SHARP ERROR BOUNDS FOR APPROXIMATE EIGENVALUES AND SINGULAR VALUES FROM SUBSPACE METHODS *

IRINA-BEATRICE HAAS† AND YUJI NAKATSUKASA†

Abstract. Subspace methods are commonly used for finding approximate eigenvalues and singular values of large-scale matrices. Once a subspace is found, the Rayleigh-Ritz method (for symmetric eigenvalue problems) and Petrov-Galerkin projection (for singular values) are the de facto method for extraction of eigenvalues and singular values. In this work we derive quadratic error bounds for approximate eigenvalues of symmetric matrices obtained via the Rayleigh-Ritz process. Our bounds take advantage of the fact that extremal eigenpairs tend to converge faster than the rest, hence having smaller residuals $\|A\widehat{x}_i - \theta_i \widehat{x}_i\|_2$, where $(\theta_i, \widehat{x}_i)$ is a Ritz pair (approximate eigenpair). The proof uses the structure of the perturbation matrix underlying the Rayleigh-Ritz method to bound the components of its eigenvectors. In this way, we obtain a bound of the form $c\frac{\|A\widehat{x}_i-\theta_i\widehat{x}_i\|_2^2}{\mathrm{Gap}_i}$, where Gap_i is roughly the gap between the ith Ritz value and the eigenvalues that are not approximated by the Ritz process, and $c > 1$ is a modest scalar. Our bound is adapted to each Ritz value and is robust to clustered Ritz values, which is a key improvement over existing results. We further show that the bound is asymptotically sharp, and generalize it to singular values of arbitrary real matrices. Finally, we apply these bounds to several methods for computing eigenvalues and singular values, and illustrate the sharpness of our bounds in a number of computational settings, including Krylov methods and randomized algorithms.

Key words. Matrix perturbation theory, Ritz values, Rayleigh-Ritz process, singular values, Petrov-Galerkin process

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

1. Introduction.

The symmetric eigenvalue problem and the Singular value decomposition (SVD) are key computational tasks in many numerical methods for engineering and data analysis applications. Classical algorithms perform these decompositions in polynomial time (cubic in the dimension, for an ϵ-accurate solution), but the matrices that are considered nowadays are often of very large scale, sometimes even too large to fit into memory. Subspace methods (including Krylov subspace methods [20] and randomized algorithms [5]) form a leading class of methods that allow us to tackle such problems. Once a subspace $\mathrm{Span}(Q)$ is identified, where $Q \in \mathbb{C}^{n\times k}$ is a matrix with orthonormal columns, the most common approach to extracting eigenvalues is the Rayleigh-Ritz (RR) process [20, Ch. 11], which outputs the eigenvalues of Q^TAQ as approximations to those of A. The quality of the Ritz eigenpairs as an approximation to the actual eigenpairs of the original matrix has been extensively studied in the literature: for the eigenvectors, a bound on the canonical (or *principal*) angle between the subspace spanned by the Ritz vectors and the actual eigenspace is provided in the classical Davis-Kahan sin θ theorem [3][20, Chap.11] and was improved in [16]. For the Ritz values, error bounds have been derived in [7, 12, 20][6, Cor.7.3.5][22, Cor.1.4.31]. A similar process is applicable to approximate the (usually leading) singular values of $A \in \mathbb{C}^{m\times n}$ from subspace(s) spanned by $Q_1 \in \mathbb{C}^{m\times k_1}, Q_2 \in \mathbb{C}^{n\times k_2}$, as the singular values of the projected matrix $Q_1^TAQ_2$; this is called a *Petrov-Galerkin (PG) projection* method. More recently, randomized algorithms such as randomized SVD by Halko, Martinson and Tropp (HMT) [5] and (Generalized) Nyström [2, 17, 24] have been receiving significant attention, as they

---

[*] Submitted to the editors on June 1st, 2025.

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

provide an efficient way of computing a low-rank approximation of the original matrix that is nearly optimal in the Frobenius norm [5]. Similarly to RR and PG, randomized SVD is also based on (one-sided) projection, $A \approx QQ^T A$, where $\mathrm{Span}(Q) = \mathrm{Span}(A\Omega)$ for a random (sketch) matrix Ω. Error bounds on the resulting singular values for these methods are derived in [4, 11, 21].

In this paper we first examine the accuracy of the output approximate eigenvalues from the RR process. Our work is primarily motivated by the fact that existing error bounds for the Ritz eigenvalues can be improved greatly, assuming that the norms of the residuals corresponding to each eigenvalue differ significantly, which is common in practice. That is why we derive sharp upper bounds on the difference between the Ritz values and the exact eigenvalues by exploiting the perturbation structure underlying the Rayleigh-Ritz process. We then generalize the result for singular values of arbitrary matrices, and compare the bound to the approximate singular values obtained from HMT in our numerical experiments, along with other approximation methods including subspace iteration and Krylov methods.

To make the problem precise, let us fix the notation and consider the structure afforded by the RR and PG processes.

Notation : We use MATLAB notation for matrix indexing, in which $X(:,i)$ denotes the ith column of X and $X(:,i:j)$ is the matrix consisting of the ith to the jth columns of X. If not stated otherwise, we also note X_i the ith column of the matrix X. We let $\|X\|_2$ denote the spectral norm of a matrix X, which reduces to the Euclidean norm when X is a vector. For simplicity we focus on the real symmetric case (for the eigenvalue problem). Our analysis holds more generally verbatim for complex Hermitian matrices; in this case, one replaces ^{T} with the Hermitian transpose * in what follows. Finally, the error bounds presented below often involve spectral gaps that we denote $\mathrm{Gap}_i, \mathrm{gap}_i$ or Γ_i. By definition Gap_i and Γ_i will involve the part of the spectrum that is not approximated and the eigenvalue whose error is bounded, while gap_i will involve (roughly) the approximated eigenvalue and the next closest exact eigenvalue. Therefore typically $\mathrm{Gap}_i > \mathrm{gap}_i$, and we emphasize the case where $\mathrm{Gap}_i \gg \mathrm{gap}_i$. We also use the letter δ to denote available lower bounds on the spectral gaps.

Perturbed matrix and available information from the Rayleigh-Ritz process. Suppose that k eigenpairs of an $n \times n$ symmetric matrix A are sought (typically the smallest or largest eigenpairs), and we have a k-dimensional[1] trial subspace spanned by an $n \times k$ matrix Q_1 with orthonormal columns that approximates the corresponding eigenspace. To extract an approximation to the desired eigenpairs from Q_1 with RR one computes the $k \times k$ symmetric eigendecomposition $Q_1^T A Q_1 = Y\widehat{\Lambda}Y^T$ where $\widehat{\Lambda} = \mathrm{diag}(\theta_1, \ldots, \theta_k)$ is the matrix of Ritz values and Y is unitary (see e.g., [20, Ch. 11]). The matrix of Ritz vectors $\widehat{X} = [\widehat{x}_1, \ldots, \widehat{x}_k]$ is defined as $\widehat{X} = Q_1 Y$. Then let $\widehat{X}_\perp$ be such that $[\widehat{X}\ \widehat{X}_\perp]$ is a unitary matrix, and form the *perturbation matrix*

$$\bar{A} = [\widehat{X}\ \widehat{X}_\perp]^T A [\widehat{X}\ \widehat{X}_\perp] = \begin{bmatrix} \widehat{\Lambda} & E^T \\ E & A_2 \end{bmatrix} \tag{1.1}$$

where $E = \widehat{X}_\perp^T A \widehat{X} = [E_1, E_2, \ldots, E_k]$. The matrix $\bar{A}$ is similar to A so it has the same eigenvalues as A. Throughout the paper we let λ_i (resp. σ_i) denote the ith eigenvalue (resp. singular value) of the matrix A. The overarching goal is to bound $|\theta_i - \lambda_i|$, the error in the Ritz values as approximations to the exact eigenvalues λ_i.

---

[1] In practice one often uses a subspace whose dimension is larger than the desired number of eigenpairs k, but this does not affect what follows.

In practice, the information available after RR are the Ritz pairs $(\theta_i, \widehat{x}_i)$ for $i = 1, \ldots, k$ and the norms of the columns of E, because they are equal to the norm of the residual: $\|E_i\|_2 = \|A\widehat{x}_i - \theta_i\widehat{x}_i\|_2$. In addition, since Q_1 is designed to approximate a k-dimensional eigenspace, a rough estimate of the eigenvalues of A_2 is usually available. For example, when Q_1 approximates the smallest eigenspace we can reasonably expect $\lambda(A_2) \gtrsim \max_i \theta_i$ (this must hold in the limit $E \to 0$), where $\lambda(A_2)$ denotes the spectrum of A_2. Another option is to use oversampling, i.e. to approximate $k+\ell$ eigenvalues (or singular values) so that the ℓ additional Ritz values provide a good approximation of $\mathrm{Gap}_i := \min_j |\theta_i - \lambda_j(A_2)| \approx \min_{k<j\leq k+\ell} |\theta_i - \theta_j|$. More generally to obtain an error bound on θ_i we only need to assume the knowledge of a lower bound $\widehat{\delta}_i$ on Gap_i.

