ANALYSIS OF RANDOMIZED BLOCK KRYLOV METHODS

JOEL A. TROPP

Technical Report No. 2018-02
June 2021

APPLIED & COMPUTATIONAL MATHEMATICS
CALIFORNIA INSTITUTE OF TECHNOLOGY
mail code 9-94 · pasadena, ca 91125

CALIFORNIA INSTITUTE OF TECHNOLOGY

Department of Computing & Mathematical Sciences, Annenberg Hall
14 June 2021

Re: Caltech ACM Technical Report 2018-02

This report comprises three manuscripts, written in 2017–2018, that provide an analysis of randomized block Krylov methods for spectral computations and low-rank matrix approximation:

- Analysis of randomized block Krylov methods. Part I: Spectral norm estimation
- Analysis of randomized block Krylov methods. Part II: Eigenvalue approximation
- Analysis of randomized block Krylov methods. Part III: Low-rank approximation of matrices

These works were produced under a subcontract to ONR Award 00014-16-C-2009 to UTRC, and they have been cleared for public release.

This report is intended to provide a permanent reference for these manuscripts, since they have been cited in the literature but are not publicly available.

I hope that this report provides a more systematic way of analyzing randomized block Krylov methods, that it sheds light on their behavior, and advances research on randomized matrix computations.

Please be aware that the existing manuscripts have a number of shortcomings. The most serious point is that **the pseudocode is not numerically sound, and it is not recommended for use in practice**. Second, these manuscripts lack numerical work that interrogates whether the theoretical bounds are informative. Third, the treatment of related work is incomplete and out of date. These issues will be addressed in subsequent versions of the papers.

The first manuscript, on spectral norm estimation, has already been revised for publication. This report contains the original version. The updated version will be released on arXiv upon acceptance to a journal, likely in mid-2021.

The second and third manuscripts need to be combined for publication. This consolidated paper will be released when it has been completed, likely in late 2021.

Please send comments or feedback to jtropp@cms.caltech.edu.

Joel A. Tropp
Steele Family Professor of Applied & Computational Mathematics
California Institute of Technology
Pasadena, California

ANALYSIS OF RANDOMIZED BLOCK KRYLOV METHODS
PART I: SPECTRAL NORM ESTIMATION

JOEL A. TROPP

ABSTRACT. Randomized block Krylov subspace methods form a powerful class of algorithms for computing the (extreme) eigenvalues and singular values of a matrix. The purpose of this paper is to develop new theoretical bounds on the performance of randomized block Krylov subspace methods for these problems. The results demonstrate that, for many matrices, it is possible to obtain an accurate spectral norm estimate using only a *constant* number of steps of the randomized block Krylov method. Furthermore, the analysis reveals that the behavior of the algorithm depends in a delicate way on the block size.

1. MOTIVATION AND MAIN RESULTS

Randomized block Krylov methods have emerged as a powerful tool for spectral computation and matrix approximation [HMT11, MRT11, RST09, HMST11, MM15, WZZ15, DIKMI16]. At present, our understanding of the performance of these methods is more rudimentary than our understanding of simple Krylov subspace methods in either the deterministic or random setting. The goal of this paper and its companions is to develop refined bounds that help explain the remarkable numerical performance [HMST11, MM15] of randomized block Krylov methods.

1.1. Project Overview.

Part I, this paper, focuses on the most basic questions. How well can we estimate the maximum (or minimum) eigenvalue of a symmetric matrix using a randomized block Krylov method? How well can we estimate the maximum (or minimum) singular value of a general matrix? We develop the first detailed theory for these problems. The analysis streamlines and improves the work [KW92] of Kuczyński & Woźniakowski on the simple Krylov method with a randomized starting vector.

Part II of this paper [Tro18a] studies the problem of estimating the largest (or smallest) k eigenvalues of a symmetric matrix. It also touches on the analogous questions for singular values. We have chosen to treat these methods separately because the analysis requires more difficult tools from randomized matrix approximation and random matrix theory [HMT11].

Part III of this paper [Tro18b] turns to the problem of computing a rank-k approximation of a general matrix. This problem has already been studied in some depth [MM15, DIKMI16]. We are able to obtain more precise results by exploiting insights from the earlier chapters of our research.

We have tried to make each of these papers self-contained, while minimizing repetition of material. In future work, we may also treat the problem of estimating invariant subspaces of a symmetric matrix associated with the largest k eigenvalues. This problem is closer in spirit to low-rank matrix approximation than eigenvalue estimation, but it requires independent arguments.

1.2. Block Krylov Methods for Computing the Maximum Eigenvalue.

Let us begin with a mathematical description of a block Krylov method for estimating the maximum eigenvalue of a symmetric matrix. See Section 1.2.6 for a brief discussion about implementations.

Date: 30 September 2017. Revised 26 February 2018.

2010 *Mathematics Subject Classification.* Primary: 65F30. Secondary: 68W20, 60B20.

Key words and phrases. Eigenvalue computation; Krylov subspace method; Lanczos method; numerical analysis; randomized algorithm; singular value computation.

1.2.1. *Block Krylov Subspaces.* Suppose that we are given a symmetric matrix $\boldsymbol{A} \in \mathbb{R}^{n\times n}$. Choose a test matrix $\boldsymbol{B} \in \mathbb{R}^{n\times \ell}$, where the number ℓ is called the *block size*. Select a parameter $q \in \mathbb{N}$ that controls the *depth* of the Krylov subspace. In concept, the block Krylov method constructs the matrix

$$\boldsymbol{S}_q(\boldsymbol{A};\boldsymbol{B}) := \begin{bmatrix} \boldsymbol{B} & \boldsymbol{A}\boldsymbol{B} & \boldsymbol{A}^2\boldsymbol{B} & \dots & \boldsymbol{A}^q\boldsymbol{B} \end{bmatrix} \in \mathbb{R}^{n\times(q+1)\ell}. \tag{1.1}$$

The range K_q of the matrix $\boldsymbol{S}_q$ is called a *block Krylov subspace*:

$$K_q(\boldsymbol{A};\boldsymbol{B}) := \operatorname{range}(\boldsymbol{S}_q) \subset \mathbb{R}^n. \tag{1.2}$$

The block Krylov subspace admits an alternative representation in terms of polynomials:

$$K_q(\boldsymbol{A};\boldsymbol{B}) = \bigoplus\nolimits_{\varphi\in\mathscr{P}_q} \operatorname{range}\big(\varphi(\boldsymbol{A})\boldsymbol{B}\big) \tag{1.3}$$

where $\mathscr{P}_q$ is the linear space of real polynomials with degree at most q. The $\oplus$ operator refers to the ordinary subspace sum.

1.2.2. *Invariance Properties of Krylov Subspaces.* Krylov subspaces have remarkable invariance properties that help explain their computational value.

- The block Krylov subspace only depends on the range of the test matrix:

$$K_q(\boldsymbol{A};\boldsymbol{B}\boldsymbol{T}) = K_q(\boldsymbol{A};\boldsymbol{B}) \quad \text{for all nonsingular } \boldsymbol{T} \in \mathbb{R}^{\ell\times\ell}.$$

- The block Krylov subspace co-varies with the orientation of the matrices:

$$K_q(\boldsymbol{U}\boldsymbol{A}\boldsymbol{U}^*;\boldsymbol{U}\boldsymbol{B}) = \boldsymbol{U}K_q(\boldsymbol{A};\boldsymbol{B}) \quad \text{for all orthogonal } \boldsymbol{U} \in \mathbb{R}^{n\times n}.$$

- The block Krylov subspace is invariant under affine transformations of the spectrum of the input matrix:

$$K_q(\alpha\boldsymbol{A} + \beta\mathbf{I};\boldsymbol{B}) = K_q(\boldsymbol{A};\boldsymbol{B}) \quad \text{for all } \alpha,\beta \in \mathbb{R}.$$

These facts follow directly from the definition (1.1)–(1.2) of the Krylov subspace and the representation (1.3) using polynomials. For example, see [Par98, Sec. 12.2.2] for the simple case $\ell = 1$.

1.2.3. *Computing Maximum Eigenvalues.* Block Krylov subspaces support a wide range of matrix computations. The core idea is to compress the input matrix to the Krylov subspace and to perform calculations on the (small) compressed matrix. In other words, Krylov methods belong to the class of Ritz–Galerkin methods; see [Lan50, Pai71, Saa80, Par98, Saa11].

In particular, we can obtain an estimate $\xi_{\max}(\boldsymbol{A};\boldsymbol{B};q)$ for the maximum eigenvalue $\lambda_{\max}(\boldsymbol{A})$ of the input matrix by maximizing the Rayleigh quotient of $\boldsymbol{A}$ over the block Krylov subspace $K_q(\boldsymbol{A};\boldsymbol{B})$:

$$\xi_{\max}(\boldsymbol{A};\boldsymbol{B};q) := \max_{\boldsymbol{v}\in K_q(\boldsymbol{A};\boldsymbol{B})} \frac{\boldsymbol{v}^*\boldsymbol{A}\boldsymbol{v}}{\boldsymbol{v}^*\boldsymbol{v}}. \tag{1.4}$$

The symbol * denotes the transpose of a matrix of vector, and we instate the convention that $0/0 = 0$. We may suppress the dependence of $\xi_{\max}$ on $\boldsymbol{A}$, $\boldsymbol{B}$, or q when they are clear from context.

The Rayleigh–Ritz theorem [Bha97, Cor. III.2.1] implies that the maximum eigenvalue estimate (1.4) always satisfies

$$\lambda_{\min}(\boldsymbol{A}) \le \xi_{\max}(\boldsymbol{A};\boldsymbol{B};q) \le \lambda_{\max}(\boldsymbol{A}). \tag{1.5}$$

The goal of our analysis is to understand how well $\xi_{\max}(\boldsymbol{A};\boldsymbol{B};q)$ *approximates* the maximum eigenvalue of the input matrix $\boldsymbol{A}$ as a function of the block size ℓ and the depth q of the Krylov space.

Warning 1.1 (Eigenvectors)**.** The vector $\boldsymbol{v}$ that maximizes the Rayleigh quotient in (1.4) need not provide a good approximation for any maximum eigenvector of $\boldsymbol{A}$. See the companion paper [Tro18b] for some discussion about estimating invariant subspaces.

1.2.4. *Invariance Properties of the Eigenvalue Estimate.* The eigenvalue estimate $\xi_{\max}(\boldsymbol{A};\boldsymbol{B};q)$ inherits some invariances from the block Krylov subspace. These properties help us develop effective implementations and to analyze their performance.

- For fixed depth q, the estimate only depends on the range of the test matrix $\boldsymbol{B}$:

$$\xi_{\max}(\boldsymbol{A};\boldsymbol{BT}) = \xi_{\max}(\boldsymbol{A};\boldsymbol{B}) \quad \text{for all nonsingular } \boldsymbol{T} \in \mathbb{R}^{\ell\times\ell}. \tag{1.6}$$

- For fixed depth q, the estimate does not depend on the orientation of $\boldsymbol{A}$ and $\boldsymbol{B}$ in the sense that

$$\xi_{\max}(\boldsymbol{UAU}^*;\boldsymbol{UB}) = \xi_{\max}(\boldsymbol{A};\boldsymbol{B}) \quad \text{for all orthogonal } \boldsymbol{U} \in \mathbb{R}^{n\times n}. \tag{1.7}$$

- For fixed depth q, the estimate covaries with *increasing* affine transformations of $\boldsymbol{A}$:

$$\xi_{\max}(\alpha\boldsymbol{A}+\beta\mathbf{I};\boldsymbol{B}) = \alpha\xi_{\max}(\boldsymbol{A};\boldsymbol{B}) + \beta \quad \text{for all } \alpha \geq 0 \text{ and } \beta \in \mathbb{R}. \tag{1.8}$$

These results all follow immediately from the invariance properties of the Krylov subspace (Section 1.2.2) and the definition (1.4) of the eigenvalue estimate. See [KW92] for the simple case $\ell = 1$.

1.2.5. *A Random Test Matrix.* To ensure that we can estimate the maximum eigenvalue of an *arbitrary* input matrix $\boldsymbol{A}$, we draw the test matrix $\boldsymbol{B}$ that generates the Krylov subspace *at random*.

How should we select the distribution? Observe that the eigenvalue estimate $\xi_{\max}(\boldsymbol{A};\boldsymbol{B};q)$ only depends on the range of the test matrix $\boldsymbol{B}$ because of the property (1.6). Furthermore, the property (1.7) shows that the eigenvalue estimate is invariant under rotations. Therefore, we can choose any random test matrix whose range has a rotationally invariant distribution. This idea is extracted from [KW92]. See Appendix A.1 for further justification.

In this series of papers, we will consider a standard normal test matrix $\boldsymbol{\Omega} \in \mathbb{R}^{n\times\ell}$. That is, the entries of $\boldsymbol{\Omega}$ are statistically independent Gaussian random variables, each with mean zero and variance one. It is well known that the range of this random matrix has a rotationally invariant distribution. The goal of this paper is to study the behavior of the random eigenvalue estimate $\xi_{\max}(\boldsymbol{A};\boldsymbol{\Omega};q)$.

Remark 1.2 (Other Test Matrices)**.** The analysis and detailed results in this paper depend heavily on the choice of a standard normal test matrix. In practice, we can achieve similar empirical behavior from test matrices with other types of distributions. For randomized Krylov subspace methods, the computational benefits of changing the test matrix are limited because we need to perform repeated multiplications with the input matrix to generate the Krylov matrix, cf. (1.1). See [HMT11, Secs. 4.6 and 7.4] for some discussion about other random test matrices.

1.2.6. *Implementation.* For completeness, we describe the simplest implementation of a block Krylov method for computing the largest eigenvalue of a symmetric matrix. See Algorithm 1 for pseudocode with MATLAB-like notation; this approach is adapted from [RST09, HMST11]. Let us emphasize that this algorithm is not suitable for numerical computations when the depth q is moderate or large or when $\boldsymbol{A}$ is ill-conditioned.

Here is a summary of the costs of Algorithm 1:

- A total of q matrix–matrix multiplies between $\boldsymbol{A}$ and an $n\times\ell$ matrix, plus another multiplication between $\boldsymbol{A}$ and an $n\times(q+1)\ell$ matrix. The arithmetic cost depends on whether the matrix $\boldsymbol{A}$ supports a fast multiplication operation. For example, the algorithm is more efficient when $\boldsymbol{A}$ is sparse.
- Orthogonalization of $(q+1)\ell$ vectors of length n at a cost of $\mathcal{O}(q^2\ell^2 n)$ operations.
- Solution of a dense maximum eigenvalue problem at a cost of $\mathcal{O}(q^3\ell^3)$ operations.
- Storage of the matrix $\boldsymbol{S}_q$, which requires $\mathcal{O}(q\ell n)$ units of storage.

More refined algorithms can reduce the resource usage somewhat, but these requirements are fairly typical.

As noted, Krylov subspace methods are particularly valuable when we have an efficient procedure for computing matrix–vector products with $\boldsymbol{A}$. On contemporary computer architectures, the cost of performing a product with several vectors is similar to the cost of a product with a single vector. Therefore,

Algorithm 1 A block Krylov method for computing the largest eigenvalue of a symmetric matrix

Input: Symmetric $n \times n$ matrix $\boldsymbol{A}$; block size ℓ; depth q
Output: Estimate ξ of largest eigenvalue

function BLOCKKRYLOVMAXEIG($\boldsymbol{A}$, ℓ, q)
  $\boldsymbol{Y}_0 \leftarrow \texttt{randn}(\texttt{size}(\boldsymbol{A}, 1), \ell)$ ▷ Draw $n \times \ell$ standard normal test matrix
  for $t \leftarrow 1, 2, 3, \dots, q$ **do**
    $\boldsymbol{Y}_t \leftarrow \boldsymbol{A}\boldsymbol{Y}_{t-1}$ ▷ Form blocks of Krylov matrix by repeated multiplication
  $\boldsymbol{Q} \leftarrow \texttt{orth}([\boldsymbol{Y}_0, \boldsymbol{Y}_1, \dots, \boldsymbol{Y}_q])$ ▷ Find orthonormal basis for block Krylov space
  $\boldsymbol{H} \leftarrow \boldsymbol{Q}^*(\boldsymbol{A}\boldsymbol{Q})$ ▷ Form Rayleigh matrix
  $(\sim, \xi) \leftarrow \texttt{eigs}(\boldsymbol{H}, 1)$ ▷ Compute maximum eigenvalue via *dense* linear algebra

block methods may offer practical benefits. We refer to the books [Par98, BDD+00, Saa11, GVL13] and the paper [HMST11] for more discussion and references.

1.3. **The Role of the Spectrum.** Owing to invariance, the theoretical behavior of the eigenvalue estimate $\xi_{\max}(\boldsymbol{A}; \boldsymbol{\Omega}; q)$ depends only on the spectrum of the input matrix $\boldsymbol{A}$. In this section, we develop this idea further and introduce some spectral features that affect the performance of the eigenvalue estimate.

Warning 1.3 (Numerical Behavior)**.** The numerical performance of a (block) Krylov method depends on other properties of the input matrix aside from the spectrum; for example, see [Par98, Chaps. 12, 13]. The current paper does not address numerical issues.

1.3.1. *Invariance Properties of the Random Eigenvalue Estimate.* The random estimate $\xi_{\max}(\boldsymbol{A}; \boldsymbol{\Omega}; q)$ of the maximum eigenvalue has several invariance properties that allow us to simplify the analysis.

First, the rotation invariance (1.7) of the eigenvalue estimate and the rotational invariance of the range of $\boldsymbol{\Omega}$ imply that

$$\xi_{\max}(\boldsymbol{A}; \boldsymbol{\Omega}) \sim \xi_{\max}(\boldsymbol{\Lambda}; \boldsymbol{\Omega}) \quad \text{where } \boldsymbol{A} = \boldsymbol{U}\boldsymbol{\Lambda}\boldsymbol{U}^* \text{ is an eigenvalue factorization.} \tag{1.9}$$

The symbol $\sim$ signifies equality of distribution for two random variables. In other words, the maximum eigenvalue estimate depends only on the eigenvalues of the input matrix—but not the eigenvectors.

Second, owing to the affine covariance property (1.8), the eigenvalue estimate $\xi_{\max}(\boldsymbol{A}; \boldsymbol{\Omega}; q)$ only depends on the "shape" of the spectrum of $\boldsymbol{A}$, but not its location or scale. As a consequence, we must assess the behavior of the eigenvalue estimate in terms of spectral features that are affine invariant.

1.3.2. *Spectral Features of the Input Matrix.* To express the results of our analysis, we introduce terminology for some spectral features of the symmetric matrix $\boldsymbol{A} \in \mathbb{R}^{n \times n}$. First, let us instate compact notation for the eigenvalues of $\boldsymbol{A}$:

$$a_i := \lambda_i(\boldsymbol{A}) \quad \text{for } i = 1, \dots, n, \quad \text{and} \quad a_{\max} := a_1 \geq a_2 \geq \cdots \geq a_n =: a_{\min}.$$

The map $\lambda_i(\cdot)$ returns the ith largest eigenvalue of a symmetric matrix.

Let us define some functions of the eigenvalue spectrum:

- The *spectral range* ρ of the input matrix is the distance between the extreme eigenvalues. That is, $\rho := a_{\max} - a_{\min}$.
- The *spectral gap* γ is the relative difference between the maximum eigenvalue and the next distinct eigenvalue:

$$\gamma := \frac{a_{\max} - a_{m+1}}{a_{\max} - a_{\min}} \quad \text{where} \quad a_{\max} = a_m > a_{m+1}. \tag{1.10}$$

  If $\boldsymbol{A}$ is a multiple of the identity, then $\gamma = 0$. Note that $\gamma \in [0, 1]$.

- Let ν be a nonnegative number. The *ν-stable rank* is a continuous measure of the "dimension" of the range of $\boldsymbol{A} - a_{\min}\mathbf{I}$ that reflects how quickly the spectrum decays. It is defined as

$$\operatorname{srk}(\nu) := \sum_{i=1}^{n} \left(\frac{a_i - a_{\min}}{a_{\max} - a_{\min}} \right)^{2\nu}. \tag{1.11}$$

  If $\boldsymbol{A}$ is a multiple of the identity, then $\operatorname{srk}(\nu) = 0$. Otherwise, $1 \le \operatorname{srk}(\nu) \le \operatorname{rank}(\boldsymbol{A} - a_{\min}\mathbf{I}) \le n - 1$. When the eigenvalues of $\boldsymbol{A}$ decay at a polynomial rate, the stable rank can be much smaller than the rank for an appropriate choice of ν.
- Let ξ be any estimate for the largest eigenvalue $a_{\max}$ of the input matrix $\boldsymbol{A}$. We measure the error in the estimate relative to the spectral range:

$$\operatorname{err}(\xi) := \frac{a_{\max} - \xi}{a_{\max} - a_{\min}}. \tag{1.12}$$

  The relative error in the Krylov estimate $\xi = \xi_{\max}(\boldsymbol{A}; \boldsymbol{B}; q)$ falls in the interval $[0, 1]$ because of (1.5).

The spectral gap, the stable rank, and the error measure are all invariant under increasing affine transformations of the spectrum of $\boldsymbol{A}$. We suppress the dependence of these quantities on the input matrix $\boldsymbol{A}$, unless emphasis is required.

1.4. **Matrices with Few Distinct Eigenvalues.** Before continuing, we must address an important special case. When the input matrix has few distinct eigenvalues, the block Krylov method computes the maximum eigenvalue of the matrix perfectly.

Proposition 1.4 (Randomized Block Krylov: Matrices with Few Eigenvalues)**.** *Let $\boldsymbol{A} \in \mathbb{R}^{n\times n}$ be a symmetric matrix. Fix the block size $\ell \ge 1$ and the depth $q \ge 0$ of the block Krylov subspace. Draw a standard normal matrix $\boldsymbol{\Omega} \in \mathbb{R}^{n\times \ell}$. If $\boldsymbol{A}$ has $q+1$ or fewer distinct eigenvalues, then $\operatorname{err}(\xi_{\max}(\boldsymbol{A}; \boldsymbol{\Omega}; q)) = 0$ with probability one.*

This type of result is well known (e.g., see [KW92]), but we include a short proof in Section 3.

1.5. **Matrices without a Spectral Gap.** Our first result gives probabilistic bounds for the maximum eigenvalue estimate $\xi_{\max}(\boldsymbol{A}; \boldsymbol{\Omega}; q)$ without any additional assumptions. In particular, it does not require a lower bound on the spectral gap γ.

Theorem 1.5 (Randomized Block Krylov: Maximum Eigenvalue Estimate)**.** *Instate the following hypotheses.*

- *Let $\boldsymbol{A} \in \mathbb{R}^{n\times n}$ be a symmetric input matrix.*
- *Draw a standard normal test matrix $\boldsymbol{\Omega} \in \mathbb{R}^{n\times \ell}$ with block size ℓ.*
- *Fix the depth parameter $q \ge 0$, and let $q = q_1 + q_2$ be an arbitrary nonnegative integer partition.*

We have the following probability bounds for the estimate $\xi_{\max}(\boldsymbol{A}; \boldsymbol{\Omega}; q)$, defined in (1.4), of the maximum eigenvalue of the input matrix.

(1) The relative error (1.12) in the eigenvalue estimate satisfies the probability bound

$$\mathbb{P}\left\{ \operatorname{err}(\xi_{\max}(\boldsymbol{A}; \boldsymbol{\Omega}; q)) \ge \varepsilon \right\} \le 1 \wedge \sqrt{2} \left[8 \operatorname{srk}(q_1) \cdot \mathrm{e}^{-2(2q_2+1)\sqrt{\varepsilon}} \right]^{\ell/2} \quad \text{for } \varepsilon \in [0, 1].$$

(2) The expectation of the relative error satisfies

$$\mathbb{E} \operatorname{err}(\xi_{\max}(\boldsymbol{A}; \boldsymbol{\Omega}; q)) \le 1 \wedge \left[\frac{2.70\ell^{-1} + \log(8 \operatorname{srk}(q_1))}{2(2q_2 + 1)} \right]^2.$$

The symbol $\wedge$ denotes the minimum, and the stable rank $\operatorname{srk}(\cdot)$ is defined in (1.11).

The proof of Theorem 1.5 appears in Sections 4 and 5.

Let us take a moment to explain the content of this result. We begin with a discussion about the role of the second depth parameter q_2, and then we explain the role of the first depth parameter q_1. Let us emphasize that the user does not choose the partition $q = q_1 + q_2$; the block Krylov method automatically makes the optimal selection.

For the moment, we fix q_1. The key message of Theorem 1.5 is that the relative error satisfies the bound

$$\mathbb{E}\,\mathrm{err}(\xi_{\max}(\boldsymbol{A};\boldsymbol{\Omega};q_1+q_2)) \le \varepsilon$$

once the depth parameter q_2 exceeds

$$q_2(\varepsilon) := -\frac{1}{2} + \frac{2.70\ell^{-1} + \log(8\,\mathrm{srk}(q_1))}{4\sqrt{\varepsilon}}.$$

Once the depth q_2 attains this level, the probability of error drops off exponentially fast:

$$q_2 = q_2(\varepsilon) + k\varepsilon^{-1/2} \quad \text{implies} \quad \mathbb{P}\left\{\mathrm{err}(\xi_{\max}(\boldsymbol{A};\boldsymbol{\Omega};q_1+q_2)) \ge \varepsilon\right\} \le \mathrm{e}^{-2k\ell}.$$

In fact, we need $q_2 \ge q_2(\varepsilon)$ just to ensure that the probability bound is nontrivial.

The most important aspect of this result is that the depth $q_2(\varepsilon)$ scales with $\varepsilon^{-1/2}$, so it is possible to achieve moderate relative error using a block Krylov space with limited depth. In contrast, the power method [KW92] requires depth q to be proportional to ε^{-1} to achieve a relative error of ε.

The second thing to notice is that the depth $q_2(\varepsilon)$ scales with $\log(\mathrm{srk}(q_1))$. The stable rank is never larger than the ambient dimension n, but it can be significantly smaller—even constant—when the spectrum of the matrix decays. We believe this is the first result that identifies conditions under which the block Krylov method has dimensionless behavior.

Here is another way to look at these facts. As we increase the depth parameter q, the block Krylov method exhibits a burn-in period whose length $q_1 + q_2(1)$ depends on $\mathrm{srk}(q_1)$. While the depth $q_2 \le q_2(1)$, the algorithm does not make any progress in estimating the maximum eigenvalue. Once the depth satisfies $q_2 \ge q_2(1)$, the algorithm the expected relative error decreases in proportion to q_2^{-2}. In contrast, the power method [KW92] reduces the expected relative error in proportion to q_2^{-1}.

We can now appreciate the role of the first depth parameter q_1. When the spectrum of the input matrix exhibits polynomial decay, $\mathrm{srk}(q_1)$ is *constant* for an appropriate value of q_1. For many matrices, this situation already occurs for $q_1 = 1$ or $q_1 = 2$. The analysis shows that the total burn-in period $q_1 + q_2(1)$ can be as short as $\mathcal{O}(1)$ steps when the eigenvalues of the input matrix decay.

The block size ℓ does not play a significant role in determining the average error. But changing the block size has a large effect on the probability of failure, i.e., the event that the relative error exceeds ε. For example, suppose that we increase the block size ℓ from one to three. For each increment of $\varepsilon^{-1/2}$ in the depth q_2, the failure probability with block size $\ell = 3$ is a factor of 403× smaller than the failure probability with block size $\ell = 1$!

Remark 1.6 (Prior Work)**.** The simple case $\ell = 1$ in Theorem 1.5 has been studied in the paper [KW92]. Our work introduces two major innovations. First, we obtain bounds in terms of the stable rank, which allows us to mitigate the dimensional dependence that appears in [KW92]. Second, we have obtained precise results for larger block sizes ℓ, which indicate potential benefits of using block Krylov methods. Our proof strategy is motivated by the work in [KW92], but we have been able to streamline and extend their arguments by using a more transparent random model for the test matrix.

1.6. **Matrices with a Spectral Gap.** Our second result gives probabilistic bounds for the maximum eigenvalue estimate $\xi_{\max}(\boldsymbol{A};\boldsymbol{\Omega};q)$ when we have a lower bound for the spectral gap γ of the input matrix.

Theorem 1.7 (Randomized Block Krylov: Maximum Eigenvalue Estimate with Spectral Gap)**.** *Instate the following hypotheses.*

- *Let $\boldsymbol{A} \in \mathbb{R}^{n\times n}$ be a symmetric input matrix.*
- *Draw a standard normal test matrix $\boldsymbol{\Omega} \in \mathbb{R}^{n\times\ell}$ with block size ℓ.*

- *Fix the depth parameter $q \geq 0$, and let $q = q_1 + q_2$ be an arbitrary nonnegative integer partition.*

We have the following probability bounds for the estimate $\xi_{\max}(\boldsymbol{A}; \boldsymbol{\Omega}; q)$, defined in (1.4), of the maximum eigenvalue of the input matrix.

(1) The relative error (1.12) in the eigenvalue estimate satisfies the probability bound

$$\mathbb{P}\left\{\operatorname{err}(\xi_{\max}(\boldsymbol{A}; \boldsymbol{\Omega}; q)) \geq \varepsilon\right\} \leq \sqrt{2}\left[\frac{8 \operatorname{srk}(q_1)}{\varepsilon} \cdot \mathrm{e}^{-4q_2\sqrt{\gamma}}\right]^{\ell/2} \quad \text{for } \varepsilon \in (0, 1].$$

(2) The expectation of the relative error satisfies

$$\begin{aligned}
\mathbb{E} \operatorname{err}(\xi_{\max}(\boldsymbol{A}; \boldsymbol{\Omega})) &\leq \frac{1}{\ell - 2}\big(4 \operatorname{srk}(q_1)\big) \cdot \mathrm{e}^{-4q_2\sqrt{\gamma}} && (\ell \geq 3); \\
\mathbb{E} \operatorname{err}(\xi_{\max}(\boldsymbol{A}; \boldsymbol{\Omega})) &\leq \sqrt{4 \operatorname{srk}(q_1)} \cdot \mathrm{e}^{-2q_2\sqrt{\gamma}} && (\ell = 2); \\
\mathbb{E} \operatorname{err}(\xi_{\max}(\boldsymbol{A}; \boldsymbol{\Omega})) &\leq 1.13\ \sqrt{4 \operatorname{srk}(q_1)} \cdot \mathrm{e}^{-2q_2\sqrt{\gamma}} && (\ell = 1).
\end{aligned}$$

The spectral gap is defined in (1.10), and the stable rank is defined in (1.11).

The proof of Theorem 1.7 appears in Sections 4 and 6.

As before, we fix the first depth parameter q_1. One implication of this result is that the relative error satisfies

$$\mathbb{P}\left\{\operatorname{err}(\xi_{\max}(\boldsymbol{A}; \boldsymbol{\Omega}; q)) \leq \varepsilon\right\} > 0$$

when the second depth parameter q_2 exceeds

$$q_2(\varepsilon; \gamma) := \frac{0.70\ell^{-1} + \log(8\varepsilon^{-1} \operatorname{srk}(q_1))}{4\sqrt{\gamma}}.$$

In this case, the depth q_2 must scale with $\log(1/\varepsilon)$, so the block Krylov method can achieve a very small relative error—provided that the spectral gap γ exceeds the desired error tolerance ε. As before, the depth q_2 scales with $\log(\operatorname{srk}(q_1))$, so the dimensional dependence is weak—or even nonexistent if the spectrum has polynomial decay and q_1 is sufficiently large.

When the depth $q_2 \geq q_2(\varepsilon; \gamma)$, the error probability drops off quickly:

$$q_2 = q_2(\varepsilon; \gamma) + k\gamma^{-1/2} \quad \text{implies} \quad \mathbb{P}\{\operatorname{err}(\xi_{\max}(\boldsymbol{A}; \boldsymbol{\Omega})) \geq \varepsilon\} \leq \mathrm{e}^{-2k\ell}.$$

This bound indicates that $\gamma^{-1/2}$ is the scale on which the depth q_2 needs to increase to reduce the failure probability by a constant multiple (which depends on the block size).

We discover another phenomenon when we examine the expectation of the error. On average, to achieve a relative error of ε, it is sufficient that the depth $q_2 \geq q_2'(\varepsilon; \gamma)$, where

$$\begin{aligned}
\ell = 3: \quad q_2'(\varepsilon; \gamma) &:= \frac{-\log(\ell - 2) + \log(4 \operatorname{srk}(q_1)) + \log(1/\varepsilon)}{4\sqrt{\gamma}} \\
\ell = 1, 2: \quad q_2'(\varepsilon; \gamma) &:= \frac{0.25 + \log(4 \operatorname{srk}(q_1)) + \log(1/\varepsilon^2)}{4\sqrt{\gamma}}.
\end{aligned}$$

In other words, the depth q_2 of the block Krylov space needs to be about $\log(4 \operatorname{srk}(q_1))/(4\sqrt{\gamma})$ before we obtain an average relative error less than one; we can reduce this requirement slightly by increasing the block size ℓ. But once the depth q_2 reaches this level, Theorem 1.7 suggests that the block Krylov method with $\ell \geq 3$ reduces the average error *twice as fast* as the block Krylov method with $\ell = 1, 2$.

Remark 1.8 (Prior Work)**.** The simple case $\ell = 1$ in Theorem 1.7 has been studied in the paper [KW92]. See Remark 1.6 for a discussion of the improvements we have achieved.

1.7. **Estimating Minimum Eigenvalues.** We can also use Krylov subspace methods to obtain an estimate $\xi_{\min}(\boldsymbol{A};\boldsymbol{B};q)$ for the *minimum* eigenvalue of a symmetric matrix $\boldsymbol{A}$. Conceptually, the simplest way to accomplish this task is to apply the Krylov subspace method to the negation $-\boldsymbol{A}$. The minimum eigenvalue estimate takes the form

$$\xi_{\min}(\boldsymbol{A};\boldsymbol{B};q) := -\xi_{\max}(-\boldsymbol{A};\boldsymbol{B};q).$$

Owing to (1.5), this estimate is never smaller than $\lambda_{\min}(\boldsymbol{A})$.

It is straightforward to adapt Theorems 1.5 to obtain bounds for the minimum eigenvalue estimate with a random test matrix $\boldsymbol{\Omega}$. In particular, we always have the bound

$$\mathbb{E}\left[\frac{\xi_{\min}(\boldsymbol{A};\boldsymbol{\Omega};q) - a_{\min}}{a_{\max} - a_{\min}}\right] \le \left[\frac{2.70\ell^{-1} + \log(8\,\mathrm{srk}(-\boldsymbol{A};q_1))}{2(2q_2+1)}\right]^2.$$

In this context, the stable rank takes the form

$$\mathrm{srk}(-\boldsymbol{A};\nu) = \sum_{i=1}^{n}\left(\frac{a_{\max} - a_i}{a_{\max} - a_{\min}}\right)^{2\nu}.$$

See Section 1.5 for discussion of this type of bound.

We can also use Theorem 1.7 obtain results in terms of the spectral gap. The spectral gap $\gamma(-\boldsymbol{A})$ is the magnitude of the difference between the smallest two eigenvalues of $\boldsymbol{A}$, relative to the spectral range. For example, when the block size $\ell = 3$, we have the bound

$$\ell = 3: \qquad \mathbb{E}\left[\frac{\xi_{\min}(\boldsymbol{A};\boldsymbol{\Omega};q) - a_{\min}}{a_{\max} - a_{\min}}\right] \le 4\,\mathrm{srk}(-\boldsymbol{A};q_1)\cdot \mathrm{e}^{-4q_2\sqrt{\gamma(-\boldsymbol{A})}}.$$

See Section 1.6 for discussion of this type of bound.

1.8. **Estimating Singular Values.** We now arrive at the problem of estimating the spectral norm of a general matrix $\boldsymbol{C} \in \mathbb{R}^{m\times n}$ using Krylov subspace methods. To do so, we just apply the block Krylov method to the square $\boldsymbol{C}\boldsymbol{C}^*$. This yields an estimate $\xi_{\max}(\boldsymbol{C}\boldsymbol{C}^*;\boldsymbol{B};q)$ for the *square* of the spectral norm of $\boldsymbol{C}$.

