper bounds, one from the true spectra $\Sigma \in \mathbb{C}^{r \times r}$ and the other from the l approximated singular values from Algorithm 1. For the latter, we pad the approximated spectrum $\tilde{\Sigma}_l = \text{diag}(\tilde{\sigma}_1, \dots, \tilde{\sigma}_l)$ with $r-l$ copies of $\tilde{\sigma}_l$ and evaluate the canonical angle bounds and estimates with $\tilde{\Sigma} = \text{diag}(\tilde{\sigma}_1, \dots, \tilde{\sigma}_l, \dots, \tilde{\sigma}_l) \in \mathbb{C}^{r \times r}$.

From Figure 1 to Figure 11,

1. red lines and dashes (**Thm1 w/ σ and $\hat{\sigma}$**) represent the space-agnostic probabilistic bounds in Theorem 1 evaluated with the true (lines) and approximated (dashes) singular values, Σ , and $\tilde{\Sigma}$, respectively, where we simply ignore tail decay and suppress constants for the distortion factors and set $\epsilon_1 = \sqrt{\frac{k}{l}}$ and $\epsilon_2 = \sqrt{\frac{l}{r-k}}$ in (3.1) and (3.2);
2. blue lines and dashes (**Prop1 w/ σ and $\hat{\sigma}$**) represent the unbiased space-agnostic estimates in Proposition 1 (averages of $N = 3$ independent trials with

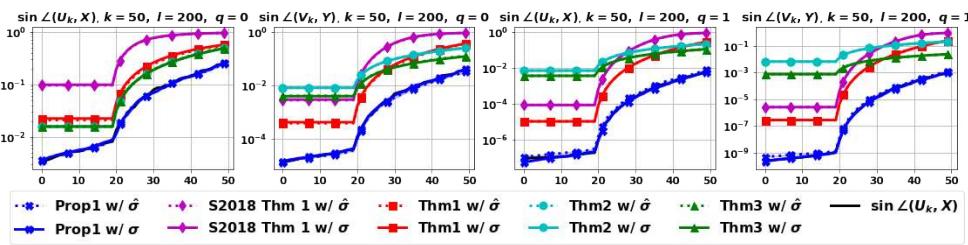


FIG. 1. Synthetic Gaussian with the slower spectral decay. $k = 50, l = 200, q = 0, 1$.

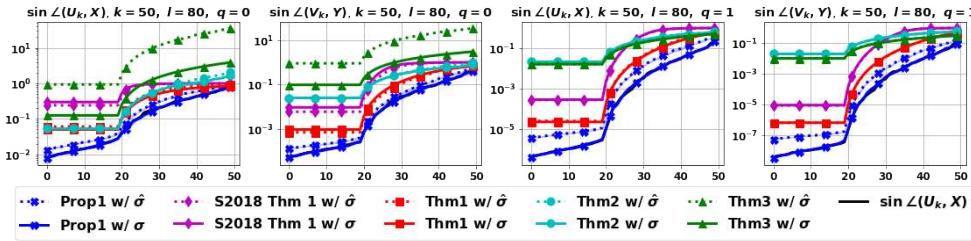


FIG. 2. Synthetic Gaussian with the slower spectral decay. $k = 50, l = 80, q = 0, 1$.

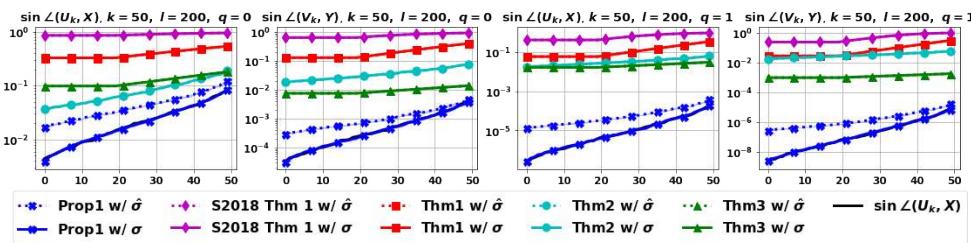


FIG. 3. Synthetic Gaussian with the faster spectral decay. $k = 50, l = 200, q = 0, 1$.

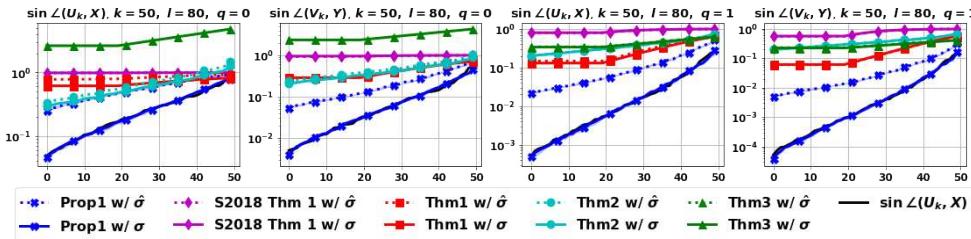


FIG. 4. Synthetic Gaussian with the faster spectral decay. $k = 50, l = 80, q = 0, 1$.