A key observation here is that the residuals $\|E_i\|_2$ typically vary with i by several orders of magnitude. For example, if the smallest k eigenpairs are sought and the Ritz values θ_i are arranged in increasing order, then we typically have a *graded* structure $\|E_1\|_2 \lesssim \|E_2\|_2 \lesssim \ldots \lesssim \|E_k\|_2$, with $\|E_1\|_2 \ll \|E_k\|_2$, because extremal eigenvalues tend to converge faster than interior ones by widely used methods such as Krylov subspace methods (like Lanczos) and LOBPCG [1, 9] (and HMT for singular values); this is related to the convergence of the power method. In this situation we observe that the error $|\theta_i - \lambda_i|$ is $O(\|E_i\|_2^2)$. Our goal is to derive bounds that accurately reflect this for each eigenpair.

Turning to the general case, to approximate k singular values of a general matrix $A \in \mathbb{R}^{m\times n}$, using the PG method, assume we have a left trial subspace spanned by $Q_1 \in \mathbb{R}^{m\times k}$ and a right trial subspace spanned by $Q_2 \in \mathbb{R}^{n\times k_2}$, with Q_1, Q_2 orthogonal, and $k_2 \geq k$ without loss of generality. Take the SVD decomposition of $Q_1^T A Q_2 = X\widehat{\Sigma}Y^T$ with $\widehat{\Sigma} = \mathrm{diag}(\theta_1, \theta_2, \ldots, \theta_k)$ containing the approximate singular values, then multiply A by the appropriate orthogonal matrices on both sides as follows to obtain

$$\bar{A} = \begin{bmatrix} X^T & \\ & I_{m-r} \end{bmatrix} \begin{bmatrix} Q_1^T \\ Q_{1\perp}^T \end{bmatrix} A \begin{bmatrix} Q_2 & Q_{2\perp} \end{bmatrix} \begin{bmatrix} Y & \\ & I_{n-r} \end{bmatrix} = \begin{bmatrix} \widehat{\Sigma} & E^T \\ F & A_2 \end{bmatrix}. \tag{1.2}$$

Note that when $k_2 > k$, the first $k_2 - k$ columns of E^T are zero; indeed $k_2 = n$ is allowed (as will be in HMT as we will see), in which case $E = 0$. The perturbation matrix $\bar{A}$ is equivalent to A up to orthogonal multiplication, so it has the same singular values as A. Similarly to the symmetric case, we assume that the available information is the residual norms $\|E_i\|_2 = \|A^T\hat{u}_i - \theta_i\hat{v}_i\|_2$ and $\|F_i\|_2 = \|A\hat{v}_i - \theta_i\hat{u}_i\|_2$, where F_i and E_i are the columns of F and E respectively, and some gap information on the approximate singular values θ_i and the singular values of A_2.

Below we review existing results that use similar information, namely the residuals and some (approximate) gap information, to derive error bounds for singular values or eigenvalues.

Literature review. Fundamentally, the problem set out in Equation (1.1) is that of eigenvalue perturbation: how do the eigenvalues of $\begin{bmatrix} \widehat{\Lambda} & \\ & A_2 \end{bmatrix}$ change by the off-diagonal perturbation E? A number of results are available in this direction, which we review here. A similar problem was considered in [13], which investigates the special case of a multiple Ritz value, ie. when $\widehat{\Lambda} = \mu I_k$, but the problem considered in this paper is for a general set of Ritz values.

The simplest and most general (i.e. making no use of the perturbation structure) symmetric eigenvalue or singular value perturbation bound is provided by Weyl's theorem [6, Cor.7.3.5][22, Cor.I.4.31], which states that

$$|\theta_i - \lambda_i| \leq \|E_i\|_2, \tag{1.3}$$

in the symmetric case, and in the general case

$$|\theta_i - \sigma_i| \leq \max\{\|E_i\|_2, \|F_i\|_2\}. \tag{1.4}$$

However this only gives a linear bound (with respect to the residual $\|E_i\|_2$) although usually in practice the accuracy of θ_i is much higher. An extensive line of work such as [26] and the references therein offer improved linear bounds. Other authors [11, 12, 15] have derived quadratic bounds that depend on the spectral norm of the error blocks from the perturbation matrix. Improving the bound from Mathias [14], C.-K. Li and R.-C. Li [12] derive an error bound for the case where the perturbation matrix is as in Equations (1.1) and (1.2). For symmetric matrices, using again $\mathrm{Gap}_i = \min_j |\theta_i - \lambda_j(A_2)|$ as above, they prove the bound:

$$|\lambda_i - \theta_i| \leq \frac{2\|E\|_2^2}{\mathrm{Gap}_i + \sqrt{\mathrm{Gap}_i^2 + 4\|E\|_2^2}} \tag{1.5}$$

and similarly, for singular values of general matrices:

$$|\sigma_i - \theta_i| \leq \frac{2\max\{\|E\|_2, \|F\|_2\}^2}{\mathrm{Gap}_i + \sqrt{\mathrm{Gap}_i^2 + 4\max\{\|E\|_2, \|F\|_2\}^2}}, \tag{1.6}$$

with $\mathrm{Gap}_i = \min_j |\theta_i - \sigma_j(A_2)|$.

Alternatively, [11, 15] give more general bounds since their perturbation matrix also potentially has nonzero components in the diagonal blocks. In [11] the authors apply such bound to the structure afforded from Generalized Nyström, HMT and RR and compare the accuracy of these methods for singular value estimation. With our notation and considering only off diagonal perturbations, the bound from [11, Theorem 4.1] is :

$$|\sigma_i - \theta_i| \leq \frac{2\max\{\|E\|_2, \|F\|_2\}^2}{\min_k |\sigma_i - \sigma_k(A_2)| - 2\max\{\|E\|_2, \|F\|_2\}}. \tag{1.7}$$

Equations (1.5)–(1.7) are relevant in cases where the individual residuals are not available, and [12] showed that their bound is asymptotically sharp as $\|E\|_2 \longrightarrow 0$ when Gap_i and $\|E\|_2$ are the only available information.

Recalling from above that we expect $|\theta_i - \lambda_i|$ to be $O(\|E_i\|_2^2)$, the above bounds clearly overestimate the error for the extremal Ritz values.

Other existing quadratic bounds on the accuracy of Ritz values that make use of the individual residuals include the bounds in [7], which uses the angle between Q_1 and an exact eigenspace, and the classical result

$$|\theta_i - \lambda_i| \leq \frac{\|E_i\|_2^2}{\widetilde{\mathrm{gap}}_i} \tag{1.8}$$

from [20, Thm. 11.7.1], where $\widetilde{\mathrm{gap}}_i$ is the gap between θ_i and the eigenvalues of A excluding the one closest to θ_i. When the Ritz values are well separated, the bound Equation (1.8) is much sharper than the previous ones, especially assuming we have graded residuals, but this bound is not usable for clustered singular values and, similarly to [16], we will show in this paper that the denominator in Equation (1.8) can be made larger.

One can also use Equation (1.5) from [12] (or alternatively the bound from [15]) with a 1-by-1 block partitioning as follows: the perturbation matrix $\bar{A}$ can be rearranged as

$$\begin{bmatrix} \theta_i & & E_i^T \\ & \widehat{\Lambda}_{\setminus i} & E_{\setminus i}^T \\ E_i & E_{\setminus i} & A_2 \end{bmatrix} =: \left[\begin{array}{c|c} \theta_i & 0_{1\times(k-1)} \quad E_i^T \\ \hline \begin{matrix} 0_{(k-1)\times 1} \\ E_i \end{matrix} & B \end{array}\right], \tag{1.9}$$

where $\setminus i$ indicates that we removed only the elements corresponding to the i-th Ritz value from the blocks E and $\widehat{\Lambda}$. This gives the bound

$$|\theta_i - \lambda_i| \leq \frac{\|E_i\|_2^2}{\mathrm{gap}_i + \sqrt{\mathrm{gap}_i^2 + 4\|E_i\|_2^2}} \tag{1.10}$$

where the denominator gap_i is the gap between θ_i and the eigenvalues of the $(n-1) \times (n-1)$ matrix B defined above. The drawback of this approach is that it puts the remaining Ritz values in the second diagonal block, and thus it reduces Gap_i in Equations (1.5) and (1.6) to approximately the same denominator as in Equation (1.8). Hence, another way of looking at the contribution made in this paper is that we improve the spectral gap in the error bounds of Ritz values from (roughly) $\mathrm{gap}_i = \min_{j \neq i} |\theta_i - \lambda_j(A)|$ to Gap_i, while keeping the denominator as small as $\|E_i\|_2^2$.

Key contributions. Motivated by the above discussion, in this paper we derive bounds of the form

$$|\theta_i - \lambda_i| \leq c \frac{\|E_i\|_2^2}{\mathrm{Gap}_i}, \tag{1.11}$$

where c is a modest constant such that $c \to 1$ as $E \to 0$. Our proposed bounds are particularly good in the case where the extremal eigenvalues are sought and the residuals are graded (but these conditions are not required for the bounds to be applicable). In this context, we show that our bounds are significantly sharper than those in the literature, including [12, 20], accurately reflecting the error observed in practice. Moreover, contrarily to Equation (1.8), our bound Equation (1.11) is also suitable for clustered eigenvalues, as Gap_i is much larger than the denominator $\widetilde{\mathrm{gap}}_i$ in Equation (1.8).

Since the constant c above converges to 1 as $E \to 0$, we have

$$\lim_{E \to 0} \frac{|\theta_i - \lambda_i|}{\|E_i\|_2^2} \leq \frac{1}{\min_j |\theta_i - \lambda_j(A_2)|} \leq \frac{1}{\widehat{\delta}_i},$$

and we show that this asymptotic bound is sharp, that is, it cannot be improved without more information about the eigenvalues of A_2.