Theorems 1.5 and 1.7 immediately yield error bounds for the random test matrix $\boldsymbol{\Omega}$. In particular, we always have the bound

$$\mathbb{E}\left[\frac{\|\boldsymbol{C}\|^2 - \xi_{\max}(\boldsymbol{C}\boldsymbol{C}^*;\boldsymbol{\Omega};q)}{\|\boldsymbol{C}\|^2}\right] \le \left[\frac{2.70\ell^{-1} + \log(8\,\mathrm{srk}(\boldsymbol{C}\boldsymbol{C}^*;q_1))}{2(2q_2+1)}\right]^2.$$

We also have a bound in terms of the spectral gap $\gamma(\boldsymbol{C}\boldsymbol{C}^*)$, which is the difference between the squares of the largest two distinct singular values. For block size $\ell = 3$, we have

$$\mathbb{E}\left[\frac{\|\boldsymbol{C}\|^2 - \xi_{\max}(\boldsymbol{C}\boldsymbol{C}^*;\boldsymbol{\Omega})}{\|\boldsymbol{C}\|^2}\right] \le 4\,\mathrm{srk}(\boldsymbol{C}\boldsymbol{C}^*;q_1)\cdot \mathrm{e}^{-4q_2\sqrt{\gamma(\boldsymbol{C}\boldsymbol{C}^*)}}.$$

In this case, it is natural to bound the stable rank as

$$\mathrm{srk}(\boldsymbol{C}\boldsymbol{C}^*;\nu) \le \left(\frac{\|\boldsymbol{C}\|_{4\nu}}{\|\boldsymbol{C}\|}\right)^{4\nu}.$$

We have written $\|\cdot\|$ for the spectral norm and $\|\cdot\|_p$ for the Schatten p-norm.

Remark 1.9 (Other Approaches)**.** Identical results hold if we apply the block Krylov method to the other square $\boldsymbol{C}^*\boldsymbol{C}$. It is also possible to work with "odd" Krylov subspaces $K_q(\boldsymbol{C}\boldsymbol{C}^*;\boldsymbol{C}\boldsymbol{B})$ or $K_q(\boldsymbol{C}^*\boldsymbol{C};\boldsymbol{C}^*\boldsymbol{B})$, but the analysis requires some modifications.

Remark 1.10 (Minimum Singular Values)**.** The quantity $\zeta_q(\boldsymbol{C}\boldsymbol{C}^*;\boldsymbol{B})$ gives an estimate for the minimum squared singular value of $\boldsymbol{C}$. It is straightforward to derive results for the estimate using the principles outlined above. We omit the details.

1.9. **Extensions.** The approach in this paper extends easily to the complex setting. The bounds for a complex Gaussian test matrix are similar in spirit, but they are even better than the bounds in the real case. For brevity, we omit the details.

1.10. **Summary of Results.** Theorems 1.5 and 1.7 provide information about the behavior of block Krylov methods with a random test matrix for computing the extreme eigenvalues of a symmetric matrix.

First, the results demonstrate that the depth of the Krylov subspace only needs to increase with the logarithm of the stable rank. For many input matrices, we can achieve satisfactory eigenvalue estimates even when the depth of the block Krylov space is a constant independent of the dimension.

Second, our results indicate that the block Krylov method may accrue significant performance benefits when the block size $\ell \geq 3$. Indeed, the probability that the method fails to achieve a relative error ε at a depth q may be orders of magnitude smaller when $\ell \geq 3$ as compared with $\ell = 1, 2$. Moreover, when the matrix has a spectral gap, the block Krylov method with $\ell \geq 3$ appears to reduce the average relative error twice as fast as when $\ell = 1, 2$.

On contemporary computer architectures, the cost of implementing the block Krylov method for a moderate block size $\ell = 3, 4$ may be similar to the cost when $\ell = 1, 2$. This observation suggests that block Krylov methods may be more suitable for estimating extreme eigenvalues and spectral norms than the simple Krylov method with block size $\ell = 1$.

This paper does not implement block Krylov methods or attempt to verify the error bounds. One main reason is that it takes care to develop algorithms that are both efficient and numerically stable. We leave this investigation for future research.

2. History, Related Work, and Contributions

Krylov subspace methods are a wide class of algorithms that use matrix–vector products ("Krylov information") to compute eigenvalues and eigenvectors and to solve linear systems. These methods are especially appropriate in situations where we can only interact with a matrix through its action on vectors. In this treatment, we only discuss Krylov methods for spectral computations. Some of the basic algorithms in this area are the power method, the inverse power method, subspace iteration, the Lanczos method, the block Lanczos method, and the band Lanczos method. See the books [Par98, BDD+00, Saa11, GVL13] for more background and details.

2.1. **Simple Krylov Methods.** Simple Krylov methods are algorithms based on a Krylov subspace $K_q(\boldsymbol{A}; \boldsymbol{b})$ constructed from a single starting vector $\boldsymbol{b}$. That is, the block size $\ell = 1$.

The power method, which dates to the 19th century, is probably the earliest algorithm that relies on Krylov information to compute eigenvalues and eigenvectors of a symmetric matrix. The power method is degenerate in the sense that it only keeps the highest-order term in the Krylov subspace.

In the late 1940s, Lanczos [Lan50] developed a sophisticated Krylov subspace method for solving the symmetric eigenvalue problem. (More precisely, the Lanczos method uses a three-term recurrence to compute a basis for the Krylov subspace so that the compression of the input matrix to the Krylov subspace is tridiagonal.) In exact arithmetic, the Lanczos estimate of the maximum eigenvalue of a symmetric matrix coincides with $\xi_{\max}(\boldsymbol{A}; \boldsymbol{b}; q)$ for a fixed vector $\boldsymbol{b}$. On the other hand, the Lanczos method has complicated behavior in finite-precision arithmetic.

The first analysis of the Lanczos method with a deterministic starting vector $\boldsymbol{b}$ dates to the work of Lanczos [Lan50]. Kaniel, Paige, and Saad also made major theoretical contributions in the 1970s and 1980s; see [Par98, Saa11] for details and references. In the 1980s, Nemirovsky, Yudin, and Chou showed that Krylov subspace methods are the optimal deterministic algorithms for solving the symmetric eigenvalue problem, assuming we only have access to the matrix via matrix–vector multiplication; see [NY83, Cho87, Nem91].

2.2. **Random Starting Vectors.** Practitioners have often suggested using randomized variants of Krylov subspace methods. That is, the starting vector $\boldsymbol{b}$ is chosen at random. Historically, randomness was just used to avoid the situation where the starting vector $\boldsymbol{b}$ is orthogonal to the leading invariant subspace of the matrix.

Later, deeper justifications for random starting vectors appeared. The first substantive theoretical analysis of a randomized Krylov method appears in Dixon's paper [Dix83] on the power method with a

random starting vector. We believe that this is the first paper to recognize that Krylov methods can be successful without the presence of a spectral gap.

In 1992, Kuczyński & Woźniakowski published an analysis [KW92] of the Lanczos method with a random starting vector. Their work highlighted the benefits of randomization, and it provided a clear explanation of the advantages of using full Krylov information instead of the power method. See the papers [KW94, DCM97, LW98] for further work in this direction.

The recent paper [SAR17] contains lower bounds on the performance of randomized algorithms for the symmetric eigenvalue problem that use Krylov information.

2.3. **Block Krylov Methods.** Block Krylov subspace methods use multiple starting vectors to generate the Krylov subspace, instead of just one. In other words, the algorithms form a Krylov subspace $K_q(\boldsymbol{A}; \boldsymbol{B})$, where $\boldsymbol{B}$ is a matrix. These methods were developed in the late 1960s and 1970s in an effort to resolve multiple eigenvalues more reliably.

The block analog of the power method is called subspace iteration; see the books [Par98, Saa11] for discussion.

There are also block versions of the Lanczos method, which were developed by Cullum & Donath [CD74] and Golub & Underwood [GU77]. (More precisely, the block Lanczos method uses a recurrence to compute a basis for the block Krylov subspace so that the compression of the input matrix to the block Krylov subspace is block tridiagonal.) In exact arithmetic, the block Lanczos estimate of the maximum eigenvalue of a symmetric matrix coincides with $\xi_{\max}(\boldsymbol{A}; \boldsymbol{B}; q)$ for a fixed matrix $\boldsymbol{B}$.

Most of the early work on block Krylov subspace methods focuses on the case where the block size ℓ is small, while the depth q of the Krylov space is moderately large. This leads to significant problems with numerical stability, especially in the case where we use a recurrence to perform orthogonalization. Furthermore, most of the existing analysis of block Krylov methods is deterministic; for example, see [Saa80, LZ15].

2.4. **Randomized Block Krylov Methods.** Over the last decade, randomized block Krylov subspace methods have emerged as a powerful tool for spectral computations on large matrices. These algorithms use a Krylov subspace $K_q(\boldsymbol{A}; \boldsymbol{B})$ generated by a *random* test matrix $\boldsymbol{B}$.

In contrast with earlier block Krylov algorithms, contemporary methods use a much larger block size ℓ and a much smaller depth q. Furthermore, the randomness plays a central role in supporting performance guarantees for the algorithms.

Most of the recent literature concerns the problem of computing a low-rank approximation to a large matrix, rather than estimating eigenvalues or invariant subspaces. Some of the initial work on randomized algorithms for matrix approximation appears in [PRTV98, FKV98, DKM06, MRT11]. Randomized subspace iteration was proposed in [RST09] and further developed in [HMT11]. Randomized block Krylov methods that use the full block Krylov subspace were proposed in the papers [RST09, HMST11]; see also [DIKMI16]. See [HMT11] for more background and history.

There is some theoretical and empirical evidence [HMST11, MM15] that randomized block Krylov methods can produce low-rank matrix approximations with higher accuracy and with less computation than randomized subspace iteration.

2.5. **Analysis of Randomized Block Krylov Methods.** There is a growing body of literature that develops theoretical performance guarantees for randomized block Krylov methods. The papers [RST09, HMT11, WC15, Gu15, MM15] contain theoretical analyses of randomized subspace iteration. The papers [MM15, WZZ15, DIKMI16] contain theoretical analysis of randomized methods that use the full block Krylov space. These works all focus on low-rank matrix approximation.

Remark 2.1 (Contemporary Work)**.** After this paper was completed, we learned about a related contemporary manuscript [DI18]. We hope to include a discussion in a future version of this work.

2.6. **Contributions.** We set out to develop highly refined bounds for the behavior of randomized block Krylov methods that use the full Krylov subspace. Our aim is to present useful and informative results in the spirit of Saad [Saa80], Kucziński & Woźniakowski [KW92], and Halko et al. [HMT11]. Our research has a number of specific benefits over prior work.

- We have shown that randomized block Krylov methods have exceptional performance for matrices with spectral decay. In fact, we can often obtain accurate results even when the block Krylov subspace has *constant* depth. We believe our work provides the first convincing explanation for the remarkable empirical performance [HMST11, MM15] of these algorithms.
- We have obtained detailed information about the role of the block size ℓ. It is possible to estimate the largest k eigenvalues (singular values) reliably by choosing the block size $\ell = k + 2$. For estimating the maximum eigenvalue (singular value), block size $\ell = 3$ is sufficient.
- Our work gives the first results on the performance of randomized block Krylov methods for the symmetric eigenvalue problem. Earlier work only addresses singular value estimates and matrix approximation.
- Our bounds have explicit and modest constants, which gives the bounds some predictive power.

We hope that these results help clarify the benefits of randomized block Krylov methods. We also hope that our work encourages researchers to develop new implementations of these algorithms that fully exploit contemporary computer architectures.

3. The Error in the Block Krylov Subspace Method

In this section, we initiate the proof of Theorems 1.5 and 1.7. Along the way, we establish Proposition 1.4. First, we show how to replace the block Krylov subspace by a simple Krylov subspace (with block size one). Afterward, we develop an explicit representation for the error in the eigenvalue estimate derived from the simple Krylov subspace. Finally, we explain how to construct the simple Krylov subspace so that we preserve the benefits of computing a block Krylov subspace.

The ideas in this section are drawn from several sources. The strategy of reducing a block Krylov subspace to a simple Krylov subspace already appears in [Saa80], but we use a different technique that is adapted from [HMT11]. The kind of analysis we perform for the simple Krylov method is standard; we have closely followed the presentation in [KW92].

3.1. **Simplifications.** Suppose that the input matrix $\boldsymbol{A}$ is a multiple of the identity matrix. From the definitions (1.1), (1.2), and (1.4), it is straightforward to check that the eigenvalue estimate $\xi_{\max}(\boldsymbol{A}; \boldsymbol{\Omega}; q) = \lambda_{\max}(\boldsymbol{A})$ with probability one for each $q \geq 0$. Therefore, we may as well assume that $\boldsymbol{A}$ is not a multiple of the identity.

In view of (1.9), we may also assume that the input matrix is diagonal with weakly decreasing entries:

$$\boldsymbol{A} = \operatorname{diag}(a_1, a_2, \dots, a_n) \quad \text{where} \quad a_1 \geq a_2 \geq \cdots \geq a_n. \tag{3.1}$$

Since $\boldsymbol{A}$ has at least two distinct eigenvalues, we may normalize the extreme eigenvalues of $\boldsymbol{A}$:

$$a_1 = 1 \quad \text{and} \quad a_n = 0. \tag{3.2}$$

The main results are all stated in terms of affine invariant quantities, so we have not lost any generality. These choices help to streamline the proof.

3.2. **Block Krylov Subspaces and Simple Krylov Subspaces.** The first key step in the argument is to reduce the block Krylov subspace to a Krylov subspace with block size one. This idea allows us to avoid any computations involving matrices. To that end, recall that the block Krylov subspace takes the form

$$K_q(\boldsymbol{A}; \boldsymbol{\Omega}) = \operatorname{range}\begin{bmatrix} \boldsymbol{\Omega} & \boldsymbol{A}\boldsymbol{\Omega} & \boldsymbol{A}^2\boldsymbol{\Omega} & \dots & \boldsymbol{A}^q\boldsymbol{\Omega} \end{bmatrix}.$$

In particular, for any vector $\boldsymbol{x} \in \operatorname{range}(\boldsymbol{\Omega})$,

$$K_q(\boldsymbol{A}; \boldsymbol{x}) = \operatorname{range}\begin{bmatrix} \boldsymbol{x} & \boldsymbol{A}\boldsymbol{x} & \boldsymbol{A}^2\boldsymbol{x} & \dots & \boldsymbol{A}^q\boldsymbol{x} \end{bmatrix} \subset K_q(\boldsymbol{A}; \boldsymbol{\Omega}).$$

Later, we will make a careful choice of the vector $\boldsymbol{x}$ so that we do not abandon the benefits of computing the block Krylov subspace.

3.3. **Representation of the Error Using Polynomials.** The next step in the argument is to exploit the close relationship between Krylov subspaces and polynomial filtering to obtain an explicit representation of the error in the eigenvalue estimate. Using (1.3), we may rewrite the last display in the form

$$K_q(\boldsymbol{A};\boldsymbol{x}) = \operatorname{span}\{\varphi(\boldsymbol{A})\boldsymbol{x} : \varphi \in \mathscr{P}_q\} \subset K_q(\boldsymbol{A};\boldsymbol{\Omega}).$$

As a consequence of this containment,

$$\xi_{\max}(\boldsymbol{A};\boldsymbol{\Omega};q) = \max_{\boldsymbol{v}\in K_q(\boldsymbol{A};\boldsymbol{\Omega})} \frac{\boldsymbol{v}^*\boldsymbol{A}\boldsymbol{v}}{\boldsymbol{v}^*\boldsymbol{v}} \geq \max_{\boldsymbol{v}\in K_q(\boldsymbol{A};\boldsymbol{x})} \frac{\boldsymbol{v}^*\boldsymbol{A}\boldsymbol{v}}{\boldsymbol{v}^*\boldsymbol{v}} = \max_{\varphi\in\mathscr{P}_q} \frac{(\varphi(\boldsymbol{A})\boldsymbol{x})^*\boldsymbol{A}(\varphi(\boldsymbol{A})\boldsymbol{x})}{(\varphi(\boldsymbol{A})\boldsymbol{x})^*(\varphi(\boldsymbol{A})\boldsymbol{x})}.$$

Owing to the normalization (3.2), the relative error (1.12) in the eigenvalue estimate satisfies

$$\operatorname{err}(\xi_{\max}(\boldsymbol{A};\boldsymbol{\Omega};q)) = 1 - \xi_{\max}(\boldsymbol{A};\boldsymbol{\Omega};q) \leq \min_{\varphi\in\mathscr{P}_q} \frac{(\varphi(\boldsymbol{A})\boldsymbol{x})^*(\mathbf{I}-\boldsymbol{A})(\varphi(\boldsymbol{A})\boldsymbol{x})}{(\varphi(\boldsymbol{A})\boldsymbol{x})^*(\varphi(\boldsymbol{A})\boldsymbol{x})}.$$

The fraction is invariant under scaling of the polynomial φ, so we may normalize $\varphi(1) = 1$. With the definition $\mathscr{P}_q(1) := \{\varphi \in \mathscr{P}_q : \varphi(1) = 1\}$, we arrive at

$$\operatorname{err}(\xi_{\max}(\boldsymbol{A};\boldsymbol{\Omega};q)) \leq \inf_{\varphi\in\mathscr{P}_q(1)} \frac{(\varphi(\boldsymbol{A})\boldsymbol{x})^*(\mathbf{I}-\boldsymbol{A})(\varphi(\boldsymbol{A})\boldsymbol{x})}{(\varphi(\boldsymbol{A})\boldsymbol{x})^*(\varphi(\boldsymbol{A})\boldsymbol{x})}.$$

We remark that this inequality becomes an equality in the case where the block size $\ell = 1$.

Invoke (3.1) to rewrite this bound in terms of the eigenvalues of $\boldsymbol{A}$:

$$\operatorname{err}(\xi_{\max}(\boldsymbol{A};\boldsymbol{\Omega};q)) \leq \inf_{\varphi\in\mathscr{P}_q(1)} \frac{\sum_{i=1}^n X_i^2\varphi^2(a_i)(1-a_i)}{\sum_{i=1}^n X_i^2\varphi^2(a_i)} = \inf_{\varphi\in\mathscr{P}_q(1)} \frac{\sum_{i>1} X_i^2\varphi^2(a_i)(1-a_i)}{X_1^2 + \sum_{i>1} X_i^2\varphi^2(a_i)}. \tag{3.3}$$

We have introduced the coordinates of the vector $\boldsymbol{x} = (X_1,\dots,X_n)$, and we have used normalization (3.2) to simplify the expression.

Remark 3.1 (Multiplicity)**.** When the maximum eigenvalue has multiplicity greater than one, we can refine this argument to obtain stronger results.

3.4. **When the Matrix has Few Distinct Eigenvalues.** Suppose that the distinct eigenvalues of the input matrix $\boldsymbol{A}$ are $\mu_1 = 1$ and $\mu_2,\dots,\mu_r$ where $r \leq q+1$. Consider the polynomial

$$\varphi_0(s) = \prod_{i=2}^r \frac{s-\mu_i}{1-\mu_i} \in \mathscr{P}_{r-1}(1) \subset \mathscr{P}_q(1).$$

This polynomial annihilates each point in the spectrum, except for μ_1. Setting $\varphi = \varphi_0$ in (3.3), we discover that $\operatorname{err}(\xi_{\max}(\boldsymbol{A};\boldsymbol{\Omega};q)) = 0$ with probability one. This fact appears as Proposition 1.4.

3.5. **Choosing the Simple Krylov Space.** The next step in the argument is to select a particular vector $\boldsymbol{x} \in \operatorname{range}(\boldsymbol{\Omega})$. Let $\boldsymbol{\omega}_1^* \in \mathbb{R}^\ell$ denote the first row of the matrix $\boldsymbol{\Omega}$. Set

$$\boldsymbol{x} = \frac{\boldsymbol{\Omega}\boldsymbol{\omega}_1}{\|\boldsymbol{\omega}_1\|} \in \operatorname{range}(\boldsymbol{\Omega}).$$

The rows of $\boldsymbol{\Omega}$ are statistically independent standard normal vectors, which are rotationally invariant. It follows that the entries of $\boldsymbol{x}$ are also independent. Moreover,

$$X_1 \sim \textsc{chi}(\ell) \quad \text{and} \quad X_i \sim \textsc{normal}(0,1) \quad \text{for } i > 1. \tag{3.4}$$

We write $\textsc{chi}(\ell)$ for the chi distribution with ℓ degrees of freedom. This choice of $\boldsymbol{x}$ ensures that X_1^2 is large relative to the other X_i^2. The argument here is inspired by the analysis in Halko et al. [HMT11, Secs. 9, 10].

4. Probabilistic Bounds for the Error

This section contains the proof of the probability bounds that appear in Theorem 1.5 and 1.7. The argument is based on the bound (3.3) for the error and the distributional properties (3.4) of the random vector $\boldsymbol{x}$.

The proof is inspired by the argument in Kuczyński & Woźniakowski [KW92], but our approach is technically easier. Indeed, they work with a random vector that is uniformly distributed on the Euclidean unit sphere, which leads to a difficult multivariate integration. In contrast, our random vector $\boldsymbol{x}$ has independent entries, which means that we only have to compute one-dimensional integrals.

4.1. **A Bound for the Probability.** Let $\varepsilon \in (0,1)$ be an error tolerance. Our goal is to control the probability P_ε that the relative error in the eigenvalue estimate (1.4) is at least ε. In other words, we wish to bound

$$P_\varepsilon := \mathbb{P}\left\{\operatorname{err}(\xi_{\max}(\boldsymbol{A};\boldsymbol{\Omega};q)) \ge \varepsilon\right\}$$

In view of the upper bound (3.3) for the relative error, we obtain the estimate

$$P_\varepsilon \le \mathbb{P}\left\{ \inf_{\varphi\in\mathscr{P}_q(1)} \frac{\sum_{i>1} X_i^2 \varphi^2(a_i)(1-a_i)}{X_1^2 + \sum_{i>1} X_i^2 \varphi^2(a_i)} \ge \varepsilon \right\}.$$

Fix a polynomial $\varphi \in \mathscr{P}_q(1)$ to be determined later. Then rearrange the inequality in the event:

$$\begin{aligned} P_\varepsilon &\le \mathbb{P}\left\{\textstyle\sum_{i>1} X_i^2\varphi^2(a_i)(1-a_i) \ge \varepsilon X_1^2 + \varepsilon\sum_{i>1} X_i^2\varphi^2(a_i)\right\} \\ &= \mathbb{P}\left\{-\varepsilon X_1^2 + \textstyle\sum_{i>1} X_i^2\varphi^2(a_i)(1-\varepsilon-a_i) \ge 0\right\} \\ &\le \mathbb{P}\left\{-\varepsilon X_1^2 + \textstyle\sum_{a_i<1-\varepsilon} X_i^2\varphi^2(a_i)(1-\varepsilon-a_i) \ge 0\right\} \\ &=: \mathbb{P}\left\{-\varepsilon X_1^2 + \textstyle\sum_{i\in I} c_i X_i^2 \ge 0\right\}. \end{aligned} \tag{4.1}$$

To reach the second line, we have dropped the nonpositive terms in the sum. Then we introduced the compact notation

$$I := \{i : a_i < 1-\varepsilon\} \quad\text{and}\quad c_i := \varphi^2(a_i)(1-\varepsilon-a_i) > 0.$$

To continue the argument, we apply some elementary notions from the theory of concentration of measure.

4.2. **The Laplace Transform Argument.** We invoke the Laplace transform method to convert the probability bound into an expectation bound. Introduce a parameter $\theta > 0$, to be chosen later. Continuing from (4.1), we write the probability as the expectation of an indicator function:

$$\begin{aligned} P_\varepsilon &\le \mathbb{E}\,\mathbb{1}\left\{-\varepsilon X_1^2 + \textstyle\sum_{i\in I} c_i X_i^2 \ge 0\right\} \\ &\le \mathbb{E}\exp\left(-\theta\varepsilon X_1^2 + \textstyle\sum_{i\in I}\theta c_i X_i^2\right). \end{aligned}$$

To reach the second line, we bound the indicator $\mathbb{1}\{s \ge 0\}$ above by the function $s \mapsto \mathrm{e}^{\theta s}$. Write the exponential as a product, and invoke the independence of the family $\{X_i\}$ to obtain

$$P_\varepsilon \le \mathbb{E}\left[\mathrm{e}^{-\theta\varepsilon X_1^2}\prod_{i\in I}\mathrm{e}^{\theta c_i X_i^2}\right] = \left(\mathbb{E}\,\mathrm{e}^{-\theta\varepsilon X_1^2}\right)\left(\prod_{i\in I}\mathbb{E}\,\mathrm{e}^{\theta c_i X_i^2}\right).$$

The distributional property (3.4) implies that X_1^2 is a chi-squared variable with ℓ degrees of freedom, while X_i^2 is a chi-squared variable with one degree of freedom. Computing the remaining expectations is a standard exercise, which results in the bound

$$P_\varepsilon \le (1+2\theta\varepsilon)^{-\ell/2}\left(\textstyle\prod_{i\in I}(1-2\theta c_i)\right)^{-1/2} \quad \text{when } \theta < (2c_i)^{-1} \text{ for each } i \in I.$$

We make a coarse estimate to arrive at

$$P_\varepsilon \le (1+2\theta\varepsilon)^{-\ell/2}\left(1-2\theta\textstyle\sum_{i\in I} c_i\right)^{-1/2} \quad \text{when } \theta < \left(2\sum_{i\in I} c_i\right)^{-1}.$$

The last bound follows from repeated application of the numerical inequality $(1-s)(1-t) \ge 1-(s+t)$, which is valid when $st \ge 0$.

Next, we must identify a suitable value for θ. It is possible to minimize the probability bound with respect to θ, but it is more expedient to select $\theta^{-1} = 4\sum_{i\in I} c_i$. This choice yields

$$P_\varepsilon \le \sqrt{2}\left(1 + \frac{\varepsilon/2}{\sum_{i\in I} c_i}\right)^{-\ell/2} = \sqrt{2}\left[1 + \frac{\varepsilon/2}{\sum_{a_i<1-\varepsilon} \varphi^2(a_i)(1-\varepsilon-a_i)}\right]^{-\ell/2}. \tag{4.2}$$

It remains to choose a good polynomial φ. Constructions and analysis of polynomials appear in Appendix A.2, and we simply present the information that is relevant for each argument.

4.3. **Probability Bound without a Spectral Gap.** In this section, we establish the probability bound that appears in Theorem 1.5. Introduce the quantity $\beta := 1-\varepsilon$. For a partition $q = q_1 + q_2$, consider the polynomial

$$\varphi_1(s) := \frac{s^{q_1} U_{2q_2}\big(\sqrt{s/\beta}\big)}{U_{2q_2}\big(\sqrt{1/\beta}\big)} \in \mathscr{P}_q(1).$$

According to (A.5) and (A.6), this polynomial satisfies

$$\varphi_1^2(s)(\beta - s) \le \frac{4\varepsilon s^{2q_1}\delta^{2q_2+1}}{(1-\delta^{2q_2+1})^2} \quad \text{for } 0 \le s \le \beta \quad \text{where} \quad \delta \le \mathrm{e}^{-2\sqrt{\varepsilon}}. \tag{4.3}$$

See Appendix A.2.5 for a more detailed discussion.

Using these facts, we may estimate

$$\begin{aligned}\sum_{a_i<1-\varepsilon} \varphi_1^2(a_i)(1-\varepsilon-a_i) &= \sum_{a_i<\beta} \varphi_1^2(a_i)(\beta - a_i) \\ &\le 4\varepsilon\left(\sum_{a_i<\beta} a_i^{2q_1}\right)\delta^{2q_2+1}\big(1-\delta^{2q_2+1}\big)^{-2} \\ &\le 4\varepsilon\,\mathrm{srk}(q_1)\,\delta^{2q_2+1}\big(1-\delta^{2q_2+1}\big)^{-2}.\end{aligned}$$

In the last step, we bound the sum in terms of the stable rank (1.11). We rely on the normalization (3.2) to recognize the stable rank.

Select $\varphi = \varphi_1$ in our probability bound (4.2). Substitute in the last display to arrive at

$$P_\varepsilon \le \sqrt{2}\left[1 + \frac{\big(1-\delta^{2q_2+1}\big)^2}{8\,\mathrm{srk}\,\delta^{2q_2+1}}\right]^{-\ell/2}.$$

We can develop a lower bound on the bracket as follows.

$$1 + \frac{\big(1-\delta^{2q_2+1}\big)^2}{8\,\mathrm{srk}(q_1)\delta^{2q_2+1}} \ge 1 + \frac{1-2\delta^{2q_2+1}}{8\,\mathrm{srk}(q_1)\,\delta^{2q_2+1}} \ge \frac{1}{8\,\mathrm{srk}(q_1)\,\delta^{2q_2+1}}.$$

The last inequality follows from the fact that $\mathrm{srk}(q_1) \ge 1$ because $\boldsymbol{A}$ is not a multiple of the identity. Combining the last two displays, we obtain

$$P_\varepsilon \le \sqrt{2}\left[8\,\mathrm{srk}(q_1)\,\delta^{2q_2+1}\right]^{\ell/2} \le \sqrt{2}\left[8\,\mathrm{srk}(q_1)\cdot \mathrm{e}^{-2(2q_2+1)\sqrt{\varepsilon}}\right]^{\ell/2}. \tag{4.4}$$

The final relation is a consequence of the bound for δ in (4.3). This is the required statement.

4.4. **Probability Bound with a Spectral Gap.** Now, we establish the probability bound that appears in Theorem 1.7. Recall that the spectral gap γ is defined in (1.10). This time, set $\beta := 1-\gamma$. For a partition $q = q_1 + q_2$, construct the polynomial

$$\varphi_2(s) := \frac{s^{q_1} T_{q_2}((2/\beta)s - 1)}{T_{q_2}((2/\beta) - 1)} \in \mathscr{P}_q(1). \tag{4.5}$$

According to (A.5) and (A.6), this polynomial satisfies

$$\varphi_2^2(s) \le 4 s^{2q_1}\mathrm{e}^{-4q_2\sqrt{\gamma}} \quad \text{for } 0 \le s \le \beta. \tag{4.6}$$

See Appendix A.2.3 for more details.

Using these facts, we calculate that

$$\begin{aligned}\sum_{a_i\le 1-\varepsilon}\varphi_2^2(a_i)(1-\varepsilon-a_i) &\le \sum_{a_i\le 1-\gamma}\varphi_2^2(a_i)\\ &\le 4\left(\sum_{a_i\le 1-\gamma} a_i^{2q_1}\right)\mathrm{e}^{-4q_2\sqrt{\gamma}} \le 4\,\mathrm{srk}(q_1)\cdot\mathrm{e}^{-4q_2\sqrt{\gamma}}.\end{aligned}$$

In the first step, we apply $1-\varepsilon-a_i\le 1$, and then we observe that $1-\gamma$ is the first eigenvalue smaller than one. Last, we use (3.2) to identify the stable rank (1.11).

Instantiate the probability estimate (4.2) with $\varphi=\varphi_2$, and substitute in the last display to obtain

$$P_\varepsilon \le \sqrt{2}\left[1+\frac{\varepsilon}{8\,\mathrm{srk}(q_1)\cdot\mathrm{e}^{-4q_2\sqrt{\gamma}}}\right]^{-\ell/2} \le \sqrt{2}\left[\frac{8\,\mathrm{srk}(q_1)}{\varepsilon}\cdot\mathrm{e}^{-4q_2\sqrt{\gamma}}\right]^{\ell/2}. \tag{4.7}$$

We have used the numerical inequality $(1+1/s)^{-1}\le s$, valid for $s>0$. This is the advertised result.

5. A Bound for the Expected Error without a Spectral Gap

In this section, we establish the expectation bound that appears in Theorem 1.5. To obtain this result, we simply integrate the probability bound (4.4). Surprisingly, this approach appears to be more accurate than a direct computation of the expected error. This insight yields a better expected error bound than the one obtained in [KW92] by a direct argument.

5.1. **Computing the Expectation.**

We may express the expectation of the relative error as an integral:

$$E := \mathbb{E}\,\mathrm{err}(\xi_{\max}(\boldsymbol{A};\boldsymbol{\Omega};q)) = \int_0^1 \mathbb{P}\left\{\mathrm{err}(\xi_{\max}(\boldsymbol{A};\boldsymbol{\Omega};q))\ge\varepsilon\right\}\mathrm{d}\varepsilon = \int_0^1 P_\varepsilon\,\mathrm{d}\varepsilon.$$

The limits of the integral follow from the fact that the relative error falls in the interval $[0,1]$. We split the integral at a value $c>0$, to be determined later. Then make the estimates

$$E \le c + \int_c^1 P_\varepsilon\,\mathrm{d}\varepsilon \le c + \sqrt{2}\,(8\,\mathrm{srk}(q_1))^{\ell/2}\int_c^\infty \mathrm{e}^{-(2q_2+1)\ell\sqrt{\varepsilon}}\,\mathrm{d}\varepsilon.$$

To obtain the first inequality, we use the trivial bound $P_\varepsilon\le 1$. The second inequality is a consequence of (4.4). The remaining integral can be calculated by changing the variable and integrating by parts. Indeed,

$$\int_c^\infty \mathrm{e}^{-p\sqrt{\varepsilon}}\,\mathrm{d}\varepsilon = 2\left(\frac{\sqrt{c}}{p}+\frac{1}{p^2}\right)\mathrm{e}^{-p\sqrt{c}} \quad \text{for } p>0.$$

Together, the last two displays yield

$$E \le c + 2\sqrt{2}(8\,\mathrm{srk}(q_1))^{\ell/2}\left(\frac{\sqrt{c}}{(2q_2+1)\ell}+\frac{1}{(2q_2+1)^2\ell^2}\right)\mathrm{e}^{-(2q_2+1)\ell\sqrt{c}}.$$

Now, select the (optimal) value

$$c = \left(\frac{\ell^{-1}\log 2+\log(8\,\mathrm{srk}(q_1))}{2(2q_2+1)}\right)^2.$$

Combine the last two displays to reach

$$E \le \left(\frac{\ell^{-1}(2+\log 2)+\log(8\,\mathrm{srk}(q_1))}{2(2q_2+1)}\right)^2$$

Bound the numerical constant by 2.70 to complete the proof.

6. A Bound for the Expected Error with a Spectral Gap

In this section, we prove the expectation bounds for the relative error that appear in Theorem 1.7. In this case, we achieve better results by a direct computation, rather than by integrating the probability bound (4.7). These arguments are inspired by the approach in [KW92], but our computations are technically easier because the random vector $\boldsymbol{x}$ has independent entries.

6.1. **Form of the Expected Error.** Fix a polynomial $\varphi \in \mathscr{P}_q(1)$. Take the expectation of the error bound (3.3):

$$\mathbb{E}\,\mathrm{err}(\xi_{\max}(\boldsymbol{A};\boldsymbol{\Omega};q)) \le \mathbb{E}\left[\frac{\sum_{i>1} X_i^2\varphi^2(a_i)(1-a_i)}{X_1^2+\sum_{i>1} X_i^2\varphi^2(a_i)}\right]. \tag{6.1}$$

By independence, we may compute the expectation with respect to X_1, holding X_i fixed for $i>1$. The computation of this integral depends on the block size ℓ.

6.2. **Error Bound for Block Size $\ell \ge 3$.** We begin with the case $\ell \ge 3$, which is technically simplest. Use the fact that $X_1 \sim \mathrm{CHI}(\ell)$ to compute that

$$\mathbb{E}\left[\frac{1}{X_1^2+c}\right] = \frac{1}{2}\mathrm{e}^{c/2}\int_1^\infty s^{-\ell/2}\mathrm{e}^{-cs/2}\,\mathrm{d}s \le \frac{1}{(\ell-2)+c} \quad \text{for } \ell \ge 2 \text{ and } c \ge 0.$$

The first relation depends on a standard identity for the partial gamma function [OLBC10, Sec. 8.6.4]. The second relation is a classic bound for the exponential integral due to Hopf [Hop34, p. 26]; see [Gau59] or [AS64, Sec. 5.1.19].