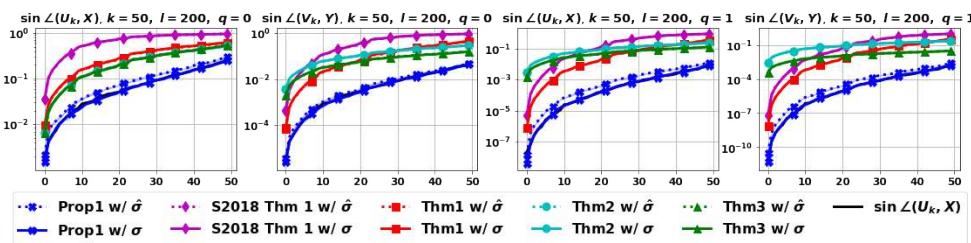
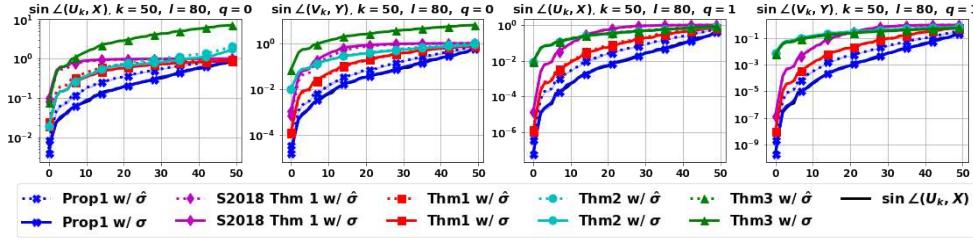
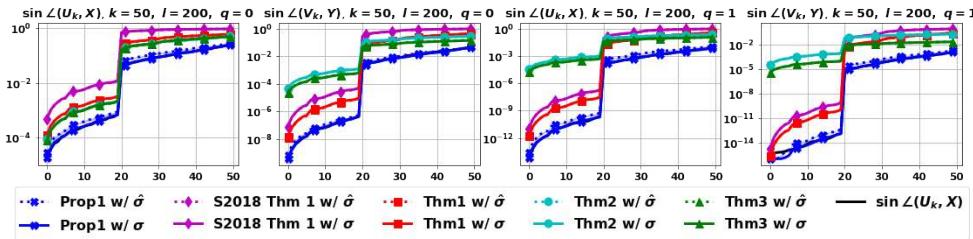
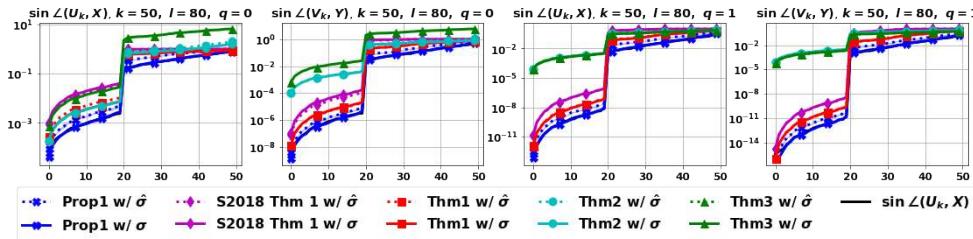
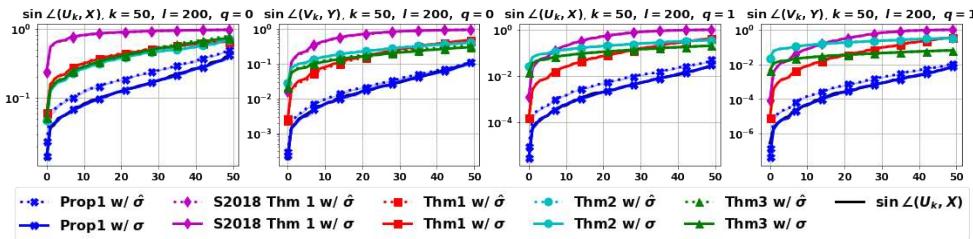


FIG. 5. SNN with $r_1 = 20, a = 1$. $k = 50, l = 200, q = 0, 1$.

blue shades marking the corresponding minima and maxima in the trials) evaluated with the true (lines) and approximated (dashes) singular values, Σ and $\tilde{\Sigma}$, respectively;

3. cyan lines and dashes (**Thm2 w/ σ and $\hat{\sigma}$**) represent the posterior residual-based bounds in Theorem 2 evaluated with the true (lines) and approximated (dashes) singular values, Σ , and $\tilde{\Sigma}$, respectively;

FIG. 6. SNN with $r_1 = 20$, $a = 1$. $k = 50$, $l = 80$, $q = 0, 1$.FIG. 7. SNN with $r_1 = 20$, $a = 100$. $k = 50$, $l = 200$, $q = 0, 1$.FIG. 8. SNN with $r_1 = 20$, $a = 100$. $k = 50$, $l = 80$, $q = 0, 1$.FIG. 9. 800 randomly sampled images from the MNIST training set. $k = 50$, $l = 200$, $q = 0, 1$.

4. green lines and dashes (**Thm3 w/ σ and $\hat{\sigma}$**) represent the posterior residual-based bounds (5.3) and (5.4) in Theorem 4 evaluated with the true (lines) and approximated (dashes) singular values, Σ , and $\tilde{\Sigma}$, respectively;
5. magenta lines and dashes (**S2018 Thm1 w/ σ and $\hat{\sigma}$**) represent the upper bounds in [27, Theorem 1] (i.e., (3.6) and (3.7)) evaluated with the true (lines) and approximated (dashes) singular values, Σ and $\tilde{\Sigma}$, respectively, and the unknown true singular subspace such that $\Omega_1 = \mathbf{V}_k^* \Omega$ and $\Omega_2 = \mathbf{V}_{r \setminus k}^* \Omega$;
6. black lines mark the true canonical angles $\sin \angle(\mathbf{U}_k, \mathbf{\widehat{U}}_l)$.

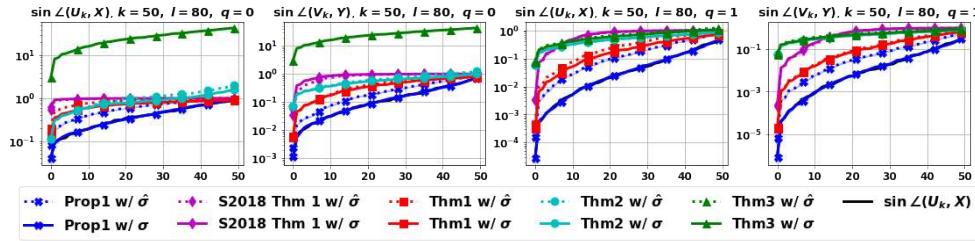


FIG. 10. 800 randomly sampled images from the MNIST training set. $k = 50$, $l = 80$, $q = 0, 1$.

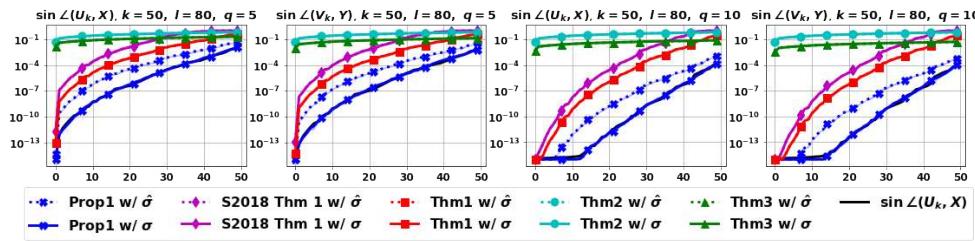


FIG. 11. 800 randomly sampled images from the MNIST training set. $k = 50$, $l = 80$, $q = 5, 10$.

We recall from Remark 3 that, by the algorithmic construction of Algorithm 1, for given q , canonical angles of the right singular spaces $\sin \angle(\mathbf{V}_k, \widehat{\mathbf{V}}_l)$ are evaluated with half more power iterations than those of the left singular spaces $\sin \angle(\mathbf{U}_k, \widehat{\mathbf{U}}_l)$. That is, $\sin \angle(\mathbf{U}_k, \widehat{\mathbf{U}}_l)$, $\sin \angle(\mathbf{V}_k, \widehat{\mathbf{V}}_l)$ with $q = 0, 1$ in Figures 1–10 can be viewed as canonical angles of randomized subspace approximation with $q = 0, 0.5, 1, 1.5$ power iterations, respectively; while Figure 11 corresponds to randomized subspace approximations constructed with $q = 5, 5.5, 10, 10.5$ power iterations analogously.

