Les Cahiers du GERAD ISSN: 0711–2440

A general error analysis for randomized low-rank approximation with application to data assimilation

A. Scotto Di Perrotolo, Y. Diouane, S. Gürol, X. Vasseur

G–2024–31

Mai 2024

GERAD HEC Montréal
3000, chemin de la Côte-Sainte-Catherine
Montréal (Québec) Canada H3T 2A7

Tél. : 514 340-6053
Téléc. : 514 340-5665
info@gerad.ca
www.gerad.ca

A general error analysis for randomized low-rank approximation with application to data assimilation

Alexandre Scotto Di Perrotolo [a]

Youssef Diouane [b]

Selime Gürol [c]

Xavier Vasseur [d]

[a] *IRT Saint Exupéry, 31400 Toulouse, France*

[b] *GERAD & Department of Mathematics and Industrial Engineering, Montréal (Qc), Canada, H3T 1J4*

[c] *CERFACS, F–31057 Toulouse, France*

[d] *F–31700 Blagnac, France*

alexandre.scotto@irt-saintexupery.com
youssef.diouane@polymtl.ca
selime.gurol@cerfacs.fr
xavier.vasseur@gmail.com

Mai 2024
Les Cahiers du GERAD
G–2024–31

Abstract : Randomized algorithms have proven to perform well on a large class of numerical linear algebra problems. Their theoretical analysis is critical to provide guarantees on their behaviour, and in this sense, the stochastic analysis of the randomized low-rank approximation error plays a central role. Indeed, several randomized methods for the approximation of dominant eigen- or singular modes can be rewritten as low-rank approximation methods. However, despite the large variety of algorithms, the existing theoretical frameworks for their analysis rely on a specific structure for the covariance matrix that is not adapted to all the algorithms. We propose a general framework for the stochastic analysis of the low-rank approximation error in Frobenius norm for centered and non-standard Gaussian matrices. Under minimal assumptions on the covariance matrix, we derive accurate bounds both in expectation and probability. Our bounds have clear interpretations that enable us to derive properties and motivate practical choices for the covariance matrix resulting in efficient low-rank approximation algorithms. The most commonly used bounds in the literature have been demonstrated as a specific instance of the bounds proposed here, with the additional contribution of being tighter. Numerical experiments related to data assimilation further illustrate that exploiting the problem structure to select the covariance matrix improves the performance as suggested by our bounds.

Keywords : Low-rank approximation methods, randomized algorithms, Singular Value Decomposition, non-standard Gaussian error analysis, data assimilation

Acknowledgements: This work was funded by ISAE-SUPAERO, France.

1 Introduction

Let $\mathbf{A} \in \mathbb{R}^{m\times n}$ be an arbitrary real matrix and consider the following minimization problem,

$$\min_{\mathbf{B}\in\mathbb{R}^{m\times n}} \|\mathbf{A}-\mathbf{B}\|_F \quad \text{subject to} \quad \mathsf{rank}(\mathbf{B}) = k \leq \min\{m,n\}, \tag{1}$$

where $\|\cdot\|_F$ denotes the Frobenius norm. This problem, referred to as the *low-rank approximation* problem, is a key ingredient in numerous applications in data analysis and scientific computing including principal component analysis [23], data compression [19] and approximation algorithms for partial differential and integral equations [13], to name a few. Its solution [7] is obtained from the order k truncated singular value decomposition (SVD) of $\mathbf{A}$.

Let $\mathbf{U} \in \mathbb{R}^{m\times m}$ and $\mathbf{V} \in \mathbb{R}^{n\times n}$ be the real orthogonal matrices containing the left and right singular vectors of $\mathbf{A}$ respectively, and $\mathbf{\Sigma} = \mathsf{diag}(\sigma_1, \ldots, \sigma_{\min\{m,n\}}) \in \mathbb{R}^{m\times n}$ the matrix containing the singular values of $\mathbf{A}$ with the convention $\sigma_1 \geq \cdots \geq \sigma_{\min\{m,n\}}$. For a given integer $k \leq \min\{m,n\}$, we consider the following partitioning of the SVD of $\mathbf{A}$,

$$\mathbf{A} = \begin{bmatrix}\mathbf{U}_k & \overline{\mathbf{U}}_k\end{bmatrix}\begin{bmatrix}\mathbf{\Sigma}_k & \\ & \overline{\mathbf{\Sigma}}_k\end{bmatrix}\begin{bmatrix}\mathbf{V}_k & \overline{\mathbf{V}}_k\end{bmatrix}^\top, \tag{2}$$

where $\mathbf{U}_k \in \mathbb{R}^{m\times k}$, $\overline{\mathbf{U}}_k \in \mathbb{R}^{m\times(m-k)}$, $\mathbf{V}_k \in \mathbb{R}^{n\times k}$, $\overline{\mathbf{V}}_k \in \mathbb{R}^{n\times(n-k)}$, $\mathbf{\Sigma}_k \in \mathbb{R}^{k\times k}$ and $\overline{\mathbf{\Sigma}}_k \in \mathbb{R}^{(m-k)\times(n-k)}$. Defining $\mathbf{A}_k = \mathbf{U}_k\mathbf{\Sigma}_k\mathbf{V}_k^\top$ and $\overline{\mathbf{A}}_k = \overline{\mathbf{U}}_k\overline{\mathbf{\Sigma}}_k\overline{\mathbf{V}}_k^\top$, one can rewrite $\mathbf{A} = \mathbf{A}_k + \overline{\mathbf{A}}_k$ where the matrix $\mathbf{A}_k$ is precisely the solution of (1). The corresponding optimal value reads $\|\mathbf{A}-\mathbf{A}_k\|_F = \left\|\overline{\mathbf{\Sigma}}_k\right\|_F$.

The matrix $\mathbf{A}_k$ can be computed with any SVD algorithm enabling a truncation mechanism [8, Section 9.6]. However, for large-scale problems, the classical approaches become prohibitively expensive or are even inapplicable if $\mathbf{A}$ is not stored explicitly. In these situations, one is rather interested in computing $\widehat{\mathbf{A}}_k$, a reasonably accurate approximation of $\mathbf{A}_k$, which is typically expected to be optimal up to a small factor $\varepsilon > 0$, that is, $\widehat{\mathbf{A}}_k$ satisfies,

$$\left\|\mathbf{A}-\widehat{\mathbf{A}}_k\right\|_F \leq (1+\varepsilon)\|\mathbf{A}-\mathbf{A}_k\|_F.$$

Over the past decade, randomized algorithms for computing such approximations have been proposed. Following [30], the authors in [14] proposed an efficient algorithm now widely known as the Randomized SVD (RSVD) algorithm. Noticing that $\mathbf{A}_k = \mathbf{U}_k\mathbf{U}_k^\top\mathbf{A} = \pi(\mathbf{U}_k)\mathbf{A}$, the RSVD computes a low-rank approximation of $\mathbf{A}$ using a randomized procedure for estimating $\mathcal{R}(\mathbf{U}_k)$, where $\mathcal{R}(\cdot)$ denotes the column space. Several so-called *randomized range-finder* algorithms have been proposed and studied in the literature like the randomized subspace iteration method [23] or the randomized (block) Krylov subspace methods [21, 28, 31]. Regardless of the range-finder method, the RSVD can be theoretically analyzed by studying the following quantity

$$\left\|[\mathbf{I}_m - \pi(\mathbf{Z})]\mathbf{A}\right\|_F, \tag{3}$$

where $\mathbf{Z} \in \mathbb{R}^{m\times\ell}$ is a random matrix satisfying $\ell \geq k$ assumed to be ideally such that $\mathcal{R}(\mathbf{Z})$ approximates $\mathcal{R}(\mathbf{U}_k)$.

In this manuscript, we will focus on the randomized subspace iteration method which is the most commonly used method to approximate the dominant singular modes (and sometimes the only practicable one). The theoretical analysis of the RSVD based on the randomized subspace iteration has first been proposed in [14], where the authors have derived bounds for (3) both in expectation and in probability. Their analysis in Frobenius norm was limited to the case where $\mathbf{Z} = \mathbf{AG}$ with $\mathbf{G}$ a matrix whose columns are independently sampled from a n-variate standard Gaussian distribution. In a subsequent study, the author in [12] extended the analysis to the full randomized subspace iteration

and provided error bounds when $\mathbf{Z} = (\mathbf{A}\mathbf{A}^\top)^q\mathbf{A}\mathbf{G}$ with $q \geq 0$. More recently, the authors in [2] have proposed a generalization of the RSVD to the infinite-dimensional case, where $\mathbf{A}$ is replaced by a linear differential operator. From their theoretical analysis in the infinite-dimensional setting, they were able to propose error bounds [1] for (3) in the more general setting where $\mathbf{Z} = \mathbf{A}\mathbf{G}'$ and the columns of $\mathbf{G}'$ are non-standard Gaussian vectors with covariance matrix $\mathbf{C} \in \mathbb{R}^{n\times n}$. Their approach showed that a relevant choice for $\mathbf{C}$ could improve the performance.

In its generality, for randomized low-rank approximation, a structure is imposed to $\mathbf{Z}$, namely $\mathbf{Z} = f(\mathbf{A})\,\mathbf{G}$, where f is a function in $\mathbf{A}$ and $\mathbf{G}$ a Gaussian matrix whose columns are sampled from a Gaussian distribution with zero mean and covariance matrix $\mathbf{C}$. Here, the columns $\mathbf{z}_i = f(\mathbf{A})\mathbf{g}_i$ of $\mathbf{Z}$ also follow a Gaussian distribution with zero mean and covariance matrix

$$\mathbf{K} = \mathbb{E}\left[\, f(\mathbf{A})\,\mathbf{g}_i\mathbf{g}_i^\top\, f(\mathbf{A})^\top \,\right] = f(\mathbf{A})\,\mathbb{E}\left[\,\mathbf{g}_i\mathbf{g}_i^\top\,\right]\, f(\mathbf{A})^\top = f(\mathbf{A})\,\mathbf{C}\,f(\mathbf{A})^\top. \tag{4}$$

Consequently, the stochastic analysis of (3) is determined by the covariance matrix $\mathbf{K}$ of $\mathbf{z}_i$, and existing results all rely on particular structures for $\mathbf{K}$.

Beyond the RSVD, various randomized algorithms were proposed to address more elaborated eigenvalue and singular value problems. In [24, 27], the authors proposed methods to compute dominant eigenpairs of generalized Hermitian eigenvalue problems, while in [25], algorithms for computing truncated generalized singular value decomposition [29] were derived. Considering the appropriate matrices $\mathbf{A}$ and $\mathbf{Z}$, these algorithms all try, by design, to minimize a quantity of the form (3). However, their theoretical analysis cannot rely on the existing frameworks because the resulting covariance matrix of $\mathbf{Z}$ has a structure that differs from (4). In this regard, the analysis proposed by the authors was, if not incomplete, at least not entirely satisfying. This is for instance the case in [26], where no analysis in expectation are provided and only one algorithm out of three has an analysis in probability. Filling this gap in the literature is our main motivation.

We propose a stochastic analysis of the low-rank approximation error (3) which holds for any matrix $\mathbf{Z}$ whose columns are independently sampled from a multivariate Gaussian distribution with zero mean and covariance matrix $\mathbf{K}$. In particular, we do not assume any particular structure for $\mathbf{K}$. The proposed expectation and probability bounds allow one to obtain key properties of the covariance matrix $\mathbf{K}$ resulting in improved performance. This analysis can then be used to increase the efficiency of the current algorithms in learning the matrix $\mathbf{A}$. The main advantage of our general framework is that it allows us to analyze any randomized algorithms that can be rewritten as a low-rank approximation problem. In addition, all the prior results can be recovered by using our general framework unifying a large variety of error analyses.

The outline is the following. We first introduce key background material that will be helpful throughout the manuscript in Section 2. Section 3 details our main result on the general error analysis stated in Theorem 3.1. Section 3.1 illustrates our proposed analysis in three practical contexts: the classical power iteration scheme, the generalized RSVD [1] and the case where the covariance matrix is constructed out of a priori information. For the first two contexts, a comparison of our bounds with the ones from [12] and [1] is proposed. In Section 4, we propose numerical illustrations on a data assimilation problem [5, 18] and show that exploiting the particular structure of the problem to define the covariance matrix $\mathbf{K}$ improves the performance of the RSVD.

2 Preliminaries

In this section, we recall well-known key results and definitions from the literature.

Submultiplicativity Let $\|\cdot\|_2$ and $\|\cdot\|_F$ denote the spectral and the Frobenius norm respectively. The strong submultiplicativity property [15, Relation (B.7)] reads

$$\forall\ \mathbf{M} \in \mathbb{R}^{n\times p}, \mathbf{N} \in \mathbb{R}^{p\times q}, \quad \|\mathbf{M}\mathbf{N}\|_F \leq \|\mathbf{M}\|_F\|\mathbf{N}\|_2\,. \tag{5}$$

Partial ordering on the set of symmetric matrices Let $\mathbf{M} \in \mathbb{R}^{m\times m}, \mathbf{N} \in \mathbb{R}^{m\times m}$ be two symmetric matrices. The notation $\mathbf{M} \preccurlyeq \mathbf{N}$ means that $\mathbf{N} - \mathbf{M}$ is positive semi-definite. This relation defines a partial ordering on the set of symmetric matrices [16, Section 7.7]. An important property [16, Theorem 7.7.2] is that the partial ordering is preserved under the conjugation rule, i.e.,

$$\mathbf{M} \preccurlyeq \mathbf{N} \implies \mathbf{Q}^\top \mathbf{M}\mathbf{Q} \preccurlyeq \mathbf{Q}^\top \mathbf{N}\mathbf{Q}, \quad \forall\, \mathbf{Q} \in \mathbb{R}^{m\times n}. \tag{6}$$

We note that as a consequence of [16, Corollary 7.7.4(c)], the trace is monotonic with respect to the partial ordering, i.e.,

$$\mathbf{M} \preccurlyeq \mathbf{N} \implies \mathsf{tr}(\mathbf{M}) \leq \mathsf{tr}(\mathbf{N}). \tag{7}$$

Projection matrices Suppose that $\mathbf{M} \in \mathbb{R}^{m\times n}$ has full column rank with column range space denoted by $\mathcal{R}(\mathbf{M})$. We denote by $\mathbf{M}^\dagger$ the left multiplicative inverse of $\mathbf{M}$, i.e., the Moore-Penrose inverse of $\mathbf{M}$, see, e.g., [16]. The orthogonal projection on $\mathcal{R}(\mathbf{M})$ is then given by $\pi(\mathbf{M}) = \mathbf{M}\mathbf{M}^\dagger$, in particular, one has $\mathcal{R}(\pi(\mathbf{M})) = \mathcal{R}(\mathbf{M})$.

Sherman-Morrison formula Let $\mathbf{M} \in \mathbb{R}^{m\times n}$ and $\mathbf{N} \in \mathbb{R}^{n\times m}$ such that $\mathbf{I}_n + \mathbf{N}\mathbf{M}$ is non-singular. Then, $\mathbf{I}_m + \mathbf{M}\mathbf{N}$ is also non-singular and one has [9, Section 2.1.4],

$$(\mathbf{I}_m + \mathbf{M}\mathbf{N})^{-1} = \mathbf{I}_m - \mathbf{M}(\mathbf{I}_n + \mathbf{N}\mathbf{M})^{-1}\mathbf{N}. \tag{8}$$

Principal angles between subspaces Let $\mathcal{M}, \mathcal{N} \subset \mathbb{R}^m$ be two k-dimensional subspaces, and $\mathbf{M}, \mathbf{N} \in \mathbb{R}^{m\times k}$ be two associated matrices with orthogonal columns satisfying $\mathcal{R}(\mathbf{M}) = \mathcal{M}$ and $\mathcal{R}(\mathbf{N}) = \mathcal{N}$ respectively. If $\sigma_i(\cdot)$ denotes the i-th largest singular value of a given matrix, then the principal angles between $\mathcal{M}$ and $\mathcal{N}$, denoted by $\theta_1, \ldots, \theta_k$ are defined as follow [17],

$$\theta_i = \arccos(\sigma_i(\mathbf{M}^\top \mathbf{N})), \quad 1 \leq i \leq k.$$

Further details on this notion can be found in [22].