The main message of this work is that when E is sufficiently small, the error in a Ritz value θ_i is bounded by the square of the corresponding residual divided by the gap between θ_i and the eigenvalues of A_2, which are roughly the eigenvalues that are not sought; they converge when $E \to 0$. A practical implication is that in large-scale symmetric eigenvalue problems, working with a subspace of dimension larger than k (which is often done for the purpose of avoiding missing eigenvalues) helps improve the accuracy of the desired ones, because doing so increases the denominator in the error bound Equation (1.11), which implies that a larger tolerance becomes acceptable for the residual $\|E_i\|_2 = \|A\widehat{x}_i - \theta_i\widehat{x}_i\|_2$ to ensure a given required accuracy of θ_i.

The paper is organized as follows: in Section 2 we describe the method that we use to obtain our theoretical error bounds for the eigenvalues of symmetric matrices in Section 3 and for singular values of arbitrary matrices in Section 4. We then illustrate our theoretical bounds with numerical experiments in Section 5 and apply them for estimates obtained via RR, HMT, PG and Lanczos.

2. Basic approach.

In this section we explain the main ideas that we use to derive the theorems in Section 3. We first recall a well-known result on the derivative of simple eigenvalues [23].

LEMMA 2.1. *Let A_0 and F be symmetric matrices. Denote by $\lambda_i(t)$ the ith eigenvalue of $A_0 + tF$ such that $(A_0 + tF)x(t) = \lambda_i(t)x(t)$ where $\|x(t)\|_2 = 1$ for $t \in [0, 1]$. If $\lambda_i(t)$ is simple, then*

$$\frac{d\lambda_i(t)}{dt} = x(t)^* F x(t). \tag{2.1}$$

In our analysis below we will have F of the form $F = \left[\begin{smallmatrix} 0 & F_i^T \\ F_i & 0 \end{smallmatrix}\right]$, where $F_i = [0_{(n-k)\times(i-1)} \;\; E_i \;\; 0_{(n-k)\times(n-i)}]$ has just one (ith) nonzero column. Then with the partitioning $x(t) = \begin{bmatrix} x_{1:k}(t) \\ y(t) \end{bmatrix}$ (where for a vector x, we denote by x_j its jth entry and by $x_{j:\ell}$ with $\ell \geq j$ the vector $[x_j, x_{j+1}, \ldots, x_\ell]^T$), if $\lambda_i(t)$ is simple for $0 \leq t \leq 1$ then by the above lemma we get

$$\begin{aligned}
|\lambda_i(0) - \lambda_i(1)| &= \left| \int_0^1 x(t)^T F x(t) dt \right| && (2.2)\\
&\leq 2 \left| \int_0^1 y(t)^T F_i x_{1:k}(t) dt \right| \\
&= 2 \left| \int_0^1 y(t)^T E_i x_i(t) dt \right| \\
&\leq 2\|E_i\|_2 \left| \int_0^1 \|y(t)\|_2 dt \right|. \qquad (\|x_i(t)\|_2 \leq 1) && (2.3)
\end{aligned}$$

Note that in the setting of (1.1), we have $\lambda_i(0) = \theta_i$ and $\lambda_i(1)$ is an eigenvalue of A, so (2.3) provides an error bound for the Ritz value θ_i. The key observation here is that (2.3) is small if $\|y(t)\|_2$ is small for all $0 \leq t \leq 1$. In view of this, our approach below is to obtain sharp bounds for $\|y(t)\|_2$, from which we get sharp bounds for $|\lambda_i(0) - \lambda_i(1)|$ by (2.3).

The idea of obtaining eigenvalue perturbation bounds via bounding eigenvector components was introduced in [15]. Moreover, it is shown there that (2.3) holds even in the presence of multiple eigenvalues, in which case $x(t)$ can be taken as any of the (many possible) eigenvectors associated with $\lambda_i(t)$, and the bounds hold as long as the bound on $\|y(t)\|_2$ holds for any eigenvector corresponding to the eigenvalue, which is the case in the forthcoming analysis. Hence in what follows we are not concerned with whether $\lambda_i(t)$ is simple or not.

3. Error bounds for Ritz values.

We are now ready to derive error bounds on the Ritz values of symmetric matrices. First we derive an error bound for a single Ritz value that depends on the corresponding residual norm $\|E_i\|_2$, then we use the same approach to derive a bound for a set of clustered eigenvalues based on the spectral norm of the corresponding submatrix of E.

3.1. Error bound for well separated Ritz values.

In this subsection we prove the following theorem, which is the main result of this paper.

THEOREM 1. *Let*

$$A = \begin{bmatrix} \widehat{\Lambda} & E^* \\ E & A_2 \end{bmatrix}$$

be a symmetric matrix where $\widehat{\Lambda} = diag(\theta_1, \dots, \theta_k)$ *and* $E = [E_1, E_2, \dots, E_k]$. *For* $i = 1, 2, \dots, k$, *suppose that*

$$\delta_i = \min_j |\theta_i - \lambda_j(A_2)| - \|E_i\|_2 \ (= \mathrm{Gap}_i - \|E_i\|_2) > 0, \tag{3.1}$$

$$\delta_{i,j} = |\theta_i - \theta_j| - \|E_i\|_2 > 0, \quad \textit{for} \quad 1 \le j \le k,\ j \ne i \tag{3.2}$$

and further that

$$d_i = \frac{1}{\delta_i - \displaystyle\sum_{\substack{1 \le j \le k \\ j \ne i}} \frac{\|E_j\|_2^2}{\delta_{i,j}}} > 0. \tag{3.3}$$

Then there exists an eigenvalue λ_i *of* A *satisfying*

$$|\lambda_i - \theta_i| \le d_i \|E_i\|_2^2. \tag{3.4}$$

The Ritz values θ_i *and eigenvalues* λ_i *have a one-to-one correspondence. Moreover, let* $x(t), \lambda_i(t)$ *be continuous functions of* $t \in [0, 1]$ *defined by*

$$\left[\begin{array}{ccc|c} \widehat{\Lambda}_{(1:i-1)} & & & E_{1:i-1}^T \\ & \theta_i & & tE_i^T \\ & & \widehat{\Lambda}_{(i+1:k)} & E_{i+1:k}^T \\ \hline E_{1:i-1} & tE_i & E_{i+1:k} & A_2 \end{array}\right] x(t) = \lambda_i(t) x(t)$$

and $\lambda_i(0) = \theta_i$, *where we note* $\widehat{\Lambda}_{(a:b)} = \widehat{\Lambda}(a : b, a : b)$. *Then the bound* (3.4) *holds for any* $\delta_i, \delta_{i,j}$ *such that* $0 < \delta_i \le \min_{j, 0 \le t \le 1} |\lambda_i(t) - \lambda_j(A_2)|$ *and* $0 < \delta_{i,j} \le \min_{0 \le t \le 1} |\lambda_i(t) - \theta_j|$.

Proof. Let $i \le k$ be an integer. First, we decompose the matrix A from the theorem as

$$A = \begin{bmatrix} \widehat{\Lambda} & \\ & A_2 \end{bmatrix} + \begin{bmatrix} & E^T \\ E & \end{bmatrix} =: A_0 + \widehat{E}. \tag{3.5}$$

As above let $F = \left[\begin{smallmatrix} 0 & F_i^T \\ F_i & 0 \end{smallmatrix}\right]$, where $F_i = [0_{(n-k)\times(i-1)}\ E_i\ 0_{(n-k)\times(n-i)}]$ is the matrix obtained by taking the ith row and column of $\widehat{E}$. Let $\widetilde{E}$ be a matrix such that $\widehat{E} = F + \widetilde{E}$, so that $A = A_0 + \widehat{E} = A_0 + \widetilde{E} + F$. Note that θ_i is an eigenvalue of $A_0 + \widetilde{E}$. Suppose that $(A_0 + \widetilde{E} + tF)x(t) = \lambda_i(t)x(t)$ for $0 \le t \le 1$, and define $y(t) = x_{k+1:n}(t)$ where $\lambda_i(t)$ is a continuous function of t satisfying $\lambda_i(0) = \theta_i$. Our goal is to bound $|\lambda_i - \theta_i| = |\lambda_i(1) - \lambda_i(0)|$ using (2.3), so we aim to derive a bound for $\|y(t)\|_2$.