For each set of upper bounds/unbiased estimates, we observe the following.

1. The space-agnostic probabilistic bounds (**Thm1 w/ σ and $\widehat{\sigma}$**) in Theorem 1 provide tighter statistical guarantees for the canonical angles of all the tested target matrices in comparison to those from [27, Theorem 1] (**S2018 Thm1 w/ σ and $\widehat{\sigma}$**), as explained in Remark 2.
2. The unbiased estimators (**Prop1 w/ σ and $\widehat{\sigma}$**) in Proposition 1 yield accurate approximations for the true canonical angles on all the tested target matrices with as few as $N = 3$ trials, while enjoying good empirical concentrate. As a potential drawback, the accuracy of the unbiased estimates may be compromised when approaching the machine epsilon (as observed in Figure 7, $\sin \angle(\mathbf{V}_k, \mathbf{Y})$, $q = 1$).
3. The posterior residual-based bounds (**Thm2 w/ σ and $\widehat{\sigma}$**) in Theorem 2 are relatively tighter among the compared bounds in the setting with larger oversampling ($l = 4k$), and no power iterations ($\sin \angle(\mathbf{U}_k, \widehat{\mathbf{U}}_l)$ with $q = 0$) or exponential spectral decay (Figure 3).
4. The posterior residual-based bounds (5.3) and (5.4) (**Thm3 w/ σ and $\widehat{\sigma}$**) in Theorem 4 have relative tightness similar to that of the posterior residual-based bounds in Theorem 2 but are slightly more sensitive to power iterations. As shown in Figure 3, on a target matrix with exponential spectral decay and large oversampling ($l = 4k$), Theorem 4 gives tighter posterior guarantees when $q > 0$. However, with the additional assumptions $\sigma_k > \widehat{\sigma}_{k+1}$ and $\sigma_k >$

$\|\mathbf{E}_{33}\|_2$, Theorem 4 usually requires large oversampling ($l = 4k$) in order to provide nontrivial (i.e., within the range $[0, 1]$) bounds.

For target matrices with various patterns of spectral decay, with different combinations of oversampling ($l = 1.6k, 4k$) and power iterations ($q = 0, 1$), we make the following observations on the relative tightness of upper bounds in Theorems 1, 2, and 4.

1. For target matrices with subexponential spectral decay, the space-agnostic bounds in Theorem 1 are relatively tighter in most tested settings, except for the setting in Figure 3 with larger oversampling ($l = 200$) and no power iterations ($q = 0$).
2. For target matrices with exponential spectral decay (Figures 3 and 4), the posterior residual-based bounds in Theorems 2 and 4 tend to be relatively tighter, especially with large oversampling (Figure 3 with $l = 4k$). Meanwhile, with power iterations $q > 0$, Theorem 4 tends to be tighter than Theorem 2.

Furthermore, considering the scenario with an unknown true spectrum Σ , we plot estimations for the upper bounds in Theorems 1, 2, and 4, and the unbiased estimates in Proposition 1, evaluated with a padded approximation of the spectrum $\tilde{\Sigma} = \text{diag}(\hat{\sigma}_1, \dots, \hat{\sigma}_l, \dots, \hat{\sigma}_l)$, which leads to mild overestimations, as marked in dashes from Figure 1 to Figure 11.

6.3. Balance between oversampling and power iterations. To illustrate the insight cast by Theorem 1 on the balance between oversampling and power iterations, we consider the following synthetic example.

Example 1. Given a target rank $k \in \mathbb{N}$, we consider a simple synthetic matrix $\mathbf{A} \in \mathbb{C}^{r \times r}$ of size $r = (1 + \beta)k$, consisting of random singular subspaces (generated by orthonormalizing Gaussian matrices) and a step spectrum:

$$\sigma(\mathbf{A}) = \text{diag}(\underbrace{\sigma_1, \dots, \sigma_1}_{\sigma_i = \sigma_1 \forall i \leq k}, \underbrace{\sigma_{k+1}, \dots, \sigma_{k+1}}_{\sigma_i = \sigma_{k+1} \forall i \geq k+1}).$$

We fix a budget of $N = \alpha k$ matrix-vector multiplications with \mathbf{A} in total. The goal is to distribute the computational budget between the sample size l and the number of power iterations q for the smaller canonical angles $\angle(\mathbf{U}_k, \widehat{\mathbf{U}}_l)$.

Leveraging Theorem 1, we start by fixing $\gamma > 1$ associated with the constants $\epsilon_1 = \gamma\sqrt{k/l}$ and $\epsilon_2 = \gamma\sqrt{l/(r-k)}$ in (3.1) such that $l \geq \gamma^2 k$ and $2q + 1 < \alpha/\gamma^2$. Characterized by γ , the right-hand side of (3.1) under fixed budget N (i.e., $N \geq l(2q + 1)$) is defined as

$$\begin{aligned} (6.2) \quad \phi_\gamma(q) &\triangleq \left(1 + \frac{1 - \epsilon_1}{1 + \epsilon_2} \cdot \frac{l}{r - k} \left(\frac{\sigma_1}{\sigma_{k+1}} \right)^{4q+2} \right)^{-\frac{1}{2}} \\ &= \left(1 + \frac{\alpha - \gamma\sqrt{\alpha(2q+1)}}{\beta(2q+1) + \gamma\sqrt{\alpha\beta(2q+1)}} \left(\frac{\sigma_1}{\sigma_{k+1}} \right)^{4q+2} \right)^{-\frac{1}{2}}. \end{aligned}$$

With the synthetic step spectrum, the dependence of (3.1) on $\sigma(\mathbf{A})$ is reduced to the spectral gap σ_1/σ_{k+1} in (6.2).

As a synopsis, Table 1 summarizes the relevant parameters that characterize the problem setup.

With $k = 10$, $\alpha = 32$, $\beta = 64$, and $\gamma \in \{1.05, 2.00\}$, Figures 12 and 13 illustrate

- (i) how the balance between oversampling and power iterations is affected by the spectral gap σ_1/σ_{k+1} , and more importantly,

TABLE 1

Given $\mathbf{A} \in \mathbb{C}^{r \times r}$ with a spectral gap σ_1/σ_{k+1} , a target rank k , and a budget of N matrix-vector multiplications, we consider applying Algorithm 1 with a sample size l and q power iterations.