Tangent matrix of the principal angles Let $[\,\mathbf{M},\, \overline{\mathbf{M}}\,] \in \mathbb{R}^{n\times n}$ be an orthogonal matrix with $\mathbf{M} \in \mathbb{R}^{n\times k}$, and $\mathbf{N} \in \mathbb{R}^{n\times k}$ a full column rank matrix. Then the following singular value decomposition [32, Theorem 3.1] yields

$$\mathbf{tan}\,(\mathbf{M},\, \mathbf{N}) \equiv \overline{\mathbf{M}}^\top \mathbf{N}(\mathbf{M}^\top \mathbf{N})^\dagger = \mathbf{U}\,\mathsf{diag}\left(\tan(\theta_1), \ldots, \tan(\theta_k)\right)\mathbf{V}^\top, \tag{9}$$

where $\mathbf{U} \in \mathbb{R}^{(n-k)\times k}$ has orthonormal columns, $\mathbf{V} \in \mathbb{R}^{k\times k}$ is orthogonal and scalars $\theta_1, \ldots, \theta_k$ are the principal angles between $\mathcal{R}(\mathbf{M})$ and $\mathcal{R}(\mathbf{N})$.

3 General randomized low-rank approximation error bounds

A generalized RSVD is given in Algorithm 1. In the first stage, one searches for an approximation of $\mathcal{R}(\mathbf{A})$ by sampling $\mathbf{A}$ on ℓ independently drawn Gaussian vectors with zero mean and covariance matrix $\mathbf{K}$. From a practical viewpoint, if a factorization $\mathbf{K} = \mathbf{L}^\top \mathbf{L}$ is available, then drawing the columns of $\mathbf{Z}$ reduces to compute $\mathbf{Z} = \mathbf{L}\mathbf{\Omega}$, where the columns of $\mathbf{\Omega}$ are standard Gaussian vectors. The second stage extracts the low-rank approximation of $\mathbf{A}$ from the orthonormal basis $\mathbf{Q}$ of $\mathcal{R}(\mathbf{Z})$.

We provide the general stochastic analysis of Algorithm 1 in Theorem 3.1 This theorem, which is our main contribution, extends the randomized low-rank approximation error analysis to a general covariance matrix $\mathbf{K}$, i.e. without assuming any specific structure. As will be shown in Section 3.1, our result unifies several existing error analysis bounds provided in the literature. For the sake of readability, we postpone the proof to Appendix A.

Algorithm 1: A general randomized singular value decomposition framework.

Input: Matrix $\mathbf{A} \in \mathbb{R}^{m\times n}$, number of samples $1 \leq \ell \leq \mathsf{rank}(\mathbf{A})$, target rank $1 \leq k \leq \ell$ to provide, covariance matrix $\mathbf{K} \in \mathbb{R}^{m\times m}$.

`// Stage 1: Randomized range-finder`

1 Draw $\mathbf{Z} \in \mathbb{R}^{m\times \ell}$ whose columns are centered Gaussian with covariance matrix $\mathbf{K}$.

2 Perform the thin QR factorization $\mathbf{Z} = \mathbf{QR}$.

`// Stage 2: Low-rank approximation`

3 Compute $\mathbf{Y} = \mathbf{Q}^\top \mathbf{A}$.

4 Perform the order k truncated SVD of $\mathbf{Y}$, i.e. $\mathbf{Y} = \mathbf{U}_k \widehat{\mathbf{\Sigma}}_k \widehat{\mathbf{V}}_k^\top$.

5 Compute $\widehat{\mathbf{U}}_k = \mathbf{Q}\mathbf{U}_k$.

Output: Matrices $\widehat{\mathbf{U}}_k \in \mathbb{R}^{m\times k}$, $\widehat{\mathbf{\Sigma}}_k \in \mathbb{R}^{k\times k}$ and $\widehat{\mathbf{V}}_k \in \mathbb{R}^{n\times k}$ such that $\mathbf{A} \approx \widehat{\mathbf{U}}_k \widehat{\mathbf{\Sigma}}_k \widehat{\mathbf{V}}_k^\top$.

Theorem 3.1. Let $\mathbf{A} \in \mathbb{R}^{m\times n}$ be an arbitrary matrix and $\mathbf{Z} \in \mathbb{R}^{m\times \ell}$ be a matrix whose columns are independently sampled from a multivariate Gaussian distribution with covariance matrix $\mathbf{K} \in \mathbb{R}^{m\times m}$, such that $\ell \leq \min(\mathsf{rank}(\mathbf{A}), \mathsf{rank}(\mathbf{K}))$. Let us further consider $\mathbf{A} = \mathbf{U\Sigma V}^\top$, the singular value decomposition of $\mathbf{A}$, and the related matrices $\mathbf{U}_k$, $\mathbf{\Sigma}_k$ and $\overline{\mathbf{\Sigma}}_k$ as given by (2).

For any integer $k \leq \ell - 2$, if $\mathbf{K}_k \equiv \mathbf{U}_k^\top \mathbf{K} \mathbf{U}_k$ is non-singular, then one has

$$\mathbb{E}\left[\left\|[\mathbf{I}_n - \pi(\mathbf{Z})]\mathbf{A}\right\|_F\right] \leq \left(1 + \tau_k(\mathbf{K})^2 + \frac{\rho_k(\mathbf{K})^2}{\ell - k - 1}\right)^{\frac{1}{2}} \left\|\overline{\mathbf{\Sigma}}_k\right\|_F.$$

Moreover, if $k \leq \ell - 4$, then for all $u, t \geq 1$,

$$\left\|[\mathbf{I}_n - \pi(\mathbf{Z})]\mathbf{A}\right\|_F \leq \left(1 + \tau_k(\mathbf{K}) + \sqrt{3}ut \cdot \frac{\rho_k(\mathbf{K})}{\sqrt{\ell - k + 1}}\right)\left\|\overline{\mathbf{\Sigma}}_k\right\|_F.$$

holds with probability at least $1 - e^{-u^2/2} - t^{-(\ell-k)}$ with

$$\tau_k(\mathbf{K}) = \frac{\left\|\mathbf{tan}(\mathbf{U}_k, \mathbf{K}\mathbf{U}_k)\mathbf{\Sigma}_k\right\|_F}{\left\|\overline{\mathbf{\Sigma}}_k\right\|_F},$$

$$\rho_k(\mathbf{K}) = \left\|[\mathbf{I}_m - \pi(\mathbf{K}^{\frac{1}{2}}\mathbf{U}_k)]\mathbf{K}^{\frac{1}{2}}\right\|_F \frac{\sqrt{\mathsf{tr}(\mathbf{\Sigma}_k^2 \mathbf{K}_k^{-1})}}{\left\|\overline{\mathbf{\Sigma}}_k\right\|_F}.$$

Theorem 3.1 states that the deviation from the optimal error of the randomized low-rank approximation is monitored by two coefficients $\tau_k(\mathbf{K})$ and $\rho_k(\mathbf{K})$, respectively. The error is optimal whenever they both equal zero. The first one depends on the tangent of the principal angles (9) between $\mathcal{R}(\mathbf{U}_k)$ and $\mathcal{R}(\mathbf{K}\mathbf{U}_k)$. Consequently, $\tau_k(\mathbf{K})$ will be small when the principal angles are small, that is when $\mathcal{R}(\mathbf{U}_k)$ is close to $\mathcal{R}(\mathbf{K}\mathbf{U}_k)$. This suggests to choose $\mathbf{K}$ so that $\mathcal{R}(\mathbf{U}_k)$ is an approximate invariant subspace under the action of $\mathbf{K}$. The second coefficient $\rho_k(\mathbf{K})$ depends on two quantities. The term $\left\|[\mathbf{I}_m - \pi(\mathbf{K}^{\frac{1}{2}}\mathbf{U}_k)]\mathbf{K}^{\frac{1}{2}}\right\|_F$ is the low-rank approximation error of $\mathbf{K}^{\frac{1}{2}}$ induced by $\mathcal{R}(\mathbf{K}^{\frac{1}{2}}\mathbf{U}_k)$. This quantity is small if $\mathcal{R}(\mathbf{K}^{\frac{1}{2}}\mathbf{U}_k)$ is close to $\mathcal{R}(\mathbf{K}^{\frac{1}{2}})$, and the minimum value is reached whenever $\mathcal{R}(\mathbf{K}^{\frac{1}{2}}\mathbf{U}_k)$ is equal to the dominant eigenspace of $\mathbf{K}$.

For the second term, one observes that

$$\mathsf{tr}\left(\mathbf{\Sigma}_k^2 \mathbf{K}_k^{-1}\right) \leq \frac{1}{\lambda_{\min}(\mathbf{K}_k)} \|\mathbf{\Sigma}_k\|_F. \tag{10}$$

This term can thus be arbitrarily small when the eigenvalues of $\mathbf{K}_k$ (or equivalently $\pi(\mathbf{U}_k)\mathbf{K}$) are large.

Altogether, Theorem 3.1 highlights the main property that $\mathbf{K}$ must satisfy in order to obtain an efficient low-rank approximation: $\mathcal{R}(\mathbf{U}_k)$ should be well aligned with the dominant eigenspace of $\mathbf{K}$. Consequently, choosing the matrix $\mathbf{K}$ such that

$$\mathbf{K} = \mathbf{U}_k \mathbf{\Delta} \mathbf{U}_k^\top, \tag{11}$$

with $\mathbf{\Delta} \in \mathbb{R}^{k\times k}$ a symmetric positive definite matrix, leads to optimal values for the first error term $\tau_k(\mathbf{K})$ and the first coefficient of the second error term $\rho_k(\mathbf{K})$ for any choice of $\mathbf{\Delta}$. Note that, with the choice of the $\mathbf{K}$ matrix (11), the bound (10) depends on the choice of $\mathbf{\Delta}$, i.e. $\lambda_{\min}(\mathbf{K}_k) = \lambda_{\min}(\mathbf{\Delta})$.

In practice, since $\mathbf{U}_k$ is precisely the quantity that needs to be determined, the choice of the matrix $\mathbf{K}$ expressed by (11) is irrelevant. Instead, one rathers consider $\mathbf{K} = \mathbf{A}\mathbf{A}^\top$ as in the RSVD algorithm. Observing that $\mathbf{A}\mathbf{A}^\top = \mathbf{A}_k\mathbf{A}_k^\top + \overline{\mathbf{A}}_k\overline{\mathbf{A}}_k^\top$ and that $\mathbf{A}_k\mathbf{A}_k^\top = \mathbf{U}_k\mathbf{\Sigma}_k^2\mathbf{U}_k^\top$ is precisely of the form (11), then the RSVD is particularly efficient if $\mathbf{A}\mathbf{A}^\top \approx \mathbf{A}_k\mathbf{A}_k^\top$, that is whenever $\overline{\mathbf{A}}_k\overline{\mathbf{A}}_k^\top$ is negligible compared to $\mathbf{A}_k\mathbf{A}_k^\top$. We recover here a property well documented in the randomized linear algebra literature, that is Algorithm 1 is expected to perform well when the tail of the singular value distribution is negligible compared to the first k singular values. Further discussions on practical choices for $\mathbf{K}$ are proposed in Section 3.1.

3.1 Particular choices for the covariance matrix

In this section, we particularize Theorem 3.1 for certain covariance matrices. As we will show, several prior bounds, that have been independently derived, can be recovered from Theorem 3.1. In Section 3.1.1, we first address the power iteration scheme, where we can compare our bounds to the ones proposed in [12]. Secondly, in Section 3.1.2 we propose an analysis of the Generalized RSVD introduced in [1]. Finally, in Section 3.1.3, we discuss relevant cases where the covariance matrix is constructed using available a priori information.

3.1.1 The power iteration scheme

The power iteration scheme considers $\mathbf{Z} = f(\mathbf{A})\mathbf{G}$ with $f(\mathbf{A}) = \mathbf{A}(\mathbf{A}^\top\mathbf{A})^q$, $q \geq 0$ and $\mathbf{G}$ being a matrix whose columns are independently sampled from a standard multivariate Gaussian distribution. Equivalently, one considers a matrix $\mathbf{Z}$ with columns sampled from a centered multivariate Gaussian distribution with covariance matrix $\mathbf{K}^{(\mathrm{p})} = f(\mathbf{A})f(\mathbf{A})^\top = \mathbf{A}(\mathbf{A}^\top\mathbf{A})^{2q}\mathbf{A}^\top$. Exploiting this structure in Theorem 3.1 yields the following corollary.

Corollary 3.2. Let $\mathbf{A} \in \mathbb{R}^{m\times n}$ be an arbitrary matrix, $\mathbf{G} \in \mathbb{R}^{n\times \ell}$ a matrix whose columns are independently sampled from a standard multivariate Gaussian distribution with covariance matrix such that $\ell \leq \mathsf{rank}(\mathbf{A})$, and $\mathbf{Z} = \mathbf{A}(\mathbf{A}^\top\mathbf{A})^q\mathbf{G}$ with $q \geq 0$. Let us further consider $\sigma_1 \geq \cdots \geq \sigma_n$ the singular values of $\mathbf{A}$ and the matrix $\overline{\mathbf{\Sigma}}_k$ as given by (2).

For any integer $k \leq \ell - 2$, one has

$$\mathbb{E}\left[\left\|[\mathbf{I}_n - \pi(\mathbf{Z})]\mathbf{A}\right\|_F\right] \leq \left(1 + \left(\frac{\sigma_{k+1}}{\sigma_k}\right)^{2q}\sqrt{\frac{k}{\ell-k-1}}\right)\left\|\overline{\mathbf{\Sigma}}_k\right\|_F.$$

Moreover, if $k \leq \ell - 4$, then for all $u, t \geq 1$,

$$\left\|[\mathbf{I}_n - \pi(\mathbf{Z})]\mathbf{A}\right\|_F \leq \left(1 + \sqrt{3}ut\cdot\left(\frac{\sigma_{k+1}}{\sigma_k}\right)^{2q}\sqrt{\frac{k}{\ell-k+1}}\right)\left\|\overline{\mathbf{\Sigma}}_k\right\|_F,$$

holds with probability at least $1 - e^{-u^2/2} - t^{-(\ell-k)}$.

Proof. Let us first define the covariance matrix $\mathbf{K}^{(\mathrm{p})}$ of the columns of $\mathbf{Z}$:

$$\begin{aligned}\mathbf{K}^{(\mathrm{p})} &= \mathbf{A}(\mathbf{A}^\top\mathbf{A})^{2q}\mathbf{A}^\top \\ &= \mathbf{U}\boldsymbol{\Sigma}(\boldsymbol{\Sigma}^\top\boldsymbol{\Sigma})^{2q}\boldsymbol{\Sigma}^\top\mathbf{U}^\top \\ &= \mathbf{U}_k\boldsymbol{\Sigma}_k^{4q+2}\mathbf{U}_k^\top + \overline{\mathbf{U}}_k(\overline{\boldsymbol{\Sigma}}_k\overline{\boldsymbol{\Sigma}}_k^\top)^{2q+1}\overline{\mathbf{U}}_k^\top.\end{aligned}$$

From this, it readily follows that $\mathbf{U}_k^\top\mathbf{K}^{(\mathrm{p})}\mathbf{U}_k = \boldsymbol{\Sigma}_k^{4q+2}$ which is non-singular since $k \le \ell \le \mathsf{rank}(\mathbf{A})$ by assumption, and Theorem 3.1 can be applied. To compute the two coefficients, let us first observe that $\mathbf{K}^{(\mathrm{p})}\mathbf{U}_k = \mathbf{U}_k\boldsymbol{\Sigma}_k^{4q+2}$, that is $\mathcal{R}(\mathbf{K}^{(\mathrm{p})}\mathbf{U}_k) = \mathcal{R}(\mathbf{U}_k)$. One thus deduces that

$$\mathbf{tan}(\mathbf{U}_k, \mathbf{K}^{(\mathrm{p})}\mathbf{U}_k) = \mathbf{0}.$$

Then, one has

$$\begin{aligned}\left\|[\mathbf{I}_m - \pi([\mathbf{K}^{(\mathrm{p})}]^{\frac{1}{2}}\mathbf{U}_k)][\mathbf{K}^{(\mathrm{p})}]^{\frac{1}{2}}\right\|_F^2 &= \mathsf{tr}\left(\mathbf{K}^{(\mathrm{p})} - \mathbf{K}^{(\mathrm{p})}\mathbf{U}_k(\mathbf{U}_k^\top\mathbf{K}^{(\mathrm{p})}\mathbf{U}_k)^{-1}\mathbf{U}_k^\top\mathbf{K}^{(\mathrm{p})}\right) \\ &= \mathsf{tr}\left(\mathbf{K}^{(\mathrm{p})} - \mathbf{U}_k\boldsymbol{\Sigma}_k^{4q+2}\mathbf{U}_k^\top\right) \\ &= \mathsf{tr}\left(\overline{\mathbf{U}}_k(\overline{\boldsymbol{\Sigma}}_k\overline{\boldsymbol{\Sigma}}_k^\top)^{2q+1}\overline{\mathbf{U}}_k^\top\right) \\ &= \mathsf{tr}\left((\overline{\boldsymbol{\Sigma}}_k\overline{\boldsymbol{\Sigma}}_k^\top)^{2q+1}\right) \\ &= \mathsf{tr}\left(\overline{\boldsymbol{\Sigma}}_k(\overline{\boldsymbol{\Sigma}}_k^\top\overline{\boldsymbol{\Sigma}}_k)^{2q}\overline{\boldsymbol{\Sigma}}_k^\top\right) = \left\|(\overline{\boldsymbol{\Sigma}}_k^\top\overline{\boldsymbol{\Sigma}}_k)^q\overline{\boldsymbol{\Sigma}}_k^\top\right\|_F^2\end{aligned}$$