Together, the last two displays imply that

$$\mathbb{E}\,\mathrm{err}(\xi_{\max}(\boldsymbol{A};\boldsymbol{\Omega};q)) \le \mathbb{E}\left[\frac{\sum_{i>1} X_i^2\varphi^2(a_i)(1-a_i)}{(\ell-2)+\sum_{i>1} X_i^2\varphi^2(a_i)}\right].$$

Set the parameter $\beta = 1-\gamma$. By the definition (1.10) of the spectral gap γ, each eigenvalue of $\boldsymbol{A}$ that exceeds β equals the maximum eigenvalue $a_1 = 1$. Therefore,

$$\begin{aligned}\mathbb{E}\,\mathrm{err}(\xi_{\max}(\boldsymbol{A};\boldsymbol{\Omega})) &\le \mathbb{E}\left[\frac{\sum_{a_i\le\beta} X_i^2\varphi^2(a_i)(1-a_i)}{(\ell-2)+\sum_{i>1} X_i^2\varphi^2(a_i)}\right] \\ &\le \mathbb{E}\left[\frac{1}{\ell-2}\sum_{a_i\le\beta} X_i^2\varphi^2(a_i)\right] = \frac{1}{\ell-2}\sum_{a_i\le\beta}\varphi^2(a_i).\end{aligned}$$

To reach the second inequality, use (3.2) to bound $1-a_i \le 1$, and drop the second term from the denominator. Last, we compute the expectation using the fact that $X_i \sim \mathrm{NORMAL}(0,1)$ for each $i>1$.

Introduce the polynomial $\varphi = \varphi_2$ from (4.5) into the last display. Using the estimate (4.6), we arrive at

$$\mathbb{E}\,\mathrm{err}(\xi_{\max}(\boldsymbol{A};\boldsymbol{\Omega})) \le \frac{4\,\mathrm{srk}(q_1)}{\ell-2}\cdot\mathrm{e}^{-4q_2\sqrt{\gamma}}.$$

This is the advertised result for $\ell = 3$ in Theorem 1.7.

6.3. **Error Bound for Block Size $\ell = 2$.** Now, assume that the block size $\ell = 2$. We begin with the bound (6.1):

$$\begin{aligned}\mathbb{E}\,\mathrm{err}(\xi_{\max}(\boldsymbol{A};\boldsymbol{\Omega})) &\le \mathbb{E}\left[\frac{\sum_{i>1} X_i^2\varphi^2(a_i)(1-a_i)}{X_1^2+\sum_{i>1} X_i^2\varphi^2(a_i)}\right] \\ &\le \mathbb{E}\left[\frac{\left(\sum_{i>1} X_i^2\varphi^2(a_i)\right)^{1/2}\left(\sum_{i>1} X_i^2\varphi^2(a_i)(1-a_i)^2\right)^{1/2}}{X_1^2+\sum_{i>1} X_i^2\varphi^2(a_i)}\right].\end{aligned} \tag{6.2}$$

To reach the second line, apply the Cauchy–Schwarz inequality to the sum in the numerator. The reason for this extra maneuver will become clear in a moment.

Next, we use statistical independence to take the expectation with respect to X_1. Since $X_1 \sim \mathrm{CHI}(2)$, it holds that

$$\mathbb{E}\left[\frac{1}{X_1^2+c}\right] = \frac{1}{2}\mathrm{e}^{c/2}\int_1^\infty s^{-1}\mathrm{e}^{-cs/2}\,\mathrm{d}s \le \frac{1}{2}\log\left(1+\frac{2}{c}\right) \le \frac{1}{\sqrt{c}}\wedge\frac{1}{c} \quad \text{for } c>0.$$

The first relation is [OLBC10, Sec. 8.6.4]. The second relation is a classical bound for the exponential integral; see [AS64, Sec. 5.1.20]. The third relation is a numerical inequality.

With $c = \sum_{i>1} X_i^2 \varphi^2(a_i)$, the last two displays imply that

$$\mathbb{E}\,\mathrm{err}(\xi_{\max}(\boldsymbol{A};\boldsymbol{\Omega})) \le \mathbb{E}\left[\frac{\sum_{i>1} X_i^2 \varphi^2(a_i)(1-a_i)^2}{1 \vee \left(\sum_{i>1} X_i^2 \varphi^2(a_i)\right)}\right]^{1/2}.$$

The rest of the argument follows the same path as in the case $\ell \ge 3$.

As before, set the parameter $\beta = 1-\gamma$. Since each eigenvalue of the input matrix that exceeds β equals the maximum eigenvalue $a_1 = 1$, we have

$$\mathbb{E}\,\mathrm{err}(\xi_{\max}(\boldsymbol{A};\boldsymbol{\Omega})) \le \mathbb{E}\left[\sum\nolimits_{a_i \le \beta} X_i^2 \varphi^2(a_i)(1-a_i)^2\right]^{1/2} \le \left[\sum\nolimits_{a_i \le \beta} \varphi^2(a_i)\right]^{1/2}.$$

This calculation just combines the spectral gap argument for $\ell = 3$ with the bounds we used in the spectral-gap-free argument for $\ell = 2$. Select the polynomial $\varphi = \varphi_2$ from (4.5), and invoke the bound (4.6) to arrive at

$$\mathbb{E}\,\mathrm{err}(\xi_{\max}(\boldsymbol{A};\boldsymbol{\Omega})) \le \sqrt{4\,\mathrm{srk}(q_1)} \cdot \mathrm{e}^{-2q_2\sqrt{\gamma}}.$$

This is the desired outcome for block size $\ell = 2$ in Theorem 1.7.

6.4. **Error Bound for Block Size $\ell = 1$.** Finally, we outline the changes required to handle the case where the block size $\ell = 1$. As in the case $\ell = 2$, we start with the bound (6.2). To continue, use the fact that $X_1 \sim \mathrm{CHI}(1)$ to see that

$$\mathbb{E}\left[\frac{1}{X_1^2 + c}\right] = \frac{1}{2\sqrt{c}}\mathrm{e}^{-c/2}\int_{\sqrt{c}}^{\infty} \mathrm{e}^{-s^2/2}\,\mathrm{d}s \le \frac{\sqrt{\pi/2}}{\sqrt{c}} \wedge \frac{1}{c}.$$

The first relation is [OLBC10, Sec. 8.6.4], followed by a change of variable. The second relation depends on two well-known tail bounds for a standard normal random variable. The remaining steps in the proof are identical with the case $\ell = 2$. We omit the details, except to note that the numerical constant $(\pi/2)^{1/4} \le 1.13$. The final statement appears in Theorem 1.7.

Appendix A. Complements

This appendix contains some supplementary material. In Appendix A.1, we provide further justification for requiring that the test matrix has a uniformly random range. In Appendix A.2, we sketch the properties of Chebyshev polynomials that we use in the body of the paper.

A.1. **Rotationally Invariant Distributions.** This series of papers studies problems about estimating spectral properties of a matrix from randomized block Krylov information. In particular, we wish to obtain probabilistic upper bounds on the error in these spectral estimates for any input matrix. This section contains a general argument that explains why we ought to use a random test matrix with a rotationally invariant range in these applications.

Proposition A.1 (Uniformly Random Range)**.** *Consider any bivariate function $f : \mathbb{R}^{n\times n} \times \mathbb{R}^{n\times \ell} \to \mathbb{R}$ that is orthogonally invariant:*

$$f(\boldsymbol{A};\boldsymbol{U}\boldsymbol{B}) = f(\boldsymbol{U}\boldsymbol{A}\boldsymbol{U}^*;\boldsymbol{B}) \quad \text{for each orthogonal } \boldsymbol{U} \in \mathbb{R}^{n\times n}.$$

Fix a symmetric $n \times n$ matrix $\boldsymbol{\Lambda}$, and consider the orthogonal orbit

$$\mathscr{A} := \mathscr{A}(\boldsymbol{\Lambda}) := \left\{\boldsymbol{U}\boldsymbol{\Lambda}\boldsymbol{U}^* : \boldsymbol{U} \in \mathbb{R}^{n\times n} \text{ orthogonal}\right\}.$$

Let $\boldsymbol{B} \in \mathbb{R}^{n\times \ell}$ be a random matrix. Let $\boldsymbol{V} \in \mathbb{R}^{n\times n}$ be a uniformly random orthogonal matrix, drawn independently from $\boldsymbol{B}$. Then

$$\max_{\boldsymbol{A}\in\mathscr{A}} \mathbb{E}_{\boldsymbol{V},\boldsymbol{B}}\, f(\boldsymbol{A};\boldsymbol{V}\boldsymbol{B}) \le \max_{\boldsymbol{A}\in\mathscr{A}} \mathbb{E}_{\boldsymbol{B}}\, f(\boldsymbol{A};\boldsymbol{B}).$$

Proof. By rotation invariance of f,

$$\begin{aligned}\max_{\boldsymbol{A}\in\mathscr{A}} \mathbb{E}_{\boldsymbol{V},\boldsymbol{B}}\, f(\boldsymbol{A};\boldsymbol{V}\boldsymbol{B}) &= \max_{\boldsymbol{A}\in\mathscr{A}} \mathbb{E}_{\boldsymbol{V},\boldsymbol{B}}\, f(\boldsymbol{V}^*\boldsymbol{A}\boldsymbol{V};\boldsymbol{B}) \\ &\le \mathbb{E}_{\boldsymbol{V}} \max_{\boldsymbol{A}\in\mathscr{A}} \mathbb{E}_{\boldsymbol{B}}\, f(\boldsymbol{V}^*\boldsymbol{A}\boldsymbol{V};\boldsymbol{B}) \\ &= \max_{\boldsymbol{A}\in\mathscr{A}} \mathbb{E}_{\boldsymbol{B}}\, f(\boldsymbol{A};\boldsymbol{B}).\end{aligned}$$

The inequality is Jensen's, and the last identity follows from the definition of the class $\mathscr{A}$. □

As a specific application, we study the problem of estimating the maximum eigenvalue of the worst matrix with eigenvalue spectrum $\boldsymbol{\Lambda}$ using the block Krylov method. Define the rotationally invariant[1] function $f(\boldsymbol{A};\boldsymbol{B}) = \operatorname{err}(\xi_{\max}(\boldsymbol{A};\boldsymbol{B};q))$. Proposition A.1 states that

$$\max_{\boldsymbol{A}\in\mathscr{A}} \mathbb{E}\operatorname{err}(\xi_{\max}(\boldsymbol{A};\boldsymbol{U}\boldsymbol{B};q)) \le \max_{\boldsymbol{A}\in\mathscr{A}} \mathbb{E}\operatorname{err}(\xi_{\max}(\boldsymbol{A};\boldsymbol{B};q)).$$

That is, the symmetrized test matrix $\boldsymbol{U}\boldsymbol{B}$ is better than the test matrix $\boldsymbol{B}$ if we want to minimize the worst-case expectation of the error; the same kind of bound holds for tail probabilities. We surmise that the test matrix $\boldsymbol{B}$ should have a uniformly random range. Moreover, because of (1.6), we can select *any* distribution with uniformly random range, such as the standard normal matrix $\boldsymbol{\Omega}$.

A.2. **Chebyshev Polynomials.** The computations in this paper depend on the properties of some polynomials derived from the Chebyshev polynomials. This appendix collects the required results. We have drawn some of this material from the paper [KW92] by Kuczyński & Woźniakowski. For general information about Chebyshev polynomials, we refer the reader to [AS64, OLBC10].

A.2.1. *Chebyshev Polynomials of the First Kind.* We can define the Chebyshev polynomials of the first kind via the formula

$$T_p(s) := \frac{1}{2}\left[\left(s+\sqrt{s^2-1}\right)^p + \left(s-\sqrt{s^2-1}\right)^p\right] \quad \text{for } s\in\mathbb{R} \text{ and } p\in\mathbb{Z}_+. \tag{A.1}$$

Using the binomial theorem, it is easy to check that this expression coincides with a polynomial of degree p with real coefficients.

We require two properties of the Chebyshev polynomial T_p. First, it satisfies a uniform bound on the unit interval:

$$\left|T_p(s)\right| \le 1 \quad \text{for } |s|\le 1. \tag{A.2}$$

This result is an immediate consequence of the representation

$$T_p(s) = \cos\left(p\cos^{-1}(s)\right) \quad \text{for } |s|\le 1.$$

The latter formula follows from (A.1) after we apply de Moivre's theorem for complex exponentiation. Second, the polynomial grows quickly outside of the unit interval:

$$T_p\left(\frac{1+s}{1-s}\right) \ge \frac{1}{2}\left(\frac{1+\sqrt{s}}{1-\sqrt{s}}\right)^p \quad \text{for } 0\le s<1. \tag{A.3}$$

This estimate is a direct consequence of the definition (A.1).

[1] The relative error (1.12) is orthogonally invariant because the eigenvalues of $\boldsymbol{A}$ are rotationally invariant and (1.7) states that the eigenvalue estimate (1.4) is also rotationally invariant.

A.2.2. *The Attenuation Factor.* Let $\beta \in [0,1]$ be a parameter, and define the quantity

$$\delta := \delta(\beta) = \frac{1-\sqrt{1-\beta}}{1+\sqrt{1-\beta}}. \tag{A.4}$$

This definition is closely connected with the growth properties of the T_p. We can bound the attenuation factor in two ways:

$$\delta \le \mathrm{e}^{-2\sqrt{1-\beta}} \quad \text{and} \quad \delta \le \beta \cdot 2^{-2\sqrt{1-\beta}}. \tag{A.5}$$

These numerical inequalities can be justified using basic calculus. The first is very accurate for $\beta \approx 1$, while the second is better across the full range $\beta \in [0,1]$.

A.2.3. *First Polynomial Construction.* Choose a nonnegative integer partition $q = q_1 + q_2$. Consider the polynomial

$$\varphi_{\beta,q_1,q_2}(s) := \frac{s^{q_1} T_{q_2}((2/\beta)s-1)}{T_{q_2}((2/\beta)-1)} \quad \text{for } s \in \mathbb{R}.$$

The polynomial $\varphi_{\beta,q_1,q_2} \in \mathscr{P}_q(1)$ because it has degree q and it takes the value one at $s=1$. It holds that

$$\varphi^2_{\beta,q_1,q_2}(s) \le \frac{s^{2q_1}}{T^2_{q_2}((2/\beta)-1)} \le 4s^{2q_1}\left(\frac{1-\sqrt{1-\beta}}{1+\sqrt{1-\beta}}\right)^{2q_2} = 4s^{2q_1}\delta^{2q_2} \quad \text{for } 0 \le s \le \beta. \tag{A.6}$$

The first inequality follows from (A.2), and the second follows from (A.3). Last, we instate the definition (A.4).

Remark A.2 (The Monomial)**.** In contrast to [KW92] and other prior work, we use products of Chebyshev polynomials with low-degree monomials. This seemingly minor change leads to better results, phrased in terms of the stable rank, rather than the ambient dimension.

A.2.4. *Chebyshev Polynomials of the Second Kind.* We can define the Chebyshev polynomials of the second kind via the formula

$$U_p(s) := \frac{1}{2\sqrt{s^2-1}}\left[\left(s+\sqrt{s^2-1}\right)^{p+1} - \left(s-\sqrt{s^2-1}\right)^{p+1}\right] \quad \text{for } s \in \mathbb{R} \text{ and } p \in \mathbb{Z}_+. \tag{A.7}$$

Using the binomial theorem, it is easy to check that this expression coincides with a polynomial of degree p with real coefficients. Moreover, when p is an even number, the polynomial U_p is an even function.

We require two properties of the Chebyshev polynomial U_p. First, it satisfies a weighted uniform bound on the unit interval:

$$\left|\sqrt{1-s^2}\,U_p(s)\right| \le 1 \quad \text{for } |s| \le 1. \tag{A.8}$$

This result is an immediate consequence of the representation

$$U_p(s) = \frac{\sin\left((p+1)\cos^{-1}(s)\right)}{\sqrt{1-s^2}} \quad \text{for } |s| \le 1.$$

The latter formula follows from (A.7) after we apply de Moivre's theorem for complex exponentiation. Second, we can evaluate the polynomial at a specific point:

$$U^2_{2p}\left(\sqrt{1/\beta}\right) = \frac{\beta\left(1-\delta^{2p+1}\right)^2}{4(1-\beta)\delta^{2p+1}} \quad \text{where } 0 < \beta \le 1. \tag{A.9}$$

We defined $\delta = \delta(\beta)$ above in (A.4). This formula is a direct—but unpleasant—consequence of the definition (A.7).

A.2.5. *Second Polynomial Construction.* As before, introduce a parameter $\beta \in [0,1]$. Choose a nonnegative integer partition $q = q_1 + q_2$. Consider the polynomial

$$\psi_{\beta,q_1,q_2}(s) := \frac{s^{q_1} U_{2q_2}\big(\sqrt{s/\beta}\big)}{U_{2q_2}\big(\sqrt{1/\beta}\big)} \quad \text{for } s \in \mathbb{R}.$$

Since U_{2q_2} is an even polynomial, this expression defines a polynomial $\psi_{\beta,q_1,q_2} \in \mathscr{P}_q(1)$. We have the bound

$$(\beta - s)\,\psi^2_{\beta,q_1,q_2}(s) \le \frac{s^{2q_1}\beta}{U^2_{2q_2}(\sqrt{1/\beta})} = \frac{4(1-\beta)s^{2q_1}\delta^{2q_2+1}}{\big(1-\delta^{2q_2+1}\big)^2} \quad \text{for } 0 \le s \le \beta.$$

The inequality follows from (A.8), and the equality follows from (A.9). Last, we instate the definition (A.4).

Acknowledgments

This paper was produced under a subcontract to ONR Award 00014-16-C-2009. The author thanks Shaunak Bopardikar for helpful discussions and feedback.

References

[AS64] M. Abramowitz and I. A. Stegun. *Handbook of mathematical functions with formulas, graphs, and mathematical tables*, volume 55 of *National Bureau of Standards Applied Mathematics Series*. For sale by the Superintendent of Documents, U.S. Government Printing Office, Washington, D.C., 1964.

[BDD+00] Z. Bai, J. Demmel, J. Dongarra, A. Ruhe, and H. van der Vorst, editors. *Templates for the solution of algebraic eigenvalue problems*, volume 11 of *Software, Environments, and Tools*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 2000. A practical guide.

[Bha97] R. Bhatia. *Matrix analysis*, volume 169 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1997.

[CD74] J. Cullum and W. E. Donath. A block generalization of the s-step Lanczos algorithm. Technical Report Report RC 4845 (21570), IBM Thomas J. Watson Research Center, Yorktown Heights, New York, 1974.

[Cho87] A. W. Chou. On the optimality of Krylov information. *J. Complexity*, 3(1):26–40, 1987.

[DCM97] G. M. Del Corso and G. Manzini. On the randomized error of polynomial methods for eigenvector and eigenvalue estimates. *J. Complexity*, 13(4):419–456, 1997.

[DI18] P. Drineas and I. Ipsen. Low-rank matrix approximations do not need a singular value gap. Available at `http://arxiv.org/abs/1801.00670`, Jan. 2018.

[DIKMI16] P. Drineas, I. Ipsen, E.-M. Kontopoulou, and M. Magdon-Ismail. Structural convergence results for approximation of dominant subspaces from block krylov spaces. Available at `http://arXiv.org/abs/1609.00671`, Sep. 2016.

[Dix83] J. D. Dixon. Estimating extremal eigenvalues and condition numbers of matrices. *SIAM J. Numer. Anal.*, 20(4):812–814, 1983.

[DKM06] P. Drineas, R. Kannan, and M. W. Mahoney. Fast Monte Carlo algorithms for matrices. II. Computing a low-rank approximation to a matrix. *SIAM J. Comput.*, 36(1):158–183, 2006.

[FKV98] A. Frieze, R. Kannan, and S. Vempala. Fast Monte Carlo algorithms for finding low-rank approximations. In *Proc. 39th Ann. IEEE Symp. Foundations of Computer Science (FOCS)*, pages 370–378, 1998.

[Gau59] W. Gautschi. Exponential integral $\int_1^\infty e^{-xt} t^{-n}\,dt$ for large values of n. *J. Res. Nat. Bur. Standards*, 62:123–125, 1959.

[GU77] G. H. Golub and R. Underwood. The block Lanczos method for computing eigenvalues. In *Mathematical software, III (Proc. Sympos., Math. Res. Center, Univ. Wisconsin, Madison, Wis., 1977)*, pages 361–377. Publ. Math. Res. Center, No. 39. Academic Press, New York, 1977.

[Gu15] M. Gu. Subspace iteration randomization and singular value problems. *SIAM J. Sci. Comput.*, 37(3):A1139–A1173, 2015.

[GVL13] G. H. Golub and C. F. Van Loan. *Matrix computations*. Johns Hopkins Studies in the Mathematical Sciences. Johns Hopkins University Press, Baltimore, MD, fourth edition, 2013.

[HMST11] N. Halko, P.-G. Martinsson, Y. Shkolnisky, and M. Tygert. An algorithm for the principal component analysis of large data sets. *SIAM J. Sci. Comput.*, 33(5):2580–2594, 2011.

[HMT11] N. Halko, P. G. Martinsson, and J. A. Tropp. Finding structure with randomness: probabilistic algorithms for constructing approximate matrix decompositions. *SIAM Rev.*, 53(2):217–288, 2011.

[Hop34] E. Hopf. *Mathematical problems of radiative equilibrium*. Number 31 in Cambridge Tracts in Mathematics and Mathematical Physics. Cambridge Univ. Press, 1934.

[KW92] J. Kuczyński and H. Woźniakowski. Estimating the largest eigenvalue by the power and Lanczos algorithms with a random start. *SIAM J. Matrix Anal. Appl.*, 13(4):1094–1122, 1992.

[KW94] J. Kuczyński and H. Woźniakowski. Probabilistic bounds on the extremal eigenvalues and condition number by the Lanczos algorithm. *SIAM J. Matrix Anal. Appl.*, 15(2):672–691, 1994.

[Lan50] C. Lanczos. An iteration method for the solution of the eigenvalue problem of linear differential and integral operators. *J. Research Nat. Bur. Standards*, 45:255–282, 1950.

[LW98] Z. Leyk and H. Woźniakowski. Estimating a largest eigenvector by Lanczos and polynomial algorithms with a random start. *Numer. Linear Algebra Appl.*, 5(3):147–164, 1998.

[LZ15] R.-C. Li and L.-H. Zhang. Convergence of the block Lanczos method for eigenvalue clusters. *Numer. Math.*, 131(1):83–113, 2015.

[MM15] C. Musco and C. Musco. Stronger and faster approximate singular value decomposition via the block Lanczos method. In *Proc. Adv. Neural Information Processing Systems 28 (NIPS 2015)*, 2015. Available at http://arXiv.org/abs/1504.05477.

[MRT11] P.-G. Martinsson, V. Rokhlin, and M. Tygert. A randomized algorithm for the decomposition of matrices. *Appl. Comput. Harmon. Anal.*, 30(1):47–68, 2011.

[Nem91] A. S. Nemirovsky. On optimality of Krylov's information when solving linear operator equations. *J. Complexity*, 7(2):121–130, 1991.

[NY83] A. S. Nemirovsky and D. B. a. Yudin. *Problem complexity and method efficiency in optimization*. A Wiley-Interscience Publication. John Wiley & Sons, Inc., New York, 1983. Translated from the Russian and with a preface by E. R. Dawson, Wiley-Interscience Series in Discrete Mathematics.

[OLBC10] F. W. J. Olver, D. W. Lozier, R. F. Boisvert, and C. W. Clark, editors. *NIST handbook of mathematical functions*. U.S. Department of Commerce, National Institute of Standards and Technology, Washington, DC; Cambridge University Press, Cambridge, 2010. With 1 CD-ROM (Windows, Macintosh and UNIX).

[Pai71] C. C. Paige. *The computation of eigenvalues and eigenvectors of very large sparse matrices*. PhD thesis, Univ. London, 1971.

[Par98] B. N. Parlett. *The symmetric eigenvalue problem*, volume 20 of *Classics in Applied Mathematics*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 1998. Corrected reprint of the 1980 original.

[PRTV98] C. H. Papadimitriou, P. Raghavan, H. Tamaki, and S. Vempala. Latent semantic indexing: A probabilistic analysis. In *Proc. 17th ACM Symp. Principles of Database Systems (PODS)*, 1998.

[RST09] V. Rokhlin, A. Szlam, and M. Tygert. A randomized algorithm for principal component analysis. *SIAM J. Matrix Anal. Appl.*, 31(3):1100–1124, 2009.

[Saa80] Y. Saad. On the rates of convergence of the Lanczos and the block-Lanczos methods. *SIAM J. Numer. Anal.*, 17(5):687–706, 1980.

[Saa11] Y. Saad. *Numerical methods for large eigenvalue problems*, volume 66 of *Classics in Applied Mathematics*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 2011. Revised edition of the 1992 original [1177405].

[SAR17] M. Simchowitz, A. E. Alaoui, and B. Recht. On the gap between strict-saddles and true convexity: An $\omega(\log d)$ lower bound for eigenvector approximation. Available at http://arXiv.org/abs/1704.04548, Apr. 2017.

[Tro18a] J. A. Tropp. Analysis of randomized block Krylov methods. Part II: Eigenvalue estimation. Available at http://users.cms.caltech.edu/~jtropp/, Feb. 2018.

[Tro18b] J. A. Tropp. Analysis of randomized block Krylov methods. Part III: Matrix approximation. Available at http://users.cms.caltech.edu/~jtropp/, Feb. 2018.

[WC15] R. Witten and E. Candès. Randomized algorithms for low-rank matrix factorizations: sharp performance bounds. *Algorithmica*, 72(1):264–281, 2015.

[WZZ15] S. Wang, Z. Zhang, and T. Zhang. Improved analyses of the randomized power method and block Lanczos method. Available at http://arXiv.org/abs/1508.06429, Aug. 2015.

ANALYSIS OF RANDOMIZED BLOCK KRYLOV METHODS
PART II: EIGENVALUE APPROXIMATION

JOEL A. TROPP

ABSTRACT. Randomized block Krylov subspace methods are a powerful class of algorithms for computing information about the spectrum of a matrix. The purpose of this note is to develop new theoretical bounds on the performance of randomized block Krylov subspace methods for estimating a number of extreme eigenvalues. The results demonstrate that, for many matrices, it is possible to obtain accurate approximations using only a *constant* number of steps of the randomized block Krylov method.

1. MOTIVATION AND MAIN RESULTS

Randomized block Krylov methods are a powerful tool for spectral computation and matrix approximation [MRT11, RST09, HMT11, HMST11, MM15, WZZ15, DIKMI16]. Our understanding of the performance of these methods is less complete than our understanding of simple Krylov subspace methods. The goal of this paper, and its companions, is to develop detailed bounds that help explain the striking empirical behavior [HMST11, MM15] of randomized block Krylov methods.

1.1. Project Overview.

This paper, Part II of the project, studies how well we can estimate the largest (or smallest) k eigenvalues of a symmetric matrix using a randomized block Krylov method. It also discusses the analogous questions for singular values. We develop the first detailed theory for these problems. The analysis uses ideas from the papers [KW92, HMT11, MM15].

Part I of this project [Tro18a] focuses on the simpler problem of estimating the largest (or smallest) eigenvalue (or singular value) of a symmetric matrix. We have treated these problems separately because the proofs are elementary, and we can elucidate the precise role of the block size.

Part III of this project [Tro18b] turns to the problem of computing a rank-k approximation of a general matrix. This problem has already been investigated [MM15, DIKMI16]. We have been able to obtain more precise results by taking advantage of ideas from the other parts of this project.

Our expository goal is to make each work self-contained, while limiting repetition of material. In future work, we may also treat the problem of estimating invariant subspaces of a symmetric matrix associated with the largest k eigenvalues. This problem is more like low-rank matrix approximation than eigenvalue estimation, but it requires an independent argument.

1.2. Eigenvalue Approximation via the Block Krylov Method.

We begin with a mathematical description of a block Krylov method for approximating the largest eigenvalues of a symmetric matrix. See Section 1.2.5 for some discussion about the implementation. See Section 1.7 for some discussion of related computational problems.

1.2.1. *Block Krylov Subspaces.* Suppose that we are given a symmetric *input matrix* $\boldsymbol{A} \in \mathbb{R}^{n\times n}$. Fix a *test matrix* $\boldsymbol{B} \in \mathbb{R}^{n\times \ell}$, where ℓ is the block size. Select a *depth* parameter q that controls the amount of information we collect about the input matrix. Implicitly, the block Krylov method constructs the matrix

$$\boldsymbol{S}_q := \boldsymbol{S}_q(\boldsymbol{A};\boldsymbol{B}) := \begin{bmatrix} \boldsymbol{B} & \boldsymbol{AB} & \boldsymbol{A}^2\boldsymbol{B} & \dots & \boldsymbol{A}^q\boldsymbol{B} \end{bmatrix} \in \mathbb{R}^{n\times(q+1)\ell}.$$

Date: 21 December 2017. Revised 26 February 2018.

2010 *Mathematics Subject Classification.* Primary: 65F30. Secondary: 68W20, 60B20.

Key words and phrases. Eigenvalue computation; Krylov subspace method; Lanczos method; numerical analysis; randomized algorithm; singular value computation.

The range of this matrix is called a *block Krylov subspace*:

$$K_q(\boldsymbol{A};\boldsymbol{B}) := \operatorname{range}(\boldsymbol{S}_q) \subset \mathbb{R}^n.$$

The block Krylov space can also be expressed using polynomials:

$$K_q(\boldsymbol{A};\boldsymbol{B}) = \bigoplus\nolimits_{\varphi\in\mathscr{P}_q} \operatorname{range}\big(\varphi(\boldsymbol{A})\boldsymbol{B}\big), \tag{1.1}$$

where $\mathscr{P}_q$ is the space of polynomials with real coefficients and degree at most q. Here, $\oplus$ denotes the ordinary subspace sum. The block Krylov subspace has remarkable invariance properties; see [Tro18a, Sec. 1.2.2].

1.2.2. *Eigenvalue Approximation.* We can compute spectral information about the input matrix by compressing it to the block Krylov subspace and performing spectral computations on the compressed matrix. For example, see [Par98, Saa11].

Let k be the number of eigenvalues that we wish to compute, and select a block size ℓ that is (slightly) larger than k. Next, construct a matrix $\boldsymbol{Q}$ whose columns form an orthonormal basis for the block Krylov subspace $K_q(\boldsymbol{A};\boldsymbol{B})$. Introduce the *Rayleigh matrix*

$$\boldsymbol{H} := \boldsymbol{H}_q(\boldsymbol{A};\boldsymbol{B};\boldsymbol{Q}) := \boldsymbol{Q}^*\boldsymbol{A}\boldsymbol{Q}. \tag{1.2}$$

The * denotes the transpose operation on matrices and vectors. The eigenvalues of $\boldsymbol{H}$ are called *Ritz values*:

$$\xi_r := \xi_r(\boldsymbol{A};\boldsymbol{B};q) := \lambda_r(\boldsymbol{H}) \quad \text{for } r = 1,2,3,\dots,\dim(K_q). \tag{1.3}$$

The map $\lambda_r(\cdot)$ returns the rth largest eigenvalue of a symmetric matrix. Even though the matrix $\boldsymbol{Q}$ is not unique, the Ritz values are a function of the Krylov subspace $K_q(\boldsymbol{A};\boldsymbol{B})$.

The Ritz values tend to approximate the extreme eigenvalues of the input matrix $\boldsymbol{A}$. Let us define a measure of the relative error in estimating the rth eigenvalue by means of the corresponding Ritz value:

$$\operatorname{err}_r := \operatorname{err}_r(\boldsymbol{A};\boldsymbol{B};q) := \frac{\lambda_r(\boldsymbol{A}) - \xi_r(\boldsymbol{A};\boldsymbol{B};q)}{\lambda_r(\boldsymbol{A}) - \lambda_n(\boldsymbol{A})} \quad \text{for } r = 1,2,3,\dots,\dim(K_q). \tag{1.4}$$

It is always the case that $\operatorname{err}_r \in [0,1]$ because of the Cauchy interlacing theorem [Bha97, Cor. III.1.5]. To handle degenerate cases, we use the conventions $0/0 = 0$ and $c/0 = +\infty$ for $c > 0$. As we will explain, this error measure is natural because of invariance properties of the block Krylov method.

Remark 1.1 (Invariant Subspaces)**.** Let us emphasize that the eigenvectors of the Rayleigh matrix $\boldsymbol{H}$ are not always related to the eigenvectors of $\boldsymbol{A}$. We present some results on the quality of invariant subspace approximations in the companion paper [Tro18b], but a detailed treatment requires its own argument.

1.2.3. *Invariance Properties of the Ritz Values.* We can derive some important invariance properties of the block Krylov method for estimating the rth eigenvalue of the input matrix using a fixed depth q.

- For a fixed input matrix $\boldsymbol{A}$, the Ritz values only depend on the range of the test matrix $\boldsymbol{B}$:

$$\xi_r(\boldsymbol{A};\boldsymbol{B}) = \xi_r(\boldsymbol{A};\boldsymbol{B}\boldsymbol{T}) \quad \text{for all nonsingular } \boldsymbol{T} \in \mathbb{R}^{\ell\times\ell}. \tag{1.5}$$

- The Ritz values are invariant to the orientation of $\boldsymbol{A}$ and $\boldsymbol{B}$ in the sense that

$$\xi_r(\boldsymbol{A};\boldsymbol{B}) = \xi_r(\boldsymbol{U}\boldsymbol{A}\boldsymbol{U}^*;\boldsymbol{U}\boldsymbol{B}) \quad \text{for all orthogonal } \boldsymbol{U} \in \mathbb{R}^{n\times n}. \tag{1.6}$$

- The Ritz values co-vary with increasing affine maps of the spectrum of the input matrix:

$$\xi_r(\alpha\boldsymbol{A} + \beta\mathbf{I};\boldsymbol{B}) = \alpha\xi_r(\boldsymbol{A};\boldsymbol{B}) + \beta\mathbf{I} \quad \text{for all } \alpha \geq 0 \text{ and } \beta \in \mathbb{R}. \tag{1.7}$$

In particular, the error (1.4) is *invariant* to increasing affine transformations of the spectrum of $\boldsymbol{A}$.

These facts follow in a straightforward way from the invariance properties of the Krylov subspace [Tro18a, Sec. 1.2.2] and the definitions of the Rayleigh matrix (1.2) and the Ritz values (1.3).

Algorithm 1 A block Krylov method for computing the largest eigenvalues of a symmetric matrix

Input: Symmetric $n \times n$ matrix $\boldsymbol{A}$; number k of eigenvalues; block size ℓ; depth q
Output: Vector $\boldsymbol{\xi} \in \mathbb{R}^k$ with estimates of k largest eigenvalues

```
function BlockKrylovEigs(A, k, ℓ, q)
    Y_0 ← randn(size(A, 2), ℓ)          ▷ Draw n × ℓ standard normal test matrix
    for t ← 1, 2, 3, ..., q do
        Y_t ← A Y_{t-1}                 ▷ Form blocks of Krylov matrix by repeated multiplication
    Q ← orth([Y_0, Y_1, ..., Y_q])      ▷ Find orthonormal basis for block Krylov space
    H ← Q*(AQ)                          ▷ Form Rayleigh matrix
    (~, ξ) ← eigs(H, k)                 ▷ Compute Ritz values via dense linear algebra
```

1.2.4. *A Random Test Matrix.* To ensure that the block Krylov method is effective for estimating the largest eigenvalues of an arbitrary input matrix $\boldsymbol{A}$, we may draw the test matrix *at random*. In this work, we always consider a test matrix $\boldsymbol{\Omega} \in \mathbb{R}^{n \times \ell}$ drawn from the standard normal distribution. We will assess how well the Ritz values $\xi_r(\boldsymbol{A}; \boldsymbol{\Omega}; q)$ approximate the eigenvalues of $\boldsymbol{A}$ by studying the distribution of the relative error (1.4).