α	budget parameter		$N = \alpha k$
β	size parameter		$r = (1 + \beta)k$
γ	oversampling parameter	$l \geq \gamma^2 k$ and $2q + 1 \leq \frac{\alpha}{\gamma^2}$	

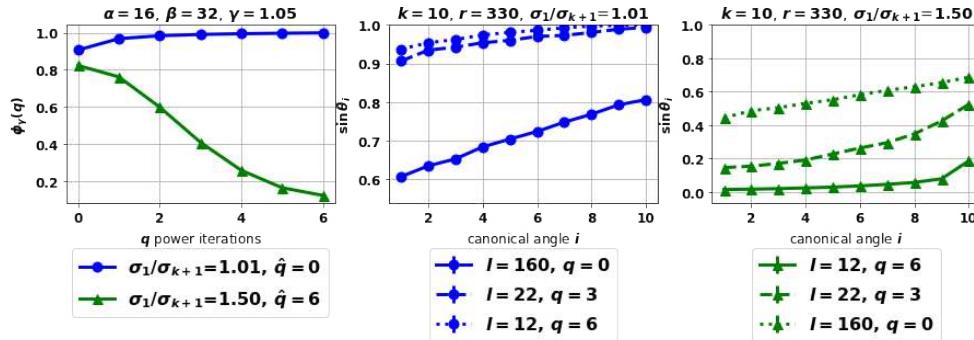


FIG. 12. For $k = 10$, $\alpha = 16$, $\beta = 32$, $\gamma = 1.05$, the left figure marks $\phi_\gamma(q)$ (i.e., the right-hand side of (3.1) under the fixed budget N) with two different spectral gaps ($\hat{q} = \arg\min_{1 \leq 2q+1 \leq \alpha/\gamma^2} \phi_\gamma(q)$), while the middle and the right figures demonstrate how the relative magnitudes of canonical angles $\sin \angle_i(\mathbf{U}_k, \hat{\mathbf{U}}_l)$ ($i \in [k]$) under different configurations (i.e., choices of (l, q) , showing the averages and ranges of 5 trials) align with the trends in $\phi_\gamma(q)$.

(ii) how (6.2) unveils the trend in true canonical angles $\sin \angle_i(\mathbf{U}_k, \hat{\mathbf{U}}_l)$ among different configurations $\{(l, q) \mid l \geq \gamma^2 k, 2q + 1 \leq \alpha/\gamma^2\}$.

Concretely, both Figure 12 and Figure 13 imply that more oversampling (e.g., $q = 0$) is preferred when the spectral gap is small (e.g., $\sigma_1/\sigma_{k+1} = 1.01$), while more power iterations (e.g., $q = \lfloor \frac{\alpha/\gamma^2 - 1}{2} \rfloor$) are preferred when the spectral gap is large (e.g., $\sigma_1/\sigma_{k+1} = 1.5$). Such trends are both observed in the true canonical angles $\sin \angle_i(\mathbf{U}_k, \hat{\mathbf{U}}_l)$ ($i \in [k]$) and well reflected by $\phi_\gamma(q)$.

7. Discussion, limitations, and future directions. We presented prior and posterior bounds and estimates that can be computed efficiently for canonical angles of the randomized subspace approximation. Under moderate multiplicative oversampling, our prior probabilistic bounds are space-agnostic (i.e., independent of the unknown true subspaces) and asymptotically tight and can be computed in linear ($O(\text{rank}(\mathbf{A}))$) time, while casting a clear guidance on the balance between oversampling and power iterations for a fixed budget of matrix-vector multiplications. As corollaries of the prior probabilistic bounds, we introduce a set of unbiased canonical angle estimates that are efficiently computable and applicable to arbitrary choices of oversampling with good empirical concentrations. In addition to the prior bounds and estimates, we further discuss two sets of posterior bounds that provide deterministic guarantees for canonical angles given the computed low-rank approximations. With numerical experiments, we compare the empirical tightness of different canonical angle bounds and estimates on various data matrices under a diverse set of algorithmic choices for the randomized subspace approximation.

As a major limitation of our space-agnostic bounds and estimates, we note that Theorem 1 and Proposition 1 rely crucially on the isotropic property of Gaussian

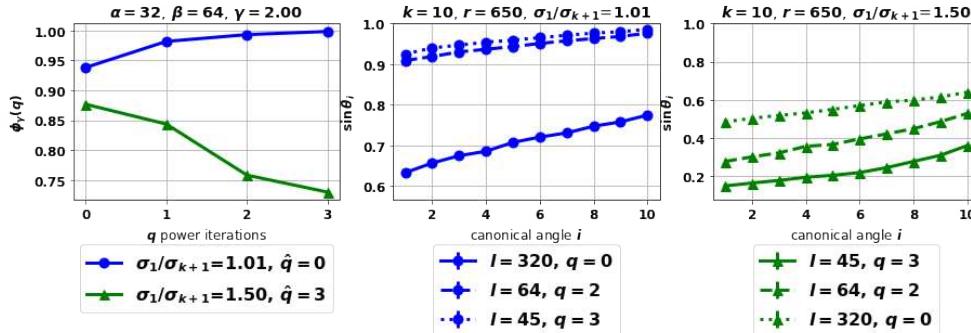


FIG. 13. Under the same setup as Figure 12, for $k = 10$, $\alpha = 32$, $\beta = 64$, $\gamma = 2.00$, the trend in $\phi_\gamma(q)$ also aligns well with that in true canonical angles $\sin \angle_i(\mathbf{U}_k, \tilde{\mathbf{U}}_l)$ ($i \in [k]$).

random matrices and therefore are only applicable to Gaussian embeddings. Although Gaussian embedding is one of the most commonly used Johnson–Lindenstrauss transforms (JLT) both theoretically and in practice, in light of the appealing empirical performance [14, 20, 6] of various alternatives (e.g., the subsampled randomized trigonometric transforms [38, 26, 31, 2] and sparse sign matrices [22, 25, 37, 4, 33]), it would be interesting to explore relaxations of the isotropic assumption and extend the analysis to these fast JLTs.

Appendix A. Technical lemmas.

LEMMA 1. Let $\mathbf{x} \in \mathbb{R}^d$ be a random vector with $\mathbb{E}[\mathbf{x}] = \mathbf{0}$, $\mathbb{E}[\mathbf{x}\mathbf{x}^\top] = \Sigma$, and $\bar{\mathbf{x}} = \Sigma^{-1/2}\mathbf{x}^7$ being ρ^2 -sub-Gaussian⁸. Consider a set of n i.i.d. samples of \mathbf{x} , $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_n]^* \in \mathbb{R}^{n \times d}$, and a diagonal weight matrix $\mathbf{W} = \text{diag}(w_1, \dots, w_n)$ with $w_i > 0$ corresponding to each sample $i \in [n]$. If $n \geq \frac{\text{tr}(\mathbf{W})^2}{\text{tr}(\mathbf{W}^2)} \geq \frac{20736\rho^4 d}{\epsilon^2} + \frac{10368\rho^4 \log(1/\delta)}{\epsilon^2}$, then with probability at least $1 - \delta$,

$$(1 - \epsilon) \text{tr}(\mathbf{W}) \Sigma \preceq \mathbf{X}^\top \mathbf{W} \mathbf{X} \preceq (1 + \epsilon) \text{tr}(\mathbf{W}) \Sigma.$$

Concretely, with $\mathbf{W} = \mathbf{I}_n$, $n = \frac{\text{tr}(\mathbf{W})^2}{\text{tr}(\mathbf{W}^2)} = \Omega(\rho^4 d)$, and $\epsilon = \Theta(\rho^2 \sqrt{\frac{d}{n}})$, $(1 - \epsilon) \Sigma \preceq \frac{1}{n} \mathbf{X}^\top \mathbf{X} \preceq (1 + \epsilon) \Sigma$ with high probability (at least $1 - e^{-\Theta(d)}$).