Altogether one obtains

$$\tau_k\left(\mathbf{K}^{(\mathrm{p})}\right) = 0 \qquad \text{and} \qquad \rho_k\left(\mathbf{K}^{(\mathrm{p})}\right) = \left\|\boldsymbol{\Sigma}_k^{-2q}\right\|_F \frac{\left\|(\overline{\boldsymbol{\Sigma}}_k^\top\overline{\boldsymbol{\Sigma}}_k)^q\overline{\boldsymbol{\Sigma}}_k^\top\right\|_F}{\left\|\overline{\boldsymbol{\Sigma}}_k\right\|_F}.$$

Applying the strong submultiplicativity (5) and using norm equivalence yields $\left\|(\overline{\boldsymbol{\Sigma}}_k^\top\overline{\boldsymbol{\Sigma}}_k)^q\overline{\boldsymbol{\Sigma}}_k^\top\right\|_F \le \sigma_{k+1}^{2q}\left\|\overline{\boldsymbol{\Sigma}}_k\right\|_F$ and $\left\|\boldsymbol{\Sigma}_k^{-2q}\right\|_F \le \sqrt{k}\sigma_k^{-2q}$ respectively, from which we finally obtain

$$\rho_k\left(\mathbf{K}^{(\mathrm{p})}\right) \le \sqrt{k}\left(\frac{\sigma_{k+1}}{\sigma_k}\right)^{2q}. \qquad \square$$

If $q = 0$, we recover the expectation bound proposed in [14, Theorem 10.5]. For easier comparison of the bounds in probability, we recall [14, Theorem 10.7] which states that under the same assumptions as Corollary 3.2, then

$$\left\|[\mathbf{I}_n - \pi(\mathbf{Z})]\mathbf{A}\right\|_F \le \left(1 + t\cdot\sqrt{\frac{3k}{\ell-k+1}}\right)\left\|\overline{\boldsymbol{\Sigma}}_k\right\|_F + ut\frac{e\sqrt{\ell}}{\sqrt{\ell-k+1}}\sigma_{k+1}$$

holds with probability at least $1 - e^{-u^2/2} - 2t^{-(\ell-k)}$, where $k \le \ell - 4$ and $u, t \ge 1$. Our bound is almost identical to the first term in $\left\|\overline{\boldsymbol{\Sigma}}_k\right\|_F$ except for the extra factor u. However, since our bound does not have the second term in σ_{k+1}, then for reasonably small values of u (e.g. 2 or 3) one can expect our probability bound to have comparable accuracy.

For $q \ge 1$, we recover the convergence property of the power iteration scheme, i.e. the error approaches the optimal value as $(\sigma_{k+1}/\sigma_k)^{2q} \to 0$. For comparison, using the property that $\sqrt{1+x} \le 1 + \sqrt{x}$ for $x \ge 0$, Theorem 5.7 in [12] provides the following upper bound for $p \ge 2$:

$$\mathbb{E}\left[\left\|[\mathbf{I}_n - \pi(\mathbf{Z})]\mathbf{A}\right\|_F\right] \leq \left(1 + \sqrt{k}\left(\frac{\sigma_{k+1}}{\sigma_k}\right)^{2q} \frac{4e\sqrt{\ell}}{\ell - k + 1}\left[\sqrt{n-k} + \sqrt{\ell} + 7\right] \frac{\sigma_{k+1}}{\left\|\overline{\mathbf{\Sigma}}_k\right\|_F}\right)\left\|\overline{\mathbf{\Sigma}}_k\right\|_F.$$

This bound is similar to our bound except for the multiplying factor

$$\mathsf{c} = 4e\sqrt{\ell}\,\frac{\sqrt{\ell - k - 1}}{\ell - k + 1}\left[\sqrt{n-k} + \sqrt{\ell} + 7\right]\frac{\sigma_{k+1}}{\left\|\overline{\mathbf{\Sigma}}_k\right\|_F}.$$

Since $\left\|\overline{\mathbf{\Sigma}}_k\right\|_F \leq \sigma_{k+1}\sqrt{n-k}$ from classical norm equivalence, and both $\ell \geq \ell - k - 1$ and $\ell - k \geq 2$ by assumption, it yields

$$\mathsf{c} \geq \frac{4e}{3}(\ell - k - 1).$$

This suggests that Corollary 3.2 provides tighter bounds than the ones proposed in [12], especially when the oversampling $\ell - k$ is large. Nevertheless, the refinement might not be significant and only illustrates that bounds from [12] are slightly over-pessimistic.

3.1.2 The generalized RSVD

Let us now consider $\mathbf{Z} = f(\mathbf{A})\mathbf{G}'$ with $f(\mathbf{A}) = \mathbf{A}$ and, $\mathbf{G}'$ is a matrix whose columns are independently sampled from a multivariate Gaussian distribution with covariance matrix $\mathbf{C} \in \mathbb{R}^{n\times n}$. Again, this can be alternatively viewed as drawing the columns of $\mathbf{Z}$ from a centered multivariate Gaussian distribution with covariance matrix $\mathbf{K}^{(\mathrm{g})} = \mathbf{ACA}^\top$. This configuration has been for instance studied in [1], to improve the accuracy of low-rank approximations of differential operators. Lemma 3.3 provides simplified forms for both $\tau_k(\mathbf{K}^{(\mathrm{g})})$ and $\rho_k(\mathbf{K}^{(\mathrm{g})})$.

Lemma 3.3. Let $\mathbf{A} \in \mathbb{R}^{m\times n}$ be an arbitrary matrix, $\mathbf{G} \in \mathbb{R}^{n\times\ell}$ a matrix whose columns are independently sampled from a multivariate Gaussian distribution with covariance matrix $\mathbf{C} \in \mathbb{R}^{n\times n}$ such that $\ell \leq \mathsf{rank}(\mathbf{A})$, and $\mathbf{Z} = \mathbf{AG} \in \mathbb{R}^{m\times\ell}$. Let us further consider $\mathbf{A} = \mathbf{U\Sigma V}^\top$, the singular value decomposition of $\mathbf{A}$, and the related matrices $\mathbf{V}_k$ and $\overline{\mathbf{\Sigma}}_k$ as given by (2).

For any integer $k \leq \ell - 2$, one has

$$\tau_k\left(\mathbf{K}^{(\mathrm{g})}\right) = \frac{\left\|\overline{\mathbf{\Sigma}}_k \tan(\mathbf{V}_k, \mathbf{CV}_k)\right\|_F}{\left\|\overline{\mathbf{\Sigma}}_k\right\|_F},$$

$$\text{and} \quad \rho_k\left(\mathbf{K}^{(\mathrm{g})}\right) = \left\|[\mathbf{I}_n - \pi(\mathbf{C}^{\frac{1}{2}}\mathbf{V}_k)]\mathbf{C}^{\frac{1}{2}}\mathbf{A}^\top\right\|_F \frac{\sqrt{\mathsf{tr}((\mathbf{V}_k^\top\mathbf{CV}_k)^{-1})}}{\left\|\overline{\mathbf{\Sigma}}_k\right\|_F},$$

where $\mathbf{K}^{(\mathrm{g})} = \mathbf{ACA}^\top$ is the covariance matrix of the columns of $\mathbf{Z}$. The coefficients $\tau_k(\cdot)$ and $\rho_k(\cdot)$ are defined in Theorem 3.1.

Proof. The standard properties of Gaussian vectors yield $\mathbf{K}^{(\mathrm{g})} = \mathbf{ACA}^\top$, from which we obtain

$$\mathbf{U}^\top\mathbf{K}^{(\mathrm{g})}\mathbf{U} = \mathbf{\Sigma V}^\top\mathbf{CV\Sigma}^\top.$$

Straightforward algebraic manipulations then yield

$$\begin{aligned}\tan(\mathbf{U}_k, \mathbf{K}^{(\mathrm{g})}\mathbf{U}_k)\mathbf{\Sigma}_k &= \overline{\mathbf{U}}_k^\top\mathbf{K}^{(\mathrm{g})}\mathbf{U}_k(\mathbf{U}_k^\top\mathbf{K}^{(\mathrm{g})}\mathbf{U}_k)^{-1}\mathbf{\Sigma}_k \\ &= \overline{\mathbf{\Sigma}}_k\overline{\mathbf{V}}_k^\top\mathbf{CV}_k(\mathbf{V}_k^\top\mathbf{CV}_k)^{-1}\end{aligned}$$

$$= \overline{\mathbf{\Sigma}}_k \tan(\mathbf{V}_k, \mathbf{C}\mathbf{V}_k).$$

Likewise, one gets

$$\begin{aligned}
\left\| [\mathbf{I}_m - \pi([\mathbf{K}^{(\mathrm{g})}]^{\frac{1}{2}}\mathbf{U}_k)][\mathbf{K}^{(\mathrm{g})}]^{\frac{1}{2}} \right\|_F^2 &= \mathsf{tr}\left([\mathbf{K}^{(\mathrm{g})}]^{\frac{1}{2}} [\mathbf{I}_m - \pi([\mathbf{K}^{(\mathrm{g})}]^{\frac{1}{2}}\mathbf{U}_k)][\mathbf{K}^{(\mathrm{g})}]^{\frac{1}{2}} \right) \\
&= \mathsf{tr}\left(\mathbf{K}^{(\mathrm{g})} [\mathbf{I}_m - \pi([\mathbf{K}^{(\mathrm{g})}]^{\frac{1}{2}}\mathbf{U}_k)] \right) \\
&= \mathsf{tr}\left(\mathbf{K}^{(\mathrm{g})} [\mathbf{I}_m - \pi([\mathbf{K}^{(\mathrm{g})}]^{\frac{1}{2}}\mathbf{U}_k)] \right) \\
&= \mathsf{tr}\left(\mathbf{A}\mathbf{C}^{\frac{1}{2}} \left[\mathbf{I}_n - \pi(\mathbf{C}\mathbf{V}_k)\right] \mathbf{C}^{\frac{1}{2}}\mathbf{A}^\top \right) \\
&= \left\| [\mathbf{I}_n - \pi(\mathbf{C}^{\frac{1}{2}}\mathbf{V}_k)]\mathbf{C}^{\frac{1}{2}}\mathbf{A}^\top \right\|_F^2,
\end{aligned}$$

where we have used the circular invariance of the trace and the properties of projectors. Finally, observing that $\mathbf{K}_k = \mathbf{\Sigma}_k \mathbf{V}_k^\top \mathbf{C}\mathbf{V}_k \mathbf{\Sigma}_k$ concludes the proof. □

The structure of $\tau_k(\mathbf{K}^{(\mathrm{g})})$ and $\rho_k(\mathbf{K}^{(\mathrm{g})})$ resembles the general form stated in Theorem 3.1. As a consequence, one can readily draw similar conclusions: $\mathbf{C}$ must be chosen so that $\mathbf{V}_k$ approximates its dominant eigenspace. For instance $\mathbf{C} = \mathbf{A}_k^\top \mathbf{A}_k = \mathbf{V}_k \mathbf{\Sigma}_k^2 \mathbf{V}_k^\top$, although this is not practical since matrices $\mathbf{A}_k$, $\mathbf{V}_k$ and $\mathbf{\Sigma}_k$ may not be available.

Let us now use Lemma 3.3 and Theorem 3.1 to provide bounds comparable to [1].

Corollary 3.4. Let $\mathbf{A} \in \mathbb{R}^{m\times n}$ be an arbitrary matrix, $\mathbf{G} \in \mathbb{R}^{n\times \ell}$ a matrix whose columns are independently sampled from a multivariate Gaussian distribution with covariance matrix $\mathbf{C} \in \mathbb{R}^{n\times n}$ such that $\ell \leq \mathsf{rank}(\mathbf{A})$, and $\mathbf{Z} = \mathbf{A}\mathbf{G} \in \mathbb{R}^{m\times \ell}$. Let further $\mathbf{A} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^\top$ be the singular value decomposition of $\mathbf{A}$, and the related matrices $\mathbf{V}_k$ and $\overline{\mathbf{\Sigma}}_k$ as given by (2).

For any integer $k \leq \ell - 2$, if $\mathbf{V}_k^\top \mathbf{C}\mathbf{V}_k$ is non-singular, then one has

$$\mathbb{E}\left[\left\| [\mathbf{I}_n - \pi(\mathbf{Z})]\mathbf{A} \right\|_F \right] \leq \left(1 + \sqrt{\ell - k} \sqrt{\frac{k}{\ell - k - 1} \frac{\beta_k(\mathbf{C})}{\gamma_k(\mathbf{C})}} \right) \left\| \overline{\mathbf{\Sigma}}_k \right\|_F.$$

Moreover, if $k \leq \ell - 4$, then for all $u, t \geq 1$,

$$\left\| [\mathbf{I}_n - \pi(\mathbf{Z})]\mathbf{A} \right\|_F \leq \left(1 + \sqrt{3}\left(\sqrt{\ell - k + 1} + 1 \right) \sqrt{\frac{k}{\ell - k + 1} \frac{\beta_k(\mathbf{C})}{\gamma_k(\mathbf{C})}} \cdot ut \right) \left\| \overline{\mathbf{\Sigma}}_k \right\|_F,$$

holds with probability at least $1 - e^{-u^2/2} - t^{-(\ell - k)}$.

Here,

$$\beta_k(\mathbf{C}) = \frac{\mathsf{tr}\left(\overline{\mathbf{\Sigma}}_k^\top \overline{\mathbf{\Sigma}}_k \overline{\mathbf{V}}_k^\top \mathbf{C} \overline{\mathbf{V}}_k \right)}{\lambda_1(\mathbf{C}) \left\| \overline{\mathbf{\Sigma}}_k \right\|_F^2} \quad \text{and} \quad \gamma_k(\mathbf{C}) = \frac{k}{\lambda_1(\mathbf{C})\ \mathsf{tr}\left((\mathbf{V}_k^\top \mathbf{C}\mathbf{V}_k)^{-1} \right)}.$$

Proof. Let us denote $\mathbf{K}^{(\mathrm{g})} = \mathbf{A}\mathbf{C}\mathbf{A}^\top$ the covariance matrix of $\mathbf{Z}$. First, one has $\mathbf{U}_k^\top \mathbf{K}^{(\mathrm{g})} \mathbf{U}_k = \mathbf{\Sigma}_k \mathbf{V}_k^\top \mathbf{C}\mathbf{V}_k \mathbf{\Sigma}_k$ which is non-singular since both $\mathbf{V}_k^\top \mathbf{C}\mathbf{V}_k$ and $\mathbf{\Sigma}_k$ are non-singular by assumption, so Theorem 3.1 can be applied. Using the submultiplicativity of the Frobenius norm along with (9), we have

$$\left\| \overline{\mathbf{\Sigma}}_k \tan(\mathbf{V}_k, \mathbf{C}\mathbf{V}_k) \right\|_F^2 = \left\| \overline{\mathbf{\Sigma}}_k \overline{\mathbf{V}}_k^\top \mathbf{C}\mathbf{V}_k (\mathbf{V}_k^\top \mathbf{C}\mathbf{V}_k)^{-1} \right\|_F^2$$

$$
\begin{aligned}
&\leq \left\| \overline{\mathbf{\Sigma}}_k^\top \overline{\mathbf{\Sigma}}_k \, \overline{\mathbf{V}}_k^\top \mathbf{C}^{\frac{1}{2}} \right\|_F^2 \left\| \mathbf{C}^{\frac{1}{2}} \mathbf{V}_k (\mathbf{V}_k^\top \mathbf{C} \mathbf{V}_k)^{-1} \right\|_F^2 \\
&= \mathsf{tr}\left(\overline{\mathbf{\Sigma}}_k^\top \overline{\mathbf{\Sigma}}_k \overline{\mathbf{V}}_k^\top \mathbf{C} \overline{\mathbf{V}}_k \right) \cdot \mathsf{tr}\left((\mathbf{V}_k^\top \mathbf{C} \mathbf{V}_k)^{-1} \right) \\
&= k \left\| \overline{\mathbf{\Sigma}}_k \right\|_F^2 \, \frac{\beta_k(\mathbf{C})}{\gamma_k(\mathbf{C})}.
\end{aligned}
$$