To achieve this, we next exploit the structure of the perturbation matrix $A_0+\tilde{E}+tF$ as follows. For any j such that $j \le k$ and $i \neq j$, the jth row of $(A+\widetilde{E}+tF)x(t) = \lambda_i(t)x(t)$ is

$$\theta_j x_j(t) + E_j^* y(t) = \lambda_i(t) x_j(t),$$

hence

$$|x_j(t)| = \frac{|E_j^* y(t)|}{|\theta_j - \lambda_i(t)|} \le \frac{\|E_j\|_2 \|y(t)\|_2}{|\theta_j - \lambda_i(t)|}. \tag{3.6}$$

The last $n-k$ rows of $(A_0 + \widetilde{E} + tF)x(t) = \lambda_i(t)x(t)$ give

$$A_2 y(t) + \sum_{\substack{1\le j\le k\\ j\neq i}} E_j x_j(t) + tE_i x_i(t) = \lambda_i(t) y(t).$$

Hence

$$tE_i x_i(t) = (\lambda_i(t)I - A_2)y(t) - \sum_{\substack{1\le j\le k\\ j\neq i}} E_j x_j(t),$$

so

$$\begin{aligned} t\|E_i\|_2 |x_i(t)| &\ge \|(\lambda_i(t)I - A_2)y(t) - \sum_{\substack{1\le j\le k\\ j\neq i}} E_j x_j(t)\|_2 \\ &\ge \|(\lambda_i(t)I - A_2)y(t)\|_2 - \sum_{\substack{1\le j\le k\\ j\neq i}} \|E_j x_j(t)\|_2 \\ &\ge \|(\lambda_i(t)I - A_2)y(t)\|_2 - \sum_{\substack{1\le j\le k\\ j\neq i}} \frac{\|E_j\|_2^2 \|y(t)\|_2}{|\theta_j - \lambda_i(t)|}, \end{aligned}$$

where we used (3.6) to get the last inequality. Defining $\Gamma_i(t) = \min_j |\lambda_i(t) - \lambda_j(A_2)|$ (note that $\Gamma_i(t)$ is closely related to Gap_i; in particular $\Gamma_i(0) = \mathrm{Gap}_i$) we have $\|(\lambda_i(t)I - A_2)y(t)\|_2 \ge \Gamma_i(t)\|y(t)\|_2$, so it follows that

$$\left(\Gamma_i(t) - \sum_{\substack{1\le j\le k\\ j\neq i}} \frac{\|E_j\|_2^2}{|\theta_j - \lambda_i(t)|}\right) \|y(t)\|_2 \le t\|E_i\|_2 |x_i(t)|. \tag{3.7}$$

Suppose that there exist positive scalars $\delta_i, \delta_{i,j}$ $(i \neq j)$ such that

$$\delta_i \le \min_{0\le t\le 1} \Gamma_i(t), \quad \delta_{i,j} \le \min_{0\le t\le 1} |\theta_j - \lambda_i(t)|.$$

For example, by Weyl's theorem we can take $\delta_i = \Gamma_i(0) - \|E_i\|_2 = \min_j |\theta_i - \lambda_j(A_2)| - \|E_i\|_2$ and $\delta_{i,j} = |\theta_j - \lambda_i(0)| - \|E_i\|_2 = |\theta_j - \theta_i| - \|E_i\|_2$, provided that they are both positive. Thus $\Gamma_i(t) - \sum_{\substack{1\le j\le k\\ j\neq i}} \frac{\|E_j\|_2^2}{|\theta_j - \lambda_i(t)|} \ge \delta_i - \sum_{\substack{1\le j\le k\\ j\neq i}} \frac{\|E_j\|_2^2}{\delta_{i,j}}$ for $0 \le t \le 1$, and if the right-hand side is positive then defining

$$d_i := \frac{1}{\delta_i - \sum_{\substack{1\le j\le k\\ j\neq i}} \frac{\|E_j\|_2^2}{\delta_{i,j}}} \ (>0), \tag{3.8}$$

from (3.7) we have

$$\|y(t)\|_2 \le t d_i \|E_i\|_2 |x_i(t)| \le t d_i \|E_i\|_2,$$

where we used $|x_i(t)| \leq \|x(t)\|_2 = 1$. Thus we have a bound for $\|y(t)\|_2$. Plugging this into (2.3) we obtain

$$\begin{aligned} |\lambda_i(1) - \lambda_i(0)| &\leq 2\|E_i\|_2 \left| \int_0^1 \|y(t)\|_2 dt \right| \\ &\leq 2\|E_i\|_2 \left| \int_0^1 t d_i \|E_i\|_2 dt \right| \\ &\leq d_i \|E_i\|_2^2. \end{aligned}$$

Moreover the bound (3.4) holds for any $\delta_i, \delta_{i,j}$ such that $0 < \delta_i \leq \min_{j, 0 \leq t \leq 1} |\lambda_i(t) - \lambda_j(A_2)|$ and $0 < \delta_{i,j} \leq \min_{0 \leq t \leq 1} |\lambda_i(t) - \theta_j|$. ∎

The fact that θ_i and λ_i have a one-to-one correspondence was not explicitly stated nor necessary in our derivation of Theorem 1, but this can be verified by noting that $\lambda_i(t)$ can be taken as the ith eigenvalue of $A_0 + \tilde{E} + tF$ and by defining an appropriate ordering, e.g. by keeping the non-increasing order. Moreover this question is irrelevant when the residuals are small enough as (for well separated Ritz values) the error intervals become disjoint. The possible presence of multiple eigenvalues does not affect these arguments for the reason discussed at the end of Section 2.

From Theorem 1 we see that our bound is of the form (1.11), as announced in the introduction. In a typical application where the extremal eigenvalues are sought, in the limit $E \to 0$ we have $\theta_i \to \lambda_i$ for $i = 1, \ldots, k$, where λ_i is arranged in appropriate (increasing or decreasing) order. Furthermore $\min(\lambda(A_2)) = \lambda_{k+1}$, hence we have $\mathrm{Gap}_i \to |\lambda_i - \lambda_{k+1}|$, and d_i in (3.4) approaches $1/(\lambda_{k+1} - \lambda_i)$. It follows that (3.4) gives the asymptotic bound

$$\lim_{E \to 0} \frac{|\lambda_i - \theta_i|}{\|E_i\|_2^2} \leq \frac{1}{|\lambda_{k+1} - \lambda_i|}. \tag{3.9}$$

This asymptotic bound is tighter than any of the known bounds; for example the bound in [12] is $\lim_{E \to 0} \frac{|\lambda_i - \theta_i|}{\|E\|_2^2} \leq \frac{1}{|\lambda_{k+1} - \lambda_i|}$ (the left-hand side is smaller than in (3.9)), and the classical bound on a single Ritz value [20, Thm. 11.7.1] gives $\lim_{E \to 0} \frac{|\lambda_i - \theta_i|}{\|E_i\|_2^2} \leq \frac{1}{\min_j |\lambda_j - \lambda_i|}$ (the right-hand side is larger than in (3.9)). The lower bounds $\delta_i, \delta_{i,j}$ of $\min_j |\lambda_i(t) - \lambda_j(A_2)|, \min_{0 \leq t \leq 1} |\lambda_i(t) - \theta_j|$ as defined in (3.1), (3.2) use Weyl's theorem and so are crude bounds. However, as confirmed in our numerical experiments, this does not affect the asymptotic sharpness of the theorem in the limit $E \to 0$, because we still have $d_i \to \frac{1}{\min_j |\theta_i - \lambda_j(A_2)|}$.

Finally, the asymptotic bound (3.9) is sharp, provided that no further information on A_2 is available. In order to prove this we can show that for the specific case where $A_2 = cI$ with $c > \theta_i$ for all $i = 1, \ldots, k$, in the limit $E \longrightarrow 0$ we have

$$|\lambda_i - \theta_i| = \frac{\|E_i\|_2^2}{\mathrm{Gap}_i} + O(\|E_i\|_2^4) \tag{3.10}$$

To see this, take a fixed i and rearrange the columns and rows of $A - \theta_i I$, with A from Theorem 1, such that the (1,1) coefficient is zero. This leads to the matrix

$$B = \left[\begin{array}{c|cccc|c} 0 & 0 & \cdots & & 0 & E_i^T \\ \hline 0 & \theta_1 - \theta_i & & & & E_1^T \\ \vdots & & & \ddots & & \vdots \\ 0 & & & & \theta_k - \theta_i & E_k^T \\ \hline E_i & E_1 & \cdots & & E_k & (c - \theta_i)I \end{array} \right].$$

Define $\tilde{A}_2 = B(2:n, 2:n)$. Then [13, Thm. 3.1] states that the smallest eigenvalue of the above matrix is $|\lambda_i - \theta_i| = \left|\begin{bmatrix}0_{1\times(k-1)} & E_i^T\end{bmatrix} \tilde{A}_2^{-1} \begin{bmatrix}0_{(k-1)\times 1}\\ E_i\end{bmatrix}\right| + O(\|E_i\|_2^4/\mathrm{Gap}_i^2)$. Assuming that the residuals $E_k, k \neq i$, are small enough the leading terms arising from this expression (e.g. by Neumann series) yields Equation (3.10).

3.2. Error bound for a cluster of Ritz values. The bound derived above is sharp in the limit $E \to 0$, but for finite E we might not have $\delta_i - \sum_{\substack{1\le j\le k\\ j\neq i}} \frac{\|E_j\|_2^2}{\delta_{i,j}} > 0$ or $\delta_{i,j} > 0$ for some j, in which case Theorem 1 is not applicable. Assuming E is not too large, one sees that this can happen only if $\delta_{i,j} \lesssim \min_{0\le t\le 1} |\theta_j - \lambda_i(t)| \approx 0$ for some j. This means there is a cluster of Ritz values and θ_i belongs to it. In Theorem 2 we give a bound that is applicable in such situations. Much of the analysis is the same as above.