Proof. We first denote $\mathbf{P}_\mathcal{X} \triangleq \Sigma \Sigma^\dagger$ as the orthogonal projector onto the subspace $\mathcal{X} \subseteq \mathbb{R}^d$ supported by the distribution of \mathbf{x} . With the assumptions $\mathbb{E}[\mathbf{x}] = \mathbf{0}$ and $\mathbb{E}[\mathbf{x}\mathbf{x}^\top] = \Sigma$, we observe that $\mathbb{E}[\bar{\mathbf{x}}] = \mathbf{0}$ and $\mathbb{E}[\bar{\mathbf{x}}\bar{\mathbf{x}}^\top] = \mathbb{E}[(\Sigma^{-1/2}\mathbf{x})(\Sigma^{-1/2}\mathbf{x})^\top] = \mathbf{P}_\mathcal{X}$. Given the sample set \mathbf{X} of size $n \gg \rho^4(d + \log(1/\delta))$ for any $\delta \in (0, 1)$, we let

$$\mathbf{U} = \frac{1}{\text{tr}(\mathbf{W})} \sum_{i=1}^n w_i \left(\Sigma^{-1/2} \mathbf{x}_i \right) \left(\Sigma^{-1/2} \mathbf{x}_i \right)^\top - \mathbf{P}_\mathcal{X}.$$

Then the problem can be reduced to showing that, for any $\epsilon > 0$, with probability at least $1 - \delta$, $\|\mathbf{U}\|_2 \leq \epsilon$. For this, we leverage the ϵ -net argument as follows.

⁷In the case where Σ is rank-deficient, we slightly abuse the notation such that $\Sigma^{-1/2}$ and Σ^{-1} refer to the respective pseudo-inverses.

⁸A random vector $\mathbf{v} \in \mathbb{R}^d$ is ρ^2 -sub-Gaussian if for any unit vector $\mathbf{u} \in \mathbb{S}^{d-1}$, $\mathbf{u}^\top \mathbf{v}$ is ρ^2 -sub-Gaussian; $\mathbb{E}[\exp(s \cdot \mathbf{u}^\top \mathbf{v})] \leq \exp(s^2 \rho^2/2)$ for all $s \in \mathbb{R}$.

For an arbitrary $\mathbf{v} \in \mathcal{X} \cap \mathbb{S}^{d-1}$, we have

$$\begin{aligned}\mathbf{v}^\top \mathbf{U} \mathbf{v} &= \frac{1}{\text{tr}(\mathbf{W})} \sum_{i=1}^n w_i \left(\mathbf{v}^\top (\Sigma^{-1/2} \mathbf{x}) (\Sigma^{-1/2} \mathbf{x})^\top \mathbf{v} - 1 \right) \\ &= \frac{1}{\text{tr}(\mathbf{W})} \sum_{i=1}^n w_i \left((\mathbf{v}^\top \bar{\mathbf{x}}_i)^2 - 1 \right),\end{aligned}$$

where, given $\bar{\mathbf{x}}_i$ being ρ^2 -sub-Gaussian, $\mathbf{v}^\top \bar{\mathbf{x}}_i$ is ρ^2 -sub-Gaussian. Since

$$\mathbb{E} \left[(\mathbf{v}^\top \bar{\mathbf{x}}_i)^2 \right] = \mathbf{v}^\top \mathbb{E} [\bar{\mathbf{x}}_i \bar{\mathbf{x}}_i^\top] \mathbf{v} = 1,$$

we know that $(\mathbf{v}^\top \bar{\mathbf{x}}_i)^2 - 1$ is $16\rho^2$ -subexponential.⁹ With $\beta_i \triangleq \frac{w_i}{\text{tr}(\mathbf{W})}$ for all $i \in [n]$ such that $\boldsymbol{\beta} = [\beta_1, \dots, \beta_n]^*$, we recall Bernstein's inequality [34, Theorem 2.8.2], [35, section 2.1.3],

$$\begin{aligned}\mathbb{P} \left[\left| \mathbf{v}^\top \mathbf{U} \mathbf{v} \right| = \left| \sum_{i=1}^n \beta_i \cdot ((\mathbf{v}^\top \bar{\mathbf{x}}_i)^2 - 1) \right| > t \right] \\ \leq 2 \exp \left(-\frac{1}{2} \min \left(\frac{t^2}{(16\rho^2)^2 \|\boldsymbol{\beta}\|_2^2}, \frac{t}{16\rho^2 \|\boldsymbol{\beta}\|_\infty} \right) \right),\end{aligned}$$

where $\|\boldsymbol{\beta}\|_2^2 = \frac{\text{tr}(\mathbf{W}^2)}{\text{tr}(\mathbf{W})^2}$ and $\|\boldsymbol{\beta}\|_\infty = \frac{\max_{i \in [n]} w_i}{\text{tr}(\mathbf{W})}$.

Let $N \subset \mathcal{X} \cap \mathbb{S}^{d-1}$ be an ϵ_1 -net such that $|N| = (1 + \frac{2}{\epsilon_1})^d$. Then for some $0 < \epsilon_2 \leq 16\rho^2 \frac{\|\boldsymbol{\beta}\|_2^2}{\|\boldsymbol{\beta}\|_\infty}$, by the union bound,

$$\begin{aligned}\mathbb{P} \left[\max_{\mathbf{v} \in N} : \left| \mathbf{v}^\top \mathbf{U} \mathbf{v} \right| > \epsilon_2 \right] &\leq 2 |N| \exp \left(-\frac{1}{2} \min \left(\frac{\epsilon_2^2}{(16\rho^2)^2 \|\boldsymbol{\beta}\|_2^2}, \frac{\epsilon_2}{16\rho^2 \|\boldsymbol{\beta}\|_\infty} \right) \right) \\ &\leq \exp \left(d \log \left(1 + \frac{2}{\epsilon_1} \right) - \frac{1}{2} \cdot \frac{\epsilon_2^2}{(16\rho^2)^2 \|\boldsymbol{\beta}\|_2^2} \right) \leq \delta\end{aligned}$$

whenever $\frac{1}{\|\boldsymbol{\beta}\|_2^2} \geq \frac{\text{tr}(\mathbf{W})^2}{\text{tr}(\mathbf{W}^2)} = 2(\frac{16\rho^2}{\epsilon_2})^2 (d \log(1 + \frac{2}{\epsilon_1}) + \log \frac{1}{\delta})$, where $1 < \frac{\text{tr}(\mathbf{W})^2}{\text{tr}(\mathbf{W}^2)} \leq n$.