The choice of a standard normal test matrix is justified by the invariance properties (1.5) and (1.6) of the Ritz values. The companion paper [Tro18a, Sec. 1.2.5 and App. A.1] contains some additional support for this choice.

Remark 1.2 (Other Test Matrices)**.** The detailed results in this paper depend heavily on the standard normal distribution of the test matrix. Nevertheless, the overall strategy of the proof extends to other types of random test matrices. See [HMT11, Secs. 4.6 and 7.4] for discussion of other test matrices. For block Krylov methods, fancy test matrices offer very limited benefits.

1.2.5. *Implementation of the Block Krylov Method.* This paper focuses on mathematical analysis of block Krylov methods, but it is worth a moment to comment on possible implementations.

Algorithm 1 provides pseudocode for a simple variant of the block Krylov method for computing the largest eigenvalues of a matrix. This approach is adapted from [HMST11]. We do not expect this algorithm to be numerically stable for moderate or large q or when $\boldsymbol{A}$ is poorly conditioned.

Here is a short breakdown of the computational resources required.

- We multiply $\boldsymbol{A} \in \mathbb{R}^{n \times n}$ by an $n \times \ell$ matrix a total of q times and by an $n \times (q+1)\ell$ matrix once. In case $\boldsymbol{A}$ is dense, this requires $\mathcal{O}(q\ell n^2)$ arithmetic operations, but it may be far more efficient when $\boldsymbol{A}$ is sparse or structured.
- Orthogonalization of the block Krylov matrix requires $\mathcal{O}(q^2\ell^2 n)$ arithmetic operations.
- Computation of Ritz values uses $\mathcal{O}(q^3\ell^3)$ arithmetic operations.
- Storage of the block Krylov matrix requires $\mathcal{O}(q\ell n)$ units of storage.

It is possible to achieve some further efficiencies using block Lanczos algorithms [CD74, GU77].

1.3. **The Role of the Spectrum.** Since the range of the random test matrix $\boldsymbol{\Omega}$ is rotationally invariant, property (1.6) implies that the Ritz values depend only on the eigenvalue spectrum of $\boldsymbol{A}$. Indeed, for each index r, we have

$$\xi_r(\boldsymbol{A}; \boldsymbol{\Omega}; q) \sim \xi_r(\boldsymbol{\Lambda}; \boldsymbol{\Omega}; q) \quad \text{where} \quad \boldsymbol{A} = \boldsymbol{U}\boldsymbol{\Lambda}\boldsymbol{U}^* \text{ is an eigenvalue factorization.}$$

The symbol $\sim$ denotes identical distribution. As a consequence, we may as well assume that the input matrix $\boldsymbol{A}$ is diagonal and weakly decreasing:

$$\boldsymbol{A} = \operatorname{diag}(a_1, a_2, \dots, a_n) \quad \text{where} \quad a_{\max} := a_1 \geq a_2 \geq \cdots \geq a_n =: a_{\min}. \tag{1.8}$$

This hypothesis remains in force for the rest of the paper.

On account of (1.7), the Ritz values co-vary with increasing affine transformations of the spectrum of the input matrix $\boldsymbol{A}$. Therefore, we have adopted an error measure (1.4) that shares the same invariances, and we must control this error using spectral statistics with the same invariances. Recall that we have fixed the number k of eigenvalues that we seek to approximate, and the remaining eigenvalues compose the *tail* of the spectrum.

- For each index $1 \le r \le k+1$, the *tail spectral gap* is a measure of the relative difference between the rth eigenvalue and the eigenvalues in the tail:
$$\gamma_r := \frac{a_r - a_{k+1}}{a_r - a_{\min}}. \tag{1.9}$$
If $a_r = a_{\min}$, we define $\gamma_r = 0$.
- For each index $1 \le r \le k+1$ and a number $\nu \ge 0$, the *tail content* measures how quickly the tail eigenvalues decay, relative to the rth eigenvalue:
$$\tau(r;\nu) := \sum_{j>k} \left(\frac{a_j - a_{\min}}{a_r - a_{\min}}\right)^{2\nu}. \tag{1.10}$$
If $a_r = a_{\min}$, we define $\tau(r;\nu) = 0$.

The tail content is a critical new feature of our analysis, so it merits a short discussion. A general bound for the tail content is

$$\tau(r;\nu) \le \left(\frac{a_{k+1}}{a_r}\right)^{2\nu} (n-k) \le n-k \quad \text{where } 1 \le r \le k+1.$$

We see that the tail content is never larger than the ambient dimension. Furthermore, it can be quite small when $a_{k+1} \ll a_r$. When the input matrix exhibits spectral decay, the tail content is often far smaller than this bound suggests. For example, the polynomial decay profile

$$\frac{a_j}{a_r} \le C j^{-s} \quad \text{for } j > k \quad \text{implies} \quad \tau(r;\nu) < 2C \quad \text{for } \nu \ge 1/s.$$

In many situations, $\nu = 1$ or $\nu = 2$ is already adequate to ensure that the tail content is *constant*.

1.4. **Matrices with Few Eigenvalues.** The block Krylov method is very effective for matrices that have few distinct eigenvalues.

Proposition 1.3 (Randomized Block Krylov: Matrices with Few Eigenvalues)**.** *Let $\boldsymbol{A} \in \mathbb{R}^{n\times n}$ be a symmetric matrix, and fix the number k of eigenvalues to estimate. Suppose that $\boldsymbol{A}$ has m distinct eigenvalues that are strictly smaller than $\lambda_k(\boldsymbol{A})$. Choose the depth $q \ge m$ and the block size $\ell \ge k$, and draw a standard normal matrix $\boldsymbol{\Omega} \in \mathbb{R}^{n\times\ell}$. Then, with probability one, the relative error (1.4) satisfies*

$$\mathrm{err}_r(\boldsymbol{A};\boldsymbol{\Omega};q) = 0 \quad \textit{for each } r = 1,\dots,k.$$

In particular, when $m \in \{0,1\}$, depth $q = 1$ is sufficient to estimate the largest k eigenvalues of $\boldsymbol{A}$ exactly. Moreover, the relative error (1.4) is always finite when $m \ge 2$.

This type of result is well-known, but we include a short proof for completeness. See Section 3.4.

1.5. **Estimating Eigenvalues without a Spectral Gap.** Our first result describes the performance of the block Krylov method for eigenvalue estimation. This result does not require that *any* of the leading eigenvalues of the input matrix are separated from the tail eigenvalues.

Theorem 1.4 (Randomized Block Krylov: Eigenvalue Estimation)**.** *Instate the following hypotheses.*

- *Let $\boldsymbol{A} \in \mathbb{R}^{n\times n}$ be a symmetric matrix.*
- *Fix the desired number k of eigenvalues, the block size ℓ, and the oversampling $p := \ell - k$.*
- *Draw a standard normal test matrix $\boldsymbol{\Omega} \in \mathbb{R}^{n\times\ell}$.*
- *Set the depth parameter $q \ge 2$, and choose an integer partition $q = q_1 + q_2$ with $0 < q_1 < q$.*

For each index $1 \le r \le k$, we have the following bounds for the Ritz value approximations (1.3) to the eigenvalues of $\boldsymbol{A}$.

(1) For oversampling $p \geq 2$, the relative error (1.4) satisfies the probability bound

$$\mathbb{P}\left\{\max_{i\leq k} \mathrm{err}_i(\boldsymbol{A};\boldsymbol{\Omega};q) \geq \varepsilon\right\} \leq 1 \wedge 2\left[\frac{11\max\{1, k\tau(k;q_1)\}}{\varepsilon}\cdot \mathrm{e}^{-4q_2\sqrt{\varepsilon/2}}\right]^{p/2} \quad \text{for each } \varepsilon \in (0,1].$$

(2) For oversampling $p \geq 3$, the expected relative error satisfies

$$\mathbb{E}\max_{i\leq k} \mathrm{err}_i(\boldsymbol{A};\boldsymbol{\Omega};q) \leq 1 \wedge \left[\frac{p/(p-2) + 0.36\cdot\log\left(11q_2^2\max\{1, k\tau(k;q_1)\}\right)}{q_2}\right]^2.$$

The operator $\wedge$ returns the minimum, and the tail content τ is defined in (1.10).

The proof of Theorem 1.4 appears in Section 2.

Let us take a few paragraphs to explain the content of this result. For the moment, fix the first depth parameter q_1. Note, however, that the user does not select the partition $q = q_1 + q_2$; the block Krylov algorithm automatically makes an optimal choice.

Suppose that we wish to achieve relative error ε. The probability bound in Theorem 1.4 is vacuous unless the second depth parameter q_2 exceeds

$$q_2(\varepsilon) := \frac{\log\left(11\varepsilon^{-1}\max\{1, k\tau(k;q_1)\}\right)}{4\sqrt{\varepsilon}} \quad \text{for } \varepsilon \in (0,1].$$

If we select $q_2 \geq q_2(\varepsilon)$, we also ensure that the expectation of the maximum relative error in the first k eigenvalue estimates has order ε. Once $q_2 \geq q_2(\varepsilon)$, small increases in q_2 result in rapid decreases in the failure probability. The rate of decrease is driven by the oversampling parameter p.

This observation has two important consequences. First, the threshold $q_2(\varepsilon)$ scales with $\varepsilon^{-1/2}\log(1/\varepsilon)$, so we can obtain a modest relative error with a shallow Krylov subspace. In contrast, randomized subspace iteration requires the total depth q to increase with ε^{-1} to achieve relative error ε in spectral computations; for related results, see [KW92, HMT11, MM15].

Second, the threshold $q_2(\varepsilon)$ scales with $\log(k\tau(k;q_1))$. As a consequence, we only need q_2 to grow with $\log k$ to estimate each of the k largest eigenvalues to fixed relative error. Furthermore, the appearance of the tail content $\tau(k;q_1)$ suggests that the block Krylov method is more effective when the input matrix has a decaying spectrum. The paper [HMT11] identified the benefits of spectral decay for randomized subspace iteration, but we are not aware of an analogous result for a block Krylov method.

This discussion points toward the role of the first depth parameter q_1. When the spectrum of the input matrix decays, the tail content $\tau(k;q_1)$ may be constant even when q_1 is quite small. For many matrices, $q_1 = 1$ or $q_1 = 2$ is sufficient.

In the best situation, where the tail has polynomial decay, $\tau(k;q_1)$ is constant for constant q_1. In this case, the theorem shows that the total depth q of the Krylov subspace should satisfy

$$\text{Polynomial tail decay:} \qquad q = q_1 + q_2(\varepsilon) \approx \frac{\log(k/\varepsilon)}{\sqrt{\varepsilon}}$$

to estimate each of the largest k eigenvalues of the matrix to relative accuracy ε.

In the worst situation, where the tail does not decay at all, we have $\tau(k;q_1) = n - k$ for each q_1. In this case, the theorem suggests that the total depth q of the Krylov space should satisfy

$$\text{No tail decay:} \qquad q = 1 + q_2(\varepsilon) \approx \frac{\log(kn/\varepsilon)}{\sqrt{\varepsilon}}$$

to estimate the largest k eigenvalues to relative accuracy ε.

Prior work [KW92, MM15] on randomized (block) Krylov methods suggests that the depth q needs to scale with $\log n$ to achieve nontrivial results. Our research shows that we can avoid the dimensional dependency when the spectrum of the input matrix exhibits polynomial decay. In case $k = 1$, the companion paper [Tro18a, Thm. 1.5] obtains similar bounds without the factor of $\log(1/\varepsilon)$; it is not clear whether this logarithmic term can be removed from the results here.

Here is another interpretation of Theorem 1.4. As we increase the depth q, the block Krylov method exhibits a burn-in period of length $q_1 + q_2(1)$. During the burn-in period, the algorithm may not make any progress on estimating the largest k eigenvalues. Afterward, as q_2 continues to increase, the relative errors start to decrease in proportion to q_2^{-2}. Randomized subspace iteration exhibits the same burn-in period, after which it decreases the errors in proportion to q_2^{-1}. See [HMT11, MM15] for related results.

We have focused on the case where the oversampling parameter $p \geq 2$ because the companion paper [Tro18a, Thm. 1.5] indicates that the behavior of the block Krylov method is qualitatively worse when $p = 0$ or $p = 1$. Once $p \geq 2$, increasing the oversampling does not reduce the expected error substantially, but it reduces the failure probability very quickly. On modern computer architectures, a modest increase of the oversampling parameter has a negligible impact on the algorithmic resources required.

Remark 1.5 (Prior Work)**.** The paper [KW92] studies the behavior of the randomized Krylov method (with block size $\ell = 1$) for estimating the largest eigenvalue of a symmetric matrix. There is a long literature, including [GU77, Saa80, LZ15], that contains deterministic bounds for eigenvalue estimation using block Krylov methods. Several recent papers [MM15, WZZ15, DIKMI16] analyze low-rank matrix approximation with randomized block Krylov methods. We are not aware of analogous work on eigenvalue estimation, so there is no clear point of comparison for Theorem 1.4.

1.6. **Estimating Eigenvalues with a Spectral Gap.** Our second result describes how well the block Krylov method can estimate eigenvalues that stand well above the tail eigenvalues. For most matrices, this result gives significantly stronger bounds than Theorem 1.4.

Theorem 1.6 (Randomized Block Krylov: Eigenvalue Estimation with Tail Spectral Gap)**.** *Instate the following hypotheses.*

- *Let $\boldsymbol{A} \in \mathbb{R}^{n\times n}$ be a symmetric matrix.*
- *Fix the desired number k of eigenvalues, the block size ℓ, and the oversampling $p := \ell - k$.*
- *Draw a standard normal test matrix $\boldsymbol{\Omega} \in \mathbb{R}^{n\times \ell}$.*
- *Set the depth parameter $q \geq 1$, and choose an integer partition $q = q_1 + q_2$ with $0 < q_1 \leq q$.*

For each index $1 \leq r \leq k$, we have the following bounds for the Ritz value approximations (1.3) to the eigenvalues of $\boldsymbol{A}$.

(1) For oversampling $p \geq 2$, the relative error (1.4) in the approximation satisfies the probability bound

$$\mathbb{P}\left\{\max_{i\leq r} \operatorname{err}_i(\boldsymbol{A};\boldsymbol{\Omega};q) \geq \varepsilon\right\} \leq \left[\frac{9r\tau(r;q_1)}{\varepsilon}\cdot \mathrm{e}^{-4q_2\sqrt{\gamma_r}}\right]^{p/2} \quad \text{for each } \varepsilon \in (0,1).$$

(2) For oversampling $p \geq 3$, the expected relative error satisfies

$$\mathbb{E}\max_{i\leq r} \operatorname{err}_i(\boldsymbol{A};\boldsymbol{\Omega};q) \leq \frac{p}{p-2}\cdot 9r\tau(r;q_1)\cdot \mathrm{e}^{-4q_2\sqrt{\gamma_r}}.$$

The tail spectral gap γ_r is defined in (1.9), and the tail content τ is defined in (1.10).

The proof of Theorem 1.6 appears in Section 2.

Let us comment briefly on this result. For a fixed relative error ε, the probability bound only has content when the second depth parameter q_2 exceeds

$$q_2(\varepsilon;\gamma_r) := \frac{\log\left(9\varepsilon^{-1}r\tau(r;q_1)\right)}{4\sqrt{\gamma_r}}.$$

This choice also yields expected error with order ε.

The key improvement over Theorem 1.4 is that the expected relative error in the first r eigenvalue estimates decreases *exponentially* with q_2 if the tail spectral gap γ_r is bounded away from zero. As such, we can estimate any large eigenvalue to exorbitant accuracy. The other features of this result are similar to Theorem 1.4. In particular, we can often make the tail content $\tau(r;q_1)$ constant when q_1 is constant. It is interesting that the threshold function $\max\{1,\cdot\}$ does not appear in this result, so we can sometimes achieve further benefits by increasing q_1 more.

1.7. **Related Problems.** We can also apply the block Krylov method to perform some other spectral computations.

- We can use the block Krylov method to approximate the smallest eigenvalues of the input matrix. Indeed, $\zeta_r(\boldsymbol{A};\boldsymbol{B};q) := -\xi_r(-\boldsymbol{A};\boldsymbol{B};q)$ is an estimate for the rth smallest eigenvalue of $\boldsymbol{A}$. We can easily adapt Theorems 1.4 and 1.6 to obtain theoretical results about this problem.
- Let $\boldsymbol{C} \in \mathbb{R}^{n\times m}$ be a rectangular matrix. If we apply the block Krylov method to the input matrix $\boldsymbol{A} = \boldsymbol{C}\boldsymbol{C}^*$, then the Ritz value $\xi_r(\boldsymbol{A};\boldsymbol{B};q)$ is an estimate for $\sigma_r^2(\boldsymbol{C})$, the *square* of the rth largest singular value of $\boldsymbol{C}$.
- Combining the last two observations, we can also obtain estimates for minimum singular values of a rectangular matrix.

See the companion paper [Tro18a, Secs. 1.7, 1.8] for some additional discussion.

1.8. **Extensions.** There are a number of ways to refine the results in this paper:

- **Oversampling.** When the oversampling parameter p is proportional to the number k of eigenvalues we wish to estimate, it is possible to obtain stronger bounds. This improvement requires a significant amount of additional argument, plus a strong dose of random matrix theory. We have also left out the cases $p = 0$ and $p = 1$.
- **Subspace Iteration.** The arguments can be adapted to study randomized subspace iteration. The literature [HMT11] already detailed results for this type of algorithm, although prior work focuses on matrix approximation rather than eigenvalue estimation.
- **Complex Matrices.** The arguments do not depend in any significant way on the fact that the matrices are real-valued. We can estimate spectral properties of complex-valued input matrices by using a complex standard normal test matrix. These results are actually stronger than the results in the real case.

We have chosen not to pursue these extensions because they bring additional complexity without adding much insight.

1.9. **Summary of Contributions.** We conclude with an overview of the contributions of this paper and some comments on related work.

- The main innovation in our work is to identify the role of the tail content in controlling the performance of randomized block Krylov methods. For matrices with spectral decay, the depth of the Krylov subspace need not depend on the ambient dimension. We believe this paper is the first to obtain dimension-independent bounds for these algorithms.
- We have shown that it is possible to estimate the largest (and smallest) k eigenvalues of a symmetric matrix using a randomized block Krylov method. In contrast, previous work on randomized methods only addresses the problem of estimating the maximum (or minimum) eigenvalue [KW92] or the largest k singular values [HMT11, MM15].
- Our results point to the benefits of implementing block Krylov methods with oversampling $p \geq 2$.
- We have obtained explicit and reasonable constants, which gives the results predictive power.

Our analysis of randomized block Krylov methods for eigenvalue approximation is inspired by several sources. The overall argument is similar in spirit to the proof of [MM15, Thm. 12], although the details are comprehensively different. In particular, the approach here relies on the research of Halko et al. [HMT11] about randomized matrix approximation using a Gaussian test matrix. We have also adapted our treatment [Tro18a] of the argument from Kuczyński & Woźniakowski [KW92]. See the companion paper [Tro18a, Sec. 2] for a more complete discussion about the history of (randomized) block Krylov methods. We provide more detailed pointers to the literature throughout the paper.

2. Main Technical Results

This section contains the main technical results, and we explain how they imply Theorems 1.4 and 1.6. The proofs of the technical results unfold in the remaining sections of the paper.

2.1. **Simplifications.** We instate the notation from the introduction. In particular, we maintain the assumption (1.8) that the input matrix $\boldsymbol{A} \in \mathbb{R}^{n\times n}$ is diagonal, with weakly decreasing entries. Throughout the discussion, $\boldsymbol{\Omega} \in \mathbb{R}^{n\times \ell}$ remains a standard normal test matrix.

Fix the number k of eigenvalues that we wish to estimate, and set the oversampling $p := \ell - k$. Proposition 1.3, which we establish below in Section 3.4, gives conditions under which $\mathrm{err}_r = 0$ for all $r \le k$. Therefore, we assume that we are in the complementary case, where at least two distinct eigenvalues are smaller than a_k.

We will translate the spectrum of the input matrix so that $a_n = 0$. In particular, we may assume that $\boldsymbol{A}$ is positive semidefinite (psd). Theorems 1.4 and 1.6 are both stated in terms of affine invariant quantities, so there is no loss of generality.

Finally, it is productive to introduce an explicit family of truncations of the input matrix:

$$\boldsymbol{A}_r = \operatorname{diag}(a_1, a_2, \dots, a_r, 0, \dots, 0) \quad \text{for each index } 1 \le r \le n.$$

Evidently, $\boldsymbol{A}_r$ lists the largest r eigenvalues of $\boldsymbol{A}$, and it has rank no greater than r.

2.2. **Probability Bounds.** The main part of the analysis consists of two lemmas that address two complementary challenges. Fix an error parameter $\varepsilon \in (0, 1]$. The first result will provide information about how well we estimate the large eigenvalues of $\boldsymbol{A}$ that exceed $(1-\varepsilon/2)^{-1} a_k$. The second result provides information about how well we estimate the small eigenvalues that are less than $(1-\varepsilon/2)^{-1} a_k$.

Lemma 2.1 (Large Eigenvalues)**.** *Instate the prevailing notation. Let $q = q_1 + q_2$ be an integer partition of the depth parameter q with $1 \le q_1 \le q$. For each index $1 \le r \le k$,*

$$\mathbb{P}\left\{ \max_{i \le r} \mathrm{err}_i \ge \varepsilon \right\} \le 1 \wedge \left[\frac{9 r \tau(r; q_1)}{\varepsilon} \cdot \mathrm{e}^{-4q_2\sqrt{\gamma_r}} \right]^{p/2}.$$

The error is defined in (1.4).

The proof of this result appears in Section 4.

Lemma 2.2 (Small Eigenvalues)**.** *Instate the prevailing notation. Let r be any index where $1 \le r \le k$ and $a_r \le (1-\varepsilon/2)^{-1} a_k$. Let $q = q_1 + q_2$ be an integer partition of the depth parameter q with $0 \le q_1 \le q$. Then*

$$\mathbb{P}\left\{ \max_{r \le i \le k} \mathrm{err}_i \ge \varepsilon \right\} \le 1 \wedge \left[11 \max\{1, k\tau(k; q_1)\} \cdot \mathrm{e}^{-2(2q_2+1)\sqrt{\varepsilon/2}} \right]^{p/2}.$$

The error is defined in (1.4).

The proof of this result appears in Section 5.

These two lemmas are analogous with [MM15, Lem. 9, Claims 1 and 2]. Nevertheless, the results and the proofs here are significantly different. We also introduced many new ideas to sharpen the analysis.

The probability bounds in the main results are immediate consequences of Lemmas 2.1 and 2.2.

Proof of Theorem 1.4(1) *and Theorem 1.6*(1). The probability bound in Theorem 1.6(1) is simply a restatement of Lemma 2.1.

To establish the probability bound in Theorem 1.4(1), we use Lemma 2.1 to handle estimates of the eigenvalues $a_r \ge (1-\varepsilon/2)^{-1} a_k$. In this case, the tail spectral gap (1.9) satisfies

$$\gamma_r = \frac{a_r - a_{k+1}}{a_r} \ge \frac{a_r - a_k}{a_r} \ge \varepsilon/2.$$

We use Lemma 2.2 to handle estimates of the eigenvalues $a_r < (1-\varepsilon/2)^{-1} a_k$. Weakening and combining the two lemmas, we arrive at the uniform bound

$$\mathbb{P}\left\{ \max_{i \le k} \mathrm{err}_i \ge \varepsilon \right\} \le 2 \left[\frac{11 \max\{1, k\tau(k; q_1)\}}{\varepsilon} \cdot \mathrm{e}^{-4q_2\sqrt{\varepsilon/2}} \right]^{p/2}.$$

Among other things, we have used the fact that $\tau(\cdot; q_1)$ is increasing. □

2.3. **Expectation Bounds.** Using the probability bounds from the last section, we can obtain bounds for the expectation of the error.

Lemma 2.3 (Expected Error)**.** *Instate the prevailing notation. Assume that the oversampling parameter $p := \ell - k \geq 3$, and fix an index $1 \leq r \leq k$. Let $q = q_1 + q_2$ be an integer partition with $1 \leq q_1 \leq q$. Then*

$$\mathbb{E} \max_{i \leq r} \mathrm{err}_i \leq \frac{p}{p-2} \cdot 9r\tau(r; q_1) \cdot \mathrm{e}^{-4q_2\sqrt{\gamma_r}}. \tag{2.1}$$

Furthermore, if $1 \leq q_2$,

$$\mathbb{E} \max_{i \leq k} \mathrm{err}_i \leq \left[\frac{p/(p-2) + 0.36 \cdot \log\left(11 q_2^2 \max\{1, k\tau(k; q_1)\}\right)}{q_2} \right]^2. \tag{2.2}$$

The relative error is defined in (1.4)*.*

Proof. To establish (2.1), define

$$c := 9r\tau(r; q_1) \cdot \mathrm{e}^{-4q_2\sqrt{\gamma_r}}.$$

The error (1.4) always falls in the range $[0, 1]$, so

$$\mathbb{E} \max_{i \leq r} \mathrm{err}_i = \int_0^1 \mathbb{P}\left\{ \max_{i \leq r} \mathrm{err}_i \geq \varepsilon \right\} \mathrm{d}\varepsilon \leq c + \int_c^\infty (c/\varepsilon)^{p/2} \, \mathrm{d}\varepsilon = \frac{pc}{p-2}. \tag{2.3}$$

The inequality depends on Lemma 2.1. This is the stated result.

To prove (2.2), define the quantities

$$d := 11 \max\{1, k\tau(k; q_1)\} \quad \text{and} \quad f := 2 \left[\frac{\log\left(q_2^2 d\right)}{4q_2} \right]^2.$$

First, consider the largest index $r \leq k$ where $f/2 \leq \gamma_r$. We can instantiate the expectation bound (2.3) to obtain

$$\mathbb{E} \max_{i \leq r} \mathrm{err}_i \leq \frac{pc}{p-2} \leq \frac{pd}{p-2} \cdot \mathrm{e}^{-4q_2\sqrt{f/2}} = \frac{p}{p-2} \cdot \frac{1}{q_2^2}. \tag{2.4}$$

The second inequality holds because $\tau(r; q_1) \leq \tau(k; q_1)$. The last relation follows from the definition of f.

Now, if $r < k$, then we must have the bound $f/2 \geq \gamma_{r+1}$. By the definition (1.9) of γ_{r+1} and the normalization $a_n = 0$, this condition implies that $a_{r+1} \leq (1 - f/2)^{-1} a_k$. Therefore, Lemma 2.2 implies that

$$\begin{aligned} \mathbb{E} \max_{r < i \leq k} \mathrm{err}_i &\leq f + \int_f^1 \left[d\mathrm{e}^{-2(2q_2+1)\sqrt{\varepsilon/2}} \right]^{p/2} \mathrm{d}\varepsilon \\ &\leq f + d^{p/2} \int_f^\infty \mathrm{e}^{-2pq_2\sqrt{\varepsilon/2}} \mathrm{d}\varepsilon = f + d^{p/2} \mathrm{e}^{-2pq_2\sqrt{f/2}} \left[\frac{4\sqrt{f/2}}{pq_2} + \frac{1}{(pq_2)^2} \right]. \end{aligned}$$

We obtain the last identity by changing variables and integrating by parts; the same calculation appears in [Tro18a, Sec. 5.1]. Introduce the value of f and make some simple bounds to reach

$$\mathbb{E} \max_{r < i \leq k} \mathrm{err}_i \leq \frac{1}{q_2^2} \left[\frac{1}{p^2} + \frac{\log(q_2^2 d)}{p} + \frac{\log^2(q_2^2 d)}{8} \right]. \tag{2.5}$$

In particular, we have used the trivial estimate $q_2^p \geq 1$.

Combine (2.4) and (2.5) to arrive at

$$\begin{aligned} \mathbb{E} \max_{i \leq k} \mathrm{err}_i \leq \mathbb{E} \max_{i \leq r} \mathrm{err}_i + \mathbb{E} \max_{r < i \leq k} \mathrm{err}_i &\leq \frac{1}{q_2^2} \left[\frac{p}{p-2} + \frac{1}{p^2} + \frac{\log\left(q_2^2 d\right)}{p} + \frac{\log^2\left(q_2^2 d\right)}{8} \right] \\ &\leq \left[\frac{p/(p-2) + 8^{-1/2} \log\left(q_2^2 d\right)}{q_2} \right]^2. \end{aligned}$$

The second inequality follows from a short computation. We make a numerical bound for the remaining constant, and recall the definition of d to complete the proof of (2.2). □

Lemma 2.3 implies the expectation bounds in the main results.

Proof of Theorem 1.4(2) and Theorem 1.6(2). The expectation bounds stated in Theorem 1.4(2) and Theorem 1.6(2) are just restatements of Lemma 2.3. □

3. Technical Preparations

In this section, we lay the foundations for the proofs of Lemmas 2.1 and 2.2. We maintain the simplifications introduced in Section 2.1. Some of the constructions in this section are adapted from the paper [MM15], where they appear in a disguised form.

3.1. **The Extended Rayleigh matrix.** Define the orthogonal projector $\mathbf{\Pi} \in \mathbb{R}^{n\times n}$ whose range is the block Krylov subspace $K_q(\boldsymbol{A};\boldsymbol{\Omega})$. Consider the *extended Rayleigh matrix*:

$$\boldsymbol{H}_\circ := \mathbf{\Pi}\boldsymbol{A}\mathbf{\Pi} \in \mathbb{R}^{n\times n}. \tag{3.1}$$

Whereas the Rayleigh matrix $\boldsymbol{H}$ is a linear operator on the Krylov subspace, the extended Rayleigh matrix $\boldsymbol{H}_\circ$ is the trivial extension of $\boldsymbol{H}$ to all of $\mathbb{R}^n$. Since $\boldsymbol{A}$ is psd, the extended Ritz matrix $\boldsymbol{H}_\circ$ is also psd.

In (1.3), we defined the Ritz values ξ_r to be the eigenvalues of the Rayleigh matrix $\boldsymbol{H}$, listed in weakly decreasing order. Since $\boldsymbol{H}$ is psd, the Ritz values coincide with the eigenvalues of the extended Rayleigh matrix:

$$\xi_r = \lambda_r(\boldsymbol{H}_\circ) \quad \text{for } r = 1,2,3,\dots,\dim(K_q). \tag{3.2}$$

3.2. **Graded Spectral Projectors.** To perform the analysis, we construct a special family of orthogonal projectors that align with the largest invariant subspaces of the extended Rayleigh matrix $\boldsymbol{H}_\circ$. Fix an eigenvalue decomposition

$$\boldsymbol{H}_\circ = \textstyle\sum_{i=1}^n \xi_i \boldsymbol{u}_i\boldsymbol{u}_i^* \quad \text{where } \{\boldsymbol{u}_i\} \subset \mathbb{R}^n \text{ is an orthonormal basis of eigenvectors.} \tag{3.3}$$

Now, define the orthogonal projector

$$\mathbf{\Pi}_r := \textstyle\sum_{i=1}^r \boldsymbol{u}_i\boldsymbol{u}_i^* \quad \text{for each index } 0 \le r \le k. \tag{3.4}$$

This sequence of projectors has several immediate properties:

(i) The sequence is graded: $\operatorname{rank}(\mathbf{\Pi}_r) = r$ for each index $0 \le r \le k$.

(ii) The ranges are nested subspaces of the block Krylov space:

$$\{\mathbf{0}\} = \operatorname{range}(\mathbf{\Pi}_0) \subset \operatorname{range}(\mathbf{\Pi}_1) \subset \cdots \subset \operatorname{range}(\mathbf{\Pi}_k) \subset \operatorname{range}(\mathbf{\Pi}) = K_q(\boldsymbol{A};\boldsymbol{\Omega}).$$

(iii) The projectors give us alternative expressions for the Ritz values:

$$\xi_r = \lambda_r(\boldsymbol{H}_\circ) = \lambda_r(\mathbf{\Pi}_r\boldsymbol{H}_\circ\mathbf{\Pi}_r) = \lambda_r(\mathbf{\Pi}_r\boldsymbol{A}\mathbf{\Pi}_r).$$

The first two relations depend on (3.2), (3.3), and (3.4). Last, we introduced the definition (3.1) of $\boldsymbol{H}_\circ$ and invoked property (ii) to see that $\mathbf{\Pi}_r = \mathbf{\Pi}_r\mathbf{\Pi}$.

3.3. **A Variational Principle.** Our argument relies on a variational principle that highlights the distinguished role of the graded spectral projectors.

Lemma 3.1 (Maximum Principle)**.** *For indices with $1 \le i \le r \le k$, the orthogonal projector $\mathbf{\Pi}_r$ satisfies*

$$\lambda_i(\mathbf{\Pi}_r\boldsymbol{A}\mathbf{\Pi}_r) = \max\{\lambda_i(\boldsymbol{PAP}) : \operatorname{rank}(\boldsymbol{P}) \le r \text{ and } \operatorname{range}(\boldsymbol{P}) \subset \operatorname{range}(\mathbf{\Pi})\}.$$

The maximum extends over all orthogonal projectors $\boldsymbol{P} \in \mathbb{R}^{n\times n}$ that meet the constraints.

Proof. It is immediate that

$$\lambda_i(\mathbf{\Pi}_r\boldsymbol{A}\mathbf{\Pi}_r) \le \max\{\lambda_i(\boldsymbol{PAP}) : \operatorname{rank}(\boldsymbol{P}) \le r \text{ and } \operatorname{range}(\boldsymbol{P}) \subset \operatorname{range}(\mathbf{\Pi})\}.$$

Indeed, $\mathbf{\Pi}_r$ is a rank-r orthogonal projector whose range is a subset of $\operatorname{range}(\mathbf{\Pi})$ because of property (ii).

On the other hand, suppose that $\boldsymbol{P}$ is a rank-r orthogonal projector with $\operatorname{range}(\boldsymbol{P}) \subset \operatorname{range}(\mathbf{\Pi})$. Then

$$\lambda_i(\boldsymbol{PAP}) = \lambda_i(\boldsymbol{P}\mathbf{\Pi}\boldsymbol{A}\mathbf{\Pi}\boldsymbol{P}) = \lambda_i(\boldsymbol{PH}_\circ\boldsymbol{P}) \le \lambda_i(\boldsymbol{H}_\circ) = \lambda_i(\mathbf{\Pi}_i\boldsymbol{A}\mathbf{\Pi}_i) \le \lambda_i(\mathbf{\Pi}_r\boldsymbol{A}\mathbf{\Pi}_r).$$

The first identity holds because $\boldsymbol{P}\boldsymbol{\Pi} = \boldsymbol{P}$, and the second is the definition (3.1) of $\boldsymbol{H}_\circ$. The third is property (iii) of the projector $\boldsymbol{\Pi}_r$. The inequalities follow from the Cauchy interlacing theorem [Bha97, Thm. III.1.5], as well as property (ii). □

3.4. **Matrices with Few Eigenvalues.** We are now prepared to present a short proof of Proposition 1.3.

Proposition 1.3: Proof Sketch. Without loss of generality, assume that $\ell = k$. Let $\mu_1, \dots, \mu_m$ be the distinct eigenvalues of $\boldsymbol{A}$ that are strictly smaller than a_k. Construct a polynomial that annihilates the tail eigenvalues:

$$\varphi_0(s) := \prod_{i=1}^{m} \frac{s - \mu_i}{a_k - \mu_i} \in \mathscr{P}_m \subset \mathscr{P}_q.$$

Note that $\varphi_0(a_i) \geq 1$ for each $i \leq k$.

Let $\boldsymbol{P}_{\boldsymbol{Y}}$ be the rank-k orthogonal projector onto the range of the matrix $\boldsymbol{Y} := \varphi_0(\boldsymbol{A})\boldsymbol{\Omega}$, which is a subset of the block Krylov subspace $K_q(\boldsymbol{A}; \boldsymbol{\Omega})$ because of (1.1). With probability one over the randomness in the standard normal matrix $\boldsymbol{\Omega}$, the projector $\boldsymbol{P}_{\boldsymbol{Y}} = \mathbf{I}_k \oplus \mathbf{0}_{n-k}$. Therefore, $\boldsymbol{P}_{\boldsymbol{Y}} \boldsymbol{A} \boldsymbol{P}_{\boldsymbol{Y}} = \boldsymbol{A}_k$.