Thus,

$$
\tau_k\left(\mathbf{K}^{(\mathrm{g})}\right) = \frac{\left\| \overline{\mathbf{\Sigma}}_k \tan(\mathbf{V}_k, \, \mathbf{C}\mathbf{V}_k) \right\|_F}{\left\| \overline{\mathbf{\Sigma}}_k \right\|_F} \leq \sqrt{k} \sqrt{\frac{\beta_k(\mathbf{C})}{\gamma_k(\mathbf{C})}}
$$

Then, let us use the following fact which will be proved in Appendix A, namely

$$
\left\| [\mathbf{I}_m - \pi([\mathbf{K}^{(\mathrm{g})}]^{\frac{1}{2}} \mathbf{U}_k)][\mathbf{K}^{(\mathrm{g})}]^{\frac{1}{2}} \right\|_F^2 = \mathsf{tr}\left(\overline{\mathbf{U}}_k^\top \mathbf{K}^{(\mathrm{g})} \overline{\mathbf{U}}_k - \overline{\mathbf{U}}_k^\top \mathbf{K}^{(\mathrm{g})} \mathbf{U}_k \mathbf{K}_k^{-1} \mathbf{U}_k^\top \mathbf{K}^{(\mathrm{g})} \overline{\mathbf{U}}_k \right).
$$

Since $\overline{\mathbf{U}}_k^\top \mathbf{K}^{(\mathrm{g})} \mathbf{U}_k \mathbf{K}_k^{-1} \mathbf{U}_k^\top \mathbf{K}^{(\mathrm{g})} \overline{\mathbf{U}}_k$ is symmetric positive semi-definite, one has,

$$
\begin{aligned}
\left\| [\mathbf{I}_m - \pi([\mathbf{K}^{(\mathrm{g})}]^{\frac{1}{2}} \mathbf{U}_k)][\mathbf{K}^{(\mathrm{g})}]^{\frac{1}{2}} \right\|_F^2 &\leq \mathsf{tr}\left(\overline{\mathbf{U}}_k^\top \mathbf{K}^{(\mathrm{g})} \overline{\mathbf{U}}_k \right) \\
&= \mathsf{tr}\left(\overline{\mathbf{\Sigma}}_k \overline{\mathbf{V}}_k^\top \mathbf{C} \overline{\mathbf{V}}_k \overline{\mathbf{\Sigma}}_k^\top \right) \\
&= \mathsf{tr}\left(\overline{\mathbf{\Sigma}}_k^\top \overline{\mathbf{\Sigma}}_k \overline{\mathbf{V}}_k^\top \mathbf{C} \overline{\mathbf{V}}_k \right).
\end{aligned}
$$

Altogether, this yields

$$
\rho_k\left(\mathbf{K}^{(\mathrm{g})}\right) \leq \sqrt{k} \sqrt{\frac{\beta_k(\mathbf{C})}{\gamma_k(\mathbf{C})}}.
$$

Plugging these inequalities in Theorem 3.1 and using that $\sqrt{1+x} \leq 1 + \sqrt{x}$ for all $x \geq 0$ end the proof. □

Corollary 3.4 provides bounds similar to the ones in [1]. There is nonetheless a difference that is worth pointing out. The probability and expectation bounds in [1], stated respectively in Theorem 2 and Proposition 6, predict a deviation from the optimal characterized by a form of $O(k)$ where k represents the target rank. On the contrary, Corollary 3.4 exhibits bounds for the low-rank approximation error with deviation from the optimal of the form $O(\sqrt{k})$. This shows that the influence of the target rank k is not as penalizing as the bounds in [1] would suggest.

3.1.3 K as a low-rank matrix update

We focus on exploring the potential benefits of incorporating an available approximate SVD within the covariance matrix to enhance the achievable bounds.

Let $\widehat{\mathbf{U}}_k \in \mathbb{R}^{m \times k}$ be an approximation of $\mathbf{U}_k$ satisfying $\widehat{\mathbf{U}}_k^\top \widehat{\mathbf{U}}_k = \mathbf{I}_k$, and $\widehat{\mathbf{\Sigma}}_k = \mathsf{diag}(\hat{\sigma}_1, \ldots, \hat{\sigma}_k)$ an approximation of $\mathbf{\Sigma}_k$. In this case, Theorem 3.1 suggests to consider a covariance matrix of the form

$$
\mathbf{K}_{\alpha,\beta} = \alpha \widehat{\mathbf{U}}_k \widehat{\mathbf{\Sigma}}_k^2 \widehat{\mathbf{U}}_k^\top + \beta(\mathbf{I}_n - \widehat{\mathbf{U}}_k \widehat{\mathbf{U}}_k^\top),
$$

with $\alpha > 0$ and $0 \leq \beta \leq \alpha \, \hat{\sigma}_k^2$. Therefore, the columns of $\widehat{\mathbf{U}}_k$ are eigenvectors of $\mathbf{K}_{\alpha,\beta}$ associated with eigenvalues $\alpha\hat{\sigma}_1^2, \ldots, \alpha\hat{\sigma}_k^2$, and the orthogonal of $\mathcal{R}(\widehat{\mathbf{U}}_k)$ is an eigenspace associated with the eigenvalue β. Note that the bounds in Theorem 3.1 are invariant under scalar multiplication of the covariance matrix. Therefore, one can consider either $\mathbf{K}_{\alpha/\beta,1}$ or $\mathbf{K}_{1,\beta/\alpha}$, depending on the most convenient form.

The parameters α and β play the roles of relative weights, allowing to balance the confidence in the approximation $\widehat{\mathbf{U}}_k$. If $\beta = 0$, then $\mathbf{K}_{\alpha,0}$ is a rank k matrix such that $\mathcal{R}(\mathbf{K}_{\alpha,0}\mathbf{U}_k) = \mathcal{R}(\widehat{\mathbf{U}}_k)$, which trivially yields

$$\tau_k(\mathbf{K}_{\alpha,0}) = \frac{\left\|\mathbf{tan}(\mathbf{U}_k, \widehat{\mathbf{U}}_k)\boldsymbol{\Sigma}_k\right\|_F}{\left\|\overline{\boldsymbol{\Sigma}}_k\right\|_F} \quad \text{and} \quad \rho_k(\mathbf{K}_{\alpha,0}) = 0.$$

In this configuration, the low-rank approximation error is only governed by how $\widehat{\mathbf{U}}_k$ accurately approximates $\mathbf{U}_k$. We note that using $\mathbf{K} = \mathbf{K}_{\alpha,0}$ in Algorithm 1 yields a deterministic algorithm since $\mathsf{rank}(\mathbf{K}_{\alpha,0}) = k \leq \ell$ so that $\mathbf{Z} \in \mathbb{R}^{m\times\ell}$ drawn with covariance matrix $\mathbf{K}_{\alpha,0}$ satisfies $\mathcal{R}(\mathbf{Z}) = \mathcal{R}(\mathbf{K}_{\alpha,0}) = \mathcal{R}(\widehat{\mathbf{U}}_k)$ with probability one.

Considering $\beta > 0$, the columns of $\mathbf{Z}$ will have components both in $\mathcal{R}(\widehat{\mathbf{U}}_k)$ and its orthogonal complement. Since $\mathbf{U}_k$ is approximated by $\widehat{\mathbf{U}}_k$, it also has non-zero intersection both with $\mathcal{R}(\widehat{\mathbf{U}}_k)$ and its orthogonal. Consequently, if $\widehat{\mathbf{U}}_k$ is a poor approximation, then the orthogonal of $\mathcal{R}(\widehat{\mathbf{U}}_k)$ still contains relevant information regarding $\mathbf{U}_k$, and β should be taken as large as possible to ensure the random exploration of the remaining space. Conversely, if $\widehat{\mathbf{U}}_k$ is an accurate approximation of $\mathbf{U}_k$, then one should choose a small (non-zero) value for β so that the columns of $\mathbf{Z}$ are mostly aligned with $\mathcal{R}(\mathbf{U}_k)$.

Alternatively, considering $\widehat{\mathbf{V}}_k \in \mathbb{R}^{n\times k}$, an approximation of $\mathbf{V}_k$, one can also define the matrix,

$$\mathbf{C}_{\alpha,\beta} = \alpha\widehat{\mathbf{V}}_k\widehat{\boldsymbol{\Sigma}}_k^2\widehat{\mathbf{V}}_k^\top + \beta(\mathbf{I}_n - \widehat{\mathbf{V}}_k\widehat{\mathbf{V}}_k^\top),$$

and consider the covariance matrix $\mathbf{K}_{\alpha,\beta} = \mathbf{A}\mathbf{C}_{\alpha,\beta}\mathbf{A}^\top$. The same analysis can be performed, replacing $\mathbf{U}_k$ by $\mathbf{V}_k$ in the arguments.

4 Numerical experiments

In recent years, randomized SVD has a received considerable attention from the field of data assimilation (DA) [4, 3, 6]. DA is an intensive field of research, with various applications in Earth system modelling, for instance, numerical weather forecast, and oceanography, to name a few. DA is a methodology used to estimate the physical state of a dynamical system by using different sources of information taking into consideration their uncertainties. For our numerical experiments, we use a three-dimensional variational data assimilation formulation. In this approach, we solve a weighted nonlinear least-squares problem for a fixed time:

$$\min_{\mathbf{x}\in\mathbb{R}^n} f(\mathbf{x}) = \frac{1}{2}\left\|\mathbf{y} - \mathcal{H}(\mathbf{x})\right\|_{\mathbf{R}^{-1}}^2 + \frac{1}{2}\|\mathbf{x} - \mathbf{x}_b\|_{\mathbf{B}^{-1}}^2, \tag{12}$$

where $\mathbf{x} \in \mathbb{R}^n$ is the state of a dynamical system, for instance temperature, $\mathbf{y} \in \mathbb{R}^m$ is a vector consisting of observations, and $\mathbf{x}_b \in \mathbb{R}^n$ is a priori information. The operator $\mathcal{H} : \mathbb{R}^n \to \mathbb{R}^m$ is the nonlinear observation operator mapping the state vector from the model space to the observation space. $\mathbf{B} \in \mathbb{R}^{n\times n}$ and $\mathbf{R} \in \mathbb{R}^{m\times m}$ represent the error covariance matrices of a priori information and observations respectively.

A common approach for solving (12) is the truncated Gauss-Newton method [10]. This method relies on the linearization of the nonlinear observation operator around the current iterate $\mathbf{x}_j$, which results in a weighted linear least-squares sub-problem whose solution can be obtained by solving a preconditioned linear system:

$$(\mathbf{I}_n + \mathbf{L}\mathbf{H}_j^\top\mathbf{R}^{-1}\mathbf{H}_j\mathbf{L})\delta\mathbf{v}_j = -\mathbf{L}\mathbf{b}_j \tag{13}$$

where $\mathbf{L} \in \mathbb{R}^{n\times n}$ is a symmetric square-root factorization of $\mathbf{B}$ used as a preconditioner, $\mathbf{H}_j \in \mathbb{R}^{m\times n}$ is the Jacobian matrix of the observation operator at the current iterate, $\delta\mathbf{v}_j = \mathbf{L}\delta\mathbf{x}_j$ is the increment (search direction) in the preconditioned space, and $\mathbf{b}_j$ is the gradient of the nonlinear cost function (12)

calculated at the current iterate. The linear system (13) consisting of the symmetric positive definite matrix,

$$\mathbf{A}^{(j)} = \mathbf{I}_n + \mathbf{L}\mathbf{H}_j^\top \mathbf{R}^{-1}\mathbf{H}_j\mathbf{L} \tag{14}$$

can then be solved by using an iterative method such as the preconditioned conjugate gradients method (PCG). While performing PCG for costly matrix-vector products with $\mathbf{A}$ like in numerical weather forecast, it is crucial to use an efficient second-level preconditioner like limited memory preconditioner [11] (LMP) to accelerate the convergence within the Gauss-Newton process. As an alternative to LMPs which obtain spectral information based on the matrix $\mathbf{A}^{(j-1)}$, one can obtain a second-level preconditioner based on the spectral information of the current matrix $\mathbf{A}^{(j)}$ by using randomized SVD [27].

In our numerical experiments, we then aim to approximate the eigenspectrum of the symmetric positive definite matrix $\mathbf{A}^{(k)}$ by using randomized SVD based on non-standard Gaussian sample matrices for the first Gauss-Newton iteration $k = 1$. For simplicity, we will use $\mathbf{A}$ referring to $\mathbf{A}^{(1)}$. The error covariance matrix $\mathbf{B}$ is obtained from the discretization of the diffusion operator on a regular grid, $\mathcal{H}$ is chosen as a selection operator ($\mathcal{H}$ is a linear operator), and $\mathbf{R} = \sigma_{\mathbf{R}}\,\mathbf{I}_m$ with $\sigma_{\mathbf{R}} > 0$ being the standard deviation. The matrix of interest defined in (14) is a rank m update of the identity. Consequently, the number of observations m has a critical influence on the eigenvalue distribution, i.e. for $n > m$ there are $n - m$ eigenvalues equal to 1. The eigenvalues of $\mathbf{A}$ also have a strongly decaying spectrum as shown in Figure 1 which is a property in favour of the randomized SVD. We choose $n = 1000$, and we consider two different scenarios, with a different number of observations: `LowObs` for $m = 200$, and `HighObs` for $m = 500$.

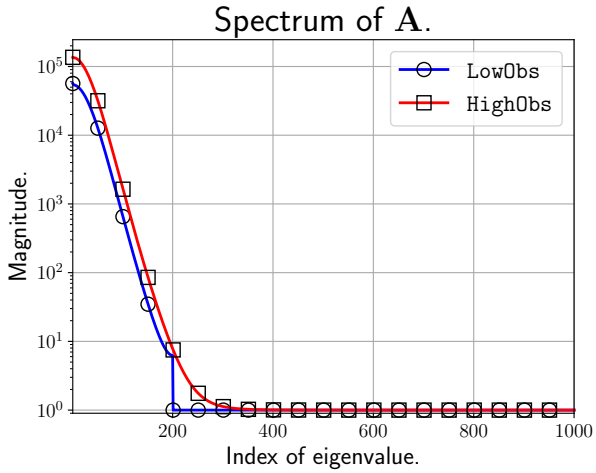

Figure 1: Eigenvalue distribution of $\mathbf{A}$ in the `LowObs` and `HighObs` scenarios.

4.1 Accuracy of the error bounds

In this section, we analyze the accuracy of the bounds in Theorem 3.1 for a particular choice of the covariance matrix $\mathbf{K}$. Note that the dominant eigenvectors of $\mathbf{A}$ correspond to the dominant eigenvectors of $\mathbf{W}$, expressed as $\mathbf{W} = \mathbf{L}\mathbf{H}_j^\top \mathbf{R}^{-1}\mathbf{H}_j\mathbf{L}$, with its rank equal to $m < n$. Using the SVD, $\mathbf{L}\mathbf{H}_j^\top \mathbf{R}^{-1/2} = \mathbf{S}_m \mathbf{\Lambda}_m^{1/2} \mathbf{F}_m^T$, where $\mathbf{S}_m \in \mathbb{R}^{n\times m}$, $\mathbf{\Lambda}_m \in \mathbb{R}^{m\times m}$ and $\mathbf{F}_m \in \mathbb{R}^{n\times m}$, the eigenspectrum of $\mathbf{W}$ is then represented as $\mathbf{W} = \mathbf{S}_m \mathbf{\Lambda}_m \mathbf{S}_m^\top$. As a result, $\mathbf{S}_m$ represents the eigenvectors of $\mathbf{A}$ corresponding to the m largest eigenvalues. Leveraging Theorem 3.1's results, we can improve the performance of Algorithm 1 by choosing the covariance matrix $\mathbf{K} = \mathbf{A}\mathbf{S}_k\mathbf{S}_k^T\mathbf{A}$, with $k \leq m$ being the target rank. Since $\mathbf{S}_k$ is not readily available and is indeed the quantity of interest, we select $\mathbf{K} = \mathbf{A}\mathbf{L}^2\mathbf{A} = \mathbf{A}\mathbf{B}\mathbf{A}$. Note that $\mathbf{W}$ is in the image of the matrix $\mathbf{L}$ (this remark is further investigated in Section 4.2). When $\mathbf{H}$ is an identity matrix, i.e. $m = n$, $\mathcal{R}(\mathbf{L}) = \mathcal{R}(\mathbf{S}_m)$. In the numerical experiments, $\mathbf{H}$ is chosen as a

selection operator. In this case, $\mathcal{R}(\mathbf{S}_m) = \mathcal{R}(\mathbf{L}_m)$ where $\mathbf{L}_m \in \mathbb{R}^{m\times n}$ is a submatrix of $\mathbf{L}$ with its m columns selected through the matrix $\mathbf{H}$. Therefore, the performance of Algorithm 1 will be improved with the choice of $\mathbf{K} = \mathbf{ABA}$, whenever the subspace spanned by the k-th leading eigenvectors of $\mathbf{L}$ gets closer to the range of $\mathbf{S}_k$.