Now for any $\mathbf{v} \in \mathcal{X} \cap \mathbb{S}^{d-1}$, there exists some $\mathbf{v}' \in N$ such that $\|\mathbf{v} - \mathbf{v}'\|_2 \leq \epsilon_1$. Therefore,

$$\begin{aligned}|\mathbf{v}^\top \mathbf{U} \mathbf{v}| &= \left| \mathbf{v}'^\top \mathbf{U} \mathbf{v}' + 2\mathbf{v}'^\top \mathbf{U} (\mathbf{v} - \mathbf{v}') + (\mathbf{v} - \mathbf{v}')^\top \mathbf{U} (\mathbf{v} - \mathbf{v}') \right| \\ &\leq \left(\max_{\mathbf{v} \in N} : |\mathbf{v}^\top \mathbf{U} \mathbf{v}| \right) + 2\|\mathbf{U}\|_2 \|\mathbf{v}'\|_2 \|\mathbf{v} - \mathbf{v}'\|_2 + \|\mathbf{U}\|_2 \|\mathbf{v} - \mathbf{v}'\|_2^2 \\ &\leq \left(\max_{\mathbf{v} \in N} : |\mathbf{v}^\top \mathbf{U} \mathbf{v}| \right) + \|\mathbf{U}\|_2 (2\epsilon_1 + \epsilon_1^2).\end{aligned}$$

Taking the supremum over $\mathbf{v} \in \mathbb{S}^{d-1}$, with probability at least $1 - \delta$,

$$\max_{\mathbf{v} \in \mathcal{X} \cap \mathbb{S}^{d-1}} : |\mathbf{v}^\top \mathbf{U} \mathbf{v}| = \|\mathbf{U}\|_2 \leq \epsilon_2 + \|\mathbf{U}\|_2 (2\epsilon_1 + \epsilon_1^2), \quad \|\mathbf{U}\|_2 \leq \frac{\epsilon_2}{2 - (1 + \epsilon_1)^2}.$$

With $\epsilon_1 = \frac{1}{3}$, we have $\epsilon = \frac{9}{2}\epsilon_2$.

⁹We abbreviate (ν, ν) -subexponential (i.e., recall that a random variable X is (ν, α) -subexponential if $\mathbb{E}[\exp(sX)] \leq \exp(s^2 \nu^2/2)$ for all $|s| \leq 1/\alpha$) simply as ν -subexponential.

Overall, if $n \geq \frac{\text{tr}(\mathbf{W})^2}{\text{tr}(\mathbf{W}^2)} \geq \frac{1024\rho^4 d}{\epsilon_2^2} + \frac{512\rho^4}{\epsilon_2^2} \log \frac{1}{\delta}$, then with probability at least $1 - \delta$, we have $\|\mathbf{U}\|_2 \leq \epsilon$.

As a concrete instance, when $\mathbf{W} = \mathbf{I}_n$ and $n = \frac{\text{tr}(\mathbf{W})^2}{\text{tr}(\mathbf{W}^2)} \geq 9^2 \cdot 1025 \cdot \rho^4 d$, by taking $\epsilon_2 = \sqrt{\frac{1025\rho^4 d}{n}}$, we have $\|\mathbf{U}\|_2 \leq \frac{1}{2} \sqrt{\frac{9^2 \cdot 1025 \cdot \rho^4 d}{n}}$ with high probability (at least $1 - \delta$ where $\delta = \exp(-\frac{d}{512})$). \square

LEMMA 2 (Cauchy interlacing theorem). *Given an arbitrary matrix $\mathbf{A} \in \mathbb{C}^{m \times n}$ and an orthogonal projection $\mathbf{Q} \in \mathbb{C}^{n \times k}$ with orthonormal columns, for all $i = 1, \dots, k$,*

$$\sigma_i(\mathbf{AQ}) \leq \sigma_i(\mathbf{A}).$$

Proof of Lemma 2. Let $\{\mathbf{v}_j \in \mathbb{C}^k \mid j = 1, \dots, k\}$ be right singular vectors of \mathbf{AQ} . By the min-max theorem (cf. [12, Theorem 8.6.1]),

$$\begin{aligned} \sigma_i(\mathbf{AQ})^2 &= \min_{\mathbf{x} \in \text{span}\{\mathbf{v}_1, \dots, \mathbf{v}_i\} \setminus \{\mathbf{0}\}} \frac{\mathbf{x}^\top \mathbf{Q}^\top \mathbf{A}^\top \mathbf{AQ} \mathbf{x}}{\mathbf{x}^\top \mathbf{x}} \\ &\leq \max_{\dim(\mathcal{V})=i} \min_{\mathbf{x} \in \mathcal{V} \setminus \{\mathbf{0}\}} \frac{\mathbf{x}^\top \mathbf{A}^\top \mathbf{A} \mathbf{x}}{\mathbf{x}^\top \mathbf{x}} = \sigma_i(\mathbf{A})^2. \end{aligned} \quad \square$$

LEMMA 3 ([15, (7.3.14)]). *For arbitrary matrices $\mathbf{A}, \mathbf{B} \in \mathbb{C}^{m \times n}$,*

$$(A.1) \quad \sigma_i(\mathbf{AB}^*) \leq \sigma_i(\mathbf{A}) \sigma_j(\mathbf{B})$$

for all $i \in [\text{rank}(\mathbf{A})]$, $j \in [\text{rank}(\mathbf{B})]$ such that $i + j - 1 \in [\text{rank}(\mathbf{AB}^)]$.*

LEMMA 4 ([15, (7.3.13)]). *For arbitrary matrices $\mathbf{A}, \mathbf{B} \in \mathbb{C}^{m \times n}$,*

$$(A.2) \quad \sigma_{i+j-1}(\mathbf{A} + \mathbf{B}) \leq \sigma_i(\mathbf{A}) + \sigma_j(\mathbf{B})$$

for all $i \in [\text{rank}(\mathbf{A})]$, $j \in [\text{rank}(\mathbf{B})]$ such that $i + j - 1 \in [\text{rank}(\mathbf{A} + \mathbf{B})]$.

Appendix B. Supplementary experiments: Lower space-agnostic bounds. In this section, we visualize and compare the upper and lower bounds in Theorem 1 under the sufficient multiplicative oversampling regime (i.e., $l = 4k$). Recall that $k < l < r = \text{rank}(\mathbf{A})$ where k is the target rank, l is the oversampled rank, and r is the full rank of the matrix \mathbf{A}).