It now follows from the maximum principle, Lemma 3.1, that

$$a_i = \lambda_i(\boldsymbol{A}_k) = \lambda_i(\boldsymbol{P}_{\boldsymbol{Y}} \boldsymbol{A} \boldsymbol{P}_{\boldsymbol{Y}}) \leq \lambda_i(\boldsymbol{\Pi}_k \boldsymbol{A} \boldsymbol{\Pi}_k) = \lambda_i(\boldsymbol{H}_\circ) \quad \text{for } 1 \leq i \leq k.$$

The last identity is property (iii). In other words, largest k eigenvalues of the Rayleigh matrix (1.2) coincide with the largest k eigenvalues of $\boldsymbol{A}$. □

3.5. **Randomized Matrix Approximation.** The most important ingredients in the argument are two results on randomized matrix approximation adapted from the paper [HMT11]. The first result reduces the problem of randomized matrix approximation to a problem about random matrix theory.

Fact 3.2 (Randomized Matrix Approximation)**.** *Let $\boldsymbol{\Sigma} \in \mathbb{R}^{n \times n}$ be a diagonal matrix. Draw a standard normal test matrix $\boldsymbol{\Omega} \in \mathbb{R}^{n \times \ell}$, and form the matrix $\boldsymbol{Y} := \boldsymbol{\Sigma}\boldsymbol{\Omega}$. Construct the orthogonal projector $\boldsymbol{P}_{\boldsymbol{Y}} \in \mathbb{R}^{n \times n}$ onto the range of $\boldsymbol{Y}$.*

To study these objects, we decompose

$$\boldsymbol{\Sigma} = \begin{bmatrix} \boldsymbol{\Sigma}_1 & \mathbf{0} \\ \mathbf{0} & \boldsymbol{\Sigma}_2 \end{bmatrix} \quad \text{and} \quad \boldsymbol{\Omega} = \begin{bmatrix} \boldsymbol{\Omega}_1 \\ \boldsymbol{\Omega}_2 \end{bmatrix} \quad \text{where} \quad \boldsymbol{\Sigma}_1 \in \mathbb{R}^{k \times k} \text{ and } \boldsymbol{\Omega}_1 \in \mathbb{R}^{k \times \ell}.$$

Note that $\boldsymbol{\Omega}_1$ and $\boldsymbol{\Omega}_2$ are statistically independent, standard normal matrices. Form a matrix

$$\boldsymbol{R} := \begin{bmatrix} \boldsymbol{R}_1 & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix} \in \mathbb{R}^{n \times n} \quad \text{where } \boldsymbol{R}_1 \in \mathbb{R}^{k \times k} \text{ is an arbitrary matrix.}$$

Then, with probability one,

$$\|(\mathbf{I} - \boldsymbol{P}_{\boldsymbol{Y}})\boldsymbol{\Sigma}\boldsymbol{R}\|_{\mathrm{F}} \leq \|\boldsymbol{\Sigma}_2 \boldsymbol{\Omega}_2 \boldsymbol{\Omega}_1^{\dagger} \boldsymbol{R}_1\|_{\mathrm{F}}.$$

This statement follows from an easy modification of the argument in [HMT11, Thm .9.1] and the fact that a standard normal matrix has full rank with probability one. See also [TYUC17, Prop. 9.2].

The second result encapsulates the random matrix theory required to invoke Fact 3.2.

Fact 3.3 (Random Matrix Bound)**.** *Assume that $p := \ell - k \geq 2$. Let $\boldsymbol{M} \in \mathbb{R}^{m \times m}$ be an arbitrary matrix, and let $\boldsymbol{R}_1 \in \mathbb{R}^{k \times k}$ be a rank-r orthogonal projector. Let $\boldsymbol{\Omega}_1 \in \mathbb{R}^{k \times \ell}$ and $\boldsymbol{\Omega}_2 \in \mathbb{R}^{m \times \ell}$ be statistically independent, standard normal matrices. For all $t > 0$,*

$$\mathbb{P}\left\{\|\boldsymbol{M}\boldsymbol{\Omega}_2 \boldsymbol{\Omega}_1^{\dagger} \boldsymbol{R}_1\|_{\mathrm{F}}^2 \geq t\right\} \leq \left[\frac{2.25 r \|\boldsymbol{M}\|_{\mathrm{F}}^2}{t}\right]^{p/2}.$$

Qualitatively different bounds hold when $p \in \{0, 1\}$.

This statement is closely related to the bound [HMT11, Thm. 10.1], but it contains some refinements that are valuable for our purposes. See Appendix A for a sketch of the proof of Fact 3.3.

Remark 3.4 (Spectral Norm)**.** We have chosen to work with the Frobenius norm for ease of argument. Related results hold for the spectral norm. The spectral norm bounds and their proofs are more complicated, and they do not offer a significant benefit unless the oversampling p is proportional to the number k of eigenvalues we wish to estimate.

4. Approximation of Large Eigenvalues

In this section, we prove Lemma 2.1, which describes how the randomized block Krylov method approximates eigenvalues of the input matrix that are far above the tail eigenvalues.

The overall argument is similar in spirit to the proof of [MM15, Lem. 9(1)], but all of the technical details are different. In particular, our approach depends on ideas from Kuczyński & Woźniakowski [KW92], Halko et al. [HMT11], and the companion paper [Tro18a].

4.1. **Normalization of Spectrum.** Lemma 2.1 is invariant under affine transformations of the spectrum of the input matrix $\boldsymbol{A}$. For the proof, it is convenient to scale the input matrix so that

$$a_{k+1} = 1 \quad \text{and} \quad a_n = 0. \tag{4.1}$$

This action is possible because of our assumption that the tail of the spectrum contains at least two distinct values.

4.2. **Setup.** Fix a parameter $\varepsilon \in (0,1]$, and choose an index $r \le k$. We will produce a bound for the probability

$$P_1 := \mathbb{P}\left\{\max_{i\le r} \mathrm{err}_i \ge \varepsilon\right\} = \mathbb{P}\left\{\max_{i\le r} \frac{a_i - \xi_i}{a_i} \ge \varepsilon\right\} \le \mathbb{P}\left\{\max_{i\le r} (a_i - \xi_i) \ge \varepsilon a_r\right\}. \tag{4.2}$$

We have used the definition (1.4) of the error and the fact $a_i \ge a_r$. We will explain how to use methods from randomized matrix approximation to accomplish this goal.

4.3. **Construction of a Projector.** Fix a filter polynomial $\varphi \in \mathscr{P}_q$ that we will describe later. Introduce the matrix $\boldsymbol{Y} := \varphi(\boldsymbol{A})\boldsymbol{\Omega}$, whose range is contained in the Krylov subspace $K_q(\boldsymbol{A};\boldsymbol{\Omega})$ because of (1.1). Let $\boldsymbol{P}_{\boldsymbol{Y}}$ be the orthogonal projector onto the range of $\boldsymbol{Y}$.

Fix an index $i \le r$. Form the rank-i orthogonal projector $\boldsymbol{P}_i$ onto the range of the matrix $\boldsymbol{P}_{\boldsymbol{Y}}\boldsymbol{A}_i$, which is also contained in the block Krylov subspace. In particular, note that $\boldsymbol{P}_i\boldsymbol{A}_i = \boldsymbol{P}_{\boldsymbol{Y}}\boldsymbol{A}_i$. We use this construction to develop a lower bound on the ith Ritz value:

$$\xi_i = \lambda_i(\boldsymbol{\Pi}_i\boldsymbol{A}\boldsymbol{\Pi}_i) \ge \lambda_i(\boldsymbol{P}_i\boldsymbol{A}\boldsymbol{P}_i) \ge \lambda_i(\boldsymbol{P}_i\boldsymbol{A}_i\boldsymbol{P}_i) = \lambda_i(\boldsymbol{P}_{\boldsymbol{Y}}\boldsymbol{A}_i\boldsymbol{P}_{\boldsymbol{Y}}) = \lambda_i\big(\boldsymbol{A}_i^{1/2}\boldsymbol{P}_{\boldsymbol{Y}}\boldsymbol{A}_i^{1/2}\big).$$

The first relation is property (iii) of the spectral projector $\boldsymbol{\Pi}_r$. The inequalities follow from the maximum principle, Lemma 3.1, and the semidefinite inequality $\boldsymbol{A} \succcurlyeq \boldsymbol{A}_i$. The next relation depends on the construction of the projector $\boldsymbol{P}_i$. The last step is the identity $\lambda_i(\boldsymbol{C}\boldsymbol{C}^*) = \lambda_i(\boldsymbol{C}^*\boldsymbol{C})$, valid for every square matrix $\boldsymbol{C}$. As usual, $(\cdot)^{1/2}$ denotes the psd square root of a psd matrix.

To control the ith error, observe that $a_i = \lambda_i(\boldsymbol{A}_i)$. Using the last display, we discover that

$$\begin{aligned} a_i - \xi_i &\le \lambda_i(\boldsymbol{A}_i) - \lambda_i\big(\boldsymbol{A}_i^{1/2}\boldsymbol{P}_{\boldsymbol{Y}}\boldsymbol{A}_i^{1/2}\big) \\ &\le \lambda_{\max}\big(\boldsymbol{A}_i - \boldsymbol{A}_i^{1/2}\boldsymbol{P}_{\boldsymbol{Y}}\boldsymbol{A}_i^{1/2}\big) = \lambda_{\max}\big(\boldsymbol{A}_i^{1/2}(\mathbf{I} - \boldsymbol{P}_{\boldsymbol{Y}})\boldsymbol{A}_i^{1/2}\big) \\ &= \lambda_{\max}\big((\mathbf{I} - \boldsymbol{P}_{\boldsymbol{Y}})\boldsymbol{A}_i(\mathbf{I} - \boldsymbol{P}_{\boldsymbol{Y}})\big) \le \lambda_{\max}\big((\mathbf{I} - \boldsymbol{P}_{\boldsymbol{Y}})\boldsymbol{A}_r(\mathbf{I} - \boldsymbol{P}_{\boldsymbol{Y}})\big). \end{aligned}$$

The second inequality follows from Lidskii's theorem [Bha97, Eqn. (III.12)]. The last inequality holds because $\boldsymbol{A}_i \preccurlyeq \boldsymbol{A}_r$. Take the maximum over $i \le r$ to achieve

$$\max_{i\le r}(a_i - \xi_i) \le \lambda_{\max}\big((\mathbf{I} - \boldsymbol{P}_{\boldsymbol{Y}})\boldsymbol{A}_r(\mathbf{I} - \boldsymbol{P}_{\boldsymbol{Y}})\big). \tag{4.3}$$

See [MM15, p. 10] for a related argument.

4.4. **Reduction to Randomized Matrix Approximation.** Let us manipulate the expression (4.3) so that we can control it using Fact 3.2, the result on randomized matrix approximation. To do so, we must replace the matrix $\boldsymbol{A}_r$ with the filtered matrix $\varphi^2(\boldsymbol{A}_r)$.

Before continuing, it is necessary to place some restrictions on the filter polynomial $\varphi \in \mathscr{P}_q$. We require that $\varphi(0) = 0$ and that the function $s \mapsto s/\varphi^2(s)$ is finite and decreasing for $s \geq 1$.

Since $r \leq k$, we have the bound $a_r \geq 1$ because of the normalization (4.1). The monotonicity of $s \mapsto s/\varphi^2(s)$ implies that

$$a_i = \frac{a_i}{\varphi^2(a_i)} \cdot \varphi^2(a_i) \leq \frac{a_r}{\varphi^2(a_r)} \cdot \varphi^2(a_i) \quad \text{for each index } 1 \leq i \leq r.$$

Owing to the facts that $\boldsymbol{A}_r = \operatorname{diag}(a_1, \dots, a_r, 0, \dots, 0)$ and $\varphi(0) = 0$, these inequalities yield

$$\boldsymbol{A}_r \preccurlyeq \frac{a_r}{\varphi^2(a_r)} \cdot \varphi^2(\boldsymbol{A}_r).$$

Introduce this semidefinite inequality into (4.3) to arrive at

$$\begin{aligned} \max_{i \leq r} (a_i - \xi_i) &\leq \frac{a_r}{\varphi^2(a_r)} \cdot \lambda_{\max}\big((\mathbf{I} - \boldsymbol{P}_{\boldsymbol{Y}})\varphi^2(\boldsymbol{A}_r)(\mathbf{I} - \boldsymbol{P}_{\boldsymbol{Y}})\big) \\ &= \frac{a_r}{\varphi^2(a_r)} \cdot \big\|(\mathbf{I} - \boldsymbol{P}_{\boldsymbol{Y}})\varphi(\boldsymbol{A}_r)\big\|^2 \leq \frac{a_r}{\varphi^2(a_r)} \cdot \big\|(\mathbf{I} - \boldsymbol{P}_{\boldsymbol{Y}})\varphi(\boldsymbol{A}_r)\big\|_{\mathrm{F}}^2. \end{aligned} \tag{4.4}$$

We have replaced the spectral norm with the Frobenius norm to facilitate the rest of the argument. This expression has the same form as the quantity that is controlled by Fact 3.2. See [HMT11, Thm. 9.2] and [MM15, p. 11] for related arguments.

4.5. **Randomized Matrix Approximation.** We proceed to study the probability (4.2) that the error is large:

$$P_1 \leq \mathbb{P}\left\{\max_{i \leq r} (a_i - \xi_i) \geq \varepsilon a_r\right\} \leq \mathbb{P}\left\{\big\|(\mathbf{I} - \boldsymbol{P}_{\boldsymbol{Y}})\varphi(\boldsymbol{A}_r)\big\|_{\mathrm{F}}^2 \geq \varepsilon\varphi^2(a_r)\right\}. \tag{4.5}$$

The second relation is a consequence of (4.4).

We plan to invoke Facts 3.2 and 3.3 to control the right-hand side of (4.5). Introduce the matrix $\boldsymbol{\Sigma} := \varphi(\boldsymbol{A})$, and recall that $\boldsymbol{Y} = \boldsymbol{\Sigma}\boldsymbol{\Omega}$. The leading $k \times k$ principal submatrix $\boldsymbol{\Sigma}_1 = \operatorname{diag}(\varphi(a_1), \dots, \varphi(a_k))$, and the complementary principal submatrix $\boldsymbol{\Sigma}_2 = \operatorname{diag}(\varphi(a_{k+1}), \dots, \varphi(a_n))$. Let $\boldsymbol{R}_1 \in \mathbb{R}^{k \times k}$ be the rank-r orthogonal projector onto the first r coordinates, and let $\boldsymbol{R} \in \mathbb{R}^{n \times n}$ be the trivial extension of $\boldsymbol{R}_1$. In particular, $\boldsymbol{\Sigma}\boldsymbol{R} = \varphi(\boldsymbol{A}_r)$.

With this notation, we can rewrite the bound (4.5) for P_1, and Fact 3.2 reveals that

$$P_1 \leq \mathbb{P}\left\{\|(\mathbf{I} - \boldsymbol{P}_{\boldsymbol{Y}})\boldsymbol{\Sigma}\boldsymbol{R}\|_{\mathrm{F}}^2 \geq \varepsilon\varphi^2(a_r)\right\} \leq \mathbb{P}\left\{\|\boldsymbol{\Sigma}_2\boldsymbol{\Omega}_2\boldsymbol{\Omega}_1^{\dagger}\boldsymbol{R}_1\|_{\mathrm{F}}^2 \geq \varepsilon\varphi^2(a_r)\right\}.$$

Now, invoke Fact 3.3 to arrive at

$$P_1 \leq \left[\frac{2.25r\,\|\boldsymbol{\Sigma}_2\|_{\mathrm{F}}^2}{\varepsilon\varphi^2(a_r)}\right]^{p/2} = \left[\frac{2.25r\sum_{i>k}\varphi^2(a_i)}{\varepsilon\varphi^2(a_r)}\right]^{p/2}. \tag{4.6}$$

It remains to identify a filter polynomial φ to make this bound effective.

4.6. **The Filter Polynomial.** We wish to design a filter polynomial that attenuates the eigenvalues of $\boldsymbol{A}$ in the tail (with index $i > k$) so that they are summable. At the same time, we want the filter polynomial to amplify the leading eigenvalues (with index $i \leq r$) so they are much larger than those in the tail.

Let $q = q_1 + q_2$ be an integer partition of the depth parameter q with $1 \leq q_1 \leq q$. We consider the filter polynomial

$$\varphi_1(s) := s^{q_1} T_{q_2}(2s - 1) \in \mathscr{P}_q, \tag{4.7}$$

where T_{q_2} is the Chebyshev polynomial of the first kind with degree q_2. As required, $\varphi_1(0) = 0$. Since $\varphi_1(1) = 1$ and φ_1 is increasing in the interval $[1, \infty)$, we discover that $s \mapsto s/\varphi_1^2(s)$ is decreasing for $s \geq 1$. The minimax property of the Chebyshev polynomial ensures that

$$\varphi_1^2(s) \leq s^{2q_1} \quad \text{for } 0 \leq s \leq 1. \tag{4.8}$$

Meanwhile, the growth properties of the Chebyshev polynomial yield the bound

$$\frac{1}{\varphi_1^2(s)} \le 4s^{-2q_1} \mathrm{e}^{-4q_2\sqrt{1-1/s}} \quad \text{for } s \ge 1. \tag{4.9}$$

The filter polynomial φ_1 also arises in the companion paper [Tro18a, App. A.2.3], which contains a short proof of its properties.

4.7. **Installing the Filter Polynomial.** We are now prepared to instantiate the error bound (4.6) with the filter polynomial (4.7). Since $a_{k+1} = 1 \le a_k$, the properties of φ_1 yield

$$\frac{\sum_{j>k} \varphi_1^2(a_j)}{\varphi_1^2(a_r)} \le 4 \cdot \left[\sum\nolimits_{j>k} (a_j/a_r)^{2q_1}\right] \cdot \mathrm{e}^{-4q_2\sqrt{1-1/a_r}} = 4\tau(r;q_1) \cdot \mathrm{e}^{-4q_2\sqrt{\gamma_r}}.$$

We have applied the bound (4.8) to the numerator and the bound (4.9) to the denominator. Then, recalling the normalization (4.1), we identified the tail spectral gap (1.9) and the tail content function (1.10). Introducing this bound into (4.6), we arrive at the estimate

$$P_1 = \mathbb{P}\left\{\max_{i\le r} \mathrm{err}_i \ge \varepsilon\right\} \le \left[\frac{9r\tau(r;q_1)}{\varepsilon} \cdot \mathrm{e}^{-4q_2\sqrt{\gamma_r}}\right]^{p/2}.$$

We have recalled the definition (4.2) of the probability P_1. This is the advertised conclusion of Lemma 2.1.

5. Approximation of Small Eigenvalues

In this section, we establish Lemma 2.2, which shows that we are unlikely to underestimate eigenvalues that are just above the tail of the input matrix.

This result is similar in flavor to [MM15, Lem. 9(2)], but the proof of completely new. The argument uses the ideas from the companion paper [Tro18a, Sec. 4] to reduce the problem to a question about randomized matrix approximation. Then we invoke the results from [HMT11] to dispatch this problem. We also rely on a polynomial construction from Kuczyński & Woźniakowski [KW92].

5.1. **Normalization of Spectrum.** Lemma 2.2 is invariant under affine transformations of the spectrum of the input matrix $\boldsymbol{A}$. For this proof, it is convenient to rescale the input matrix so that

$$a_k = 1 \quad \text{and} \quad a_n = 0. \tag{5.1}$$

As before, we have relied on the assumption that the tail of the spectrum contains at least two distinct values.

5.2. **Setup.** Fix a parameter $\varepsilon \in (0,1)$. Let $r \le k$ be the smallest index where $a_r \le (1-\varepsilon/2)^{-1}$. Suppose that

$$\mathrm{err}_k = 1 - \xi_k < \varepsilon/2.$$

The form of the error follows from (1.4) after applying the normalization (5.1). As a consequence, for any index i in the range $r \le i \le k$,

$$\xi_i \ge \xi_k > 1 - \varepsilon/2 \ge (1-\varepsilon/2)^2 a_r \ge (1-\varepsilon/2)^2 a_i$$

The first inequality holds because the Ritz values are decreasing. The next two inequalities hold by the assumptions of err_k and on a_r. The last inequality holds because $a_i \le a_r$. It follows that

$$\mathrm{err}_i = \frac{a_i - \xi_i}{a_i} < 1 - (1-\varepsilon/2)^2 < \varepsilon.$$

In summary,

$$\mathrm{err}_k < \varepsilon/2 \quad \text{implies} \quad \max_{r\le i\le k} \mathrm{err}_i < \varepsilon.$$

Therefore,

$$P_2 := \mathbb{P}\left\{\max_{r\le i\le k} \mathrm{err}_i \ge \varepsilon\right\} \le \mathbb{P}\left\{\mathrm{err}_k \ge \varepsilon/2\right\} = \mathbb{P}\left\{\xi_k \ge 1 - \varepsilon/2\right\}.$$

We will develop a bound for the latter probability.

Suppose that $\boldsymbol{V} \in \mathbb{R}^{n\times k}$ is an orthonormal matrix whose range is contained in the block Krylov subspace $K_q(\boldsymbol{A};\boldsymbol{\Omega})$. Then $\boldsymbol{P} := \boldsymbol{V}\boldsymbol{V}^*$ is a rank-k orthogonal projector with the same range as $\boldsymbol{V}$. In view of the maximum principle, Lemma 3.1,

$$\xi_k = \lambda_k(\boldsymbol{\Pi}_k \boldsymbol{A} \boldsymbol{\Pi}_k) \geq \lambda_k(\boldsymbol{P}\boldsymbol{A}\boldsymbol{P}) = \lambda_{\min}(\boldsymbol{V}^*\boldsymbol{A}\boldsymbol{V}).$$

The first relation is property (iii). It follows that

$$P_2 \leq \mathbb{P}\left\{\lambda_{\min}(\boldsymbol{V}^*\boldsymbol{A}\boldsymbol{V}) \geq 1 - \varepsilon/2\right\}. \tag{5.2}$$

It suffices to produce such a matrix $\boldsymbol{V}$ and to bound the resulting probability.

5.3. **Construction of an Orthonormal Matrix.** As in the statement of Fact 3.2, we decompose the input matrix and the test matrix:

$$\boldsymbol{A} =: \begin{bmatrix} \boldsymbol{\Lambda}_1 & \boldsymbol{0} \\ \boldsymbol{0} & \boldsymbol{\Lambda}_2 \end{bmatrix} \quad \text{and} \quad \boldsymbol{\Omega} =: \begin{bmatrix} \boldsymbol{\Omega}_1 \\ \boldsymbol{\Omega}_2 \end{bmatrix} \quad \text{where } \boldsymbol{\Lambda}_1 \in \mathbb{R}^{k\times k} \text{ and } \boldsymbol{\Omega}_1 \in \mathbb{R}^{k\times \ell}.$$

The matrix $\boldsymbol{A}$ is diagonal with weakly decreasing entries, so the matrix $\boldsymbol{\Lambda}_1$ lists the k largest eigenvalues of $\boldsymbol{A}$. Since $\boldsymbol{\Omega}$ is standard normal, $\boldsymbol{\Omega}_1$ and $\boldsymbol{\Omega}_2$ are statistically independent standard normal matrices. In particular, $\boldsymbol{\Omega}_1^\dagger$ has full row-rank with probability one.

As before, we fix a filter polynomial $\varphi \in \mathscr{P}_q$ that we will describe later. For now, we just insist that $\varphi(1) = 1$ and that φ is increasing in the interval $[1,\infty)$. In particular, φ has no roots in this interval.

Introduce the matrix $\boldsymbol{Y} := \varphi(\boldsymbol{A})\boldsymbol{\Omega} \in \mathbb{R}^{n\times\ell}$, and note that the range of $\boldsymbol{Y}$ is a subspace of the block Krylov space $K_q(\boldsymbol{A};\boldsymbol{\Omega})$ because of (1.1). Now, define the reduced matrix

$$\boldsymbol{Z} := \boldsymbol{Y}\boldsymbol{\Omega}_1^\dagger = \begin{bmatrix} \varphi(\boldsymbol{\Lambda}_1) & \boldsymbol{0} \\ \boldsymbol{0} & \varphi(\boldsymbol{\Lambda}_2) \end{bmatrix} \begin{bmatrix} \boldsymbol{I} \\ \boldsymbol{F} \end{bmatrix} \in \mathbb{R}^{n\times k} \quad \text{where } \boldsymbol{F} := \boldsymbol{\Omega}_2\boldsymbol{\Omega}_1^\dagger. \tag{5.3}$$

By construction, the range of $\boldsymbol{Z}$ is a k-dimensional subspace of the block Krylov space. Finally, introduce a matrix with orthonormal columns:

$$\boldsymbol{V} := \boldsymbol{Z}(\boldsymbol{Z}^*\boldsymbol{Z})^{-1/2} \in \mathbb{R}^{n\times k}. \tag{5.4}$$

The range of $\boldsymbol{V}$ coincides with the range of $\boldsymbol{Z}$. The construction of $\boldsymbol{V}$ is adapted from [HMT11, Sec. 9].

Using this random matrix $\boldsymbol{V}$, we will produce a bound for the right-hand side of (5.2). The approach combines ideas from [HMT11, Secs. 9–10] and [Tro18a, Sec. 4].

5.4. **Reduction to Randomized Matrix Approximation.** We will transform the probability P_2 into a quantity that we can control using tools from randomized matrix approximation. Introduce the abbreviation $\beta := 1 - \varepsilon/2$. To begin, collect everything in (5.2) on the same side of the event:

$$P_2 \leq \mathbb{P}\left\{\lambda_{\min}\big(\boldsymbol{V}^*(\boldsymbol{A} - \beta\boldsymbol{I})\boldsymbol{V}\big) \leq 0\right\}.$$

This relation depends on the fact that $\boldsymbol{V}$ has orthonormal columns. Next, substitute in the expression (5.4) for $\boldsymbol{V}$, and invoke Sylvester's law of inertia [HJ90, Thm. 4.5.8]:

$$P_2 \leq \mathbb{P}\left\{\lambda_{\min}\big((\boldsymbol{Z}^*\boldsymbol{Z})^{-1/2}\boldsymbol{Z}^*(\boldsymbol{A} - \beta\boldsymbol{I})\boldsymbol{Z}(\boldsymbol{Z}^*\boldsymbol{Z})^{-1/2}\big) \leq 0\right\} = \mathbb{P}\left\{\lambda_{\min}\big(\boldsymbol{Z}^*(\boldsymbol{A} - \beta\boldsymbol{I})\boldsymbol{Z}\big) \leq 0\right\}.$$

Indeed, conjugating by the matrix $(\boldsymbol{Z}^*\boldsymbol{Z})^{1/2}$ cannot change the sign of any eigenvalue.

We can obtain a lower bound on the minimum eigenvalue by making everything smaller in the semidefinite order. Since $\varphi(1) = 1$ and φ^2 is increasing on the interval $[1, \infty)$, we may calculate that

$$\begin{aligned}
\boldsymbol{Z}^*(\boldsymbol{A} - \beta\mathbf{I})\boldsymbol{Z} &= \begin{bmatrix} \mathbf{I} \\ \boldsymbol{F} \end{bmatrix}^* \begin{bmatrix} \varphi^2(\boldsymbol{\Lambda}_1)(\boldsymbol{\Lambda}_1 - \beta\mathbf{I}) & \mathbf{0} \\ \mathbf{0} & \varphi^2(\boldsymbol{\Lambda}_2)(\boldsymbol{\Lambda}_2 - \beta\mathbf{I}) \end{bmatrix} \begin{bmatrix} \mathbf{I} \\ \boldsymbol{F} \end{bmatrix} \\
&\succcurlyeq \begin{bmatrix} \mathbf{I} \\ \boldsymbol{F} \end{bmatrix}^* \begin{bmatrix} (1-\beta) \cdot \varphi^2(\boldsymbol{\Lambda}_1) & \mathbf{0} \\ \mathbf{0} & \varphi^2(\boldsymbol{\Lambda}_2)(\boldsymbol{\Lambda}_2 - \beta\mathbf{I})_- \end{bmatrix} \begin{bmatrix} \mathbf{I} \\ \boldsymbol{F} \end{bmatrix} \\
&\succcurlyeq \begin{bmatrix} \mathbf{I} \\ \boldsymbol{F} \end{bmatrix}^* \begin{bmatrix} (\varepsilon/2) \cdot \mathbf{I} & \mathbf{0} \\ \mathbf{0} & -\varphi^2(\boldsymbol{\Lambda}_2)(\beta\mathbf{I} - \boldsymbol{\Lambda}_2)_+ \end{bmatrix} \begin{bmatrix} \mathbf{I} \\ \boldsymbol{F} \end{bmatrix} \\
&= (\varepsilon/2) \cdot \mathbf{I} - \boldsymbol{F}^* \varphi^2(\boldsymbol{\Lambda}_2)(\beta\mathbf{I} - \boldsymbol{\Lambda}_2)_+ \boldsymbol{F}.
\end{aligned}$$

We have used the fact that the entries of $\boldsymbol{\Lambda}_1 = \operatorname{diag}(a_1, \dots, a_k)$ exceed $a_k = 1$ and that $1 - \beta = \varepsilon/2$. The function $(\cdot)_+$ returns the positive-semidefinite part of a symmetric matrix, while $(\cdot)_-$ returns the negative-semidefinite part.

Combine the last two displays to arrive at the bound

$$\begin{aligned}
P_2 &\le \mathbb{P}\left\{\lambda_{\min}\left((\varepsilon/2) \cdot \mathbf{I} - \boldsymbol{F}^* \varphi^2(\boldsymbol{\Lambda}_2)(\beta\mathbf{I} - \boldsymbol{\Lambda}_2)_+ \boldsymbol{F}\right) \le 0\right\} \\
&= \mathbb{P}\left\{\lambda_{\max}(\boldsymbol{F}^* \varphi^2(\boldsymbol{\Lambda}_2)(\beta\mathbf{I} - \boldsymbol{\Lambda}_2)_+ \boldsymbol{F}) \ge \varepsilon/2\right\}.
\end{aligned}$$

This formula has a close analogy with [Tro18a, Eqn. (4.1)]. Let us take some additional steps to bring this probability into a form that we can control using Fact 3.3.

$$\begin{aligned}
P_2 &\le \mathbb{P}\left\{\left\|\varphi(\boldsymbol{\Lambda}_2)(\beta\mathbf{I} - \boldsymbol{\Lambda}_2)_+^{1/2} \boldsymbol{F}\right\|^2 \ge \varepsilon/2\right\} \\
&\le \mathbb{P}\left\{\left\|\varphi(\boldsymbol{\Lambda}_2)(\beta\mathbf{I} - \boldsymbol{\Lambda}_2)_+^{1/2} \boldsymbol{F}\right\|_{\mathrm{F}}^2 \ge \varepsilon/2\right\} \\
&= \mathbb{P}\left\{\left\|\varphi(\boldsymbol{\Lambda}_2)(\beta\mathbf{I} - \boldsymbol{\Lambda}_2)_+^{1/2} \boldsymbol{\Omega}_2 \boldsymbol{\Omega}_1^\dagger\right\|_{\mathrm{F}}^2 \ge \varepsilon/2\right\}
\end{aligned} \tag{5.5}$$

We have replaced the spectral norm by Frobenius norm to avoid extra work. In the last step, we have recalled the definition (5.3) of the matrix $\boldsymbol{F}$.

5.5. **Randomized Matrix Approximation.** We are now prepared to invoke results on randomized matrix approximation to bound the probability (5.5). Apply Fact 3.3 with $\boldsymbol{M} := \varphi(\boldsymbol{\Lambda}_2)(\beta\mathbf{I} - \boldsymbol{\Lambda}_2)_+^{1/2}$ and with $\boldsymbol{R}_1 := \mathbf{I}_k$. It emerges that

$$P_2 \le \left[\frac{4.5k}{\varepsilon} \cdot \left\|\varphi(\boldsymbol{\Lambda}_2)(\beta\mathbf{I} - \boldsymbol{\Lambda}_2)_+^{1/2}\right\|_{\mathrm{F}}^2\right]^{p/2} = \left[\frac{4.5k}{\varepsilon} \sum\nolimits_{j>k} \varphi^2(a_j)(\beta - a_j)_+\right]^{p/2}. \tag{5.6}$$

It remains to select the filter polynomial φ.

5.6. **The Filter Polynomial.** We must choose a polynomial φ that attenuates the eigenvalues in the tail ($a_j \le \beta$). Note that the analysis has removed the influence of the eigenvalues in the range $[\beta, 1]$.

Let $q = q_1 + q_2$ be an integer partition of the depth parameter q. Consider the filter polynomial

$$\varphi_2(s) := \frac{s^{q_1} U_{2q_2}\left(\sqrt{s/\beta}\right)}{U_{2q_2}\left(\sqrt{1/\beta}\right)} \in \mathscr{P}_q.$$

In this expression, U_{2q_2} denotes the (even) Chebyshev polynomial of the second kind with degree $2q_2$. As required, $\varphi(1) = 1$ and φ_2 is increasing on the interval $[1, \infty)$. The growth properties of the Chebyshev polynomial imply that

$$\varphi_2^2(s)(\beta - s) \le 2s^{2q_1} \cdot \frac{\varepsilon\delta}{(1-\delta)^2} \quad \text{for } 0 \le s \le \beta \quad \text{where} \quad \delta := \mathrm{e}^{-2(2q_2+1)\sqrt{\varepsilon/2}}. \tag{5.7}$$

The filter polynomial φ_2 also appears in the companion paper [Tro18a, App. A.2.5], which contains a short proof of these properties.

5.7. **Installing the Filter Polynomial.** We are ready to instantiate the error bound (5.6) with the filter polynomial φ_2. Recall that the diagonal matrix $\mathbf{\Lambda}_2$ lists the eigenvalues of $\boldsymbol{A}$, beginning with the $(k+1)$th largest. The uniform bound (5.7) ensures that

$$\sum\nolimits_{j>k} \varphi_2^2(a_j)(\beta - a_j)_+ \le \frac{2\varepsilon\delta}{(1-\delta)^2} \cdot \sum\nolimits_{j>k} a_j^{2q_1} = \frac{2\varepsilon\delta}{(1-\delta)^2} \cdot \tau(k; q_1).$$

The second relation depend on the fact that $\beta \le 1 = a_k$. We have used the normalization (5.1) to identify the tail content (1.10). Introduce the last display into (5.6) to obtain

$$P_2 \le \left[k\tau(k; q_1) \cdot \frac{9\delta}{(1-\delta)^2} \right]^{p/2}.$$

To complete the proof, we just need to make a simplification.

The last part of the argument hinges on the fact that $P_2 \le 1$ because it is a probability. Thus, we can improve the last bound to

$$P_2 \le \left[1 \wedge \left(k\tau(k; q_1) \cdot \frac{9\delta}{(1-\delta)^2} \right) \right]^{p/2}.$$

It is convenient to write $\vee$ for the maximum. Observe that

$$1 \wedge \left(k\tau(k; q_1) \cdot \frac{9\delta}{(1-\delta)^2} \right) \le \big(1 \vee k\tau(k; q_1)\big) \cdot \left(1 \wedge \frac{9\delta}{(1-\delta)^2} \right) \le \big(1 \vee k\tau(k; q_1)\big) \cdot (1 \wedge 11\delta).$$

Both relations are numerical inequalities. Combine the last two displays to arrive at the bound

$$P_2 \le \big[11 \cdot (1 \vee k\tau(k; q_1)) \cdot \delta \big]^{p/2}.$$

Finally, recall the definition (5.7) of δ to arrive at the statement of Lemma 2.2.

Appendix A. Eine Kleine Random Matrix Theory

This appendix contains a sketch of the proof of Fact 3.3. The argument requires several additional ingredients.