We apply Algorithm 1 with $\mathbf{Z} = \mathbf{ALG}$, $\mathbf{G}$ being a matrix whose columns are independently drawn from a standard multivariate Gaussian distribution. We compare the bounds to the empirical mean of the randomized low-rank approximation error computed out of 20 independent runs of Algorithm 1. For the comparison, we plot the quantity

$$\frac{\left\|[\mathbf{I}_m - \pi(\mathbf{Z})]\mathbf{A}\right\|_F}{\left\|\overline{\mathbf{\Sigma}}_k\right\|_F} - 1,$$

which is as close as zero as the low-rank approximation error is close to the optimal value $\left\|\overline{\mathbf{\Sigma}}_k\right\|_F$. The coefficients u and t in Theorem 3.1 are obtained so that the probability of failure is strictly less than 10^{-3}.

Figure 2 shows the low-rank approximation error with respect to the target rank k, for a fixed oversampling parameter $p = \ell - k = 10$. Here, the bounds in expectation follow the empirical mean error trend, by a multiplicative factor of approximately 10. On Figure 2a however, the bound diverges after $k = 200$ $(= m)$. This behaviour illustrates that the use of $\mathbf{B}$ in the covariance matrix $\mathbf{K}$ may positively influence the low-rank approximation error as long as the target rank is smaller than $\mathsf{rank}(\mathbf{L}\mathbf{H}_j^\top \mathbf{R}^{-1}\mathbf{H}_j\mathbf{L}) = m$. When $k > m$, the algorithm is expected to produce approximations of eigenvectors associated with the cluster of eigenvalues at 1, which are no longer in the image of $\mathbf{L}$, thus explaining why the bound diverges. This problem is not noticeable in Figure 2b since m is equal to 500 in this scenario and k is at most 300. However, it is worth noting that the bounds diverge while the empirical error still decreases, highlighting that the bounds may overestimate the error in certain circumstances.

These considerations also apply to the probability bounds whose trends follow the expectation bounds. Given the small empirical standard deviation observed (grey halo around the empirical error), it is clear that the probability bounds severely overestimate the error by a factor of at least 10^2.

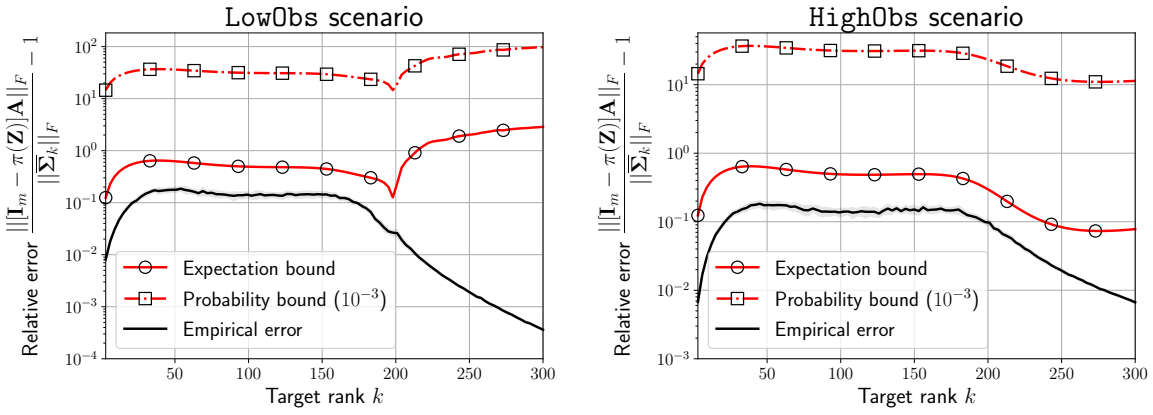

a `LowObs`: $m = 200$. **b** `HighObs`: $m = 500$.

Figure 2: Bounds for the low-rank approximation error versus the target rank k. The empirical data have been computed with $p = 10$.

Figure 3 shows the behaviour of the bounds with respect to the oversampling parameter $p = \ell - k$ for a fixed value of the target rank k $(k = 20)$. We observe no significant differences between the `LowObs` and `HighObs` scenarios. The bounds in expectation and in probability both predict a decrease rate of the error as $O(1/\sqrt{p})$, but visibly fail to capture the true decrease rate.

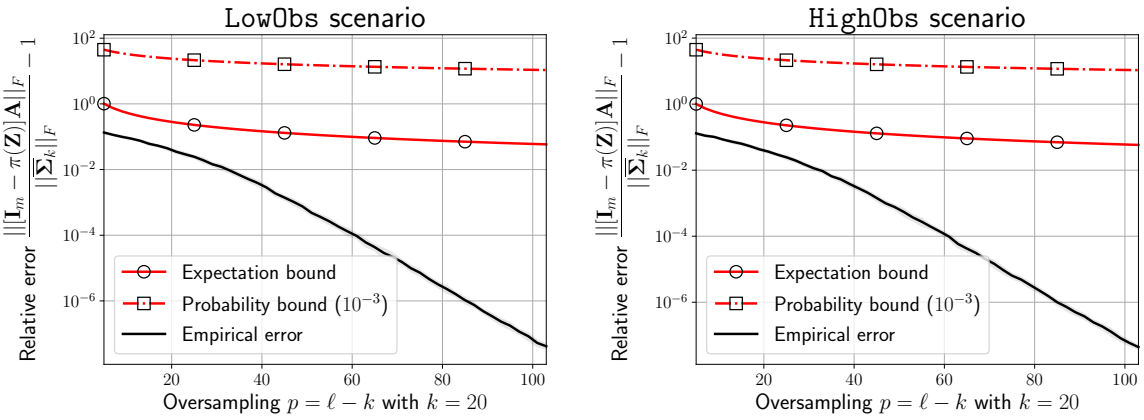

a `LowObs`: $m = 200$.

b `HighObs`: $m = 500$.

Figure 3: Bounds for the low-rank approximation error versus the oversampling parameter $p = \ell - k$. The empirical data have been computed with $k = 20$.

Finally, we propose in Figure 4 a comparison of the bounds from [1] in the `HighObs` scenario. This is meant to illustrate the discussion in Section 3.1.2, where we proposed alternative bounds to analyse the Generalized RSVD algorithm. We observe here that the bounds derived from Theorem 3.1 actually refine the ones from [1]. The impact of the extra $\sqrt{k}$ factor can be seen in Figure 4a, where the refinement brought by Theorem 3.1 is greater for large values of k: e.g. in expectation, the improvement is approximately equal to a factor of 10 for small values of k, and up to almost 100 for larger values of k. The decrease rate of the error with respect to the oversampling parameter, shown in Figure 4b, is also impacted.

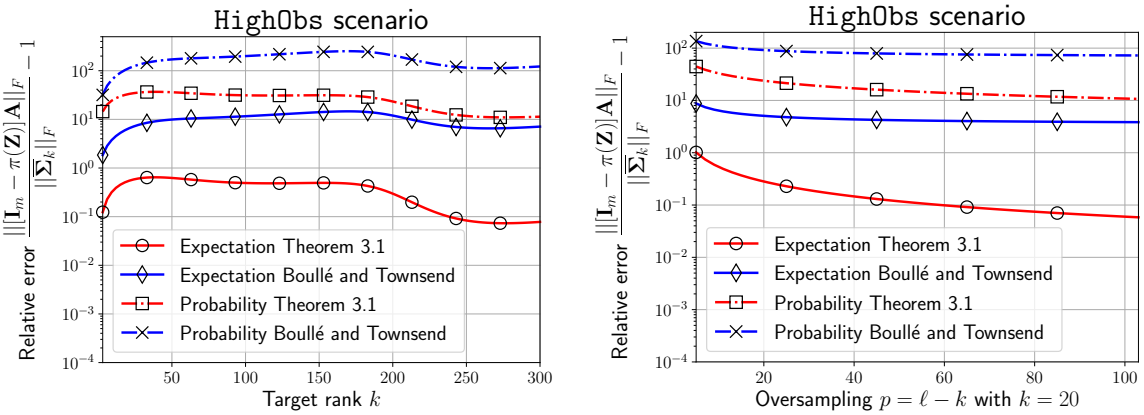

a Versus the target rank k.

b Versus the oversampling $p = \ell - k$.

Figure 4: Comparison of the bounds with [1] in the `HighObs` scenario.

4.2 Approximate subspace iteration

In this section, we consider $\mathbf{Z} = \mathbf{A}\mathbf{\Omega}$ in Algorithm 1 and we denote by $\mathbf{C}$ the covariance matrix of $\mathbf{\Omega}$ which corresponds to $\mathbf{K} = \mathbf{ACA}$. Our objective is to illustrate how Algorithm 1 behaves when $\mathbf{A}$ is replaced by an approximation of $\mathbf{A}_k$ in the power iteration. We note here that the power scheme applied to square matrices has the form of

$$\mathbf{Z} = \mathbf{A}^q\mathbf{G} \tag{15}$$

where $\mathbf{G}$ is a matrix whose columns are independently sampled from a standard multivariate Gaussian distribution. For our numerical experiments, we consider the following cases:

$\mathbf{C} = \mathbf{I}_n$: This case corresponds to the exact power iteration with $q = 1$ in Equation (15).

$\mathbf{C} = \mathbf{A}^2$: This case corresponds to the exact power iteration with $q = 2$ in Equation (15).

$\mathbf{C} = \mathbf{B}$: This case can be interpreted as an approximate power iteration with $q = 3/2$, where $\mathbf{A}\mathbf{A}^{1/2}$ is approximated by $\mathbf{A}\mathbf{L}$. Note that we assume $\mathcal{R}(\mathbf{L}) \approx \mathcal{R}(\mathbf{A}_k^{1/2})$.

$\mathbf{C} = \mathbf{B}^2$: This case can be interpreted as an approximate power iteration with $q = 2$, where $\mathbf{A}^2$ is approximated by $\mathbf{A}\mathbf{B}$.

The first case will serve as a reference, while the second one will serve as an ideal case. For the last two cases, we rely on the approximation that $\mathcal{R}(\mathbf{B})$ is an accurate approximation of $\mathcal{R}(\mathbf{L}\mathbf{H}_j^\top \mathbf{R}^{-1}\mathbf{H}_j\mathbf{L})$. Ideally, if $\mathbf{B}$ is an accurate enough approximation of $\mathbf{A}_k$, and a fortiori of $\mathbf{A}$, then the resulting low-rank approximation error should be close to the one obtained with $q = 2$ in (15). In theory, it may even yield unproved performance since using $\mathbf{A}$ is sub-optimal compared to $\mathbf{A}_k$.

Figure 5 presents the numerical results for different cases. A first important observation is that for both scenarios, using $\mathbf{L}$ or $\mathbf{B}$ improves over the reference case $\mathbf{C} = \mathbf{I}_n$, implying that corresponding covariance matrices are indeed carrying dominant eigeninformation of $\mathbf{A}$. Then, one also observes that applying Algorithm 1 with $\mathbf{C} = \mathbf{B}^2$ yields results very similar to $\mathbf{C} = \mathbf{A}^2$ and improved low-rank approximation for $k \geq 150$. These results suggest to use an approximation to $\mathbf{A}_k$ instead of $\mathbf{A}$.

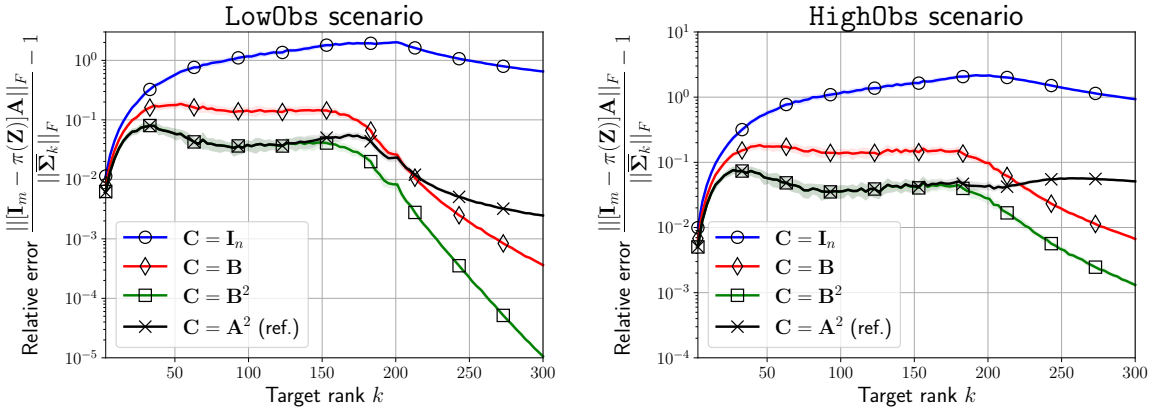

a `LowObs`: $m = 200$. **b** `HighObs`: $m = 500$.

Figure 5: Empirical low-rank approximation error with respect to the target rank k. The empirical data have been computed with $p = 10$ and for different covariance matrices.

This is also what we observe when looking at the two parameters $\tau_k(\mathbf{K})$ and $\rho_k(\mathbf{K})$ shown in Figures 6 and 7, respectively. Let us recall that $\tau_k(\mathbf{K})$ measures the stability of $\mathcal{R}(\mathbf{A}_k)$ under the action of $\mathbf{K}$ and that $\mathcal{R}(\mathbf{A}_k) = \mathcal{R}(\mathbf{S}_k)$ for $k \leq m$ in our case. From this, it is clear that $\tau_k(\mathbf{K})$ is zero for the first two cases and is thus not shown in Figure 6. For the remaining two cases, we note that $\tau_k(\mathbf{K})$ is relatively small compared to $\rho_k(\mathbf{K})$ except in the `LowObs` scenario when $k \geq m$. In this situation, the equality $\mathcal{R}(\mathbf{A}_k) = \mathcal{R}(\mathbf{S}_k)$ no longer holds, so the practical choice of using $\mathbf{L}$ or $\mathbf{B} = \mathbf{L}^2$ instead of $\mathbf{S}_k\mathbf{S}_k^\top$ becomes less relevant. Nevertheless, this limitation should not have any practical limitations since in concrete applications, the target rank k will be chosen much smaller than m the number of observations. On the other side, $\rho_k(\mathbf{K})$ is related to how the dominant eigenspectrum of $\mathbf{K}$ approximates the one of $\mathbf{A}$. Likewise, for $\tau_k(\mathbf{K})$, we observe a significant increase after $k = m$ in the `LowObs` scenario. The explanation is similar, since after this point, the target low-rank approximation $\mathbf{A}_k$ contains $m - k$ eigenmodes associated with the eigenvalue cluster at 1, while the chosen covariance matrix is only meant to improve the approximation of the first m dominant eigenmodes.

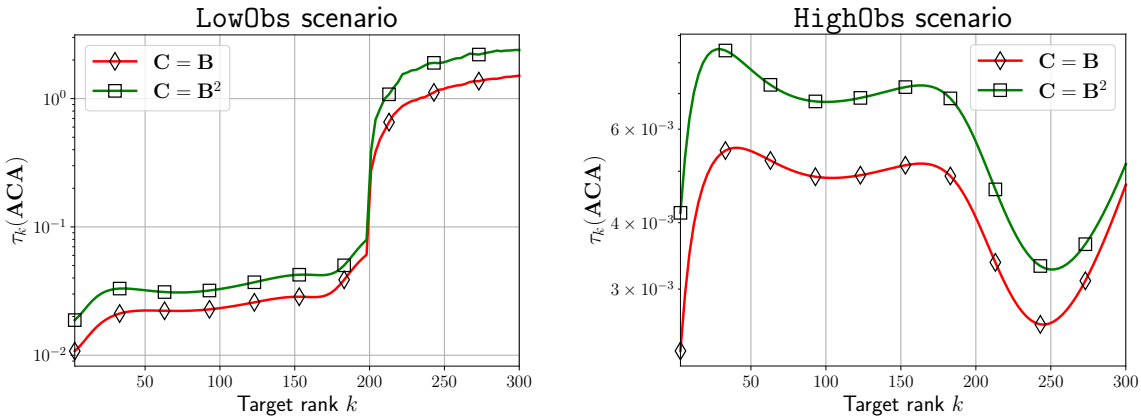

a `LowObs`: $m = 200$. **b** `HighObs`: $m = 500$.