With the same set of target matrices described in subsection 6.1, from Figure 14 to Figure 18, we observe the following:

1. Red lines and dashes (**Thm1 w/ σ and $\hat{\sigma}$**) show the upper bounds in (3.1) and (3.2) evaluated with the true (lines) and approximated (dashes) singular

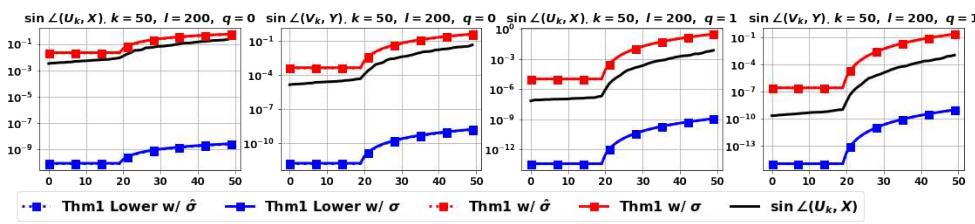


FIG. 14. *Synthetic Gaussian with the slower spectral decay. $k = 50$, $l = 200$, $q = 0, 1$.*

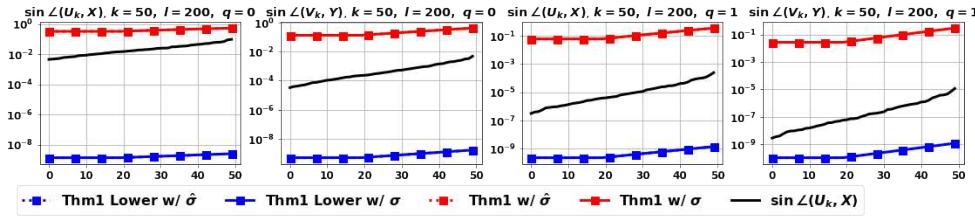


FIG. 15. *Synthetic Gaussian with the faster spectral decay. $k = 50, l = 200, q = 0, 1$.*

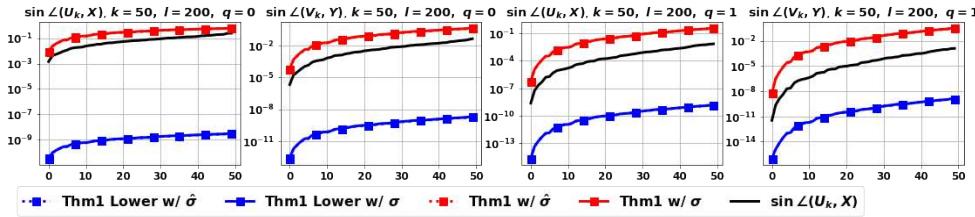


FIG. 16. *SNN with $r_1 = 20, a = 1. k = 50, l = 200, q = 0, 1$.*

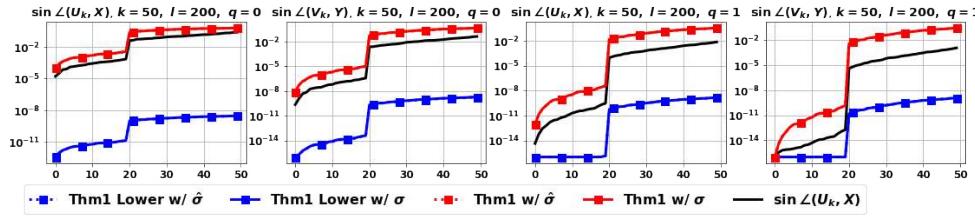


FIG. 17. *SNN with $r_1 = 20, a = 100. k = 50, l = 200, q = 0, 1$.*

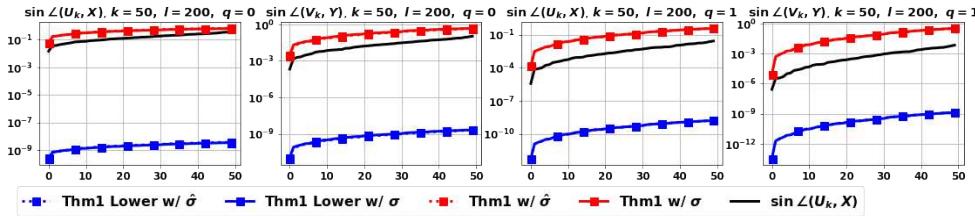


FIG. 18. *800 randomly sampled images from the MNIST training set. $k = 50, l = 200, q = 0, 1$.*

values, Σ , and $\tilde{\Sigma}$, respectively, where we simply ignore tail decay and suppress constants for the distortion factors and set

$$\epsilon_1 = \sqrt{\frac{k}{l}} \quad \text{and} \quad \epsilon_2 = \sqrt{\frac{l}{r-k}}.$$

2. Blue lines and dashes (**Thm1 Lower w/ σ and $\hat{\sigma}$**) present the lower bounds in (3.3) and (3.4) evaluated with Σ and $\tilde{\Sigma}$, respectively, and slightly larger constants associated with the distortion factors

$$\epsilon'_1 = 2\sqrt{\frac{k}{l}} \quad \text{and} \quad \epsilon'_2 = 2\sqrt{\frac{l}{r-k}}.$$

The numerical observations imply that the empirical validity of lower bounds requires more aggressive oversampling than that of upper bounds. In particular, we recall from subsection 6.2 that $l \geq 1.6k$ is usually sufficient for the upper bounds to hold numerically. In contrast, the lower bounds generally require at least $l \geq 4k$, with slightly larger constants associated with the distortion factors $\epsilon_1 = \Theta(\sqrt{k/l})$ and $\epsilon_2 = \Theta(\sqrt{l/(r-k)})$.

Acknowledgments. The authors wish to thank the editor and the referees for their valuable insights and constructive suggestions.