Fact A.1 (Moments of a Psd Gaussian Quadratic Form)**.** *Let $\boldsymbol{D} \in \mathbb{R}^{m \times m}$ be a psd matrix, and let $\boldsymbol{g} \in \mathbb{R}^m$ be a standard normal vector. For $s \ge 1$,*

$$\big[\mathbb{E} \big(\boldsymbol{g}^* \boldsymbol{D} \boldsymbol{g} \big)^p \big]^{1/s} \le s\sqrt{2} \cdot \big[\mathbb{E}\, \boldsymbol{g}^* \boldsymbol{D} \boldsymbol{g} \big].$$

Proof. The case $s = 1$ is immediate. We will establish the result for each $s \ge 2$. To treat the interval $s \in (1, 2)$, we apply Hölder's inequality to interpolate between the bounds for $s = 1$ and $s = 2$.

Suppose that $s \ge 2$. By rotational invariance, we may as well assume that $\boldsymbol{D} = \operatorname{diag}(d_1, \dots, d_m)$ with each $d_i \ge 0$. Calculate that

$$E := \big[\mathbb{E} \big(\boldsymbol{g}^* \boldsymbol{D} \boldsymbol{g} \big)^s \big]^{1/s} = \left[\mathbb{E} \left| \sum\nolimits_{i=1}^m d_i g_i^2 \right|^s \right]^{1/s} \le \sum\nolimits_{i=1}^m d_i + \left[\mathbb{E} \left| \sum\nolimits_{i=1}^m d_i (g_i^2 - 1) \right|^s \right]^{1/s}.$$

The last relation is the triangle inequality. The second term is a (homogeneous) second-order Gaussian chaos. By hypercontractivity [LT11, Sec. 3.2],

$$\begin{aligned} E &\le \operatorname{tr} \boldsymbol{D} + (s-1) \left[\mathbb{E} \left| \sum\nolimits_{i=1}^m d_i (g_i^2 - 1) \right|^2 \right]^{1/2} \\ &= \operatorname{tr} \boldsymbol{D} + (s-1) \big[2 \sum\nolimits_{i=1}^m d_i^2 \big]^{1/2} \le \Big(1 + (s-1)\sqrt{2} \Big) \operatorname{tr} \boldsymbol{D} \le s\sqrt{2} \cdot \big(\mathbb{E}\, \boldsymbol{g}^* \boldsymbol{D} \boldsymbol{g} \big). \end{aligned}$$

The second inequality follows from the fact that ℓ_s norms are monotonically decreasing as s increases. Finally, we note that $\mathbb{E}\, \boldsymbol{g}^* \boldsymbol{D} \boldsymbol{g} = \operatorname{tr} \boldsymbol{D}$. □

We also need information about the Frobenius norm of the pseudoinverse of a standard normal matrix.

Fact A.2 (Gaussian Pseudoinverse)**.** *Let $\boldsymbol{G} \in \mathbb{R}^{k\times\ell}$ be a standard normal matrix where $p := \ell - k \geq 2$. Let $\boldsymbol{R}_1 \in \mathbb{R}^{k\times k}$ be a rank-r orthogonal projector. For $s := p/2$,*

$$\left(\mathbb{E}\,\|\boldsymbol{G}^\dagger \boldsymbol{R}_1\|_{\mathrm{F}}^{2s}\right)^{1/s} \leq \frac{3.17r}{p+1}.$$

The latter result follows from a slight modification to the proof of [HMT11, Thm. A.7].

Proof Sketch. By rotational invariance of $\boldsymbol{G}^\dagger$, we may as well assume that $\boldsymbol{R}_1$ is the orthogonal projector onto the first r coordinates. Each diagonal entry $X_1, \dots, X_k$ of the matrix $(\boldsymbol{G}^\dagger)^*(\boldsymbol{G}^\dagger)$ follows the inverse chi-square distribution with $p+1$ degrees of freedom. We can write the quantity of interest as

$$\|\boldsymbol{G}^\dagger \boldsymbol{R}\|_{\mathrm{F}}^2 = \operatorname{tr}\left[\boldsymbol{R}(\boldsymbol{G}^\dagger)^*(\boldsymbol{G}^\dagger)\boldsymbol{R}\right] = \textstyle\sum_{i=1}^r X_i.$$

By the triangle inequality for the L_s norm,

$$\left(\mathbb{E}\,\|\boldsymbol{G}^\dagger \boldsymbol{R}\|_{\mathrm{F}}^{2s}\right)^{1/s} \leq \textstyle\sum_{i=1}^r \left[\mathbb{E}\, X_i^s\right]^{1/s}$$

Modifying the constants in [HMT11, Lem. A.9] to extend its validity to the range $p \geq 2$. we obtain the bound

$$\left[\mathbb{E}\, X_i^s\right]^{1/s} \leq \frac{3.17}{p+1}.$$

Combine the last two displays to complete the argument. □

Using the last two results, we quickly establish Fact 3.3.

Proof of Fact 3.3. Abbreviate $s := p/2$, and calculate that

$$\begin{aligned}
\left(\mathbb{E}_1\mathbb{E}_2\,\|\boldsymbol{M}\boldsymbol{\Omega}_2\boldsymbol{\Omega}_1^\dagger \boldsymbol{R}_1\|_{\mathrm{F}}^{2s}\right)^{1/s} &\leq s\sqrt{2}\cdot\left(\mathbb{E}_1\left[\mathbb{E}_2\,\|\boldsymbol{M}\boldsymbol{\Omega}_2\boldsymbol{\Omega}_1^\dagger \boldsymbol{R}_1\|_{\mathrm{F}}^2\right]^s\right)^{1/s} \\
&= s\sqrt{2}\cdot\left(\mathbb{E}_1\left[\|\boldsymbol{M}\|_{\mathrm{F}}^2\cdot\|\boldsymbol{\Omega}_1^\dagger \boldsymbol{R}_1\|_{\mathrm{F}}^2\right]^s\right)^{1/s} \\
&\leq s\sqrt{2}\cdot\|\boldsymbol{M}\|_{\mathrm{F}}^2\cdot\frac{3.17r}{2s+1} \leq 2.25r\,\|\boldsymbol{M}\|_{\mathrm{F}}^2.
\end{aligned}$$

In this argument, $\mathbb{E}_i$ denotes the expectation with respect to $\boldsymbol{\Omega}_i$ for $i = 1, 2$; we have invoked independence to iterate the expectation. The first inequality is a consequence of Fact A.1, which is apropos because the squared Frobenius norm is a psd quadratic form in the standard normal matrix $\boldsymbol{\Omega}_2$. The second inequality is Fact A.2. Apply Markov's inequality to complete the proof of the probability bound. □

Acknowledgments

This paper was produced under a subcontract to ONR Award 00014-16-C-2009. The author thanks Shaunak Bopardikar for helpful discussions and feedback.

References

[Bha97] R. Bhatia. *Matrix analysis*, volume 169 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1997.

[CD74] J. Cullum and W. E. Donath. A block generalization of the s-step Lanczos algorithm. Technical Report Report RC 4845 (21570), IBM Thomas J. Watson Research Center, Yorktown Heights, New York, 1974.

[DIKMI16] P. Drineas, I. Ipsen, E.-M. Kontopoulou, and M. Magdon-Ismail. Structural convergence results for approximation of dominant subspaces from block krylov spaces. Available at http://arXiv.org/abs/1609.00671, Sep. 2016.

[GU77] G. H. Golub and R. Underwood. The block Lanczos method for computing eigenvalues. In *Mathematical software, III (Proc. Sympos., Math. Res. Center, Univ. Wisconsin, Madison, Wis., 1977)*, pages 361–377. Publ. Math. Res. Center, No. 39. Academic Press, New York, 1977.

[HJ90] R. A. Horn and C. R. Johnson. *Matrix analysis*. Cambridge University Press, Cambridge, 1990. Corrected reprint of the 1985 original.

[HMST11] N. Halko, P.-G. Martinsson, Y. Shkolnisky, and M. Tygert. An algorithm for the principal component analysis of large data sets. *SIAM J. Sci. Comput.*, 33(5):2580–2594, 2011.

[HMT11] N. Halko, P. G. Martinsson, and J. A. Tropp. Finding structure with randomness: probabilistic algorithms for constructing approximate matrix decompositions. *SIAM Rev.*, 53(2):217–288, 2011.

[KW92] J. Kuczyński and H. Woźniakowski. Estimating the largest eigenvalue by the power and Lanczos algorithms with a random start. *SIAM J. Matrix Anal. Appl.*, 13(4):1094–1122, 1992.

[LT11] M. Ledoux and M. Talagrand. *Probability in Banach spaces*. Classics in Mathematics. Springer-Verlag, Berlin, 2011. Isoperimetry and processes, Reprint of the 1991 edition.

[LZ15] R.-C. Li and L.-H. Zhang. Convergence of the block Lanczos method for eigenvalue clusters. *Numer. Math.*, 131(1):83–113, 2015.

[MM15] C. Musco and C. Musco. Stronger and faster approximate singular value decomposition via the block Lanczos method. In *Proc. Adv. Neural Information Processing Systems 28 (NIPS 2015)*, 2015. Available at http://arXiv.org/abs/1504.05477.

[MRT11] P.-G. Martinsson, V. Rokhlin, and M. Tygert. A randomized algorithm for the decomposition of matrices. *Appl. Comput. Harmon. Anal.*, 30(1):47–68, 2011.

[Par98] B. N. Parlett. *The symmetric eigenvalue problem*, volume 20 of *Classics in Applied Mathematics*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 1998. Corrected reprint of the 1980 original.

[RST09] V. Rokhlin, A. Szlam, and M. Tygert. A randomized algorithm for principal component analysis. *SIAM J. Matrix Anal. Appl.*, 31(3):1100–1124, 2009.

[Saa80] Y. Saad. On the rates of convergence of the Lanczos and the block-Lanczos methods. *SIAM J. Numer. Anal.*, 17(5):687–706, 1980.

[Saa11] Y. Saad. *Numerical methods for large eigenvalue problems*, volume 66 of *Classics in Applied Mathematics*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 2011. Revised edition of the 1992 original [1177405].

[Tro18a] J. A. Tropp. Analysis of randomized block Krylov methods. Part I: Spectral norm estimation. Available at http://users.cms.caltech.edu/~jtropp/, Feb. 2018.

[Tro18b] J. A. Tropp. Analysis of randomized block Krylov methods. Part III: Matrix approximation. Available at http://users.cms.caltech.edu/~jtropp/, Feb. 2018.

[TYUC17] J. A. Tropp, A. Yurtsever, M. Udell, and V. Cevher. Randomized single-view algorithms for low-rank matrix approximation. ACM Report 2017-01, Caltech, Pasadena, CA, Jan. 2017. Available at http://users.cms.caltech.edu/~jtropp/reports/TYUC17-Randomized-Single-View-TR.pdf.

[WZZ15] S. Wang, Z. Zhang, and T. Zhang. Improved analyses of the randomized power method and block Lanczos method. Available at http://arXiv.org/abs/1508.06429, Aug. 2015.

ANALYSIS OF RANDOMIZED BLOCK KRYLOV METHODS PART III: LOW-RANK APPROXIMATION OF MATRICES

JOEL A. TROPP

ABSTRACT. Randomized block Krylov subspace methods are a powerful class of techniques for computing information about the spectrum of a matrix. The purpose of this paper is to develop new theoretical bounds on the performance of randomized block Krylov subspace methods for computing a low-rank approximation of a matrix. The results demonstrate that, for many matrices, it is possible to obtain accurate approximations using only a *constant* number of steps of the randomized block Krylov method.

1. MOTIVATION AND BACKGROUND

Randomized block Krylov methods have emerged as a powerful tool for spectral computation and matrix approximation [MRT11, RST09, HMT11, HMST11, MM15, WZZ15, DIKMI16]. At present, our understanding of these methods is less complete than our understanding of simple Krylov subspace methods. The goal of this paper and its companions is to develop detailed bounds that help explain the impressive empirical performance [HMST11, MM15] of randomized block Krylov methods.

1.1. Project Overview.

This paper, Part III of the project, addresses the problem of computing a rank-k approximation of a general matrix via a randomized block Krylov subspace method. This problem has already been investigated in several papers [MM15, WZZ15, DIKMI16], but we have been able to obtain more precise results in some settings. This paper also treats the problem of positive-semidefinite (psd) matrix approximation.

Parts I and II of the project [Tro18a, Tro18b] study how well we can estimate the largest (or smallest) eigenvalues of a symmetric matrix using a randomized block Krylov method. These papers also discuss the analogous questions for singular values. In future work, we may also address the problem of approximating invariant subspaces of a symmetric matrix by adapting the analysis of low-rank matrix approximation.

1.2. Block Krylov Subspaces.

Let us begin with a brief mathematical review of the properties of block Krylov subspaces. We describe applications to matrix approximation in Sections 2 and 3, while Section 4 contains a brief discussion about possible implementations.

Fix a symmetric matrix $\boldsymbol{A} \in \mathbb{R}^{n\times n}$. Choose a test matrix $\boldsymbol{B} \in \mathbb{R}^{n\times \ell}$, where ℓ is called the *block size*. For a depth parameter $q \geq 0$, construct the matrix

$$\boldsymbol{S}_q(\boldsymbol{A};\boldsymbol{B}) := \begin{bmatrix} \boldsymbol{B} & \boldsymbol{A}\boldsymbol{B} & \boldsymbol{A}^2\boldsymbol{B} & \dots & \boldsymbol{A}^q\boldsymbol{B} \end{bmatrix} \in \mathbb{R}^{n\times(q+1)\ell}. \tag{1.1}$$

The *block Krylov subspace* is the range of the matrix (1.1):

$$K_q := K_q(\boldsymbol{A};\boldsymbol{B}) := \operatorname{range}\big(\boldsymbol{S}_q(\boldsymbol{A};\boldsymbol{B})\big) \subset \mathbb{R}^n. \tag{1.2}$$

We can also express the block Krylov subspace using polynomials:

$$K_q(\boldsymbol{A};\boldsymbol{B}) = \bigoplus\nolimits_{\varphi\in\mathscr{P}_q} \operatorname{range}\big(\varphi(\boldsymbol{A})\boldsymbol{B}\big). \tag{1.3}$$

Date: 12 January 2018. Revised 28 February 2018.

2010 *Mathematics Subject Classification.* Primary: 65F30. Secondary: 68W20, 60B20.

Key words and phrases. Eigenvalue computation; Krylov subspace method; Lanczos method; numerical analysis; randomized algorithm; singular value computation.

where $\mathscr{P}_q$ is the linear space of polynomials with real coefficients and degree at most q. The $\oplus$ operator is the ordinary subspace sum. It is often more convenient to work with the orthogonal projector onto the Krylov subspace:

$$\boldsymbol{\Pi} := \boldsymbol{\Pi}(\boldsymbol{A}; \boldsymbol{B}; q) \in \mathbb{R}^{n \times n} \quad \text{with range}(\boldsymbol{\Pi}) = K_q(\boldsymbol{A}; \boldsymbol{B}). \tag{1.4}$$

In the notation, we may suppress the matrices $\boldsymbol{A}$ and $\boldsymbol{B}$ or the depth q if they are clear from context.

1.3. **Invariance Properties.** For any fixed depth q, the block Krylov subspace (1.2) and the orthogonal projector (1.4) have some important invariance properties:

- The block Krylov subspace only depends on the range of the test matrix:

$$\boldsymbol{\Pi}(\boldsymbol{A}; \boldsymbol{B}) = \boldsymbol{\Pi}(\boldsymbol{A}; \boldsymbol{B}\boldsymbol{T}) \quad \text{for any nonsingular } \boldsymbol{T} \in \mathbb{R}^{\ell \times \ell}. \tag{1.5}$$

- The block Krylov subspace is rotationally covariant, in the sense that

$$\boldsymbol{U}\boldsymbol{\Pi}(\boldsymbol{A}; \boldsymbol{B})\boldsymbol{U}^* = \boldsymbol{\Pi}(\boldsymbol{U}\boldsymbol{A}\boldsymbol{U}^*; \boldsymbol{U}\boldsymbol{B}) \quad \text{for any orthogonal } \boldsymbol{U} \in \mathbb{R}^{n \times n}. \tag{1.6}$$

  The * represents the (conjugate) transpose of a matrix or vector.
- The block Krylov subspace is invariant to affine transformations of the spectrum of $\boldsymbol{A}$:

$$\boldsymbol{\Pi}(\alpha\boldsymbol{A} + \beta\mathbf{I}; \boldsymbol{B}) = \boldsymbol{\Pi}(\boldsymbol{A}; \boldsymbol{B}) \quad \text{for } \alpha, \beta \in \mathbb{R}. \tag{1.7}$$

These properties are an immediate consequence of the definitions (1.1)–(1.2) and the polynomial representation (1.3). See [Tro18a, Sec. 1.2.2] for more discussion and references.

1.4. **A Random Test Matrix.** Suppose that we wish to make spectral computations on a fixed input matrix $\boldsymbol{A}$. Then it is natural to draw the test matrix $\boldsymbol{B}$ *at random* so that the Krylov subspace always captures information about the input matrix.

In this work, we will only study block Krylov subspaces built from a test matrix $\boldsymbol{\Omega} \in \mathbb{R}^{n \times \ell}$ that has the *standard normal distribution*. That is, each entry of $\boldsymbol{\Omega}$ is an independent Gaussian random variable with mean zero and variance one. A standard normal matrix has a uniformly random range. This choice is justified by the range invariance (1.5) and rotation covariance (1.6) properties. See the companion paper [Tro18a, App. A.1] for additional support.

A key advantage of choosing a standard normal matrix is that we can reduce our analysis to the situation where the input matrix $\boldsymbol{A}$ is diagonal. Indeed,

$$\boldsymbol{\Pi}(\boldsymbol{A}; \boldsymbol{\Omega}; q) \sim \boldsymbol{U}\boldsymbol{\Pi}(\boldsymbol{\Lambda}; \boldsymbol{\Omega}; q)\boldsymbol{U}^* \quad \text{where } \boldsymbol{A} = \boldsymbol{U}\boldsymbol{\Lambda}\boldsymbol{U}^* \text{ is an eigenvalue decomposition.}$$

The symbol $\sim$ denotes equality of distribution. This relation follows from the rotational invariance of range($\boldsymbol{\Omega}$) and the rotational covariance property (1.6).

2. Matrix Approximation via Randomized Block Krylov Methods

In this section, we explain how to use randomized block Krylov methods to compute a low-rank approximation of a matrix. We also present the main theoretical results on the algorithm, as well as a comparison with related work.

2.1. **Low-Rank Matrix Approximation.** Before we continue, we need to instate some notation for low-rank matrix approximation. Let $\boldsymbol{M} \in \mathbb{R}^{n \times m}$ be a general matrix. Fix a singular-value decomposition (SVD) of the matrix:

$$\boldsymbol{M} = \sum\nolimits_{i=1}^{m \wedge n} \sigma_i(\boldsymbol{M}) \boldsymbol{u}_i \boldsymbol{v}_i^*.$$

Each family $\{\boldsymbol{u}_i\} \subset \mathbb{R}^n$ and $\{\boldsymbol{v}_i\} \subset \mathbb{R}^m$ of vectors is orthonormal. The map $\sigma_i(\cdot)$ returns the ith largest singular value of its argument. We also use the standard notation $\wedge$ for the minimum.

For each natural number $k \le m \wedge n$, we can form a rank-k approximation[1] of $\boldsymbol{M}$ by truncating the SVD:

$$[\![\boldsymbol{M}]\!]_k := \sum\nolimits_{i=1}^{k} \sigma_i(\boldsymbol{M}) \boldsymbol{u}_i \boldsymbol{v}_i^*. \tag{2.1}$$

[1] We always use the adjective "rank-k" to mean "with exact rank at most k."

The Eckart–Young–Mirsky theorem [Hig89, Sec. 6] states that this construction yields a best rank-k approximation of $\boldsymbol{M}$ with respect to each unitarily invariant norm $|||\cdot|||$:

$$|||\boldsymbol{M} - [\![\boldsymbol{M}]\!]_k||| = \min_{\operatorname{rank}\boldsymbol{B}\le k} |||\boldsymbol{M} - \boldsymbol{B}|||.$$

Warning 2.1 (Non-uniqueness)**.** SVDs are not uniquely determined, so the low-rank approximation $[\![\cdot]\!]_k$ is not well-defined. We use this notation to refer to *any* low-rank approximation that has been prepared according to the foregoing recipe.

2.2. **Matrix Approximation via the Randomized Block Krylov Method.** We are now prepared to present a mathematical description of a (randomized) block Krylov method for approximating a matrix. This approach to matrix approximation was proposed in the papers [RST09, HMST11]. See Section 4.1 for some discussion about the implementation.

Consider an input matrix $\boldsymbol{C} \in \mathbb{R}^{n\times m}$. Suppose that we wish to produce a rank-k approximation of the input matrix. To do so, we set the block size ℓ to be (slightly) larger than the approximation rank k. Draw a standard normal test matrix $\boldsymbol{\Omega} \in \mathbb{R}^{m\times \ell}$. For a depth parameter $q \ge 0$, we form the block Krylov subspace and the associated orthogonal projector

$$K_q(\boldsymbol{CC}^*; \boldsymbol{C\Omega}) \subset \mathbb{R}^n \quad \text{and} \quad \boldsymbol{\Pi} := \boldsymbol{\Pi}(\boldsymbol{CC}^*; \boldsymbol{C\Omega}; q) \in \mathbb{R}^{n\times n}. \tag{2.2}$$

Then we compress the input matrix to the block Krylov subspace, and we compute a low-rank approximation $\hat{\boldsymbol{C}}_k$ of the compressed matrix:

$$\hat{\boldsymbol{C}}_k := [\![\boldsymbol{\Pi C}]\!]_k. \tag{2.3}$$

Owing to the definition (2.1) of the low-rank approximation, we can also express

$$\hat{\boldsymbol{C}}_k = \boldsymbol{\Pi}_k \boldsymbol{C} \tag{2.4}$$

for a rank-k orthogonal projector $\boldsymbol{\Pi}_k$ whose range is contained in the block Krylov subspace K_q.

The first main goal of this paper is to study how well the random matrix $\hat{\boldsymbol{C}}_k$ approximates the input matrix $\boldsymbol{C}$. In the next subsection, we discuss how to measure the approximation error and additional conclusions that we can adduce from a bound on the error.

2.3. **The Approximation Error and Its Consequences.** We will develop bounds on approximation errors with respect to the spectral norm, the Frobenius norm, and all Schatten s-norms for $s \ge 2$. All of these norms fit within a larger framework, but it is helpful to keep them in mind as concrete examples.

2.3.1. *Q-Norms.* A Q-norm is a super-quadratic unitarily invariant norm [Bha97, Sec. IV.2]. More precisely, a Q-norm $|||\cdot|||$ is a unitarily invariant norm with the property that it can be written in terms of another unitarily invariant norm $|||\cdot|||_\wedge$ as follows.

$$|||\boldsymbol{A}||| = \left[\, |||\boldsymbol{A}^*\boldsymbol{A}|||_\wedge \,\right]^{1/2}. \tag{2.5}$$

It is also convenient to introduce the symmetric gauge functions Φ and $\Phi_\wedge$ that generate these two norms:

$$|||\boldsymbol{A}||| = \Phi(\boldsymbol{\sigma}(\boldsymbol{A})) \quad \text{and} \quad |||\boldsymbol{A}|||_\wedge = \left[\Phi_\wedge\left(\boldsymbol{\sigma}^2(\boldsymbol{A})\right)\right]^{1/2}. \tag{2.6}$$

In the expression (2.6), the map $\boldsymbol{\sigma}$ returns the vector of decreasingly ordered singular values of a matrix, and $\boldsymbol{\sigma}^2$ is the componentwise square of this vector.

It may be helpful to consider how these definitions play out for the most common Q-norms. We have the following dictionary:

$\|\|\|\cdot\|\|\|$	Φ	$\Phi_\wedge$	$\|\|\|\cdot\|\|\|_\wedge$
Frobenius norm $\|\cdot\|_{\mathrm{F}}$	ℓ_2 norm	ℓ_1 norm	Schatten 1-norm
spectral norm $\|\cdot\|$	ℓ_∞ norm	ℓ_∞ norm	spectral norm
Schatten s-norm $\|\cdot\|_s$ for $s \ge 2$	ℓ_s norm	$\ell_{s/2}$ norm	Schatten $s/2$-norm

In this work, we require that every unitarily invariant norm is scaled so that $|||\mathrm{diag}(1,0,\dots,0)||| = 1$. As a consequence, the spectral norm is the smallest Q-norm, and the Frobenius norm is the largest Q-norm.

2.3.2. *The Approximation Error.* We will develop probabilistic bounds on the approximation error $|||\boldsymbol{C} - \hat{\boldsymbol{C}}_k|||$ with respect to a Q-norm. According to Mirsky's theorem [Mir60, Thm. 2], the approximation error always satisfies the bound

$$|||\boldsymbol{C} - [\![\boldsymbol{C}]\!]_k||| \le |||\boldsymbol{C} - \hat{\boldsymbol{C}}_k|||.$$

Suppose that we have obtained a reversed inequality

$$|||\boldsymbol{C} - \hat{\boldsymbol{C}}_k||| \le |||\boldsymbol{C} - [\![\boldsymbol{C}]\!]_k||| + \varepsilon =: \eta. \tag{2.7}$$

Note that both sides of the bound (2.7) depend on the choice of Q-norm, and approximation error bounds with respect to different norms are not necessarily comparable. Let us explain what other conclusions we can extract from the error bound (2.7).

2.3.3. *Linear Functionals.* Given (2.7), we immediately obtain approximation bounds for all linear functionals:

$$|\langle \boldsymbol{F},\ \boldsymbol{C}\rangle - \langle \boldsymbol{F},\ \hat{\boldsymbol{C}}_k\rangle| \le \eta\, |||\boldsymbol{F}|||_* \quad \text{for } \boldsymbol{F} \in \mathbb{R}^{n\times m}.$$

We have written $\langle \cdot,\ \cdot\rangle$ for the trace inner product, and $|||\cdot|||_*$ denotes the trace dual of the Q-norm. This is a simple consequence of duality.

2.3.4. *Singular Values.* It is always the case that the singular values of the approximation satisfy

$$0 \le \sigma_i(\boldsymbol{C}) - \sigma_i(\hat{\boldsymbol{C}}_k).$$

This point follows from the expression (2.4) for the approximation $\hat{\boldsymbol{C}}_k$ and the interlacing theorem for singular values [Bha97, Prob. III.6.5].

Given (2.7), we can reverse the latter inequality:

$$\Phi\big(\boldsymbol{\sigma}(\boldsymbol{C}) - \boldsymbol{\sigma}(\hat{\boldsymbol{C}}_k)\big) \le \eta,$$

where Φ is the symmetric gauge function that appears in (2.6). This result is a consequence of the perturbation theorem for singular values [Bha97, Prob. III.6.13]. Note, however, that we can achieve better bounds on the individual singular values using the specialized arguments in the companion paper [Tro18b].

2.3.5. *Right Invariant Subspaces.* We can also control the quality of right invariant subspace approximations. Suppose that S and $\hat{S}$ are subsets of the positive real line with

$$\min_{s\in S,\hat{s}\in\hat{S}} \big|s^2 - \hat{s}^2\big| \ge \delta.$$

Let $\boldsymbol{P}_{\mathrm{rt}}$ be the orthogonal projector onto the right invariant subspace of $\boldsymbol{C}$ associated with the singular values listed in S, and let $\hat{\boldsymbol{P}}_{\mathrm{rt}}$ be the orthogonal projector onto the right invariant subspace of $\hat{\boldsymbol{C}}_k$ associated with the singular values listed in $\hat{S}$. Let $\boldsymbol{\theta}(L, L')$ denote the vector of principal angles between two subspaces L, L'.

We have the following bounds on the principal angles:

$$\Phi_\wedge\big(\cos\boldsymbol{\theta}\big(\mathrm{range}(\boldsymbol{P}_{\mathrm{rt}}); \mathrm{range}(\hat{\boldsymbol{P}}_{\mathrm{rt}})\big)\big) = |||\boldsymbol{P}_{\mathrm{rt}}\hat{\boldsymbol{P}}_{\mathrm{rt}}|||_\wedge \le \frac{\pi}{2}\cdot\frac{\eta^2}{\delta}. \tag{2.8}$$

Indeed, the bound (2.7) implies

$$\eta^2 \ge |||(\mathbf{I} - \boldsymbol{\Pi}_k)\boldsymbol{C}|||^2 = \big|||\boldsymbol{C}^*(\mathbf{I} - \boldsymbol{\Pi}_k)\boldsymbol{C}|||_\wedge = \big|||\boldsymbol{C}^*\boldsymbol{C} - \hat{\boldsymbol{C}}_k{}^*\hat{\boldsymbol{C}}_k|||_\wedge.$$

We have used the relations (2.4) and (2.5). The right invariant subspaces of a matrix $\boldsymbol{A}^*\boldsymbol{A}$ match those of $\boldsymbol{A}$, so the claim (2.8) follows from the perturbation theorem [Bha97, Thm. VII.3.2].

Remark 2.2 (Frobenius Norm)**.** When $|||\cdot|||$ is the Frobenius norm, we can replace the constant $\pi/2$ in (2.8) with the constant 1, courtesy of [Bha97, Thm. VII.3.4].

2.3.6. *Left Invariant Subspaces.* It is also possible to control the quality of left invariant subspace approximations, but these results are less satisfactory. Suppose that S and $\hat{S}$ are subsets of the positive real line with $\operatorname{dist}(S, \hat{S}) \geq \delta$. Let $\boldsymbol{P}_{\mathrm{lt}}$ be the orthogonal projector onto the left invariant subspace of $\boldsymbol{C}$ associated with the singular values listed in S, and let $\hat{\boldsymbol{P}}_{\mathrm{lt}}$ be the orthogonal projector onto the left invariant subspace of $\hat{\boldsymbol{C}}$ associated with the singular values listed in $\hat{S}$.

Suppose that the error bound (2.7) holds *with respect to the Frobenius norm.* Then

$$\left\| \cos\boldsymbol{\theta}\big(\operatorname{range}(\boldsymbol{P}_{\mathrm{lt}}); \operatorname{range}(\hat{\boldsymbol{P}}_{\mathrm{lt}})\big) \right\|_{\ell_2} = \|\boldsymbol{P}_{\mathrm{lt}} \hat{\boldsymbol{P}}_{\mathrm{lt}}\|_{\mathrm{F}} \leq \frac{\eta}{\delta}.$$

This result follows from the perturbation theorem [Bha97, Thm. VII.5.9].

Suppose instead that the error bound (2.7) holds *with respect to the spectral norm.* Then

$$\left\| \cos\boldsymbol{\theta}\big(\operatorname{range}(\boldsymbol{P}_{\mathrm{lt}}); \operatorname{range}(\hat{\boldsymbol{P}}_{\mathrm{lt}})\big) \right\|_{\ell_\infty} = \|\boldsymbol{P}_{\mathrm{lt}} \hat{\boldsymbol{P}}_{\mathrm{lt}}\| \leq \text{const} \cdot \log(m+n) \cdot \frac{\eta}{\delta}.$$

We obtain this result by combining the proof of [Bha97, Thm. VII.5.9] with the discussion [Bha97, p. 299] and using the bound (2.7). The logarithmic factor seems to be necessary here.

Remark 2.3 (Other Q-Norms)**.** It is also possible to develop results on left invariant subspaces for the Schatten s-norm with $s \geq 2$ in the same fashion. It is not clear, however, whether there is a general theorem that holds for all Q-norms.

2.4. **Spectral Properties of the Approximation.** The spectral properties of the input matrix $\boldsymbol{C}$ play a central role in determining the quality of the approximation.

2.4.1. *Spectral Features of the Input Matrix.* In exact arithmetic, the performance of the matrix approximation scheme depends only on singular values of the input matrix. Let us introduce compact notation for the singular values of $\boldsymbol{C}$:

$$c_i := \sigma_i(\boldsymbol{C}) \quad \text{for } 1 \leq i \leq n \quad \text{and} \quad c_{\max} := c_1 \geq c_2 \geq \cdots \geq c_n =: c_{\min} \geq 0.$$

Recall that we seek a rank-k approximation, and the singular values c_j with $j > k$ compose the *tail* of the spectrum.

We will express our results in terms of several functions of the singular-value spectrum.

- For each index $1 \leq r \leq k+1$, the *tail spectral gap* measures the relative difference between the rth squared singular value and the singular values in the tail:

$$\gamma_r := \frac{c_r^2 - c_{k+1}^2}{c_r^2 - c_{\min}^2}. \tag{2.9}$$

  If the singular values of $\boldsymbol{C}$ are all identical, then we set $\gamma_r = 0$.
- Given a nonnegative number ν, the *modified tail content* measures how quickly the tail singular values decay relative to the kth singular value:

$$\tau'(k;\nu) := \sum_{j>k} \left(\frac{c_j}{c_{k+1}}\right)^2 \left(\frac{c_j^2 - c_{\min}^2}{c_k^2 - c_{\min}^2}\right)^{2\nu}. \tag{2.10}$$

  If $c_k = c_{\min}$, then we define $\tau'(k;\nu) = 0$.

We suppress the dependence of these functions on the matrix $\boldsymbol{C}$.

The appearance of the (modified) tail content function is an important new feature of our analysis, so it is worth a moment of attention. First, note the general bounds

$$\tau'(k;\nu) \leq \left(\frac{c_{k+1}}{c_k}\right)^{4\nu} (n-k) \leq n-k.$$

In the worst case, the tail content is comparable with the dimension of the matrix $\boldsymbol{C}$, but this situation is not typical when the matrix admits a good low-rank approximation. In particular, the tail content tends to be small when there is a spectral gap: $c_{k+1} \ll c_k$. Furthermore, the tail content also tends to be small

when the spectrum decays at least polynomially. For example, if $c_j / c_k \le C j^{-s}$ for $j > k$, then $\tau'(k; \nu) < 2C$ for any exponent $\nu \ge 1/(2s)$.

2.5. **Approximation of Matrices with Few Singular Values.** The block Lanczos method is very effective for approximating a matrix that has few distinct singular values.

Proposition 2.4 (Matrices with Few Singular Values)**.** *Let $\boldsymbol{C} \in \mathbb{R}^{n\times m}$ be a matrix, and fix the target rank k. Assume that $\sigma_{k+1}(\boldsymbol{C}) < \sigma_k(\boldsymbol{C})$ and that $\boldsymbol{C}$ has at most h distinct singular values smaller than $\sigma_k(\boldsymbol{C})$. Choose the depth $q \ge h$ and the block size $\ell \ge k$, and draw a standard normal matrix $\boldsymbol{\Omega} \in \mathbb{R}^{n\times \ell}$. Then, with probability one, the error in the approximation $\hat{\boldsymbol{C}}_k$ defined in (2.3) satisfies*

$$|||\boldsymbol{C} - \hat{\boldsymbol{C}}_k||| = |||\boldsymbol{C} - [\![\boldsymbol{C}]\!]_k|||.$$

Here, $|||\cdot|||$ is an arbitrary Q-norm.

See Section 5.8 for a short proof.

2.6. **Matrix Approximation without a Spectral Gap.** The first main result provides bounds for the quality of the low-rank approximation when there is a gap in the spectrum of the input matrix after the kth singular value. We focus on controlling the error with respect to the spectral norm because this is the most important case, and the proof is more transparent.

Theorem 2.5 (Randomized Block Krylov: Low-Rank Matrix Approximation)**.** *Instate the following hypotheses.*

- *Let $\boldsymbol{C} \in \mathbb{R}^{n\times m}$ be an input matrix.*
- *Fix the approximation rank k, the block size ℓ, and the oversampling $p := \ell - k$.*
- *Draw a standard normal test matrix $\boldsymbol{\Omega} \in \mathbb{R}^{n\times \ell}$.*
- *Set the depth parameter $q \ge 0$, and choose a nonnegative integer partition $q = q_1 + q_2$.*

Then the rank-k approximation $\hat{\boldsymbol{C}}_k$ computed from (2.2) and (2.3) satisfies the following bounds with respect to the spectral norm $\|\cdot\|$.