THEOREM 2. *Let*

$$A = \begin{bmatrix}\widehat{\Lambda} & E^* \\ E & A_2\end{bmatrix}$$

be a symmetric matrix where $\widehat{\Lambda} = diag(\theta_1, \dots, \theta_k)$ *and* $E = [E_1, E_2, \dots, E_k]$, *and suppose that* ℓ *Ritz values* $\theta_i, \dots, \theta_{i+\ell-1}$ *all lie in the interval* $[\lambda_0 - \Delta, \lambda_0 + \Delta]$. *Define* $\mathcal{I} = \{i, i+1, \dots, i+\ell-1\}$ *and let* $E_{\mathcal{I}} = E_{i:i+\ell-1}$ *and suppose also that*

$$\delta_{\mathcal{I}} = \min_j |\lambda_0 - \lambda_j(A_2)| - \Delta - \|E_{\mathcal{I}}\|_2 \ (\approx Gap_{\mathcal{I}} - \Delta - \|E_{\mathcal{I}}\|_2) > 0,$$

$$\delta_{\mathcal{I},j} = |\theta_j - \lambda_0| - \Delta - \|E_{\mathcal{I}}\|_2 > 0, \quad for \quad 1 \le j \le k,\ j \notin \mathcal{I}$$

and further that

$$d_{\mathcal{I}} = \frac{1}{\delta_{\mathcal{I}} - \sum\limits_{\substack{1\le j\le k\\ j\notin\mathcal{I}}} \frac{\|E_j\|_2^2}{\delta_{\mathcal{I},j}}} > 0.$$

Then there exist ℓ *eigenvalues* $\lambda_i, \dots, \lambda_{i+\ell-1}$ *of* A *satisfying*

$$|\lambda_{i+j} - \theta_{i+j}| \le d_{\mathcal{I}} \|E_{\mathcal{I}}\|_2^2, \quad j = 0, 1, 2, \dots, \ell-1. \tag{3.11}$$

Moreover, for $j = 0, 1, \dots, \ell-1$ *,let* $x(t), \lambda_{i+j}(t)$ *be continuous functions of* $t \in [0,1]$ *defined by*

$$\left[\begin{array}{ccc|c} \widehat{\Lambda}_{(1:i-1)} & & & E_{1:i-1}^T \\ & \widehat{\Lambda}_{\mathcal{I}} & & tE_{\mathcal{I}}^T \\ & & \widehat{\Lambda}_{(i+\ell:k)} & E_{i+\ell:k}^T \\ \hline E_{1:i-1} & tE_{\mathcal{I}} & E_{i+\ell:k} & A_2 \end{array}\right] x(t) = \lambda_{i+j}(t) x(t).$$

and $\lambda_{i+j}(0) = \theta_{i+j}$, *where we note* $\widehat{\Lambda}_{(a:b)} = \widehat{\Lambda}(a:b, a:b)$. *Then the bound* (3.11) *holds for any* $\delta_{\mathcal{I}}, \delta_{\mathcal{I},j}$ *such that* $0 < \delta_{\mathcal{I}} \le \min_{i\in\mathcal{I},j} \min_{0\le t\le 1} |\lambda_i(t) - \lambda_j(A_2)|$ *and* $0 < \delta_{\mathcal{I},j} \le \min_{i\in\mathcal{I}} \min_{0\le t\le 1} |\lambda_i(t) - \theta_j|$.

Proof. Suppose that ℓ Ritz values form a cluster and assume without loss of generality that their indices are also clustered in $\mathcal{I} = \{i, i+1, \dots, i+\ell-1\}$, so the cluster consists of $\theta_i, \dots, \theta_{i+\ell-1}$ all lying in the interval $[\lambda_0 - \Delta, \lambda_0 + \Delta]$ (if not, then we can apply a permutation to cluster the indices).

As before define $y(t) = x_{k+1:n}(t)$ and let $E_{\mathcal{I}} = E_{i:i+\ell-1}$.

Let $F = \left[\begin{smallmatrix} 0 & F_{\mathcal{I}}^T \\ F_{\mathcal{I}} & 0 \end{smallmatrix}\right]$, where $F_{\mathcal{I}} = [0_{(n-k)\times(i-1)}\ E_{\mathcal{I}}\ 0_{(n-k)\times(n-i-\ell+1)}]$ has ℓ nonzero columns, and let $\widetilde{E}$ be a matrix such that $E = F + \widetilde{E}$. Let $\hat{i}$ be an arbitrary index in $\mathcal{I} = \{i, i+1, \ldots, i+\ell-1\}$. For any j such that $j \le k$ and $j \notin \mathcal{I}$, the jth row of $(A + \widetilde{E} + tF)x(t) = \lambda_{\hat{i}}(t)x(t)$ is

$$\theta_j x_j(t) + E_j^* y(t) = \lambda_{\hat{i}}(t) x_j(t),$$

hence

$$|x_j(t)| = \frac{E_j^* y(t)}{|\theta_j - \lambda_{\hat{i}}(t)|} \le \frac{\|E_j\|_2 \|y(t)\|_2}{|\theta_j - \lambda_{\hat{i}}(t)|}. \tag{3.12}$$

Now, writing $x_{\mathcal{I}} = x_{i:i+\ell-1}$, from the last $n-k$ rows of $(A + \widetilde{E} + tF)x(t) = \lambda_{\hat{i}}(t)x(t)$ we get

$$tE_{\mathcal{I}} x_{\mathcal{I}}(t) + \sum_{\substack{1 \le j \le k \\ j \notin \mathcal{I}}} E_j x_j(t) + A_2 y(t) = \lambda_{\hat{i}}(t) y(t).$$

Hence

$$tE_{\mathcal{I}} x_{\mathcal{I}}(t) = (\lambda_{\hat{i}}(t) I - A_2) y(t) - \sum_{\substack{1 \le j \le k \\ j \notin \mathcal{I}}} E_j x_j(t),$$

so

$$\begin{aligned} t\|E_{\mathcal{I}}\|_2 \|x_{\mathcal{I}}(t)\|_2 &\ge \|(\lambda_{\hat{i}}(t) I - A_2) y(t) - \sum_{\substack{1 \le j \le k \\ j \notin \mathcal{I}}} E_j x_j(t)\|_2 \\ &\ge \|(\lambda_{\hat{i}}(t) I - A_2) y(t)\|_2 - \sum_{\substack{1 \le j \le k \\ j \notin \mathcal{I}}} \frac{\|E_j\|_2^2 \|y(t)\|_2}{|\theta_j - \lambda_{\hat{i}}(t)|}, \end{aligned}$$

where we used (3.12) to get the last inequality.

Defining $\Gamma_{\mathcal{I}}(t) = \min_{i \in \mathcal{I}, j} |\lambda_i(t) - \lambda_j(A_2)|$ we have $\|(\lambda_{\hat{i}}(t) I - A_2) y(t)\|_2 \ge \Gamma_{\mathcal{I}}(t) \|y(t)\|_2$, so it follows that

$$\left(\Gamma_{\mathcal{I}}(t) - \sum_{\substack{1 \le j \le k \\ j \notin \mathcal{I}}} \frac{\|E_j\|_2^2}{|\theta_j - \lambda_{\hat{i}}(t)|} \right) \|y(t)\|_2 \le t \|E_{\mathcal{I}}\|_2 \|x_{\mathcal{I}}(t)\|_2,$$

Suppose that there exist positive scalars $\delta_{\mathcal{I}}, \delta_{\mathcal{I},j}$ for $j \notin \mathcal{I}$ such that

$$\delta_{\mathcal{I}} \le \min_{0 \le t \le 1} \Gamma_{\mathcal{I}}(t), \quad \delta_{\mathcal{I},j} \le \min_{i \in \mathcal{I}} \min_{0 \le t \le 1} |\theta_j - \lambda_i(t)|.$$

For example, by Weyl's theorem we can take

$$\begin{aligned} \delta_{\mathcal{I}} &= \min_j |\lambda_0 - \lambda_j(A_2)| - \Delta - \|E_{\mathcal{I}}\|_2, \\ \delta_{\mathcal{I},j} &= |\theta_j - \lambda_{\hat{i}}(0)| - \|E_{\mathcal{I}}\|_2 \ge |\theta_j - \lambda_0| - \Delta - \|E_{\mathcal{I}}\|_2. \end{aligned}$$

Thus $\Gamma_{\mathcal{I}}(t) - \sum_{\substack{1 \le j \le k \\ j \notin \mathcal{I}}} \frac{\|E_j\|_2^2}{|\theta_j - \lambda_{\hat{i}}(t)|} \ge \delta_{\mathcal{I}} - \sum_{\substack{1 \le j \le k \\ j \notin \mathcal{I}}} \frac{\|E_j\|_2^2}{\delta_{\mathcal{I},j}}$ for $0 \le t \le 1$, and if the right-hand side is positive then defining

$$d_{\mathcal{I}} := \frac{1}{\delta_{\mathcal{I}} - \sum_{\substack{1 \le j \le k \\ j \notin \mathcal{I}}} \frac{\|E_j\|_2^2}{\delta_{\mathcal{I},j}}} \ (> 0), \tag{3.13}$$

we have

$$\|y(t)\|_2 \le td_{\mathcal{I}}\|E_{\mathcal{I}}\|_2\|x_{\mathcal{I}}\|_2 \le td_{\mathcal{I}}\|E_{\mathcal{I}}\|_2.$$

Plugging this into (2.2) yields

$$\begin{aligned}
|\lambda_{\hat{i}}(1) - \theta_{\hat{i}}| &\le 2\left|\int_0^1 x_{1:k}(t)^T E_{\mathcal{I}} y(t)dt\right| \\
&\le 2\|E_{\mathcal{I}}\|_2 \left|\int_0^1 t\|y(t)\|_2 dt\right| \qquad (\|x_{1:k}(t)\|_2 \le 1) \\
&\le 2\|E_{\mathcal{I}}\|_2 \left|\int_0^1 td_{\mathcal{I}}\|E_{\mathcal{I}}\|_2 dt\right| \\
&\le d_{\mathcal{I}}\|E_{\mathcal{I}}\|_2^2.
\end{aligned}$$

The same argument holds for any $\hat{i} \in \mathcal{I}$, which proves the bound (3.4). Moreover, the bound (3.4) holds for any $\delta_{\mathcal{I}}, \delta_{\mathcal{I},j}$ such that $0 < \delta_{\mathcal{I}} \le \min_{0\le t\le 1}|\lambda_i(t) - \lambda_j(A_2)|$ and $0 < \delta_{\mathcal{I},j} \le \min_{i\in\mathcal{I}}\min_{0\le t\le 1}|\lambda_i(t) - \theta_j|$, where $\lambda_i(t), \lambda_{i+1}(t), \ldots, \lambda_{i+\ell-1}(t)$ are continuous functions of t such that $\lambda_{i+j}(t)$ is an eigenvalue of $A + \widetilde{E} + tF$ with $\lambda_{i+j}(0) = \theta_{i+j}$ for $j = 0, 1, \ldots, \ell - 1$. ∎

Conceptually, in the definition of $\delta_{\mathcal{I}} = \min_j |\lambda_0 - \lambda_j(A_2)| - \Delta - \|E_{\mathcal{I}}\|_2$, the term $\min_j |\lambda_0 - \lambda_j(A_2)|$ corresponds to the big $\mathrm{Gap}_{\mathcal{I}}$ in previous discussions.