Figure 6: Parameter $\tau_k(\mathbf{ACA})$ versus the target rank k for the different covariance matrices $\mathbf{C}$ under study.

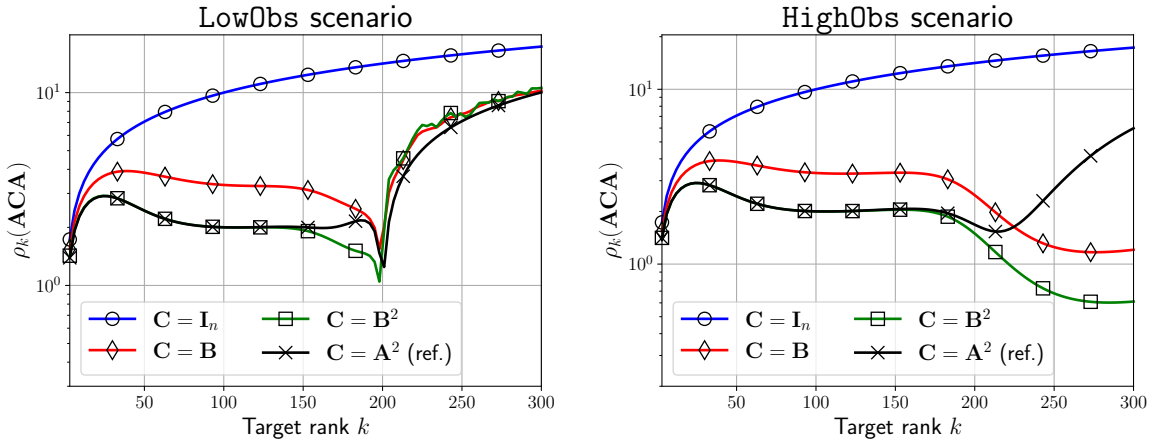

a `LowObs`: $m = 200$. **b** `HighObs`: $m = 500$.

Figure 7: Parameter $\rho_k(\mathbf{ACA})$ versus the target rank k for the different covariance matrices $\mathbf{C}$ under study.

4.3 Using available approximations

In this section, we study the case where approximations $\widehat{\mathbf{V}}_k$ and $\widehat{\mathbf{\Sigma}}_k$ of $\mathbf{V}_k$ and $\mathbf{\Sigma}_k$ respectively are already available. The objective here is to analyze the low-rank approximation error when using a covariance matrix of the specific form

$$\mathbf{C}_{\alpha,\beta} = \alpha \widehat{\mathbf{V}}_k \widehat{\mathbf{\Sigma}}_k \widehat{\mathbf{V}}_k^\top + \beta(\mathbf{I}_n - \widehat{\mathbf{V}}_k \widehat{\mathbf{V}}_k^\top).$$

For this numerical experiment, the approximations are obtained by applying Algorithm 1 to $\mathbf{A}$, with $\mathbf{K} = \mathbf{A}\mathbf{A}^\top$ ($\mathbf{C} = \mathbf{I}_n$) and for varying values of target rank k together with $p = \ell - k = 10$. Information on the quality of the resulting approximation is shown in Figure 5. We then use the resulting matrices $\widehat{\mathbf{V}}_k$ and $\widehat{\mathbf{\Sigma}}_k$ to construct the matrix $\mathbf{C}_{\alpha,\beta}$. We study the following cases:

- Different scaling: $\alpha = 0.01$, 1 and 100 for a fixed value $\beta = 1$. The scaling is expected to influence the performance as it allows to weight differently $\mathcal{R}(\widehat{\mathbf{V}}_k)$ and its orthogonal complement.
- For $\alpha = \beta = 1$, we also consider the variant $\mathbf{C}_\mathbf{L} = \mathbf{L}\,\mathbf{C}_{1,1}\,\mathbf{L}$. Regarding the results from the previous subsection, using $\mathbf{L}$ is expected to improve the performance.
- We also consider the case $\alpha = 1$ and $\beta = 0$. As discussed in Section 3.1.3, this yields a deterministic algorithm which simply reduces to perform a standard power iteration step on top of

$\widehat{\mathbf{V}}_k$. This choice will serve as a reference, and our objective is to study whether maintaining randomization (i.e. $\beta > 0$) can yield better performance. This can be interpreted as the limit of $\mathbf{C}_{\alpha,1}$ when α increases.

- Finally, we consider the case $\mathbf{C} = \mathbf{I}_n$ corresponding to a single step of randomized power iteration. This base case is here to illustrate how using existing information actually improves performance.

The obtained results for the empirical low-rank approximation error are presented in Figure 8 and have been obtained from 20 independent applications of Algorithm 1.

Regardless of the scaling, for small values of k (roughly $k \leq 100$), the resulting low-rank approximation error is almost identical and also improves over the reference case. For larger k, it seems that the larger α is, the better the performance is. This suggests that in this situation, the approximation $\widehat{\mathbf{V}}_k$ used to construct the covariance matrix is accurate enough to be trusted. Nevertheless, there is a trade-off to find since taking larger α would asymptotically yield the same results as $\mathbf{C}_{!,0}$ (the reference case). Finally, the most noticeable case is $\mathbf{C}_{\mathbf{L}}$, which outperforms the reference case by almost a factor of 10 (given the target k remains small enough). As expected, using $\mathbf{L}$ allows us to improve the performance because $\mathbf{L}$ also carries information. Of course, this also yields a more expensive algorithm since it would require an application of $\mathbf{L}$ to a $n \times \ell$ matrix, while the other scaled cases only involve thin matrices.

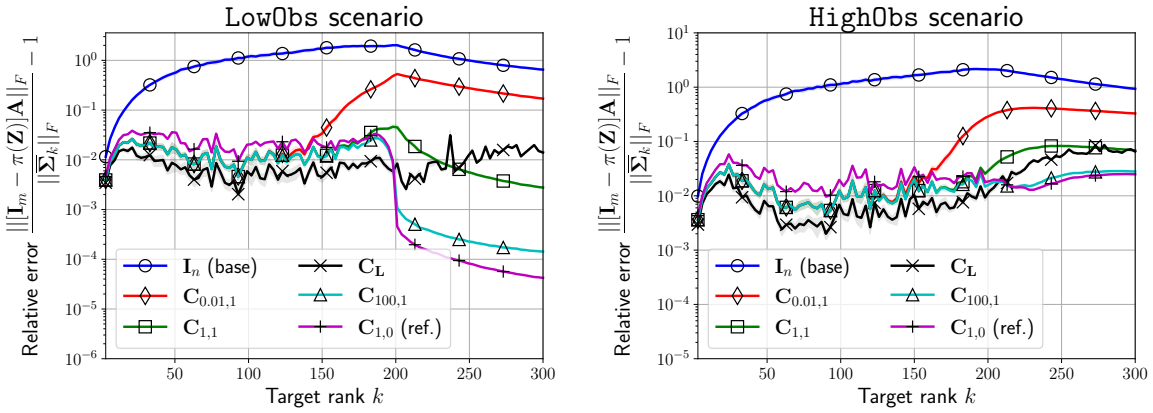

a `LowObs`: $m = 200$. **b** `HighObs`: $m = 500$.

Figure 8: Empirical low-rank approximation error with respect to the target rank k. The empirical data have been computed with $p = 10$ and for different covariance matrices.

5 Conclusions

We have proposed a general theoretical framework for the analysis of the low-rank approximation error where we have relaxed the assumptions on the covariance matrix. The obtained bounds, in expectation and in probability, have clear interpretations. We have illustrated the generality of our result by applying Theorem 3.1 to two known algorithms, and the obtained bounds have comparable accuracy as the ones derived on purpose in the literature. Finally, we have illustrated our result on a data assimilation problem. First, we have studied the accuracy of the bounds. Then, we have proposed numerical experiments in two different situations, using either an approximation of $\mathbf{A}$ or an available approximation of $\mathbf{V}_k$. In both cases, we have identified covariance matrices that enable us to improve the overall performance. This highlights that using the structure of the problem to design the covariance matrix does improve the performance of randomized low-rank approximation methods.

A Proof of Theorem 3.1

In this appendix, we prove Theorem 3.1. We begin with a lemma which is a generalization of the analysis in terms of principal angles proposed in [24].

Lemma A.1. Let $\mathbf{A} \in \mathbb{R}^{m\times n}$ be an arbitrary matrix and $\mathbf{Z} \in \mathbb{R}^{n\times \ell}$ a full column rank matrix such that $\ell \leq \mathsf{rank}(\mathbf{A})$. Let us denote $\mathbf{Z}_k = \mathbf{U}_k^\top \mathbf{Z} \in \mathbb{R}^{k\times \ell}$ and $\overline{\mathbf{Z}}_k = \overline{\mathbf{U}}_k^\top \mathbf{Z} \in \mathbb{R}^{(n-k)\times \ell}$.

For any integer $k \leq \ell$, if $\mathsf{rank}(\mathbf{Z}_k) = k$, then one has

$$\mathbf{U}_k^\top [\mathbf{I}_n - \pi(\mathbf{Z})]\mathbf{U}_k \preccurlyeq \mathbf{T}_k^\top \mathbf{T}_k, \tag{16}$$

where $\mathbf{T}_k = \overline{\mathbf{Z}}_k \mathbf{Z}_k^\dagger \in \mathbb{R}^{(n-k)\times k}$.

Proof. By assumption, $\mathbf{Z}_k$ has full row rank and thus has a right multiplicative inverse $\mathbf{Z}_k^\dagger$. Hence $\mathbf{Y}_k = \mathbf{Z}\mathbf{Z}_k^\dagger$ satisfies the two relations

$$\mathbf{U}_k^\top \mathbf{Y}_k = \mathbf{I}_k \quad \text{and} \quad \overline{\mathbf{U}}_k^\top \mathbf{Y}_k = \mathbf{T}_k.$$

Moreover, we have $\mathcal{R}(\mathbf{Y}_k) \subset \mathcal{R}(\mathbf{Z})$ which implies that $\mathbf{I}_n - \pi(\mathbf{Z}) \preccurlyeq \mathbf{I}_n - \pi(\mathbf{Y}_k)$ according to [14, Proposition 8.5]. By using the conjugation rule (6) and the identity $\mathbf{U}_k\mathbf{U}_k^\top + \overline{\mathbf{U}}_k\overline{\mathbf{U}}_k^\top = \mathbf{I}_n$, we obtain

$$\begin{aligned}
\mathbf{U}_k^\top [\mathbf{I}_n - \pi(\mathbf{Z})]\mathbf{U}_k \preccurlyeq \mathbf{U}_k^\top [\mathbf{I}_n - \pi(\mathbf{Y}_k)]\mathbf{U}_k &= \mathbf{U}_k^\top \left(\mathbf{I}_n - \mathbf{Y}_k(\mathbf{Y}_k^\top \mathbf{Y}_k)^{-1}\mathbf{Y}_k^\top\right)\mathbf{U}_k, \\
&= \mathbf{I}_k - \mathbf{U}_k^\top \mathbf{Y}_k(\mathbf{Y}_k^\top \mathbf{Y}_k)^{-1}\mathbf{Y}_k^\top \mathbf{U}_k, \\
&= \mathbf{I}_k - (\mathbf{Y}_k^\top \mathbf{Y}_k)^{-1}, \\
&= \mathbf{I}_k - \left(\mathbf{Y}_k^\top (\mathbf{U}_k\mathbf{U}_k^\top + \overline{\mathbf{U}}_k\overline{\mathbf{U}}_k^\top)\mathbf{Y}_k\right)^{-1}, \\
&= \mathbf{I}_k - \left(\mathbf{I}_k + \mathbf{T}_k^\top \mathbf{T}_k\right)^{-1}, \\
&= \mathbf{T}_k^\top \left(\mathbf{I}_{n-k} + \mathbf{T}_k\mathbf{T}_k^\top\right)^{-1}\mathbf{T}_k,
\end{aligned}$$

where the last equality is obtained using the Sherman-Morrison formula (8). To conclude, we observe that $\left(\mathbf{I}_{n-k} + \mathbf{T}_k\mathbf{T}_k^\top\right)^{-1} \preccurlyeq \mathbf{I}_{n-k}$, which implies, using the conjugation rule (6), that $\mathbf{T}_k^\top \left(\mathbf{I}_{n-k} + \mathbf{T}_k\mathbf{T}_k^\top\right)^{-1}\mathbf{T}_k \preccurlyeq \mathbf{T}_k^\top \mathbf{T}_k$. □

Remark. Let $\theta_1, \ldots, \theta_k$ denote the principal angles between $\mathcal{R}(\mathbf{U}_k)$ and $\mathcal{R}(\mathbf{Z})$, and $\widetilde{\theta}_1, \ldots, \widetilde{\theta}_k$ the ones between $\mathcal{R}(\mathbf{U}_k)$ and $\mathcal{R}(\mathbf{Y}_k)$. The statement of Lemma A.1 can be geometrically rephrased as

$$\sin(\theta_i)^2 \leq \sin(\widetilde{\theta}_i)^2 = \frac{\tan(\widetilde{\theta}_i)^2}{1 + \tan(\widetilde{\theta}_i)^2} \leq \tan(\widetilde{\theta}_i)^2, \quad 1 \leq i \leq k.$$

Less formally, the forthcoming analysis is based on how close $\mathcal{R}(\mathbf{U}_k)$ and $\mathcal{R}(\mathbf{Y}_k)$ are, while the truly computed error rather concerns how close $\mathcal{R}(\mathbf{Z})$ is from $\mathcal{R}(\mathbf{U}_k)$. Since the latter is of dimension ℓ, while $\mathcal{R}(\mathbf{Y}_k)$ is of dimension k, the looseness of the derived bounds will increase with ℓ.

The next result is a deterministic bound for the low-rank approximation error which generalizes [14, Proposition 9.6].

Proposition A.2 (Deterministic analysis)**.** Let $\mathbf{A} \in \mathbb{R}^{m\times n}$ be an arbitrary matrix and $\mathbf{Z} \in \mathbb{R}^{n\times \ell}$ a full column rank matrix such that $\ell \leq \mathsf{rank}(\mathbf{A})$. Let us further denote $\mathbf{Z}_k = \mathbf{U}_k^\top \mathbf{Z} \in \mathbb{R}^{k\times \ell}$ and $\overline{\mathbf{Z}}_k = \overline{\mathbf{U}}_k^\top \mathbf{Z} \in \mathbb{R}^{(n-k)\times \ell}$.

For any integer $k \leq \ell$, if $\mathsf{rank}(\mathbf{Z}_k) = k$, then one has

$$\left\| \left[\mathbf{I}_n - \pi(\mathbf{Z})\right] \mathbf{A} \right\|_F^2 \leq \left\| \overline{\mathbf{\Sigma}}_k \right\|_F^2 + \left\| \overline{\mathbf{Z}}_k \mathbf{Z}_k^\dagger \mathbf{\Sigma}_k \right\|_F^2.$$

Proof. By definition of the Frobenius norm, one has

$$\begin{aligned}\left\|[\mathbf{I}_n - \pi(\mathbf{Z})]\mathbf{A}\right\|_F^2 &= \mathsf{tr}\left([\mathbf{I}_n - \pi(\mathbf{Z})]\mathbf{A}\mathbf{A}^\top[\mathbf{I}_n - \pi(\mathbf{Z})]\right), \\ &= \mathsf{tr}\left([\mathbf{I}_n - \pi(\mathbf{Z})][\mathbf{A}_k\mathbf{A}_k^\top + \overline{\mathbf{A}}_k\overline{\mathbf{A}}_k^\top][\mathbf{I}_n - \pi(\mathbf{Z})]\right), \\ &= \mathsf{tr}\left([\mathbf{I}_n - \pi(\mathbf{Z})]\mathbf{A}_k\mathbf{A}_k^\top[\mathbf{I}_n - \pi(\mathbf{Z})]\right) \\ &\qquad + \mathsf{tr}\left([\mathbf{I}_n - \pi(\mathbf{Z})]\overline{\mathbf{A}}_k\overline{\mathbf{A}}_k^\top[\mathbf{I}_n - \pi(\mathbf{Z})]\right) \\ &= \left\|[\mathbf{I}_n - \pi(\mathbf{Z})]\mathbf{A}_k\right\|_F^2 + \left\|[\mathbf{I}_n - \pi(\mathbf{Z})]\overline{\mathbf{A}}_k\right\|_F^2.\end{aligned}$$

Using the unitary invariance of the Frobenius norm along with the fact that $\mathbf{I}_n - \pi(\mathbf{Z}) \preccurlyeq \mathbf{I}_n$, one readily gets $\left\|[\mathbf{I}_n - \pi(\mathbf{Z})]\overline{\mathbf{A}}_k\right\|_F^2 \leq \left\|\overline{\mathbf{\Sigma}}_k\right\|_F^2$. For the remaining term, one has

$$\left\|[\mathbf{I}_n - \pi(\mathbf{Z})]\mathbf{A}_k\right\|_F^2 = \mathsf{tr}\left(\mathbf{\Sigma}_k\mathbf{U}_k^\top[\mathbf{I}_n - \pi(\mathbf{Z})]\mathbf{U}_k\mathbf{\Sigma}_k\right).$$

Applying Lemma A.1 and conjugating by $\mathbf{\Sigma}_k$ (6) it yields

$$\mathbf{\Sigma}_k\mathbf{U}_k^\top[\mathbf{I}_n - \pi(\mathbf{Z})]\mathbf{U}_k\mathbf{\Sigma}_k \preccurlyeq \mathbf{\Sigma}_k\mathbf{T}_k^\top\mathbf{T}_k\mathbf{\Sigma}_k.$$

Using the monotonic property of the trace (7) ends the proof. □

Remark. In the particular case where $\mathbf{Z} = \mathbf{A}\mathbf{\Omega}$, we observe that Proposition A.2, up to different notation conventions, generalizes [14, Proposition 9.6].