(1) For oversampling $p \ge 2$, the approximation error satisfies the probability bound

$$\mathbb{P}\left\{\|\boldsymbol{C} - \hat{\boldsymbol{C}}_k\|^2 > (1+t)\cdot\sigma_{k+1}^2(\boldsymbol{C})\right\} \le 1 \wedge \left[\frac{18k\tau'(k;q_1)}{t}\cdot \mathrm{e}^{-4q_2\sqrt{t/(2+t)}}\right]^{p/2} \quad \text{for } t > 0.$$

(2) For oversampling $p \ge 3$, the expected approximation error satisfies

$$\mathbb{E}\,\|\boldsymbol{C} - \hat{\boldsymbol{C}}_k\|^2 \le \inf_{0<\varepsilon<1}\left[\frac{1}{1-\varepsilon} + \frac{p}{p-2}\cdot 9k\tau'(k;q_1)\cdot \mathrm{e}^{-4q_2\sqrt{\varepsilon}}\right]\cdot\sigma_{k+1}^2(\boldsymbol{C}).$$

The modified tail content τ' is defined in (2.10).

The proof of Theorem 2.5 begins in Section 5 and continues in Section 6.

Let us take a moment to explain the content of Theorem 2.5. First, Mirsky's theorem [Mir60, Thm. 2] implies that

$$\|\boldsymbol{C} - \hat{\boldsymbol{C}}_k\| \ge \sigma_{k+1}(\boldsymbol{C}).$$

In other words, our bounds for the error are on the same scale as the optimal error.

The block Krylov method automatically determines the optimal partition $q = q_1 + q_2$ of the depth parameter, but the analysis distinguishes roles for q_1 and q_2. The probability bound in Theorem 2.5 is vacuous unless the second depth parameter q_2 exceeds the level

$$q_2(t) = \frac{\log\left(18t^{-1}k\tau'(k;q_1)\right)}{4\sqrt{t/(2+t)}}.$$

Once $q_2 \ge q_2(t)$, the probability of error decreases exponentially fast, where the rate is driven by the desired relative error t and the oversampling parameter p.

The threshold $q_2(t)$ scales with $t^{-1/2}\log(1/t)$, so we can achieve a moderate relative error $1+t$ even if the Krylov subspace is fairly shallow. In contrast, randomized subspace iteration [HMT11, MM15] requires the total depth q to increase with t^{-1} to achieve relative error $1+t$.

The threshold $q_2(t)$ also scales with $\log(k\tau'(k;q_1))$. As such, the depth parameter q_2 only needs to grow with $\log k$ to compute a rank-k approximation with fixed relative error. The presence of the modified tail content $\tau'(k;q_1)$ shows that the block Krylov method is more effective when the input matrix has a decaying spectrum. The paper [HMT11] established the benefits of spectral decay for randomized subspace iteration, but we are not aware of analogous results for randomized block Krylov methods.

Now, let us discuss the role of the first depth parameter q_1. When the spectrum of the input matrix decays, the tail content $\tau'(k;q_1)$ may be constant even when $q_1 = 1$ or $q_1 = 2$. In this ideal situation, the total depth q of the Krylov subspace ought to satisfy

$$\text{Polynomial tail decay:} \qquad q = q_1 + q_2(t) \approx \frac{\log(k/t)}{\sqrt{t}}$$

to obtain a rank-k approximation with relative error $1+t$ for small t. Previous analyses [MM15] required the depth q to increase with $\log n$ in all cases.

In the worst case, the tail spectrum of the input matrix does not decay at all. In this case, $\tau'(k;q_1) = n-k$ for each choice of q_1. Theorem 2.5 suggests that the depth q should grow like

$$\text{No tail decay:} \qquad q = q_2(t) \approx \frac{\log(kn/t)}{\sqrt{t}}.$$

This result is qualitatively similar to previous worst-case analyses [MM15].

We have restricted our attention to the case where the oversampling parameter $p \geq 2$. The companion paper [Tro18a, Thm. 1.5] indicates that the performance of randomized block Krylov methods is qualitatively worse when $p = 0$ or $p = 1$. If $p \geq 2$, the oversampling does not play a role in determining the depth threshold $q_2(t)$ required to achieve relative error $1+t$. On the other hand, it plays a major role in controlling the probability that the randomized block Krylov method is successful once the depth is large enough. In contrast, the main results in [MM15] only cover the case $p = 0$.

Theorem 2.5 also contains a bound on the expectation of the error. See Section 6.3.2 for some discussion of what this formula means.

Remark 2.6 (Other Norms)**.** We were also able to establish a relative-error bound, like Theorem 2.5, with respect to other unitarily invariant norms. We have chosen to omit these results because the argument is complicated, and the spectral-norm bound is more important. A result for the Frobenius norm appears in [MM15].

Remark 2.7 (Prior Work)**.** Randomized block Krylov methods for matrix approximation were proposed in the papers [RST09, HMST11]. The papers [RST09, HMT11, Gu15] contain detailed analyses of randomized subspace iteration; see also [MM15] for qualitative results. The classic paper [KW92] contains the first results for the behavior of randomized Krylov subspace methods that do not require a spectral gap. The paper [MM15] contains the first theoretical results on randomized block Krylov subspace methods. Further analysis appears in [WZZ15, DIKMI16].

Remark 2.8 (Contemporary Work)**.** After this paper was written, we learned about a related contemporary manuscript [DI18]. We hope to discuss this paper in a future version of this work.

2.7. **Matrix Approximation with a Spectral Gap.** The second main result provides bounds for the quality of the low-rank approximation when there is a gap in the spectrum of the input matrix between the kth singular value and the $(k+1)$th singular value.

Theorem 2.9 (Randomized Block Krylov: Low-Rank Matrix Approximation with Spectral Gap)**.** *Instate the following hypotheses.*

- *Let $\boldsymbol{C} \in \mathbb{R}^{n\times m}$ be an input matrix.*

- *Fix the approximation rank k, the block size ℓ, and the oversampling $p := \ell - k$.*
- *Draw a standard normal test matrix $\boldsymbol{\Omega} \in \mathbb{R}^{m \times \ell}$.*
- *Set the depth parameter $q \geq 1$, and choose a nonnegative integer partition $q = q_1 + q_2$.*

Then the rank-k approximation $\hat{\boldsymbol{C}}_k$ computed from (2.2) and (2.3) satisfies the following bounds with respect to any Q-norm $|||\cdot|||$.

(1) *For oversampling $p \geq 2$, the approximation error satisfies the probability bound*

$$\mathbb{P}\left\{|||\boldsymbol{C} - \hat{\boldsymbol{C}}_k|||^2 > |||\boldsymbol{C} - [\![\boldsymbol{C}]\!]_k|||^2 + t \cdot \sigma_{k+1}^2(\boldsymbol{C})\right\} \leq \left[\frac{9k\tau'(k;q_1)}{t} \cdot \mathrm{e}^{-4q_2\sqrt{\gamma_k}}\right]^{p/2} \quad \text{for } t > 0.$$

(2) *For oversampling $p \geq 3$, the expected approximation error satisfies*

$$\mathbb{E}\,|||\boldsymbol{C} - \hat{\boldsymbol{C}}_k|||^2 \leq |||\boldsymbol{C} - [\![\boldsymbol{C}]\!]_k|||^2 + \left[\frac{p}{p-2} \cdot 9k\tau'(k;q_1) \cdot \mathrm{e}^{-4q_2\sqrt{\gamma_k}}\right] \cdot \sigma_{k+1}^2(\boldsymbol{C}).$$

The modified tail content τ' is defined in (2.10), and the tail spectral gap γ_k is defined in (2.9).

The proof of Theorem 2.9 begins in Section 5 and continues in Section 6.

For a fixed level t of relative error, the probability bound in Theorem 2.9 only has content when the depth q_2 exceeds

$$q_2(t;\gamma_k) = \frac{\log\left(9t^{-1}k\tau'(k;q_1)\right)}{4\sqrt{\gamma_k}}.$$

Once $q_2 \geq q_2(t)$, the probability decays exponentially fast at a rate that depends on the tail spectral gap γ_k and the oversampling p.

The critical benefit of Theorem 2.9 over Theorem 2.5 is that the expected relative error in the rank-k matrix approximation decreases *exponentially* fast with q_2 if the tail spectral gap γ_k is bounded away from zero. In this situation, we can compute matrix approximations with incredible accuracy.

Remark 2.10 (Lower-Rank Approximations)**.** Suppose that there is a significant spectral gap between $\sigma_r(\boldsymbol{C})$ and $\sigma_{k+1}(\boldsymbol{C})$ for some index $1 \leq r \leq k$. The same analysis shows that a rank-r approximation $\hat{\boldsymbol{C}}_r$ admits analogous bounds with the spectral gap γ_r in the exponent instead of γ_k. See Lemma 6.1.

Remark 2.11 (Prior Work)**.** The classic paper [KW92] contains results on the performance of randomized Krylov subspace methods for eigenvalue approximation with a spectral gap. Previous results for randomized block Krylov methods for matrix approximation with a spectral gap appear in [MM15, DIKMI16].

2.8. **Extensions.** There are a number of ways to extend the analysis in this paper:

- **Oversampling.** We can obtain related results for oversampling $p = 0$ or $p = 1$, although they are qualitatively worse.
- **Subspace iteration.** Related methods can be used to develop bounds for the performance of matrix approximation via randomized subspace iteration.
- **Complex matrices.** We can approximate a complex-valued matrix by using a complex Gaussian test matrix. The proofs are similar, and the resulting bounds turn out to be stronger.

We have omitted these developments to keep the presentation concise.

2.9. **Summary of Contributions.** We conclude with an description of the contributions of this paper relative to previous work.

- The primary advance in this work is to recognize that the modified tail content (2.10) plays a role in the performance of randomized block Krylov methods. We have shown that it is possible to approximate a matrix with spectral decay using a Krylov subspace with depth that does not depend on the ambient dimension. We believe that our work is the first to recognize this phenomenon.
- We have identified how the oversampling parameter p affects the performance of the block Krylov subspace method.

- This work contains the first detailed probability and expectation bounds for randomized block Krylov subspace methods for matrix approximation.
- For an input matrix with a spectral gap, we have obtained approximation bounds that hold with respect to any Q-norm. We have also provided a detailed explanation of the consequences of an approximation error bound with respect to a Q-norm.
- Finally, we have achieved explicit and reasonable constants, which means that the bounds have some amount of predictive power.

The previous works on this subject [MM15, WZZ15, DIKMI16] all lack one or more of these elements.

Our analysis of randomized block Krylov methods for matrix approximation is similar in design to the proof in [MM15], although many details are different. In particular, we have introduced additional ideas from the paper [HMT11] on randomized matrix approximation. The companion paper [Tro18a, Sec. 2] contains additional discussion of the history and related work. We include additional references throughout this paper.

3. Approximation of Positive-Semidefinite Matrices

The second problem that we consider is how to construct a low-rank positive-semidefinite (psd) approximation of a psd matrix by combining the randomized block Krylov method with the Nyström method. We will argue that the analysis of the approach reduces to the analysis from the last section for approximation of general matrices.

3.1. Psd Matrix Approximation via the Randomized Block Krylov Method.

We begin with a mathematical description of the randomized block Krylov method for approximating a psd matrix. See Section 4 for a discussion about implementation.

Let $\boldsymbol{A} \in \mathbb{R}^{n\times n}$ be a psd input matrix. Suppose that we wish to compute a rank-k approximation of the input matrix. To that end, select a block size ℓ that is (slightly) larger than the approximation rank k. Form the block Krylov subspace $K_q(\boldsymbol{A};\boldsymbol{\Omega})$, and construct a matrix $\boldsymbol{Q}$ with orthonormal columns that span $K_q(\boldsymbol{A};\boldsymbol{\Omega})$.

We use the Nyström method [WS00, DM05, HMT11, Git13, GM16] to produce an approximation that preserves the psd property. The Nyström approximation takes the form

$$\hat{\boldsymbol{A}}^{\mathrm{nys}} := (\boldsymbol{AQ})(\boldsymbol{Q}^*\boldsymbol{AQ})^{\dagger}(\boldsymbol{AQ})^*. \tag{3.1}$$

The dagger † denotes the Moore–Penrose pseudoinverse. The rank of this psd approximation is typically larger than the target rank k, so we pass to a rank-k approximation:

$$\hat{\boldsymbol{A}}^{\mathrm{nys}}_k := [\![\hat{\boldsymbol{A}}^{\mathrm{nys}}]\!]_k. \tag{3.2}$$

This specific approach to fixed-rank Nyström approximation was proposed independently in the recent papers [PAB16, WGM17, TYUC17].

3.2. The Approximation Error.

We will demonstrate that the analysis of the psd approximation $\hat{\boldsymbol{A}}^{\mathrm{nys}}_k$ reduces to the analysis of the general matrix approximation we studied in the last section.

Proposition 3.1 (Error in Psd Matrix Approximation)**.** *Let $\boldsymbol{A} \in \mathbb{R}^{n\times n}$ be a psd matrix, and define the psd matrix $\boldsymbol{C} := \boldsymbol{A}^{1/2} \in \mathbb{R}^{n\times n}$. Then*

$$|||\boldsymbol{A} - \hat{\boldsymbol{A}}^{\mathrm{nys}}_k||| = |||\boldsymbol{C} - \hat{\boldsymbol{C}}_k|||^2_{\wedge}.$$

The approximation $\hat{\boldsymbol{A}}^{\mathrm{nys}}_k$ is defined in (3.2), and the approximation $\hat{\boldsymbol{C}}_k$ is defined in (2.3). Here, $|||\cdot|||$ is an arbitrary Q-norm, and $|||\cdot|||_{\wedge}$ is the unitarily invariant norm that appears in (2.5).

The singular values of $\boldsymbol{C} = \boldsymbol{A}^{1/2}$ are the square roots of the eigenvalues of the psd matrix $\boldsymbol{A}$. Therefore, we immediately obtain probabilistic bounds for the error in the approximation $\hat{\boldsymbol{A}}^{\mathrm{nys}}_k$ from Theorems 2.5 and 2.9. For brevity, we omit detailed statements of these results.

The proof of this result is based on ideas from [Git13, Lem. 6.1] and [TYUC17, Thm. 4.1].

Algorithm 1 A block Krylov method for matrix approximation

Input: Input matrix $\boldsymbol{C} \in \mathbb{R}^{n\times m}$; approximation rank k; block size ℓ and depth q of Krylov space
Output: Rank-k approximation $\hat{\boldsymbol{C}}_k \in \mathbb{R}^{m\times n}$ of the input matrix

function BLOCKKRYLOVAPPROX($\boldsymbol{C}$, k, ℓ, q)
  $\boldsymbol{\Omega} \leftarrow \texttt{randn}(\texttt{size}(\boldsymbol{C},2),\ell)$ ▷ Draw $m \times \ell$ standard normal test matrix
  $\boldsymbol{Y}_0 \leftarrow \boldsymbol{C}\boldsymbol{\Omega}$
  for $t \leftarrow 1,2,3,\dots,q$ **do**
    $\boldsymbol{Y}_t \leftarrow \boldsymbol{C}(\boldsymbol{C}^*\boldsymbol{Y}_{t-1})$ ▷ Form blocks of Krylov matrix by repeated multiplication
  $\boldsymbol{Q} \leftarrow \texttt{orth}([\boldsymbol{Y}_0, \boldsymbol{Y}_1, \dots, \boldsymbol{Y}_q])$ ▷ Find orthonormal basis for block Krylov space
  $(\boldsymbol{U}, \boldsymbol{\Sigma}, \boldsymbol{V}) \leftarrow \texttt{svds}(\boldsymbol{Q}^*\boldsymbol{C}, k)$ ▷ Truncated SVD via *dense* linear algebra
  $\hat{\boldsymbol{C}}_k = (\boldsymbol{Q}\boldsymbol{U})\boldsymbol{\Sigma}\boldsymbol{V}^*$ ▷ Rank-k approximation in factored form

Proof. Abbreviate $\boldsymbol{Y} := \boldsymbol{A}^{1/2}\boldsymbol{Q}$, where $\operatorname{range}(\boldsymbol{Q}) = K_q(\boldsymbol{A};\boldsymbol{\Omega})$. Introduce the orthogonal projector $\boldsymbol{P}$ onto $\operatorname{range}(\boldsymbol{Y})$. The Nyström approximation (3.1) satisfies

$$\hat{\boldsymbol{A}}^{\mathrm{nys}} = \boldsymbol{A}^{1/2} \cdot \boldsymbol{Y}(\boldsymbol{Y}^*\boldsymbol{Y})^{\dagger}\boldsymbol{Y}^* \cdot \boldsymbol{A}^{1/2} = \boldsymbol{A}^{1/2}\boldsymbol{P}\boldsymbol{A}^{1/2}.$$

We can choose a rank-k orthogonal projector $\boldsymbol{\Pi}_k$ to realize the identity $[\![\boldsymbol{P}\boldsymbol{A}^{1/2}]\!]_k = \boldsymbol{\Pi}_k\boldsymbol{A}^{1/2}$. The construction (2.1) of a rank-k approximation implies that the rank-k approximation (3.2) satisfies

$$\hat{\boldsymbol{A}}_k^{\mathrm{nys}} = [\![(\boldsymbol{A}^{1/2}\boldsymbol{P})(\boldsymbol{P}\boldsymbol{A}^{1/2})]\!]_k = \big([\![\boldsymbol{P}\boldsymbol{A}^{1/2}]\!]_k\big)^*\big([\![\boldsymbol{P}\boldsymbol{A}^{1/2}]\!]_k\big) = \boldsymbol{A}^{1/2}\boldsymbol{\Pi}_k\boldsymbol{A}^{1/2}.$$

Therefore, we can represent the approximation error as

$$|||\boldsymbol{A} - \hat{\boldsymbol{A}}_k||| = |||\boldsymbol{A}^{1/2}(\mathbf{I} - \boldsymbol{\Pi}_k)\boldsymbol{A}^{1/2}||| = |||(\mathbf{I} - \boldsymbol{\Pi}_k)\boldsymbol{A}^{1/2}|||_{\wedge}^2 = |||\boldsymbol{A}^{1/2} - [\![\boldsymbol{P}\boldsymbol{A}^{1/2}]\!]_k|||_{\wedge}^2.$$

We have recalled the statement (2.5) of the properties of a Q-norm. Using the definition (1.1)–(1.2) of the Krylov subspace,

$$\operatorname{range}(\boldsymbol{P}) = \operatorname{range}(\boldsymbol{Y}) = \boldsymbol{A}^{1/2}K_q(\boldsymbol{A};\boldsymbol{\Omega}) = K_q(\boldsymbol{A};\boldsymbol{A}^{1/2}\boldsymbol{\Omega}).$$

Finally, we observe that the rank-k approximation (2.3) of the input matrix $\boldsymbol{C} := \boldsymbol{A}^{1/2}$ takes the form $\hat{\boldsymbol{C}}_k = [\![\boldsymbol{P}\boldsymbol{A}^{1/2}]\!]_k$. □

4. IMPLEMENTATIONS

This paper focuses on mathematical analysis of block Krylov methods, but it is worth a moment to comment on possible implementations.

4.1. **Matrix Approximation.** Algorithm 1 provides pseudocode for a simple variant of the block Krylov method for computing a rank-k approximation of a rectangular matrix $\boldsymbol{C}$. This approach is extracted from [HMST11]. Let us emphasize that this algorithm is not suitable when q is moderate or large or when $\boldsymbol{C}$ is poorly conditioned.

Here is a short breakdown of the computational resources required.

- We multiply $\boldsymbol{C} \in \mathbb{R}^{m\times n}$ by an $n \times \ell$ matrix $(q+1)$ times and $\boldsymbol{C}^*$ by an $m \times \ell$ matrix a total of q times. In case $\boldsymbol{C}$ is dense, this requires $\mathcal{O}(q\ell(m+n)^2)$ arithmetic operations, but it may be far more efficient when $\boldsymbol{C}$ is sparse or structured.
- Orthogonalization of the block Krylov matrix requires $\mathcal{O}(q^2\ell^2(m+n))$ arithmetic operations.
- The compression to $\operatorname{range}(\boldsymbol{Q})$ requires $\mathcal{O}(q\ell mn)$ arithmetic operations for dense $\boldsymbol{C}$, but it may be faster when $\boldsymbol{C}$ is sparse or structured.
- Computation of the truncated SVD uses $\mathcal{O}(q^2\ell^2 m)$ arithmetic operations.
- The block Krylov matrix requires $\mathcal{O}(q\ell n)$ units of storage.

It is possible to achieve some further improvements using block Lanczos algorithms [CD74, GU77].

Algorithm 2 A block Krylov method for psd matrix approximation

Input: Psd input matrix $\boldsymbol{A} \in \mathbb{R}^{n\times n}$; approximation rank k; block size ℓ and depth q of Krylov space
Output: Rank-k psd approximation $\hat{\boldsymbol{A}}_k \in \mathbb{R}^{n\times n}$ of the input matrix

```
function BlockKrylovNystrom(A, k, ℓ, q)
    Y_0 ← randn(size(A, 1), ℓ)          ▷ Draw n × ℓ standard normal test matrix
    for t ← 1, 2, 3, ..., q do
        Y_t ← A Y_{t-1}                 ▷ Form blocks of Krylov matrix by repeated multiplication
    Q ← orth([Y_0, Y_1, ..., Y_q])      ▷ Find orthonormal basis for block Krylov space

    F ← A Q
    S ← sqrtm(Q* F)                     ▷ Psd square root
    (U, Σ, ~) ← svds(F/S, k)            ▷ Truncated SVD via dense linear algebra
    Â_k ← (QU) Σ^2 (QU)*                ▷ Rank-k approximation in factored form
```

4.2. **Psd Matrix Approximation.** Algorithm 2 provides pseudocode for a simple variant of the block Krylov method for computing a low-rank approximation of a psd matrix. The Nyström part of the code is based on the best practices outlined in [LLS+17]; see [TYUC17] for related work. Let us emphasize that this algorithm is not suitable when q is moderate or large or when $\boldsymbol{A}$ is poorly conditioned.

Here is a short breakdown of the computational resources required.

- We multiply $\boldsymbol{A} \in R^{n\times n}$ by an $n \times \ell$ matrix q times. In case $\boldsymbol{A}$ is dense, this requires $\mathcal{O}(q\ell n^2)$ arithmetic operations, but it is faster for sparse or structured $\boldsymbol{A}$.
- Orthogonalization of the block Krylov matrix requires $\mathcal{O}(q^2\ell^2 n)$ arithmetic operations.
- The symmetric compression of $\boldsymbol{A}$ requires $\mathcal{O}(q\ell n)$ arithmetic operations when $\boldsymbol{A}$ is dense.
- The matrix square root requires $\mathcal{O}(q^3\ell^3)$ arithmetic operations.
- The least-squares problem and truncated SVD use $\mathcal{O}(q^2\ell^2 n)$ arithmetic operations.
- Storage of the block Krylov matrix, at a cost of $\mathcal{O}(q\ell m)$ units of storage.

It is possible to make some improvements with block Lanczos algorithms [CD74, GU77].

5. Technical Preparations

In this section, we initiate the proof of Theorems 2.5 and 2.9. Along the route, we establish Proposition 2.4. At a high level, the argument is similar in spirit to the proofs in [MM15], but we have isolated the key ingredients and sharpened them to obtain more detailed bounds.

5.1. **Reduction to Diagonal Form.** First, we observe that the distribution of the approximation error $|||\boldsymbol{C} - \hat{\boldsymbol{C}}_k|||$ only depends on the singular values of the input matrix $\boldsymbol{C}$. To that end, introduce the SVD

$$\boldsymbol{C} = \boldsymbol{U}\boldsymbol{\Sigma}\boldsymbol{V}^* \quad \text{where} \quad \begin{matrix} \boldsymbol{\Sigma} \in \mathbb{R}^{n\times m} \text{ is nonnegative diagonal, and} \\ \boldsymbol{U} \in \mathbb{R}^{n\times n} \text{ and } \boldsymbol{V} \in \mathbb{R}^{m\times m} \text{ are orthogonal.} \end{matrix}$$

Let $\boldsymbol{\Omega} \in \mathbb{R}^{m\times\ell}$ be a standard normal test matrix. Using the the rotational covariance (1.6) of the projector and the rotational invariance of $\boldsymbol{\Omega}$, we find that the orthogonal projector $\boldsymbol{\Pi} \in \mathbb{R}^{n\times n}$ onto the Krylov subspace $K_q(\boldsymbol{C}\boldsymbol{C}^*; \boldsymbol{C}\boldsymbol{\Omega})$ satisfies

$$\boldsymbol{\Pi}(\boldsymbol{C}\boldsymbol{C}^*; \boldsymbol{C}\boldsymbol{\Omega})\boldsymbol{C} = \boldsymbol{\Pi}(\boldsymbol{U}^*\boldsymbol{\Sigma}\boldsymbol{\Sigma}^*\boldsymbol{U}; \boldsymbol{U}\boldsymbol{\Sigma}\boldsymbol{V}^*\boldsymbol{\Omega})\boldsymbol{U}\boldsymbol{\Sigma}\boldsymbol{V}^* \sim \boldsymbol{U}\boldsymbol{\Pi}(\boldsymbol{\Sigma}^*\boldsymbol{\Sigma}; \boldsymbol{\Sigma}\boldsymbol{\Omega})\boldsymbol{\Sigma}\boldsymbol{V}^*.$$

Therefore, using the definition (2.3) of the approximation $\hat{\boldsymbol{C}}_k$ and the construction (2.1) of the low-rank approximation,

$$\hat{\boldsymbol{C}}_k := [\![\boldsymbol{\Pi}(\boldsymbol{C}\boldsymbol{C}^*; \boldsymbol{C}\boldsymbol{\Omega})\boldsymbol{C}]\!]_k \sim [\![\boldsymbol{U}\boldsymbol{\Pi}(\boldsymbol{\Sigma}\boldsymbol{\Sigma}^*; \boldsymbol{\Sigma}\boldsymbol{\Omega})\boldsymbol{\Sigma}\boldsymbol{V}^*]\!]_k = \boldsymbol{U}[\![\boldsymbol{\Pi}(\boldsymbol{\Sigma}\boldsymbol{\Sigma}^*; \boldsymbol{\Sigma}\boldsymbol{\Omega})\boldsymbol{\Sigma}]\!]_k\boldsymbol{V}^*.$$

By unitary invariance of the Q-norm $|||\cdot|||$, the approximation error satisfies

$$|||\boldsymbol{C} - \hat{\boldsymbol{C}}_k||| \sim \left|||\boldsymbol{\Sigma} - [\![\boldsymbol{\Pi}(\boldsymbol{\Sigma}\boldsymbol{\Sigma}^*; \boldsymbol{\Sigma}\boldsymbol{\Omega})\boldsymbol{\Sigma}]\!]_k\right|||.$$

It remains to check that we can pass from the rectangular matrix of singular values to a square matrix.

Since $n \le m$, the diagonal matrix $\boldsymbol{\Sigma} \in \mathbb{R}^{n\times m}$ takes the form

$$\boldsymbol{\Sigma} = \begin{bmatrix}\boldsymbol{\Sigma}_0 & \mathbf{0}\end{bmatrix} \quad \text{where } \boldsymbol{\Sigma}_0 \in \mathbb{R}^{n\times n}.$$

Let $\boldsymbol{\Omega}' \in \mathbb{R}^{n\times \ell}$ be a standard normal matrix. It is immediate that $\boldsymbol{\Sigma}\boldsymbol{\Omega} \sim \boldsymbol{\Sigma}_0\boldsymbol{\Omega}'$. It follows that

$$|||\boldsymbol{C} - \hat{\boldsymbol{C}}_k||| \sim \left|||\boldsymbol{\Sigma} - [\![\boldsymbol{\Pi}(\boldsymbol{\Sigma}\boldsymbol{\Sigma}^*; \boldsymbol{\Sigma}\boldsymbol{\Omega})\boldsymbol{\Sigma}]\!]_k\right||| = \left|||\boldsymbol{\Sigma}_0 - [\![\boldsymbol{\Pi}(\boldsymbol{\Sigma}_0^2; \boldsymbol{\Sigma}_0\boldsymbol{\Omega}')\boldsymbol{\Sigma}_0]\!]_k\right|||.$$

We see that it is sufficient to reduce our attention from a general matrix $\boldsymbol{C}$ to the square matrix $\boldsymbol{\Sigma}_0 \in \mathbb{R}^{n\times n}$ of singular values. Furthermore, we can ensure that the entries of $\boldsymbol{C}$ are nonnegative and weakly decreasing. Similar arguments apply when $n \ge m$.

5.2. **The Input Matrix.** The discussion in Section 5.1 allows us to assume that the input matrix $\boldsymbol{C}$ takes the form

$$\boldsymbol{C} = \operatorname{diag}(c_1, \dots, c_n) \in \mathbb{R}^{n\times n} \quad \text{where} \quad c_1 \ge c_2 \ge \cdots \ge c_n \ge 0.$$

It is also productive to introduce an explicit family of low-rank approximations of the input matrix:

$$\boldsymbol{C}_r := \operatorname{diag}(c_1, c_2, \dots, c_r, 0, \dots, 0) \quad \text{for each index } 1 \le r \le n.$$

Note that $\boldsymbol{C}_r$ is a valid choice for the rank-r approximation $[\![\boldsymbol{C}]\!]_r$.

5.3. **The Approximation.** The test matrix $\boldsymbol{\Omega} \in \mathbb{R}^{n\times \ell}$ is drawn from the standard normal distribution. We will work with the orthogonal projector

$$\boldsymbol{\Pi} := \boldsymbol{\Pi}(\boldsymbol{C}^2; \boldsymbol{C}\boldsymbol{\Omega}; q) \in \mathbb{R}^{n\times n} \quad \text{with} \quad \operatorname{range}(\boldsymbol{\Pi}) = K_q(\boldsymbol{C}^2; \boldsymbol{C}\boldsymbol{\Omega}).$$

Define the one-sided compression of $\boldsymbol{C}$ to the block Krylov subspace:

$$\boldsymbol{F} := \boldsymbol{\Pi}\boldsymbol{C}. \tag{5.1}$$

Recall that the desired rank-k approximation (2.3) of the input matrix takes the form $\hat{\boldsymbol{C}}_k := [\![\boldsymbol{F}]\!]_k$.

5.4. **Normalization.** The Krylov subspace $K_q(\boldsymbol{C}^2; \boldsymbol{C}\boldsymbol{\Omega})$ is invariant to the scaling of $\boldsymbol{C}$ because of (1.5) and (1.7). Therefore, the approximation $\hat{\boldsymbol{C}}_k$ defined in (2.3) is a homogeneous function of $\boldsymbol{C}$. The approximation error $|||\boldsymbol{C} - \hat{\boldsymbol{C}}_k|||$ is also homogeneous in the input matrix. Furthermore, the spectral features γ_r and $\tau'(k; v)$ that appear in the main results are scale invariant. As a consequence, we may rescale the input matrix so that

$$c_{k+1} = 1.$$

We maintain this assumption to simplify the remainder of the argument.

5.5. **Graded Spectral Projectors.** To perform the analysis, we introduce a special family of orthogonal projectors that realize rank-r approximations of the one-sided compression (5.1). Fix an SVD

$$\boldsymbol{F} = \textstyle\sum_{i=1}^n \xi_i \boldsymbol{u}_i \boldsymbol{v}_i^*$$

Now, define the orthogonal projectors

$$\boldsymbol{\Pi}_r := \textstyle\sum_{i=1}^r \boldsymbol{u}_i \boldsymbol{u}_i^* \quad \text{for each index } 0 \le r \le k.$$

This sequence of projectors has several immediate properties:

(i) The sequence is graded: $\operatorname{rank}(\boldsymbol{\Pi}_r) = r$ for each index $0 \le r \le k$.

(ii) The ranges are nested subspaces of the block Krylov space:

$$\{\mathbf{0}\} = \operatorname{range}(\boldsymbol{\Pi}_0) \subset \operatorname{range}(\boldsymbol{\Pi}_1) \subset \cdots \subset \operatorname{range}(\boldsymbol{\Pi}_k) \subset \operatorname{range}(\boldsymbol{\Pi}) = K_q(\boldsymbol{C}^2; \boldsymbol{C}\boldsymbol{\Omega}).$$

(iii) The projectors realize truncated SVDs of the compression:

$$[\![\boldsymbol{F}]\!]_r = \boldsymbol{\Pi}_r \boldsymbol{C} \quad \text{for each index } 0 \le r \le k.$$

We can—and we will—require that $\hat{\boldsymbol{C}}_k = \boldsymbol{\Pi}_k \boldsymbol{C}$. This represents no loss of generality because every rank-k approximation $[\![\boldsymbol{F}]\!]_k$ derives from some truncated SVD.

5.6. **A Variational Principle.** The graded spectral projectors admit a variational principle that is a direct consequence of the Eckart–Young Theorem. This result highlights the distinguished role that the spectral projectors play in the argument.

Lemma 5.1 (Maximum Principle)**.** *For each index $0 \le r \le k$, the orthogonal projector $\boldsymbol{\Pi}_r$ satisfies*

$$\|\boldsymbol{\Pi}_r \boldsymbol{C}\|_{\mathrm{F}}^2 = \max\big\{ \|\boldsymbol{P}\boldsymbol{C}\|_{\mathrm{F}}^2 : \operatorname{rank}(\boldsymbol{P}) \le r \text{ and } \operatorname{range}(\boldsymbol{P}) \subset \operatorname{range}(\boldsymbol{\Pi}) \big\}.$$

The maximum extends over all orthogonal projectors $\boldsymbol{P} \in \mathbb{R}^{n\times n}$ that meet the constraints.

Proof. On the one hand, it is evident that

$$\|\boldsymbol{\Pi}_r \boldsymbol{C}\|_{\mathrm{F}}^2 \le \max\big\{ \|\boldsymbol{P}\boldsymbol{C}\|_{\mathrm{F}}^2 : \operatorname{rank}(\boldsymbol{P}) \le r \text{ and } \operatorname{range}(\boldsymbol{P}) \subset \operatorname{range}(\boldsymbol{\Pi}) \big\}.$$

Indeed, $\boldsymbol{\Pi}_r$ is a rank-r orthogonal projector whose range is contained in $\operatorname{range}(\boldsymbol{\Pi})$; see properties (i) and (ii).

On the other hand, for any orthogonal projector $\boldsymbol{P}$ that satisfies the constraints, we have $\boldsymbol{P}\boldsymbol{C} = \boldsymbol{P}\boldsymbol{\Pi}\boldsymbol{C} = \boldsymbol{P}\boldsymbol{F}$. The Eckart–Young Theorem [Hig89, Sec. 6] implies that

$$\|\boldsymbol{P}\boldsymbol{C}\|_{\mathrm{F}}^2 = \|\boldsymbol{P}\boldsymbol{F}\|_{\mathrm{F}}^2 \le \|[\![\boldsymbol{F}]\!]_r\|_{\mathrm{F}}^2 = \|\boldsymbol{\Pi}_r \boldsymbol{C}\|_{\mathrm{F}}^2.$$

The last relation follows from property (iii) of the orthogonal projector $\boldsymbol{\Pi}_r$. □

5.7. **Reversed Mirsky Inequalities.** In this section, we develop a converse inequality to the Mirsky inequality on low-rank matrix approximation that allows us to extend our results to a wide family of unitarily invariant norms.

Proposition 5.2 (Reversed Mirsky)**.** *Let $\boldsymbol{M}$ and $\boldsymbol{E}$ be conformal matrices, and assume $\operatorname{rank}(\boldsymbol{E}) \le r$. Suppose that*

$$\|\boldsymbol{M} - \boldsymbol{E}\|_{\mathrm{F}}^2 \le \|\boldsymbol{M} - [\![\boldsymbol{M}]\!]_r\|_{\mathrm{F}}^2 + \eta.$$

Then, for every Q-norm $|||\cdot|||$,

$$|||\boldsymbol{M} - \boldsymbol{E}|||^2 \le |||\boldsymbol{M} - [\![\boldsymbol{M}]\!]_r|||^2 + \eta.$$

Results of this type were introduced by Gu [Gu15, Thm. 3.4]; see also [TYUC17, Prop. A.8]. The earlier work only yields a conclusion for the spectral norm.