4. Error bound for singular values.

In this section, recalling the setup of Equation (1.2), we generalize Theorem 1 to the singular values of arbitrary matrices of size m by n. To do this, we use the Jordan-Wielandt theorem from [23, Thm. I.4.2] [6, Thm. 7.3.3], which allows us to apply the bound from the symmetric case to obtain results on the singular values. This theorem states that if M is an m by n matrix with $m > n$ and singular values $\sigma_1 \ge \sigma_2 \ge \ldots \ge \sigma_n$, then the eigenvalues of the symmetric matrix $\begin{bmatrix} 0 & M \\ M^T & 0 \end{bmatrix}$ are $\sigma_1 \ge \sigma_2 \ge \ldots \ge \sigma_n \ge 0 \ge -\sigma_n \ge -\sigma_2 \ge \ldots \ge -\sigma_1$, where the eigenvalue 0 has multiplicity $m - n$. This classical result was also used in [11, 12] for a similar purpose.

This leads to the following theorem.

Theorem 3. *Let*

$$A = \begin{bmatrix} \widehat{\Sigma} & E^T \\ F & A_2 \end{bmatrix}$$

be an m-by-n matrix ($m \ne n$) where $\widehat{\Sigma} = diag(\theta_1, \ldots, \theta_k)$, $E = [E_1, E_2, \ldots, E_k]$ and $F = [F_1, F_2, \ldots, F_k]$. For $i = 1, 2, \ldots, k$, suppose that

$$\begin{aligned}
\delta_i &= \min(|\theta_i|, \min_j |\theta_i - \sigma_j(A_2)|) - \sqrt{(\|E_i\|_2^2 + \|F_i\|_2^2)/2} > 0 \\
\delta_{i,j} &= |\theta_i - \theta_j| - \sqrt{(\|E_i\|_2^2 + \|F_i\|_2^2)/2} > 0, \quad \text{for} \quad 1 \le j \le k,\ j \ne i \\
\delta'_{i,j} &= |\theta_i + \theta_j| - \sqrt{(\|E_i\|_2^2 + \|F_i\|_2^2)/2} > 0, \quad \text{for} \quad 1 \le j \le k
\end{aligned}$$

and that

$$d_i = \frac{1}{\delta_i - \sum_{\substack{1\le j\le k \\ j\ne i}} \frac{\|E_j\|_2^2 + \|F_j\|_2^2}{2\delta_{i,j}} - \sum_{1\le j\le k} \frac{\|E_j\|_2^2 + \|F_j\|_2^2}{2\delta'_{i,j}}} > 0.$$

Then there exists a singular value σ_i of A such that

$$|\sigma_i - \theta_i| \leq d_i \frac{\|E_i\|_2^2 + \|F_i\|_2^2}{2}.$$

Proof. We consider the following matrix whose eigenvalues are $\{\sigma_i, -\sigma_i\}_{i=1,\ldots,n}$ and 0, according to the Jordan-Wielandt theorem:

$$\begin{bmatrix} & & \widehat{\Sigma} & E^T \\ & & F & A_2 \\ \widehat{\Sigma} & F^T & & \\ E & A_2^T & & \end{bmatrix}.$$

As in [11, 12], by permuting the second and third block columns then the second and third block rows we obtain

$$\tilde{A} = \left[\begin{array}{cc|cc} & \widehat{\Sigma} & & E^T \\ \widehat{\Sigma} & & F^T & \\ \hline & F & & A_2 \\ E & & A_2^T & \end{array}\right] =: \begin{bmatrix} \tilde{\Sigma} & E_{tot}^T \\ E_{tot} & \tilde{A}_2 \end{bmatrix}.$$

Finally, using the unitary matrices $M = \frac{1}{\sqrt{2}} \begin{bmatrix} I_k & -I_k \\ I_k & I_k \end{bmatrix}$ and $P = \begin{bmatrix} M & 0 \\ 0 & I_{m+n-2k} \end{bmatrix}$ we have

$$P\tilde{A}P^T = \left[\begin{array}{cc|cc} -\widehat{\Sigma} & & \frac{1}{\sqrt{2}}F^T & \frac{1}{\sqrt{2}}E^T \\ & \widehat{\Sigma} & \frac{1}{\sqrt{2}}F^T & -\frac{1}{\sqrt{2}}E^T \\ \hline \frac{1}{\sqrt{2}}F & \frac{1}{\sqrt{2}}F & & A_2 \\ \frac{1}{\sqrt{2}}E & -\frac{1}{\sqrt{2}}E & A_2^T & \end{array}\right]$$

which is symmetric and of the form $B = \begin{bmatrix} \tilde{\Lambda} & \tilde{E}^T \\ \tilde{E} & \tilde{A}_2 \end{bmatrix}$, with $\tilde{E} = E_{tot}M^T = \frac{1}{\sqrt{2}} \begin{bmatrix} F & F \\ E & -E \end{bmatrix}$.

Then the Theorem 1 is directly applicable to B, which proves Theorem 3. The slight change in the coefficient δ_i (compared to Theorem 1) comes from the fact that in addition to the $\pm\sigma_j(A_2)$ the matrix $\tilde{A}_2$ also has the eigenvalue 0, and the additional terms $\delta'_{i,j}$ appear from the block $-\widehat{\Sigma}$ from $\widetilde{\Lambda}$. ◻

To compare our bound with Equation (1.6) from [12], notice that the numerator in Equation (1.6) is always larger than the one from our bound since $\max\{\|E\|_2, \|F\|_2\}^2 \geq (\|E\|_2^2 + \|F\|_2^2)/2 \geq (\|E_i\|_2^2 + \|F_i\|_2^2)/2$ for $i = 1, \ldots, k$, while the denominator tend to Gap_i in both bounds, which shows that asymptotically our bound is sharper, as announced in the introduction.

The bound from [11] with the structure considered in our Theorem 3 gives Equation (1.7) which is asymptotically looser than Equation (1.6) since we can assume $\min_k |\sigma_i - \sigma_k(A_2)| \approx \mathrm{Gap}_i = \min_j |\theta_i - \sigma_j(A_2)|$ as $E \to 0$. However we believe that [11] would provide a better tool to analyze the error when the algorithm used for the approximate SVD has perturbation terms also in the diagonal blocks. This was shown to be the case for the generalized Nyström approach in [11, Sec.3].

In practice the residuals E and F may differ significantly (e.g. $E = 0$ in HMT). However, our numerical experiments suggest that this does not affect the quality of the bound very much (see Subsection 5.2). Additionally, we derived another bound specifically for the HMT perturbation structure where $E = 0$, by starting from a singular value perturbation result similar to Lemma 2.1 and bounding the components of the singular vectors. We obtained a bound that is theoretically slightly looser than Theorem 3, so we omit this extension.

5. Numerical experiments.

In this section we present several experiments to illustrate the validity and the sharpness of our theoretical bounds compared to previous existing bounds. We first test our error bounds for the eigenvalues of a (synthetic) symmetric matrix, then for singular values of arbitrary matrices. The latter subsection also discusses the application of our bound to the singular values obtained from the randomized SVD algorithm by Halko, Martinson and Tropp [5], which is so far the most widely used randomized method for low-rank approximation of large scale matrices.

5.1. Numerical experiments on the eigenvalues of symmetric matrices.

We generate a symmetric matrix $A = VDV^T$ of size $n = 2000$ where $V \in \mathbb{R}^{n\times n}$ is orthogonal and D is a diagonal matrix of eigenvalues. Suppose that we are looking for $k = 100$ eigenpairs, and we run 40 steps of LOBPCG [8] (for *Locally Optimal Block Preconditioned Conjugate Gradient* method) without preconditioning[2] to obtain k sets of Ritz pairs $(\theta_i, \widehat{x}_i)$ for $i = 1, \dots, k$. By default LOBPCG approximates the smallest eigenvalues. We next evaluate the errors in the Ritz values $|\theta_i - \lambda_i|$ and compare their bounds both in the context of well separated eigenvalues and for a matrix D that has a cluster of eigenvalues. For the initial guess X_0 we used a randomly generated $n \times k$ Haar distributed matrix using MATLAB's `orth(randn(n,k))`.