REFERENCES

- [1] A. BJÖRCK AND G. H. GOLUB, *Numerical methods for computing angles between linear subspaces*, Math. Comp., 27 (1973), pp. 579–594.
- [2] C. BOUTSIDIS AND A. GITTENS, *Improved matrix algorithms via the subsampled randomized Hadamard transform*, SIAM J. Matrix Anal. Appl., 34 (2013), pp. 1301–1340, <https://doi.org/10.1137/120874540>.
- [3] C. BOUTSIDIS, P. KAMBADUR, AND A. GITTENS, *Spectral clustering via the power method—provably*, in International Conference on Machine Learning, PMLR, 2015, pp. 40–48.
- [4] K. L. CLARKSON AND D. P. WOODRUFF, *Low-rank approximation and regression in input sparsity time*, J. ACM, 63 (2017), pp. 1–45.
- [5] C. DAVIS AND W. M. KAHAN, *The rotation of eigenvectors by a perturbation. III*, SIAM J. Numer. Anal., 7 (1970), pp. 1–46, <https://doi.org/10.1137/0707001>.
- [6] Y. DONG AND P.-G. MARTINSSON, *Simpler is better: A comparative study of randomized pivoting algorithms for CUR and interpolative decompositions*, Adv. Comput. Math., 49 (2023), 66.
- [7] P. DRINEAS, M. MAGDON-ISMAIL, M. W. MAHONEY, AND D. P. WOODRUFF, *Fast approximation of matrix coherence and statistical leverage*, J. Mach. Learn. Res., 13 (2012), pp. 3475–3506.
- [8] P. DRINEAS, M. W. MAHONEY, AND S. MUTHUKRISHNAN, *Relative-error CUR matrix decompositions*, SIAM J. Matrix Anal. Appl., 30 (2008), pp. 844–881, <https://doi.org/10.1137/07070471X>.
- [9] C. ECKART AND G. YOUNG, *The approximation of one matrix by another of lower rank*, Psychometrika, 1 (1936), pp. 211–218.
- [10] Z. FRANGELLA, J. A. TROPP, AND M. UDELL, *Randomized Nyström preconditioning*, SIAM J. Matrix Anal. Appl., 44 (2023), pp. 718–752, <https://doi.org/10.1137/21M1466244>.
- [11] A. GITTENS AND M. MAHONEY, *Revisiting the Nyström method for improved large-scale machine learning*, in International Conference on Machine Learning, PMLR, 2013, pp. 567–575.
- [12] G. H. GOLUB AND C. F. VAN LOAN, *Matrix Computations*, 3rd ed., The Johns Hopkins University Press, 1996.
- [13] M. GU, *Subspace iteration randomization and singular value problems*, SIAM J. Sci. Comput., 37 (2015), pp. A1139–A1173, <https://doi.org/10.1137/130938700>.
- [14] N. HALKO, P.-G. MARTINSSON, AND J. A. TROPP, *Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions*, SIAM Rev., 53 (2011), pp. 217–288, <https://doi.org/10.1137/090771806>.
- [15] R. A. HORN AND C. R. JOHNSON, *Matrix Analysis*, Cambridge University Press, 2012.
- [16] J. KUCZYŃSKI AND H. WOŹNIAKOWSKI, *Estimating the largest eigenvalue by the power and Lanczos algorithms with a random start*, SIAM J. Matrix Anal. Appl., 13 (1992), pp. 1094–1122, <https://doi.org/10.1137/0613066>.
- [17] M. W. MAHONEY AND P. DRINEAS, *CUR matrix decompositions for improved data analysis*, Proc. Natl. Acad. Sci. USA, 106 (2009), pp. 697–702.
- [18] P.-G. MARTINSSON, *Randomized methods for matrix computations*, Math. Data, 25 (2019), pp. 187–231.
- [19] P.-G. MARTINSSON, V. ROKHLIN, AND M. TYGERT, *A randomized algorithm for the decomposition of matrices*, Appl. Comput. Harmon. Anal., 30 (2011), pp. 47–68.
- [20] P.-G. MARTINSSON AND J. A. TROPP, *Randomized numerical linear algebra: Foundations and algorithms*, Acta Numer., 29 (2020), pp. 403–572.
- [21] P. MASSEY, *Admissible subspaces and the subspace iteration method*, BIT, 64 (2024), pp. 1–36.

- [22] X. MENG AND M. W. MAHONEY, *Low-distortion subspace embeddings in input-sparsity time and applications to robust linear regression*, in Proceedings of the Forty-Fifth Annual ACM Symposium on Theory of Computing, 2013, pp. 91–100.
- [23] R. A. MEYER, C. MUSCO, C. MUSCO, AND D. P. WOODRUFF, *Hutch++: Optimal stochastic trace estimation*, in Symposium on Simplicity in Algorithms (SOSA), SIAM, 2021, pp. 142–155, <https://doi.org/10.1137/1.9781611976496.16>.
- [24] Y. NAKATSUKASA, *Sharp error bounds for Ritz vectors and approximate singular vectors*, Math. Comp., 89 (2020), pp. 1843–1866.
- [25] J. NELSON AND H. L. NGUYÊN, *OSNAP: Faster numerical linear algebra algorithms via sparser subspace embeddings*, in 2013 IEEE 54th Annual Symposium on Foundations of Computer Science, IEEE, 2013, pp. 117–126.
- [26] V. ROKHLIN AND M. TYGERT, *A fast randomized algorithm for overdetermined linear least-squares regression*, Proc. Natl. Acad. Sci. USA, 105 (2008), pp. 13212–13217.
- [27] A. K. SAIBABA, *Randomized subspace iteration: Analysis of canonical angles and unitarily invariant norms*, SIAM J. Matrix Anal. Appl., 40 (2019), pp. 23–48, <https://doi.org/10.1137/18M1179432>.
- [28] D. C. SORENSEN AND M. EMBREE, *A DEIM induced CUR factorization*, SIAM J. Sci. Comput., 38 (2016), pp. A1454–A1482, <https://doi.org/10.1137/140978430>.
- [29] G. W. STEWART, *Matrix Algorithms: Volume II: Eigensystems*, SIAM, 2001, <https://doi.org/10.1137/1.9780898718058>.
- [30] L. N. TREFETHEN AND D. BAU, III, *Numerical Linear Algebra*, SIAM, 2022, <https://doi.org/10.1137/1.9780898719574>.
- [31] J. A. TROPP, *Improved analysis of the subsampled randomized Hadamard transform*, Adv. Adapt. Data Anal., 3 (2011), pp. 115–126.
- [32] J. A. TROPP AND R. J. WEBBER, *Randomized Algorithms for Low-Rank Matrix Approximation: Design, Analysis, and Applications*, preprint, arXiv:2306.12418, 2023.
- [33] J. A. TROPP, A. YURTSEVER, M. UDELL, AND V. CEVHER, *Fixed-rank approximation of a positive-semidefinite matrix from streaming data*, in Advances in Neural Information Processing Systems 30, Curran Associates, 2017, pp. 1225–1234.
- [34] R. VERSHYNIN, *High-Dimensional Probability: An Introduction with Applications in Data Science*, Cambridge Ser. Stat. Probab. Math. 47, Cambridge University Press, 2018.
- [35] M. J. WAINWRIGHT, *High-Dimensional Statistics: A Non-asymptotic Viewpoint*, Cambridge Ser. Stat. Probab. Math. 48, Cambridge University Press, 2019.
- [36] P.-Å. WEDIN, *Perturbation theory for pseudo-inverses*, BIT, 13 (1973), pp. 217–232.
- [37] D. P. WOODRUFF, *Sketching as a tool for numerical linear algebra*, Found. Trends Theor. Comput. Sci., 10 (2014), pp. 1–157.
- [38] F. WOOLFE, E. LIBERTY, V. ROKHLIN, AND M. TYGERT, *A fast randomized algorithm for the approximation of matrices*, Appl. Comput. Harmon. Anal., 25 (2008), pp. 335–366.