Proposition A.3 is the final and main result from which Theorem 3.1 can be deduced. It addresses the stochastic analysis of a particular quantity that plays a central role in our proof.

Proposition A.3. Let $\mathbf{A} \in \mathbb{R}^{m\times n}$ be an arbitrary matrix, and $\mathbf{Z} \in \mathbb{R}^{m\times\ell}$ a matrix whose columns are independently sampled from a multivariate Gaussian distribution with covariance matrix $\mathbf{K} \in \mathbb{R}^{m\times m}$, such that $\ell \leq \min(\mathsf{rank}(\mathbf{A}), \mathsf{rank}(\mathbf{K}))$. Let us further denote $\mathbf{Z}_k = \mathbf{U}_k^\top\mathbf{Z} \in \mathbb{R}^{k\times\ell}$ and $\overline{\mathbf{Z}}_k = \overline{\mathbf{U}}_k^\top\mathbf{Z} \in \mathbb{R}^{(n-k)\times\ell}$.

For any integer $k \leq \ell - 2$, if $\mathbf{K}_k \equiv \mathbf{U}_k^\top\mathbf{K}\mathbf{U}_k$ is non-singular, one has

$$\mathbb{E}\left[\left\|\overline{\mathbf{Z}}_k\mathbf{Z}_k^\dagger\mathbf{\Sigma}_k\right\|_F^2\right] = \left\|\mathbf{tan}(\mathbf{U}_k, \mathbf{K}\mathbf{U}_k)\mathbf{\Sigma}_k\right\|_F^2 + \left\|[\mathbf{I}_m - \pi(\mathbf{K}^{\frac{1}{2}}\mathbf{U}_k)]\mathbf{K}^{\frac{1}{2}}\right\|_F^2 \frac{\mathsf{tr}(\mathbf{\Sigma}_k^2\mathbf{K}_k^{-1})}{\ell - k - 1}.$$

Moreover, if $k \leq \ell - 4$, then for all $u, t \geq 1$ one has

$$\left\|\overline{\mathbf{Z}}_k\mathbf{Z}_k^\dagger\mathbf{\Sigma}_k\right\|_F \leq \left\|\mathbf{tan}(\mathbf{U}_k, \mathbf{K}\mathbf{U}_k)\mathbf{\Sigma}_k\right\|_F + \sqrt{3}ut \cdot \left\|[\mathbf{I}_m - \pi(\mathbf{K}^{\frac{1}{2}}\mathbf{U}_k)]\mathbf{K}^{\frac{1}{2}}\right\|_F \sqrt{\frac{\mathsf{tr}(\mathbf{\Sigma}_k^2\mathbf{K}_k^{-1})}{\ell - k - 1}},$$

holds with probability at least $1 - e^{-u^2/2} - t^{-(\ell-k)}$.

Proof.

Expectation. Using the theorem of total expectation one has

$$\mathbb{E}\left[\left\|\overline{\mathbf{Z}}_k\mathbf{Z}_k^\dagger\mathbf{\Sigma}_k\right\|_F^2\right] = \mathbb{E}\left[\mathbb{E}\left[\left\|\overline{\mathbf{Z}}_k\mathbf{Z}_k^\dagger\mathbf{\Sigma}_k\right\|_F^2 \mid \mathbf{Z}_k\right]\right].$$

We consider the following partitioning

$$\mathbf{U}^\top \mathbf{K} \mathbf{U} = \begin{bmatrix} \mathbf{U}_k \mid \overline{\mathbf{U}}_k \end{bmatrix}^\top \mathbf{K} \begin{bmatrix} \mathbf{U}_k \mid \overline{\mathbf{U}}_k \end{bmatrix} = \begin{bmatrix} \mathbf{K}_k & \mathbf{K}_{\perp,k}^\top \\ \mathbf{K}_{\perp,k} & \overline{\mathbf{K}}_k \end{bmatrix}.$$

In the inner expectation, $\mathbf{Z}_k^\dagger$ is a fixed matrix, but unlike the standard result in [14], $\mathbf{Z}_k$ and $\overline{\mathbf{Z}}_k$ are not statistically independent. This implies that conditioned by $\mathbf{Z}_k$, the columns of $\overline{\mathbf{Z}}_k$ no longer follow a standard multivariate Gaussian distribution. Instead, if $[\mathbf{Z}_k]_{(i)}$ denotes the i-th column of $\mathbf{Z}_k$, then the i-th column of $\overline{\mathbf{Z}}_k$ follows a multivariate Gaussian distribution with mean term

$$\boldsymbol{\mu}_i \equiv \mathbf{K}_{\perp,k} \mathbf{K}_k^{-1} [\mathbf{Z}_k]_{(i)},$$

and covariance matrix

$$\mathbf{K}/\mathbf{K}_k \equiv \overline{\mathbf{K}}_k - \mathbf{K}_{\perp,k} \mathbf{K}_k^{-1} \mathbf{K}_{\perp,k}^\top.$$

Consequently, if one defines $\boldsymbol{\mu} = [\boldsymbol{\mu}_1, \ldots, \boldsymbol{\mu}_\ell]$, one can write $\overline{\mathbf{Z}}_k = \boldsymbol{\mu} + (\mathbf{K}/\mathbf{K}_k)^{\frac{1}{2}} \mathbf{G}$ where the columns of $\mathbf{G} \in \mathbb{R}^{(n-k)\times \ell}$ are sampled from a standard Gaussian distribution. Thus, conditioned by $\mathbf{Z}_k$ one has

$$\begin{aligned} \left\| \overline{\mathbf{Z}}_k \mathbf{Z}_k^\dagger \boldsymbol{\Sigma}_k \right\|_F^2 &= \left\| \left[\boldsymbol{\mu} + (\mathbf{K}/\mathbf{K}_k)^{\frac{1}{2}} \mathbf{G} \right] \mathbf{Z}_k^\dagger \boldsymbol{\Sigma}_k \right\|_F^2 \\ &= \left\| \boldsymbol{\mu}\, \mathbf{Z}_k^\dagger \boldsymbol{\Sigma}_k + (\mathbf{K}/\mathbf{K}_k)^{\frac{1}{2}} \mathbf{G} \mathbf{Z}_k^\dagger \boldsymbol{\Sigma}_k \right\|_F^2 \\ &= \left\| \boldsymbol{\mu}\, \mathbf{Z}_k^\dagger \boldsymbol{\Sigma}_k \right\|_F^2 + \left\| (\mathbf{K}/\mathbf{K}_k)^{\frac{1}{2}} \mathbf{G} \mathbf{Z}_k^\dagger \boldsymbol{\Sigma}_k \right\|_F^2 \\ &\quad + 2\,\mathsf{tr}\left(\boldsymbol{\Sigma}_k [\mathbf{Z}_k^\dagger]^\top \boldsymbol{\mu}^\top (\mathbf{K}/\mathbf{K}_k)^{\frac{1}{2}} \mathbf{G} \mathbf{Z}_k^\dagger \boldsymbol{\Sigma}_k \right). \end{aligned}$$

The linearity of the expectation allows us to deal with each term separately. Conditioned by $\mathbf{Z}_k$, the first one is a non-random constant, and the second term can be handled using [14, Proposition 10.1] since $\mathbf{G}$ has a standard multivariate Gaussian distribution. Concerning the third term, the linearity of the expectation conditioned by $\mathbf{Z}_k$ yields

$$\begin{aligned} &\mathbb{E}\left[\mathsf{tr}\left(\boldsymbol{\Sigma}_k [\mathbf{Z}_k^\dagger]^\top \boldsymbol{\mu}^\top (\mathbf{K}/\mathbf{K}_k)^{\frac{1}{2}} \mathbf{G} \mathbf{Z}_k^\dagger \boldsymbol{\Sigma}_k \right) \mid \mathbf{Z}_k \right] \\ &\quad = \mathsf{tr}\left(\boldsymbol{\Sigma}_k [\mathbf{Z}_k^\dagger]^\top \boldsymbol{\mu}^\top (\mathbf{K}/\mathbf{K}_k)^{\frac{1}{2}} \mathbb{E}\left[\mathbf{G} \mid \mathbf{Z}_k \right] \mathbf{Z}_k^\dagger \boldsymbol{\Sigma}_k \right) = 0. \end{aligned}$$

Altogether, one has

$$\begin{aligned} \mathbb{E}\left[\left\| \overline{\mathbf{Z}}_k \mathbf{Z}_k^\dagger \boldsymbol{\Sigma}_k \right\|_F^2 \mid \mathbf{Z}_k \right] &= \left\| \boldsymbol{\mu}\, \mathbf{Z}_k^\dagger \boldsymbol{\Sigma}_k \right\|_F^2 + \left\| (\mathbf{K}/\mathbf{K}_k)^{\frac{1}{2}} \right\|_F^2 \left\| \mathbf{Z}_k^\dagger \boldsymbol{\Sigma}_k \right\|_F^2 \\ &= \left\| \mathbf{K}_{\perp,k} \mathbf{K}_k^{-1} \boldsymbol{\Sigma}_k \right\|_F^2 + \left\| (\mathbf{K}/\mathbf{K}_k)^{\frac{1}{2}} \right\|_F^2 \left\| \mathbf{Z}_k^\dagger \boldsymbol{\Sigma}_k \right\|_F^2, \end{aligned}$$

where the last equality follows from the simplification of $\mathbf{Z}_k \mathbf{Z}_k^\dagger$ when substituting $\boldsymbol{\mu}$ by its expression. Taking again the expectation yields

$$\mathbb{E}\left[\left\| \overline{\mathbf{Z}}_k \mathbf{Z}_k^\dagger \boldsymbol{\Sigma}_k \right\|_F^2 \right] = \left\| \mathbf{K}_{\perp,k} \mathbf{K}_k^{-1} \boldsymbol{\Sigma}_k \right\|_F^2 + \left\| (\mathbf{K}/\mathbf{K}_k)^{\frac{1}{2}} \right\|_F^2 \cdot \mathbb{E}\left[\left\| \mathbf{Z}_k^\dagger \boldsymbol{\Sigma}_k \right\|_F^2 \right].$$

Since the columns of $\mathbf{Z}_k$ are independently sampled from a multivariate Gaussian distribution with zero mean and covariance matrix $\mathbf{K}_k$, the matrix $\mathbf{Z}_k \mathbf{Z}_k^\top$ follows a Wishart distribution of the form $\mathcal{W}_k(\ell, \mathbf{K}_k)$ [20, Definition 3.1.3]. In addition, $\left\| \mathbf{Z}_k^\dagger \boldsymbol{\Sigma}_k \right\|_F^2 = \mathsf{tr}(\boldsymbol{\Sigma}_k [\mathbf{Z}_k^\dagger]^\top \mathbf{Z}_k^\dagger \boldsymbol{\Sigma}_k) = \mathsf{tr}(\boldsymbol{\Sigma}_k [\mathbf{Z}_k \mathbf{Z}_k^\top]^{-1} \boldsymbol{\Sigma}_k)$, where the second equality holds with probability one since $\ell > k+1$. In fact, if $\mathbf{K}_k$ is non-singular,

the matrix $\mathbf{Z}_k\mathbf{Z}_k^\top$ is non-singular almost surely, see [20, Theorem 3.1.4]. In this case, according to [20, Theorem 3.2.12] for $\ell > k+1$, one has

$$\mathbb{E}\left[\,\boldsymbol{\Sigma}_k[\mathbf{Z}_k\mathbf{Z}_k^\top]^{-1}\boldsymbol{\Sigma}_k\,\right] = \boldsymbol{\Sigma}_k\mathbb{E}\left[\,[\mathbf{Z}_k\mathbf{Z}_k^\top]^{-1}\,\right]\boldsymbol{\Sigma}_k = \frac{\boldsymbol{\Sigma}_k\mathbf{K}_k^{-1}\boldsymbol{\Sigma}_k}{\ell-k-1}.$$

Hence, by the linearity of the expectation, one gets

$$\mathbb{E}\left[\left\|\mathbf{Z}_k^\dagger\boldsymbol{\Sigma}_k\right\|_F^2\right] = \mathbb{E}\left[\,\mathsf{tr}(\boldsymbol{\Sigma}_k[\mathbf{Z}_k\mathbf{Z}_k^\top]^{-1}\boldsymbol{\Sigma}_k)\,\right] = \frac{\mathsf{tr}(\boldsymbol{\Sigma}_k\mathbf{K}_k^{-1}\boldsymbol{\Sigma}_k)}{\ell-k-1} = \frac{\mathsf{tr}(\boldsymbol{\Sigma}_k^2\mathbf{K}_k^{-1})}{\ell-k-1},$$

which concludes the proof.

Probability. Let $\varphi : \mathbb{R}^{(n-k)\times\ell} \to \mathbb{R}$, $\mathbf{Y} \mapsto \left\|\,[\boldsymbol{\mu} + (\mathbf{K}/\mathbf{K}_k)^{\frac{1}{2}}\mathbf{Y}]\,\mathbf{Z}_k^\dagger\boldsymbol{\Sigma}_k\right\|_F$. Using the reverse triangular inequality, it holds for any $\mathbf{Y}_1, \mathbf{Y}_2 \in \mathbb{R}^{(n-k)\times\ell}$ that

$$\begin{aligned}\left|\varphi(\mathbf{Y}_1) - \varphi(\mathbf{Y}_2)\right| &\leq \left\|(\mathbf{K}/\mathbf{K}_k)^{\frac{1}{2}}\,[\mathbf{Y}_1 - \mathbf{Y}_2]\,\mathbf{Z}_k^\dagger\boldsymbol{\Sigma}_k\right\|_F \\ &\leq \left\|(\mathbf{K}/\mathbf{K}_k)^{\frac{1}{2}}\right\|_F\left\|\mathbf{Z}_k^\dagger\boldsymbol{\Sigma}_k\right\|_F\left\|[\mathbf{Y}_1 - \mathbf{Y}_2]\right\|_F,\end{aligned}$$

where the last inequality follows from the submultiplicativity property. Thus, φ is at worst L-Lipschitz with $L = \left\|(\mathbf{K}/\mathbf{K}_k)^{\frac{1}{2}}\right\|_F\left\|\mathbf{Z}_k^\dagger\boldsymbol{\Sigma}_k\right\|_F$. Applying [14, Proposition 10.3] then yields that if $\mathbf{G} \in \mathbb{R}^{(n-k)\times\ell}$ is a matrix whose columns are sampled from a standard Gaussian distribution, then $\forall\, u \geq 0$ one has

$$\mathbb{P}\left\{\varphi(\mathbf{G}) \leq \mathbb{E}\left[\,\varphi(\mathbf{G})\,\right] + Lu\right\} \geq 1 - e^{-u^2/2}.$$

Hölder's inequality yields $\mathbb{E}\left[\,\varphi(\mathbf{G})\,\right] \leq (\mathbb{E}\left[\,\varphi(\mathbf{G})^2\,\right])^{\frac{1}{2}}$, and combining this fact with the result proved right above and the properties of multivariate Gaussian distribution one has

$$\begin{aligned}\mathbb{E}\left[\,\varphi(\mathbf{G})\,\right] &\leq \left(\mathbb{E}\left[\left\|\overline{\mathbf{Z}}_k\mathbf{Z}_k^\dagger\boldsymbol{\Sigma}_k\right\|_F^2 \mid \mathbf{Z}_k\right]\right)^{\frac{1}{2}} \\ &= \left(\left\|\mathbf{K}_{\perp,k}\mathbf{K}_k^{-1}\boldsymbol{\Sigma}_k\right\|_F^2 + \left\|(\mathbf{K}/\mathbf{K}_k)^{\frac{1}{2}}\right\|_F^2\left\|\mathbf{Z}_k^\dagger\boldsymbol{\Sigma}_k\right\|_F^2\right)^{\frac{1}{2}} \\ &\leq \left\|\mathbf{K}_{\perp,k}\mathbf{K}_k^{-1}\boldsymbol{\Sigma}_k\right\|_F + \left\|(\mathbf{K}/\mathbf{K}_k)^{\frac{1}{2}}\right\|_F\left\|\mathbf{Z}_k^\dagger\boldsymbol{\Sigma}_k\right\|_F.\end{aligned}$$