Well-separated Ritz values. We let the eigenvalues be uniformly distributed with $D = \mathrm{diag}(1, 2, \dots, n)$. To illustrate the sharpness of our theoretical bound we first use the exact value of Gap_i, that is, taking (the usually unknown quantity) $\lambda_j(A_2)$ as known. The resulting bound is shown as "Thm. 1 exact". However in practice this distance is usually unknown, so to invoke Theorem 1 using only information that is usually available, we use the estimate $\mathrm{Gap}_i \geq \max_k |\theta_i - \theta_k|$, hence we let $\delta_i = \max_k |\theta_i - \theta_k| - \|E_i\|_2$ instead of (3.2). This assumption is valid when the eigensolver approximates the extreme eigenvalues and the approximation is high quality. The resulting bound corresponds to "Thm. 1 approximate" in Figure 1. We compare our bound to the exact error $|\theta_i - \lambda_i|$, the bound (1.5) from [12] (shown as "Li-Li (large gap)"), the bound (1.10) which is (1.5) applied to the partitioning 1.9 ("Li-Li (1-1 block)"), the "classical" bound (1.8), and $\|E_i\|_2$ which is the crude bound using Weyl's theorem. Notice from Weyl's bound that $\|E_1\|_2 \lesssim \|E_2\|_2 \lesssim \dots \lesssim \|E_k\|_2$ with $\|E_1\|_2 \ll \|E_k\|_2$, which is a typical graded behavior as mentioned in the introduction.

In Figure 1 we observe that Theorem 1 gives the sharpest error bound for the Ritz values (whether or not $\lambda(A_2)$ is known or estimated), especially for small i. The bounds Classical and Li-Li (1-1 block) are also of good quality and are nearly identical in this experiment. The latter and our theorem reflect the $\mathcal{O}(\|E_i\|_2^2)$ trend of the actual errors, unlike Weyl's theorem. For interior eigenvalues, our Theorem 1 gives results close to Li-Li and Classical (because the spectral gaps in all three bounds are about the same); but as we move to extreme eigenvalues, the term Gap_i in the denominator in Theorem 1 increases (and the residual $\|E_i\|_2$ decreases), which leads to significantly higher accuracy (here, about a factor 100 improvement). This illustrates the key idea that our bounds are asymptotically sharp as $E \to 0$.

The bound "Li-Li (large gap)" mainly serves to illustrate that, if the residual in

---

[2] An effective preconditioner can speed up the convergence of any eigenvalue, but it is usually still true that the extremal Ritz values converge first.

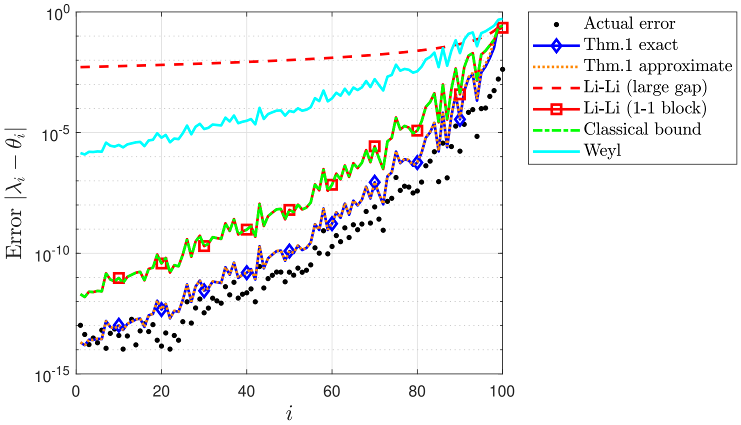

Fig. 1: Error in Ritz values $|\theta_i - \lambda_i|$ and error bounds for uniformly distributed eigenvalues: $\lambda_i = i$, $\forall i \in [1, n]$.

the numerator is $\|E\|_2$ the accuracy of the bound is significantly lower, even if the gap in the denominator is Gap_i. This is intuitive when the residuals are graded, like in this example.

Note that the two versions of our bounds (exact and approximate) are nearly identical, resulting in the two bounds being nearly superimposed in Figure 1 (as well as in the other experiments below). In color print the exact bound (blue line) and the approximate bound (orange dots) are distinguishable but in black-and-white they might appear as the same line. Hence we argue that Theorem 1 can give tight bounds using information that is usually available in practice. By contrast we note that the Classical result (1.8) uses a bound on the gap between θ_i and the eigenvalues of an $(n-1) \times (n-1)$ matrix, so it is generally not a practical bound to use (although reasonable estimates of the gap can often be obtained).

Another note of caution when using the bound $\|E_i\|_2$ (and possibly Classical) is that the Ritz values and eigenvalues of A may not have a one-to-one correspondence, as noted in [20, Sec. 11.5]. Our results and Li-Li overcome this difficulty as it explicitly specifies a one-to-one correspondence between θ_i and λ_i.

Clustered Ritz values. To illustrate the effectiveness of Theorem 2 in the presence of a clustered eigenvalue, we run the same experiment as above but now we replace the 20th to 29th eigenvalues of A by $20 + \texttt{randn}(1, 10) * \texttt{1e} - 10$ so that there are 10 eigenvalues clustered around 20. Since the clustered Ritz values lied in 20 ± 10^{-10}, we set $\mathcal{I} = \{20, \ldots, 29\}$ with $\lambda_0 = 20, \Delta = 10^{-10}$ when invoking Theorem 2 to bound the errors in θ_i for $i = 20, \ldots, 29$, and for the rest we let $\mathcal{I} = \{i\}$ with $\delta = 0$, which reduces to Theorem 1. We also performed the same experiment where 10 eigenvalues of A are clustered around $\lambda_0 = 100$ with $\Delta = 10^{-10}$, such that the Ritz values near $i = k$, ie. near $\lambda(A_2)$, are clustered. For both values of λ_0, in "Li-Li (1-1 block)" we used a partitioning similar to (1.9) that puts the clustered Ritz values in the 1-1 block and gives the residual $\|E_{\mathcal{I}}\|_2^2$ in the numerator of the bound. The results are presented in Figure 2.

For the uniformly distributed eigenvalues, the same comments as in the previous experiment apply, therefore we focus our attention on the clustered eigenvalues. In

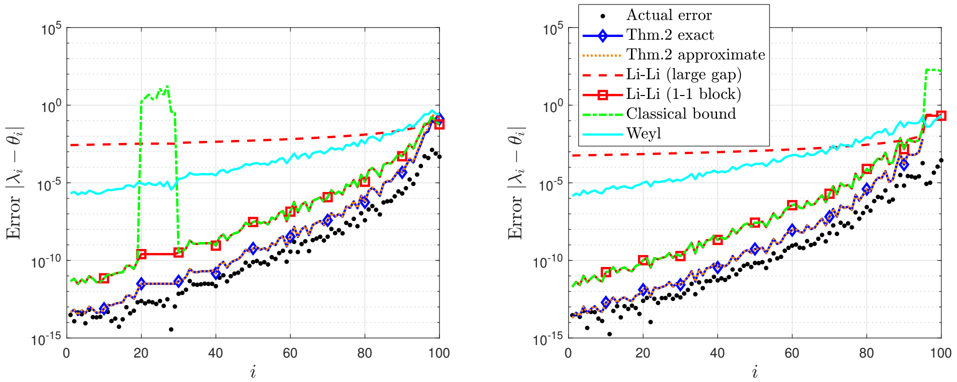

(a) Cluster: $\lambda_i = 20 + 10^{-10} \times \mathtt{randn(1)}, \forall \mathtt{i} \in [\mathtt{20}, \mathtt{29}]$.

(b) Cluster: $\lambda_i = 100 + 10^{-10} \times \mathtt{randn(1)}, \forall \mathtt{i} \in [\mathtt{96}, \mathtt{105}]$. Note that our bounds are inapplicable for $i \geq 96$, as the assumptions $\delta_{\mathcal{I}} > 0$ and $d_{\mathcal{I}} > 0$ do not hold.

Fig. 2: Error in Ritz values $|\theta_i - \lambda_i|$ and error bounds for uniformly distributed eigenvalues ($\lambda_i = i$) and a cluster of 10 eigenvalues at $\lambda_0 = 20$ (left) and $\lambda_0 = 100$ (right).

both experiments the classical bound is highly inaccurate to bound the errors on the cluster (because $\widetilde{\mathrm{gap}}_i$ in the denominator of Equation (1.8) is approximately 10^{-10}). Remarkably, Li-Li (1-1 block) gives accurate results both when $\lambda_0 = 20$ and $\lambda_0 = 100$. Our Theorem 2 is more accurate than Li-Li when the cluster is well inside the set of approximated eigenvalues (see Figure 2a), but it is inapplicable in Figure 2b, where the Ritz values cluster together with some eigenvalues of A_2. On the other hand, the bound from [12] is still applicable, so in this situation it may be necessary to use (1.10) instead of our Theorem 2.

In practice, the main assumption that we need to check in order to decide if our bounds apply is $\|E_i\|_2^2 < \mathrm{Gap}_i^2, \forall i = 1, \ldots, k$. The small gaps $|\theta_i - \theta_j|$ between neighboring Ritz values are less of a limitation since we can decide to consider certain Ritz values as clustered to ensure that $\delta_{i,j} > 0, \forall j \neq i$, as long as no cluster is not confounded with a part of $\sigma(A_2)$. Therefore, for fixed residual norms, our bounds are particularly efficient when the distribution of the eigenvalues of A (especially for the eigenvalues that we seek to approximate) is steep. Nonetheless, when E is sufficiently small relative to the gaps, our bound will give $c\dfrac{\|E_i\|_2^2}{\mathrm{Gap}_i}$ with $c \approx 1$, which is tight as shown in (3.10).