Proof. We will establish the weak majorization

$$\boldsymbol{\sigma}^2(\boldsymbol{M} - \boldsymbol{E}) \prec_w \boldsymbol{\sigma}^2(\boldsymbol{M} - [\![\boldsymbol{M}]\!]_r) + \eta \mathbf{e}_1. \tag{5.2}$$

The vector $\mathbf{e}_1$ has a one in the first coordinate and zeros elsewhere; $\boldsymbol{\sigma}$ is defined in Section 2.3.1. This weak majorization implies a result that is slightly stronger than the statement of the proposition. For example, see [Bha97, Exer. IV.2.10].

Weyl's inequalities [Bha97, Thm. III.2.1] imply that

$$\sigma_{i+r}^2(\boldsymbol{M}) \le \big[\sigma_i(\boldsymbol{M} - \boldsymbol{E}) + \sigma_{r+1}(\boldsymbol{E})\big]^2 = \sigma_i^2(\boldsymbol{M} - \boldsymbol{E}) \quad \text{for each index } i \ge 1.$$

We have used the fact that $\boldsymbol{E}$ has rank r. For each index $j \ge 1$, we may calculate that

$$\begin{aligned}
\|\boldsymbol{M} - \boldsymbol{B}\|_{\mathrm{F}}^2 &= \textstyle\sum_{i=1}^{j} \sigma_i^2(\boldsymbol{M} - \boldsymbol{E}) + \sum_{i>j} \sigma_i^2(\boldsymbol{M} - \boldsymbol{E}) \\
&\ge \textstyle\sum_{i=1}^{j} \sigma_i^2(\boldsymbol{M} - \boldsymbol{E}) + \sum_{i>j} \sigma_{i+r}^2(\boldsymbol{M}) \\
&= \textstyle\sum_{i=1}^{j} \sigma_i^2(\boldsymbol{M} - \boldsymbol{E}) + \sum_{i>j} \sigma_i^2(\boldsymbol{M} - [\![\boldsymbol{M}]\!]_r) \\
&= \textstyle\sum_{i=1}^{j} \sigma_i^2(\boldsymbol{M} - \boldsymbol{E}) - \sum_{i=1}^{j} \sigma_i^2(\boldsymbol{M} - [\![\boldsymbol{M}]\!]_r) + \|\boldsymbol{M} - [\![\boldsymbol{M}]\!]_r\|_{\mathrm{F}}^2.
\end{aligned}$$

Rearrange this inequality and invoke the hypothesis to obtain

$$\begin{aligned}\sum_{i=1}^{j} \sigma_i^2(\boldsymbol{M}-\boldsymbol{E}) &\le \sum_{i=1}^{j} \sigma_i^2(\boldsymbol{M}-[\![\boldsymbol{M}]\!]_r) + \left[\|\boldsymbol{M}-\boldsymbol{E}\|_{\mathrm{F}}^2 - \|\boldsymbol{M}-[\![\boldsymbol{M}]\!]_r\|_{\mathrm{F}}^2\right] \\ &\le \sum_{i=1}^{j} \sigma_i^2(\boldsymbol{M}-[\![\boldsymbol{M}]\!]_r) + \eta.\end{aligned}$$

This family of inequalities is equivalent with the weak majorization (5.2). □

Remark 5.3 (Reversed Mirsky II)**.** The same proof strategy yields a related result that is worth documenting. Suppose that

$$\|\boldsymbol{M}-\boldsymbol{E}\|_{S_1} \le \|\boldsymbol{M}-[\![\boldsymbol{M}]\!]_r\|_{S_1} + \eta \quad \text{where } \operatorname{rank}(\boldsymbol{E}) \le r.$$

Then, for every unitarily invariant norm $|||\cdot|||$,

$$|||\boldsymbol{M}-\boldsymbol{E}||| \le |||\boldsymbol{M}-[\![\boldsymbol{M}]\!]_r||| + \eta.$$

We have written $\|\cdot\|_{S_1}$ for the Schatten 1-norm.

5.8. **Matrices with Few Eigenvalues.** We are now prepared to present a short proof of Proposition 2.4.

Proposition 2.4: Proof Sketch. Without loss of generality, assume that $\ell = k$. The hypotheses require that $c_{k+1} < c_k$, and we write $\mu_1, \dots, \mu_h$ for the distinct singular values of $\boldsymbol{C}$ that are strictly smaller than c_k. Construct a polynomial that annihilates the tail singular values:

$$\varphi_0(s) := \prod_{i=1}^{h} \frac{s-\mu_i^2}{c_k^2-\mu_i^2} \in \mathscr{P}_h \subset \mathscr{P}_q.$$

Note that $\varphi_0(s) \ge 1$ for each $s \ge c_k$.

Let $\boldsymbol{P}_{\boldsymbol{Y}}$ be the rank-k orthogonal projector onto the range of $\boldsymbol{Y} := \varphi_0(\boldsymbol{C}^2)\boldsymbol{C}\boldsymbol{\Omega}$, which is contained within the block Krylov subspace $K_q(\boldsymbol{C}^2; \boldsymbol{C}\boldsymbol{\Omega})$ because of (1.3). With probability one over the randomness in the standard normal matrix $\boldsymbol{\Omega}$, the projector satisfies $\boldsymbol{P}_{\boldsymbol{Y}} = \mathbf{I}_k \oplus \mathbf{0}_{n-k}$. Therefore, $\boldsymbol{P}_{\boldsymbol{Y}}\boldsymbol{C} = \boldsymbol{C}_k$.

As a consequence, we may calculate that

$$\|\boldsymbol{C}-\hat{\boldsymbol{C}}_k\|_{\mathrm{F}}^2 = \|\boldsymbol{C}-\boldsymbol{\Pi}_k\boldsymbol{C}\|_{\mathrm{F}}^2 = \|\boldsymbol{C}\|_{\mathrm{F}}^2 - \|\boldsymbol{\Pi}_k\boldsymbol{C}\|_{\mathrm{F}}^2 \le \|\boldsymbol{C}\|_{\mathrm{F}}^2 - \|\boldsymbol{P}_{\boldsymbol{Y}}\boldsymbol{C}\|_{\mathrm{F}}^2 = \|\boldsymbol{C}-\boldsymbol{P}_{\boldsymbol{Y}}\boldsymbol{C}\|_{\mathrm{F}}^2 = \|\boldsymbol{C}-\boldsymbol{C}_k\|_{\mathrm{F}}^2.$$

We have used property (iii), the Pythagorean theorem, and the maximum principle, Lemma 5.1. The result follows from the reversed Mirsky inequality, Proposition 5.2. □

5.9. **Randomized Matrix Approximation.** The most important ingredient in the argument is a result on randomized matrix approximation adapted from the paper [HMT11].

Fact 5.4 (Randomized Matrix Approximation)**.** *Let $\boldsymbol{M} \in \mathbb{R}^{n\times n}$ be a diagonal matrix. Draw a standard normal test matrix $\boldsymbol{\Omega} \in \mathbb{R}^{n\times\ell}$, and construct the orthogonal projector $\boldsymbol{P}_{\boldsymbol{Y}} \in \mathbb{R}^{n\times n}$ onto the range of the matrix $\boldsymbol{Y} \in \boldsymbol{M}\boldsymbol{\Omega}$. Decompose*

$$\boldsymbol{M} = \begin{bmatrix} \boldsymbol{M}_1 & \mathbf{0} \\ \mathbf{0} & \boldsymbol{M}_2 \end{bmatrix} \quad \text{and} \quad \boldsymbol{\Omega} = \begin{bmatrix} \boldsymbol{\Omega}_1 \\ \boldsymbol{\Omega}_2 \end{bmatrix} \quad \text{where} \quad \boldsymbol{M}_1 \in \mathbb{R}^{k\times k} \text{ and } \boldsymbol{\Omega}_1 \in \mathbb{R}^{k\times\ell}.$$

Let $\boldsymbol{R} \in \mathbb{R}^{n\times n}$ be the orthogonal projector onto the first r coordinates, where $r \le k$. Writing $p := \ell - k$, we have

$$\mathbb{P}\left\{\|(\mathbf{I}-\boldsymbol{P}_{\boldsymbol{Y}})\boldsymbol{M}\boldsymbol{R}\|_{\mathrm{F}} > t\right\} \le \left[\frac{2.25k\,\|\boldsymbol{M}_2\|_{\mathrm{F}}^2}{t}\right]^{p/2}.$$

Qualitatively different bounds hold when $p = 0$ or $p = 1$.

This statement is extracted from the companion paper [Tro18b, Sec. 3.5], which explains how to derive the result from prior work.

6. Proof of Main Results

This section contains the main technical results, and we explain how they imply Theorems 2.5 and 2.9. The proofs of the technical results unfold in the remaining sections of the paper.

6.1. **Matrix Approximation Lemma.** The main part of the analysis is encapsulated in a lemma on the quality of a specific low-rank matrix approximation. This is not necessarily the same approximation that appears in the main results!

Lemma 6.1 (Approximation)**.** *Instate the prevailing notation. Fix an index $r \le k$. Let $q = q_1 + q_2$ be an integer partition with $0 \le q_1 \le q$. For each $t > 0$,*

$$\mathbb{P}\left\{|||\boldsymbol{C} - \boldsymbol{\Pi}_r \boldsymbol{C}|||^2 > |||\boldsymbol{C} - \boldsymbol{C}_r|||^2 + t\right\} \le \left[\frac{9k\tau'(k;q_1)}{t} \cdot \mathrm{e}^{-4q_2\sqrt{\gamma_r}}\right]^{p/2}.$$

The norm $|||\cdot|||$ is an arbitrary Q-norm.

The proof of Lemma 6.1 appears in Section 7.

We can establish Theorems 2.5 and 2.9 as an immediate consequence of Lemma 6.1. As noted, the overall strategy is similar to the arguments in [MM15, Thm. 10].

6.2. **Proof of Theorem 2.9.** Suppose that the there is a gap between the kth and $(k+1)$th singular values of the input matrix $\boldsymbol{C}$. This assumption manifests in the property that $\gamma_k > 0$. We can invoke Lemma 6.1 with $r = k$ to see that

$$\mathbb{P}\left\{|||\boldsymbol{C} - \hat{\boldsymbol{C}}_k|||^2 > |||\boldsymbol{C} - \boldsymbol{C}_k|||^2 + t\right\} \le \left[\frac{9k\tau'(k;q_1)}{t} \cdot \mathrm{e}^{-4q_2\sqrt{\gamma_k}}\right]^{p/2}.$$

Indeed, the approximation $\hat{\boldsymbol{C}}_k = \boldsymbol{\Pi}_k \boldsymbol{C}$. We immediately obtain an expectation bound by integration:

$$\mathbb{E}\,|||\boldsymbol{C} - \hat{\boldsymbol{C}}_k|||^2 \le |||\boldsymbol{C} - \boldsymbol{C}_k|||^2 + \frac{p}{p-2} \cdot 9k\tau'(k;q_1) \cdot \mathrm{e}^{-4q_2\sqrt{\gamma_k}}.$$

See the companion paper [Tro18b, Sec. 2.3] for a more detailed version of the latter calculation. Finally, we remove the normalization $c_{k+1} = 1$.

Remark 6.2 (Lower-Rank Approximations)**.** The same argument allows us to obtain bounds for the quality of the rank-r approximation $\hat{\boldsymbol{C}}_r := \boldsymbol{\Pi}_r \boldsymbol{C}$ where $r < k$. This modification can be useful when there is a gap in spectrum between c_r and c_{k+1}; that is, when γ_r is bounded away from zero.

6.3. **Proof of Theorem 2.5.** Next, we develop a bound that is useful even when there is no spectral gap. We have chosen to focus on the spectral norm because it is the most important case, and we can avoid complicated arguments that arise for other norms.

6.3.1. *The Probability Bound.* Fix a parameter $\varepsilon \in (0,1)$. By the definition (2.9) of the tail spectral gap γ_i, the condition $\gamma_i \le \varepsilon$ implies that $c_i^2 \le 1/(1-\varepsilon)$. We have used the normalization $c_{k+1} = 1$.

First, assume that the tail spectral gap $\gamma_i \le \varepsilon$ for each index i. In particular $\gamma_1 \le \varepsilon$. Therefore,

$$\|\boldsymbol{C} - \hat{\boldsymbol{C}}_k\|^2 = \|(\mathbf{I} - \boldsymbol{\Pi}_k)\boldsymbol{C}\|^2 \le \|\boldsymbol{C}\|^2 = c_1^2 \le \frac{1}{1-\varepsilon} = 1 + \frac{\varepsilon}{1-\varepsilon}. \tag{6.1}$$

The first relation holds because $\hat{\boldsymbol{C}}_k = \boldsymbol{\Pi}_k \boldsymbol{C}$. The inequality depends on the fact that $\gamma_1 \le \varepsilon$. This bound is deterministic.

Otherwise, we may select the first index $r \le k$ where the tail spectral gaps satisfy

$$\gamma_{r+1} \le \varepsilon \le \gamma_r.$$

This index must exist because $\gamma_1 > \varepsilon$ and $\gamma_{k+1} = 0$. We can instate Lemma 6.1 with this choice of r to obtain a bound for the spectral norm error:

$$\mathbb{P}\left\{\|\boldsymbol{C} - \boldsymbol{\Pi}_r \boldsymbol{C}\|^2 > \|\boldsymbol{C} - \boldsymbol{C}_r\|^2 + t\right\} \le \left[\frac{9k\tau'(k;q_1)}{t} \cdot \mathrm{e}^{-4q_2\sqrt{\gamma_r}}\right]^{p/2}.$$

Since $\gamma_r \ge \varepsilon$ and $\|\boldsymbol{C} - \boldsymbol{C}_r\|^2 = c_{r+1}^2 \le 1/(1-\varepsilon)$, we have

$$\mathbb{P}\left\{\|\boldsymbol{C} - \boldsymbol{\Pi}_r \boldsymbol{C}\|^2 > \frac{1}{1-\varepsilon} + t\right\} \le \left[\frac{9k\tau'(k;q_1)}{t} \cdot \mathrm{e}^{-4q_2\sqrt{\varepsilon}}\right]^{p/2}.$$

Properties (ii) and (iii) of the graded spectral projectors imply that

$$\|\boldsymbol{C} - \hat{\boldsymbol{C}}_k\|^2 = \|(\mathbf{I} - \boldsymbol{\Pi}_k)\boldsymbol{C}\|^2 \le \|(\mathbf{I} - \boldsymbol{\Pi}_r)\boldsymbol{C})\|^2 = \|\boldsymbol{C} - \boldsymbol{\Pi}_r\boldsymbol{C}\|^2,$$

which yields

$$\mathbb{P}\left\{\|\boldsymbol{C} - \hat{\boldsymbol{C}}_k\|^2 > \frac{1}{1-\varepsilon} + t\right\} \le \left[\frac{9k\tau'(k;q_1)}{t} \cdot \mathrm{e}^{-4q_2\sqrt{\varepsilon}}\right]^{p/2}. \tag{6.2}$$

Set $t = \varepsilon/(1-\varepsilon)$ to obtain the bound

$$\mathbb{P}\left\{\|\boldsymbol{C} - \hat{\boldsymbol{C}}_k\|^2 > 1 + \frac{2\varepsilon}{1-\varepsilon}\right\} \le \left[9k\tau'(k;q_1) \cdot \frac{1-\varepsilon}{\varepsilon} \cdot \mathrm{e}^{-4q_2\sqrt{\varepsilon}}\right]^{p/2}.$$

In view of (6.1), we see that the same probability bound holds trivially for the other case as well. Finally, we revert to the variable $u = 2\varepsilon/(1-\varepsilon)$, which gives

$$\mathbb{P}\left\{\|\boldsymbol{C} - \hat{\boldsymbol{C}}_k\|^2 > 1 + u\right\} \le \left[\frac{18k\tau'(k;q_1)}{u} \cdot \mathrm{e}^{-4q_2\sqrt{u/(2+u)}}\right]^{p/2}.$$

Finally, we remove the normalization $c_{k+1} = 1$. This completes the probability bound.

6.3.2. *The Expectation Bound.* To bound the expectation, we integrate the probability bound (6.2) to obtain

$$\mathbb{E}\|\boldsymbol{C} - \hat{\boldsymbol{C}}_k\|^2 \le \frac{1}{1-\varepsilon} + \frac{p}{p-2} \cdot 9k\tau'(k;q_1) \cdot \mathrm{e}^{-4q_2\sqrt{\varepsilon}}.$$

Finally, remove the normalization $c_{k+1} = 1$, and optimize this inequality over the free parameter $\varepsilon \in (0,1)$. This is the stated result.

Let us take a moment to outline the regimes of behavior that the expectation bound describes, maintaining the normalization $c_{k+1} = 1$ for this discussion. Define the quantity

$$\varepsilon_\star := \left[\frac{\log\left(9k\tau'(k;q_1)q_2^2\right)}{4q_2}\right]_+^2.$$

First, assume that $\varepsilon_\star = 0$. Then $9k\tau'(k;q_1) \le 1/q_2^2$. The choice $\varepsilon = \varepsilon_\star$ yields

$$\mathbb{E}\|\boldsymbol{C} - \hat{\boldsymbol{C}}_k\|^2 \le 1 + \frac{p}{p-2} \cdot \frac{1}{q_2^2}.$$

If $0 < \varepsilon_\star \le 1/2$, then the choice $\varepsilon = \varepsilon_\star$ yields

$$\mathbb{E}\|\boldsymbol{C} - \hat{\boldsymbol{C}}_k\|^2 \le 1 + 2\varepsilon_\star + \frac{p}{p-2} \cdot \frac{1}{q_2^2} \le 1 + \frac{p/(p-2) + 0.125\log^2\left(9k\tau'(k;q_1) \cdot q_2^2\right)}{q_2^2}.$$

The latter display should be compared with [Tro18b, Thm. 1.5]. Finally, if $\varepsilon_\star > 1/2$, we simply select $\varepsilon = 1/2$ to obtain

$$\mathbb{E}\|\boldsymbol{C} - \hat{\boldsymbol{C}}_k\|^2 \le 3 + \frac{p}{p-2} \cdot 9k\tau'(k;q_1) \cdot \mathrm{e}^{-q_2\sqrt{8}}.$$

We see that the behavior of the expected approximation error depends in a significant way on the precise value of $\varepsilon_\star$.

7. Technical Lemma on Matrix Approximation

This section contains the proof of Lemma 6.1. The argument is inspired by the proof of [MM15, Lem. 9, Claim 1].

7.1. **Setup.** We maintain the notation from Section 5. Fix an index $1 \le r \le k$. First, we observe that

$$\|\boldsymbol{C}_r\|_{\mathrm{F}}^2 - \|\boldsymbol{\Pi}_r \boldsymbol{C}\|_{\mathrm{F}}^2 \le \eta \quad \text{implies} \quad |||\boldsymbol{C} - \boldsymbol{\Pi}_r \boldsymbol{C}|||^2 \le |||\boldsymbol{C} - \boldsymbol{C}_r|||^2 + \eta. \tag{7.1}$$

Indeed, the premise ensures that

$$\begin{aligned}\|\boldsymbol{C} - \boldsymbol{\Pi}_r \boldsymbol{C}\|_{\mathrm{F}}^2 = \|(\mathbf{I} - \boldsymbol{\Pi}_r)\boldsymbol{C}\|_{\mathrm{F}}^2 = \|\boldsymbol{C}\|_{\mathrm{F}}^2 - \|\boldsymbol{\Pi}_r \boldsymbol{C}\|_{\mathrm{F}}^2 = \|\boldsymbol{C} - \boldsymbol{C}_r\|_{\mathrm{F}}^2 + \left[\|\boldsymbol{C}_r\|_{\mathrm{F}}^2 - \|\boldsymbol{\Pi}_r \boldsymbol{C}\|_{\mathrm{F}}^2\right] \\ \le \|\boldsymbol{C} - \boldsymbol{C}_r\|_{\mathrm{F}}^2 + \eta.\end{aligned}$$

The second and third identities follow from orthogonality, and the inequality depends on the premise. The conclusion of (7.1) emerges from reversed Mirsky inequality, Proposition 5.2.

Introduce the quantity

$$\Delta_r := \|\boldsymbol{C}_r\|_{\mathrm{F}}^2 - \|\boldsymbol{\Pi}_r \boldsymbol{C}\|_{\mathrm{F}}^2. \tag{7.2}$$

For a parameter $t > 0$, we will produce a bound for the probability

$$P := \mathbb{P}\left\{|||\boldsymbol{C} - \boldsymbol{\Pi}_r \boldsymbol{C}|||^2 \ge |||\boldsymbol{C} - \boldsymbol{C}_r|||^2 + t\right\} \le \mathbb{P}\{\Delta_r > t\}. \tag{7.3}$$

The second inequality follows from (7.1). To complete the estimate, we reduce the problem to a question about randomized matrix approximation.

7.2. **Construction of a Projector.** Fix a filter polynomial $\varphi \in \mathscr{P}_q$ that we will describe later. Define the matrix $\boldsymbol{Y} := \varphi(\boldsymbol{C}^2)\boldsymbol{C}\boldsymbol{\Omega}$. Construct the orthogonal projector $\boldsymbol{P}_{\boldsymbol{Y}}$ onto the range of $\boldsymbol{Y}$, which is clearly contained within the block Krylov space $K_q(\boldsymbol{C}^2; \boldsymbol{C}\boldsymbol{\Omega})$ because of (1.3). Form the rank-r orthogonal projector $\boldsymbol{P}_r$ onto the range of the matrix $\boldsymbol{P}_{\boldsymbol{Y}}\boldsymbol{C}_r$, which is also contained in the block Krylov subspace. Its key property is $\boldsymbol{P}_r \boldsymbol{C}_r = \boldsymbol{P}_{\boldsymbol{Y}} \boldsymbol{C}_r$.

The maximum principle, Lemma 5.1, implies that

$$\|\boldsymbol{\Pi}_r \boldsymbol{C}\|_{\mathrm{F}}^2 \ge \|\boldsymbol{P}_r \boldsymbol{C}\|_{\mathrm{F}}^2 = \operatorname{tr}\left[\boldsymbol{P}_r \boldsymbol{C}^2 \boldsymbol{P}_r\right] \ge \operatorname{tr}\left[\boldsymbol{P}_r \boldsymbol{C}_r^2 \boldsymbol{P}_r\right] = \operatorname{tr}\left[\boldsymbol{P}_{\boldsymbol{Y}} \boldsymbol{C}_r^2 \boldsymbol{P}_{\boldsymbol{Y}}\right].$$

The second inequality holds because of the semidefinite relation $\boldsymbol{C}_r^2 \preccurlyeq \boldsymbol{C}^2$.

As a consequence, the quantity (7.2) satisfies the bound,

$$\Delta_r \le \operatorname{tr} \boldsymbol{C}_r^2 - \operatorname{tr}\left[\boldsymbol{P}_{\boldsymbol{Y}} \boldsymbol{C}_r^2\right] = \operatorname{tr}\left[(\mathbf{I} - \boldsymbol{P}_{\boldsymbol{Y}})\boldsymbol{C}_r^2\right] = \operatorname{tr}\left[(\mathbf{I} - \boldsymbol{P}_{\boldsymbol{Y}})\boldsymbol{C}_r^2(\mathbf{I} - \boldsymbol{P}_{\boldsymbol{Y}})\right]. \tag{7.4}$$

We have used the cyclicity of the trace and the fact that orthogonal projectors are idempotent.

7.3. **Reduction to Randomized Matrix Approximation.** The next step is to massage the bound (7.4) so that we can control it using methods from randomized matrix approximation. This requires us to replace $\boldsymbol{C}_r$ with the filtered matrix $\varphi(\boldsymbol{C}_r^2)\boldsymbol{C}_r$.

Before we continue, let us instate some assumptions on the filter polynomial. Assume that $s \mapsto 1/\varphi^2(s)$ is decreasing on the interval $s \ge 1$. Then

$$c_i^2 = \frac{1}{\varphi^2(c_i^2)} \cdot \varphi^2(c_i^2) c_i^2 \le \frac{1}{\varphi^2(c_r^2)} \cdot \varphi^2(c_i^2) c_i^2 \quad \text{for each index } 1 \le i \le r.$$

Since $\boldsymbol{C}_r = \operatorname{diag}(c_1, \dots, c_r, 0, \dots 0)$, we can rewrite these inequalities as a semidefinite relation:

$$\boldsymbol{C}_r^2 \preccurlyeq \frac{1}{\varphi^2(c_r^2)} \cdot \varphi^2(\boldsymbol{C}_r^2)\boldsymbol{C}_r^2$$

Introduce this expression into (7.4) to arrive at

$$\Delta_r \le \frac{1}{\varphi^2(c_r^2)} \cdot \operatorname{tr}\left[(\mathbf{I} - \boldsymbol{P}_{\boldsymbol{Y}})\varphi^2(\boldsymbol{C}_r^2)\boldsymbol{C}_r^2(\mathbf{I} - \boldsymbol{P}_{\boldsymbol{Y}})\right] = \frac{1}{\varphi^2(c_r)} \cdot \left\|(\mathbf{I} - \boldsymbol{P}_{\boldsymbol{Y}})\varphi(\boldsymbol{C}_r^2)\boldsymbol{C}_r\right\|_{\mathrm{F}}^2. \tag{7.5}$$

Note the parallel with Fact 5.4 on randomized matrix approximation.

7.4. **Randomized Matrix Approximation.** We plan to invoke Fact 5.4, on randomized matrix approximation, to control the probability (7.3) that Δ_r is large:

$$P = \mathbb{P}\{\Delta_r > t\} \le \mathbb{P}\left\{ \left\| (\mathbf{I} - \boldsymbol{P}_{\boldsymbol{Y}}) \varphi(\boldsymbol{C}_r^2) \boldsymbol{C}_r \right\|_{\mathrm{F}}^2 > t\varphi^2(c_r^2) \right\}.$$

The second inequality is (7.5).

To align our notation with Fact 5.4, set $\boldsymbol{M} := \varphi(\boldsymbol{C}^2)\boldsymbol{C}$, and observe that $\boldsymbol{Y} = \boldsymbol{M}\boldsymbol{\Omega}$. Furthermore, with $\boldsymbol{R}$ the orthogonal projector onto the first r coordinates, we see that $\boldsymbol{M}\boldsymbol{R} = \varphi(\boldsymbol{C}_r^2)\boldsymbol{C}_r$. Meanwhile, $\boldsymbol{M}_2 = \operatorname{diag}(\varphi(c_{k+1}^2)c_{k+1}, \dots, \varphi(c_n^2)c_n)$. We determine that

$$P \le \left[\frac{2.25k \|\boldsymbol{M}_2\|_{\mathrm{F}}^2}{t\varphi^2(c_r^2)} \right]^{p/2} = \left[\frac{2.25k \sum_{j>k} \varphi^2(c_j^2) c_j^2}{t\varphi^2(c_r^2)} \right]^{p/2}. \tag{7.6}$$

It remains to select a filter polynomial $\varphi \in \mathscr{P}_q$.

7.5. **The Filter Polynomial.** We must design a filter polynomial to amplify the large singular values ($c_j > 1$) and to attenuate the singular values in the tail ($c_j \le 1$) of the matrix. Let $q := q_1 + q_2$ be an integer partition where $0 \le q_1 \le q$. Consider the polynomial

$$\varphi_1(s) := \left(\frac{s - c_n^2}{1 - c_n^2} \right)^{q_1} T_{q_2}\left(\frac{2(s - c_n^2)}{1 - c_n^2} - 1 \right). \tag{7.7}$$

We have written T_{q_2} for the Chebyshev polynomial of the first kind with degree q_2. As required, $s \mapsto 1/\varphi_1^2(s)$ is decreasing on the interval $s \ge 1$. The minimax property of the Chebyshev polynomial implies that

$$\varphi_1^2(s) \le \left(\frac{s - c_n^2}{1 - c_n^2} \right)^{2q_1} \quad \text{for } 0 \le s \le 1. \tag{7.8}$$

The growth properties of the Chebyshev polynomial yield the bound

$$\frac{1}{\varphi_1^2(s)} \le 4 \left(\frac{s - c_n^2}{1 - c_n^2} \right)^{-2q_1} \mathrm{e}^{-4q_2\sqrt{1 - (1 - c_n^2)/(s - c_n^2)}} \quad \text{for } s \ge 1. \tag{7.9}$$

The filter polynomial φ_1 is (more or less) a reparameterization of polynomials that we have considered in the companion paper [Tro18a, Sec. 4.4, Sec. 6, App. A.2.3], which contains a proof of these properties. See also [Tro18b, Sec. 5.6].

7.6. **Installing the Filter Polynomial.** We are prepared to complete the proof of Lemma 6.1. To do so, we introduce the filter polynomial (7.7) into the bound (7.6). Since $c_r \ge 1$ and $c_j \le 1$ for $j > k$, we determine that

$$\begin{aligned} \frac{\varphi_1^2(c_j^2)}{\varphi_1^2(c_r^2)} &\le 4 \left(\frac{c_j^2 - c_n^2}{1 - c_n^2} \right)^{2q_1} \left(\frac{c_r^2 - c_n^2}{1 - c_n^2} \right)^{-2q_1} \mathrm{e}^{-4q_2\sqrt{1 - (1 - c_n^2)/(c_r^2 - c_n^2)}} \\ &\le 4 \left(\frac{c_j^2 - c_n^2}{c_k^2 - c_n^2} \right)^{2q_1} \mathrm{e}^{-4q_2\sqrt{(c_r^2 - 1)/(c_r^2 - c_n^2)}} \end{aligned}$$

The bound on the numerator is (7.8), and the bound on the denominator is (7.9). We have used the fact that $c_r \le c_k$ to reach the second line. As a consequence,

$$\frac{1}{\varphi_1^2(c_r^2)} \sum_{j>k} \varphi_1^2(c_j^2) c_j^2 \le 4 \left[\sum\nolimits_{j>k} c_j^2 \left(\frac{c_j^2 - c_n^2}{c_k^2 - c_n^2} \right)^{2q_1} \right] \mathrm{e}^{-4q_2\sqrt{(c_r^2 - 1)/(c_r^2 - c_n^2)}} \le 4\tau'(k; q_1) \cdot \mathrm{e}^{-4q_2\sqrt{\gamma_r}}.$$

We have used the assumption $c_{k+1} = 1$ to identify the modified tail content (2.10) and the tail spectral gap (2.9). Substitute the last display into (7.6) to arrive at

$$P \le \left[\frac{9k\tau'(k; q_1)}{t} \cdot \mathrm{e}^{-4q_2\sqrt{\gamma_r}} \right]^{p/2}.$$

This completes the proof of Lemma 6.1.

ACKNOWLEDGMENTS

This paper was produced under a subcontract to ONR Award 00014-16-C-2009. The author thanks Shaunak Bopardikar for helpful discussions.

REFERENCES

[Bha97] R. Bhatia. *Matrix analysis*, volume 169 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1997.

[CD74] J. Cullum and W. E. Donath. A block generalization of the *s*-step Lanczos algorithm. Technical Report Report RC 4845 (21570), IBM Thomas J. Watson Research Center, Yorktown Heights, New York, 1974.

[DI18] P. Drineas and I. Ipsen. Low-rank matrix approximations do not need a singular value gap. Available at http://arxiv.org/abs/1801.00670, Jan. 2018.

[DIKMI16] P. Drineas, I. Ipsen, E.-M. Kontopoulou, and M. Magdon-Ismail. Structural convergence results for approximation of dominant subspaces from block krylov spaces. Available at http://arXiv.org/abs/1609.00671, Sep. 2016.

[DM05] P. Drineas and M. W. Mahoney. On the Nyström method for approximating a Gram matrix for improved kernel-based learning. *J. Mach. Learn. Res.*, 6:2153–2175, 2005.

[Git13] A. Gittens. *Topics in Randomized Numerical Linear Algebra*. PhD thesis, California Institute of Technology, 2013.

[GM16] A. Gittens and M. W. Mahoney. Revisiting the Nyström method for improved large-scale machine learning. *J. Mach. Learn. Res.*, 17:Paper No. 117, 65, 2016.

[GU77] G. H. Golub and R. Underwood. The block Lanczos method for computing eigenvalues. In *Mathematical software, III (Proc. Sympos., Math. Res. Center, Univ. Wisconsin, Madison, Wis., 1977)*, pages 361–377. Publ. Math. Res. Center, No. 39. Academic Press, New York, 1977.

[Gu15] M. Gu. Subspace iteration randomization and singular value problems. *SIAM J. Sci. Comput.*, 37(3):A1139–A1173, 2015.

[Hig89] N. J. Higham. Matrix nearness problems and applications. In *Applications of matrix theory (Bradford, 1988)*, volume 22 of *Inst. Math. Appl. Conf. Ser. New Ser.*, pages 1–27. Oxford Univ. Press, New York, 1989.

[HMST11] N. Halko, P.-G. Martinsson, Y. Shkolnisky, and M. Tygert. An algorithm for the principal component analysis of large data sets. *SIAM J. Sci. Comput.*, 33(5):2580–2594, 2011.

[HMT11] N. Halko, P. G. Martinsson, and J. A. Tropp. Finding structure with randomness: probabilistic algorithms for constructing approximate matrix decompositions. *SIAM Rev.*, 53(2):217–288, 2011.

[KW92] J. Kuczyński and H. Woźniakowski. Estimating the largest eigenvalue by the power and Lanczos algorithms with a random start. *SIAM J. Matrix Anal. Appl.*, 13(4):1094–1122, 1992.

[LLS+17] H. Li, G. C. Linderman, A. Szlam, K. P. Stanton, Y. Kluger, and M. Tygert. Algorithm 971: An implementation of a randomized algorithm for principal component analysis. *ACM Trans. Math. Softw.*, 43(3):28:1–28:14, Jan. 2017.

[Mir60] L. Mirsky. Symmetric gauge functions and unitarily invariant norms. *Quart. J. Math. Oxford Ser. (2)*, 11:50–59, 1960.

[MM15] C. Musco and C. Musco. Stronger and faster approximate singular value decomposition via the block Lanczos method. In *Proc. Adv. Neural Information Processing Systems 28 (NIPS 2015)*, 2015. Available at http://arXiv.org/abs/1504.05477.

[MRT11] P.-G. Martinsson, V. Rokhlin, and M. Tygert. A randomized algorithm for the decomposition of matrices. *Appl. Comput. Harmon. Anal.*, 30(1):47–68, 2011.

[PAB16] F. Pourkamali-Anaraki and S. Becker. Randomized clustered Nyström for large-scale kernel machines. Available at http://arXiv.org/abs/1612.06470, Dec. 2016.

[RST09] V. Rokhlin, A. Szlam, and M. Tygert. A randomized algorithm for principal component analysis. *SIAM J. Matrix Anal. Appl.*, 31(3):1100–1124, 2009.

[Tro18a] J. A. Tropp. Analysis of randomized block Krylov methods. Part I: Spectral norm estimation. Available at http://users.cms.caltech.edu/~jtropp/, Feb. 2018.

[Tro18b] J. A. Tropp. Analysis of randomized block Krylov methods. Part II: Eigenvalue estimation. Available at http://users.cms.caltech.edu/~jtropp/, Feb. 2018.

[TYUC17] J. A. Tropp, A. Yurtsever, M. Udell, and V. Cevher. Fixed-rank approximation of a positive-semidefinite matrix from streaming data. including supplementary appendix. In *Proc. Adv. Neural Information Processing Systems 30 (NIPS 2017)*, 2017.

[WGM17] S. Wang, A. Gittens, and M. W. Mahoney. Scalable kernel K-means clustering with Nyström approximation: relative-error bounds. Available at http://arXiv.org/abs/1706.02803, June 2017.

[WS00] C. K. I. Williams and M. Seeger. Using the Nyström method to speed up kernel machines. In *Adv. Neural Information Processing Systems 13 (NIPS)*, 2000.

[WZZ15] S. Wang, Z. Zhang, and T. Zhang. Improved analyses of the randomized power method and block Lanczos method. Available at http://arXiv.org/abs/1508.06429, Aug. 2015.