Altogether, this implies that

$$\mathbb{P}\left\{\left\|\overline{\mathbf{Z}}_k\mathbf{Z}_k^\dagger\boldsymbol{\Sigma}_k\right\|_F \leq \left\|\mathbf{K}_{\perp,k}\mathbf{K}_k^{-1}\boldsymbol{\Sigma}_k\right\|_F + (1+u)\cdot\left\|(\mathbf{K}/\mathbf{K}_k)^{\frac{1}{2}}\right\|_F\left\|\mathbf{Z}_k^\dagger\boldsymbol{\Sigma}_k\right\|_F\right\} \geq 1 - e^{-u^2/2}. \tag{17}$$

Let us now consider, for $t \geq 1$, the event

$$E_t = \left\{\left\|\mathbf{Z}_k^\dagger\boldsymbol{\Sigma}_k\right\|_F \leq \sqrt{3}t\cdot\sqrt{\frac{\mathsf{tr}(\boldsymbol{\Sigma}_k^2\mathbf{K}_k^{-1})}{\ell-k+1}}\right\}.$$

Applying [1, Lemma 8] yields for all $t \geq 1$ that

$$\mathbb{P}\left\{E_t^C\right\} \leq t^{-(\ell-k)},$$

where E_t^C denotes the complement of the event set E_t. Now, by denoting $E_{\mathbf{Z}_k}$ the event considered in (17), the law of total probabilities reads

$$\mathbb{P}\left\{E_{\mathbf{Z}_k}\right\} = \mathbb{P}\left\{E_{\mathbf{Z}_k} \mid E_t\right\}\mathbb{P}\left\{E_t\right\} + \mathbb{P}\left\{E_{\mathbf{Z}_k} \mid E_t^{\mathsf{C}}\right\}\mathbb{P}\left\{E_t^{\mathsf{C}}\right\},$$

Since one trivially has $\mathbb{P}\{E_t\} \leq 1$ and $\mathbb{P}\left\{E_{\mathbf{Z}_k} \mid E_t^{\mathsf{C}}\right\} \leq 1$, one gets

$$1 - e^{-u^2/2} \leq \mathbb{P}\left\{E_{\mathbf{Z}_k}\right\} \leq \mathbb{P}\left\{E_{\mathbf{Z}_k} \mid E_t\right\} + t^{-(\ell-k)},$$

or equivalently,

$$\mathbb{P}\left\{E_{\mathbf{Z}_k} \mid E_t\right\} \geq 1 - e^{-u^2/2} - t^{-(\ell-k)},$$

which ends the proof for the bound in probability.

Rearranging the terms. It remains to transform the terms to get more interpretable bounds. Since by definition $\left[\mathbf{U}_k \mid \overline{\mathbf{U}}_k\right]$ is orthogonal, and $\mathbf{Z}$ is full column rank, then using (9), one readily has

$$\mathbf{K}_{\perp,k}\mathbf{K}_k^{-1} = \overline{\mathbf{U}}_k^\top \mathbf{K}\mathbf{U}_k(\mathbf{U}_k^\top \mathbf{K}\mathbf{U}_k)^{-1} = \mathbf{tan}(\mathbf{U}_k, \mathbf{K}\mathbf{U}_k).$$

Then, using classical algebraic manipulations, one gets

$$\begin{aligned}
\left\|(\mathbf{K}/\mathbf{K}_k)^{\frac{1}{2}}\right\|_F^2 &= \mathsf{tr}\left(\overline{\mathbf{U}}_k^\top \mathbf{K}\overline{\mathbf{U}}_k - \overline{\mathbf{U}}_k^\top \mathbf{K}\mathbf{U}_k\mathbf{K}_k^{-1}\mathbf{U}_k^\top \mathbf{K}\overline{\mathbf{U}}_k\right) \\
&= \mathsf{tr}\left(\overline{\mathbf{U}}_k^\top \left[\mathbf{K} - \mathbf{K}\mathbf{U}_k\mathbf{K}_k^{-1}\mathbf{U}_k^\top \mathbf{K}\right]\overline{\mathbf{U}}_k\right) \\
&= \mathsf{tr}\left(\overline{\mathbf{U}}_k\overline{\mathbf{U}}_k^\top \left[\mathbf{K} - \mathbf{K}\mathbf{U}_k\mathbf{K}_k^{-1}\mathbf{U}_k^\top \mathbf{K}\right]\right) \\
&= \mathsf{tr}\left([\mathbf{I}_n - \mathbf{U}_k\mathbf{U}_k^\top]\left[\mathbf{K} - \mathbf{K}\mathbf{U}_k\mathbf{K}_k^{-1}\mathbf{U}_k^\top \mathbf{K}\right]\right) \\
&= \mathsf{tr}\left(\mathbf{K} - \mathbf{K}\mathbf{U}_k\mathbf{K}_k^{-1}\mathbf{U}_k^\top \mathbf{K}\right) - \mathsf{tr}\left(\mathbf{U}_k\mathbf{U}_k^\top \left[\mathbf{K} - \mathbf{K}\mathbf{U}_k\mathbf{K}_k^{-1}\mathbf{U}_k^\top \mathbf{K}\right]\right).
\end{aligned}$$

The second term is trivially equal to 0, which finally yields

$$\begin{aligned}
\left\|(\mathbf{K}/\mathbf{K}_k)^{\frac{1}{2}}\right\|_F^2 &= \mathsf{tr}\left(\mathbf{K} - \mathbf{K}\mathbf{U}_k\mathbf{K}_k^{-1}\mathbf{U}_k^\top \mathbf{K}\right) \\
&= \mathsf{tr}\left(\mathbf{K}^{\frac{1}{2}}\left[\mathbf{I}_n - \mathbf{K}^{\frac{1}{2}}\mathbf{U}_k\mathbf{K}_k^{-1}\mathbf{U}_k^\top \mathbf{K}^{\frac{1}{2}}\right]\mathbf{K}^{\frac{1}{2}}\right) \\
&= \mathsf{tr}\left(\mathbf{K}^{\frac{1}{2}}\left[\mathbf{I}_m - \pi(\mathbf{K}^{\frac{1}{2}}\mathbf{U}_k)\right]\mathbf{K}^{\frac{1}{2}}\right) \\
&= \left\|[\mathbf{I}_m - \pi(\mathbf{K}^{\frac{1}{2}}\mathbf{U}_k)]\mathbf{K}^{\frac{1}{2}}\right\|_F^2.
\end{aligned}$$

□

We are now ready to prove Theorem 3.1. Similarly to the proof of Proposition A.3, we prove the result in expectation and in probability, separately.

Proof of Theorem 3.1: Bound in expectation If $\mathbf{Z} \in \mathbb{R}^{n\times\ell}$ is a matrix whose columns are independently sampled from a multivariate Gaussian distribution with covariance matrix $\mathbf{K} \in \mathbb{R}^{m\times m}$, then $\mathbf{U}_k^\top \mathbf{Z}$ has full row rank with probability 1. Hence, since the expectation is monotonic and linear, taking the expectation over $\mathbf{Z}$ in Proposition A.2 yields

$$\mathbb{E}\left[\left\|\left[\mathbf{I}_n - \pi(\mathbf{Z})\right]\mathbf{A}\right\|_F^2\right] \leq \left\|\overline{\mathbf{\Sigma}}_k\right\|_F^2 + \mathbb{E}\left[\left\|\overline{\mathbf{Z}}_k\mathbf{Z}_k^\dagger\mathbf{\Sigma}_k\right\|_F^2\right].$$

Applying Proposition A.3 yields

$$\begin{aligned}
\mathbb{E}\left[\left\|\overline{\mathbf{Z}}_k\mathbf{Z}_k^\dagger\mathbf{\Sigma}_k\right\|_F^2\right] &= \left\|\mathbf{tan}(\mathbf{U}_k, \mathbf{K}\mathbf{U}_k)\mathbf{\Sigma}_k\right\|_F^2 + \left\|[\mathbf{I}_m - \pi(\mathbf{K}^{\frac{1}{2}}\mathbf{U}_k)]\mathbf{K}^{\frac{1}{2}}\right\|_F^2 \frac{\mathsf{tr}(\mathbf{\Sigma}_k^2\mathbf{K}_k^{-1})}{\ell - k - 1} \\
&= \left(\tau_k(\mathbf{K})^2 + \frac{\rho_k(\mathbf{K})^2}{\ell - k - 1}\right)\left\|\overline{\mathbf{\Sigma}}_k\right\|_F^2,
\end{aligned}$$

so that we obtain

$$\mathbb{E}\left[\left\|\left[\mathbf{I}_n - \pi(\mathbf{Z})\right]\mathbf{A}\right\|_F^2\right] \leq \left(1 + \tau_k(\mathbf{K})^2 + \frac{\rho_k(\mathbf{K})^2}{\ell - k - 1}\right)\left\|\overline{\mathbf{\Sigma}}_k\right\|_F^2.$$

Using Hölder's inequality, that is $\mathbb{E}[|X|^2] \leq \mathbb{E}[|X|]^2$ for any random variable X, concludes the proof.

Proof of Theorem 3.1: Bound in probability Since $\sqrt{a+b} \leq \sqrt{a}+\sqrt{b}$ for any $a, b \geq 0$, Proposition A.2 implies that

$$\left\|\left[\mathbf{I}_n - \pi(\mathbf{Z})\right]\mathbf{A}\right\|_F \leq \left\|\overline{\mathbf{\Sigma}}_k\right\|_F + \left\|\overline{\mathbf{Z}}_k\mathbf{Z}_k^\dagger\mathbf{\Sigma}_k\right\|_F.$$

Applying Lemma A.3 on the second term of the right-hand side concludes the proof.

References

[1] Nicolas Boullé and Alex Townsend. A generalization of the randomized singular value decomposition. arXiv:2105.13052 [cs, math, stat], January 2022.

[2] Nicolas Boullé and Alex Townsend. Learning Elliptic Partial Differential Equations with Randomized Linear Algebra. Foundations of Computational Mathematics, 23:709–739, 2023.

[3] N. Bousserez and Daven K. Henze. Optimal and scalable methods to approximate the solutions of large-scale Bayesian problems: Theory and application to atmospheric inversion and data assimilation. Quarterly Journal of the Royal Meteorological Society, 144(711):365–390, January 2018.

[4] Nicolas Bousserez, Jonathan J. Guerrette, and Daven K. Henze. Enhanced parallelization of the incremental 4D-Var data assimilation algorithm using the Randomized Incremental Optimal Technique. Quarterly Journal of the Royal Meteorological Society, 146(728):1351–1371, April 2020.

[5] Roger Daley. Atmospheric Data Analysis. Number 2 in Cambridge Atmospheric and Space Science Series. Cambridge University Press, Cambridge, first edition, 1999.

[6] Ieva Daužickaitė, Amos S. Lawless, Jennifer A. Scott, and Peter Jan Leeuwen. Randomised preconditioning for the forcing formulation of weak-constraint 4D-Var. Quarterly Journal of the Royal Meteorological Society, 147(740):3719–3734, October 2021.

[7] Carl Eckart and Gale Young. The approximation of one matrix by another of lower rank. Psychometrika, 1(3):211–218, September 1936.

[8] Gene H. Golub and Charles F. Van Loan. Matrix Computations. Johns Hopkins Studies in the Mathematical Sciences. The Johns Hopkins University Press, Baltimore, fourth edition, 2013.

[9] Gene Howard Golub and Charles F. Van Loan. Matrix Computations. Johns Hopkins Series in the Mathematical Sciences. Johns Hopkins university press, Baltimore London, third edition, 1996.

[10] S. Gratton, A. S. Lawless, and N. K. Nichols. Approximate Gauss–Newton Methods for Nonlinear Least Squares Problems. SIAM Journal on Optimization, 18(1):106–132, January 2007.

[11] S. Gratton, A. Sartenaer, and J. Tshimanga. On A Class of Limited Memory Preconditioners For Large Scale Linear Systems With Multiple Right-Hand Sides. SIAM Journal on Optimization, 21(3):912–935, July 2011.

[12] Ming Gu. Subspace Iteration Randomization and Singular Value Problems. SIAM Journal on Scientific Computing, 37(3):A1139–A1173, January 2015.

[13] Wolfgang Hackbusch. Hierarchical Matrices: Algorithms and Analysis. Number 49 in Hierarchical Matrices. Springer Berlin Heidelberg, Berlin, Heidelberg, first edition, 2015.

[14] Nathan Halko, Per-Gunnar Martinsson, and Joel A. Tropp. Finding Structure with Randomness: Probabilistic Algorithms for Constructing Approximate Matrix Decompositions. SIAM Review, 53(2):217–288, January 2011.

[15] Nicholas J. Higham. Functions of Matrices: Theory and Computation. Society for Industrial and Applied Mathematics, Philadelphia, 2008.

[16] Roger A. Horn and Charles R. Johnson. Matrix Analysis. Cambridge University Press, Cambridge ; New York, second edition, 2012.

[17] Andrew V. Knyazev and Merico E. Argentati. Principal Angles between Subspaces in an A-Based Scalar Product: Algorithms and Perturbation Estimates. SIAM Journal on Scientific Computing, 23(6):2008–2040, January 2002.

[18] William Lahoz, Boris Khattatov, and Richard Menard, editors. Data Assimilation. Springer Berlin Heidelberg, Berlin, Heidelberg, 2010.

[19] Michael W. Mahoney. Randomized Algorithms for Matrices and Data. Foundations and Trends® in Machine Learning, 3(2):123–224, 2010.

[20] Robb J. Muirhead. Aspects of Multivariate Statistical Theory. Wiley Series in Probability and Mathematical Statistics. Wiley, New York, 1982.

[21] Cameron Musco and Christopher Musco. Randomized Block Krylov Methods for Stronger and Faster Approximate Singular Value Decomposition. MIT Press, 1:1396–1404, October 2015.

[22] Christopher C. Paige and M. Wei. History and generality of the CS decomposition. Linear Algebra and its Applications, 208–209:303–326, September 1994.

[23] Vladimir Rokhlin, Arthur Szlam, and Mark Tygert. A Randomized Algorithm for Principal Component Analysis. SIAM Journal on Matrix Analysis and Applications, 31(3):1100–1124, January 2010.

[24] Arvind K. Saibaba. Randomized Subspace Iteration: Analysis of Canonical Angles and Unitarily Invariant Norms. SIAM Journal on Matrix Analysis and Applications, 40(1):23–48, January 2019.

[25] Arvind K. Saibaba, Joseph Hart, and Bart van Bloemen Waanders. Randomized algorithms for generalized singular value decomposition with application to sensitivity analysis. Numerical Linear Algebra with Applications, 28(4), August 2021.

[26] Arvind K. Saibaba, Jonghyun Lee, and Peter K. Kitanidis. Randomized algorithms for generalized Hermitian eigenvalue problems with application to computing Karhunen–Loève expansion. Numerical Linear Algebra with Applications, 23(2):314–339, March 2016.

[27] Alexandre Scotto Di perrotolo,. Randomized Numerical Linear Algebra Methods with Applications to Data Assimilation. PhD thesis, Toulouse, ISAE-SUPERO, September 2022.

[28] Joel A. Tropp. Randomized block Krylov methods for approximating extreme eigenvalues. Numerische Mathematik, 150(1):217–255, January 2022.

[29] Charles F. Van Loan. Generalizing the Singular Value Decomposition. SIAM Journal on Numerical Analysis, 13(1):76–83, March 1976.

[30] Franco Woolfe, Edo Liberty, Vladimir Rokhlin, and Mark Tygert. A fast randomized algorithm for the approximation of matrices. Applied and Computational Harmonic Analysis, 25(3):335–366, November 2008.

[31] Qiaochu Yuan, Ming Gu, and Bo Li. Superlinear Convergence of Randomized Block Lanczos Algorithm. In 2018 IEEE International Conference on Data Mining (ICDM), pages 1404–1409, Singapore, November 2018. IEEE.

[32] Peizhen Zhu and Andrew V. Knyazev. Angles between subspaces and their tangents. Journal of Numerical Mathematics, 21(4):325–340, January 2013.