Numerical experiments with the Lanczos algorithm. In this experiment, we apply the Lanczos algorithm [10, 19] (with re-orthogonalization) instead of LOBPCG to obtain a trial matrix of eigenvectors Q_1. By construction of Q_1 from the Lanczos iteration we have

$$AQ_1 = Q_1 T_k + q_{k+1}[0, \ldots, 0, t_{k+1,k}]$$

with $T_k \in \mathbb{R}^{k \times k}$ a tridiagonal matrix, $q_{k+1} \in \mathbb{R}^n$ a vector which is orthogonal to the columns of Q_1 and with $t_{k+1,k} = \|v - v^T Q_1(:,k) - v^T Q_1(:,k-1)\|_2$ where $v = AQ_1(:,k)$.

Then noting $Q_\perp$ the complement of $\begin{bmatrix} Q_1 & q_{k+1}\end{bmatrix}$ we have

$$\bar{A} = \begin{bmatrix} Q_1 & q_{k+1} & Q_\perp \end{bmatrix}^T A \begin{bmatrix} Q_1 & q_{k+1} & Q_\perp \end{bmatrix} = \left[\begin{array}{ccc|cc} & & & 0 & 0 \\ & & & \vdots & \vdots \\ & T_k & & \vdots & \vdots \\ & & & 0 & \vdots \\ & & & t_{k+1,k} & 0 \\ \hline 0 \cdots\cdots 0 & & t_{k+1,k} & \multicolumn{2}{c}{\multirow{2}{*}{$*$}} \\ 0 \cdots\cdots\cdots\cdots 0 & & & & \end{array}\right]$$

and using the eigendecomposition of the tridiagonal block $T_k = U\widehat{\Lambda}U^T$ we obtain the perturbation matrix

$$\begin{bmatrix} U^T & \\ & I \end{bmatrix} \bar{A} \begin{bmatrix} U & \\ & I \end{bmatrix} = \left[\begin{array}{c|cc} \hat{\Lambda} & e & 0 \\ \hline e^T & \multicolumn{2}{c}{\multirow{2}{*}{A_2}} \\ 0 & & \end{array}\right]$$

with $e = U^T[0,\dots,0,t_{k+1,k}]^T = t_{k+1,k}U^T(:,k)$. Therefore we have a particular structure where each residual vector E_i only has one nonzero component, namely their first component.

As in LOBPCG, with the Lanczos algorithm usually the extreme eigenvalues converge faster than the interior eigenvalues. Note that $t_{k+1,k}$ can be relatively large, leading to large residuals $\|E_i\|$ for some i, therefore we have to take a relatively large number of Lanczos iterations to ensure that a sufficient number of Ritz values have small residuals.

In our experiment, we take $n = 2000$ and we use Lanczos with full reorthogonalization to compute the trial subspace Q_1. We set the size of the Krylov subspace to $k = 400$. Some θ_i have relatively large residuals therefore, to test our bound, we rearranged the structured matrix such that only the 20 smallest approximated eigenvalues are in the $(1,1)$ diagonal block and the other ones are included in the lower diagonal block. Note that adding the larger Ritz values in the block (2,2) should not change the gap Gap_i, as those eigenvalues are far from the ones approximated in the block (1,1). The results are shown in Figure 3. Again our bounds are the sharpest for the smallest eigenvalues as $\|E_i\|_2 \to 0$, and the approximate bound is close to the exact one.

5.2. Numerical experiments for singular values obtained via the Rayleigh-Ritz and HMT method. We now illustrate the quality of the error bound derived in Section 4 for singular values of a m-by-n matrix. We defined a matrix A of size $m = 5000$ by $n = 1000$ with geometrically decaying singular values using MATLAB's command `gallery('randsvd',[m,n],1e20)`. To find approximate singular subspaces $Q_1 \in \mathbb{R}^{m\times k}, Q_2 \in \mathbb{R}^{n\times k}$ we used both the simple and the double power iteration. For instance for a single power iteration, we start from random normal matrices Ω_1, Ω_2 and take $Q_1 = \mathrm{orth}(A\Omega_1)$ and $Q_2 = \mathrm{orth}(\Omega_2^T A)$. For a double power iteration, instead of applying A one applies AA^TA. We approximated the $k = 200$ largest singular values with several methods.

To apply the error bounds from Theorem 3 for the PG method we used the structure (1.2) detailed in Section 1. We also applied this theorem to the case where the HMT algorithm is used to approximate the singular values. Let us derive the perturbation matrix associated to HMT. The main difference from PG in HMT is that only the left trial subspace Q_1 of A is used (i.e., one can think $Q_2 = I_n$ in

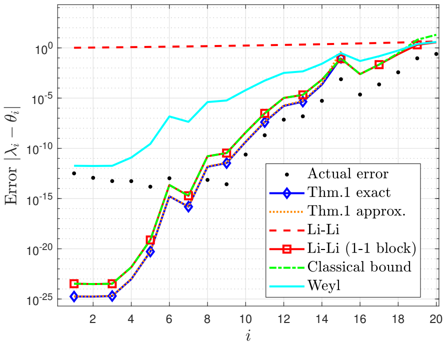

Fig. 3: Error $|\theta_i - \lambda_i|$ and bounds for approximate eigenvalues obtained with the Lanczos algorithm. The eigenvalues of A that were approximated here are $\lambda_i \in [1, 20]$. Some data points are above the bounds but only because of the limited machine precision used in the experiment.

HMT). We kept the same trial subspace to plot the RR and the HMT results (in left and right plots of Figures 4 and 5) to make sure that the results are comparable. Then let $Q_1 Q_1^T A$ be the HMT approximation of A. Taking the full SVD of the matrix $Q_1^T A = U_0 \begin{bmatrix}\Sigma_0 & 0\end{bmatrix} \begin{bmatrix} V_0^T \\ V_{0\perp}^T \end{bmatrix}$ and noting $Q_{tot} = \begin{bmatrix} Q_1 & Q_{1\perp}\end{bmatrix}$, we see that the perturbation matrix underlying the HMT approximation can be written as

$$\begin{bmatrix} U_0^T & 0 \\ 0 & I_{m-r} \end{bmatrix} Q_{tot}^T A \begin{bmatrix} V_0 & V_{0\perp} \end{bmatrix} = \begin{bmatrix} \Sigma_0 & 0 \\ Q_{1\perp}^T A V_0 & Q_{1\perp}^T A V_{0\perp} \end{bmatrix}$$

which is a particular case of the structured matrix from (3), with one of the error blocks being zero. Both in the PG method and HMT, the residual norms can be computed with the available information.

For both approximation methods, we compare our bounds (exact and approximate, as in the symmetric case) to the bound (1.6) given in [12], to (1.7) from [11] and to Weyl's theorem [6, Cor. 7.3.5] [22, Cor. I.4.31]. The results are shown in Figures 4 and 5.

Looking at the results with both single and double power iteration allows us to compare the effect of a graded and ungraded residual structure on the bounds. With single power iteration we observe that the residuals are all around 10^{-4}. Therefore the bound Li-Li (large gap) is almost as accurate as Li-Li (1-1 block), because $\|E_i\|_2 \approx \|E\|_2$, $\forall i$. In this situation the trend of the bounds is controlled by the spectral gap in the denominator.

With the double power iteration, the residuals have an exponential decay as we move towards the extreme eigenvalues, so naturally the bound Li-Li (1-1 block) and our bounds perform much better than the bounds based on $\|E\|_2$. Once again our bounds are sharper than Li-Li (1-1 block) by a factor of about 100.

Hence both for single and double power iterations, our bound is the sharpest, to the point where it even has the same fluctuations as the actual error as the index i is varied. Moreover the approximation $\mathrm{Gap}_i = |\theta_i - \theta_k|$ leads to a bound that is very close to the theoretical result from Theorem 3 and uses only available information, so this bound is computable in practice.

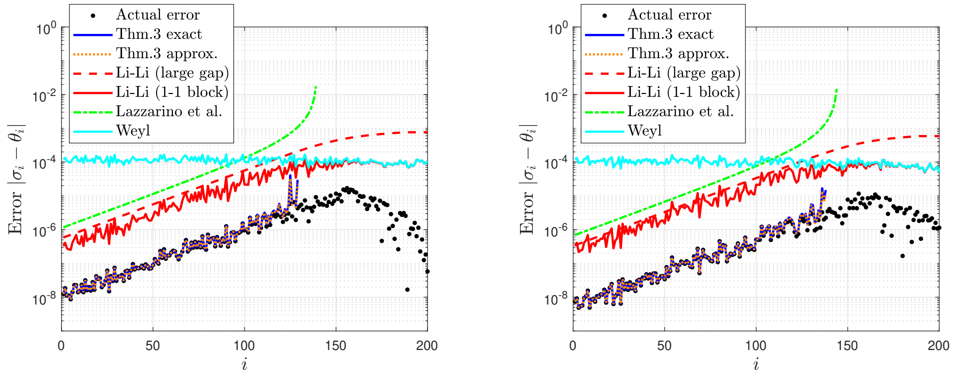

(a) Petrov-Galerkin approximation (b) Randomized SVD

Fig. 4: Error $|\sigma_i - \theta_i|$ and bounds for geometrically distributed singular values where the trial subspaces were found with a single power iteration. Left: estimation with Petrov-Galerkin approximation ; right: estimation with randomized SVD.

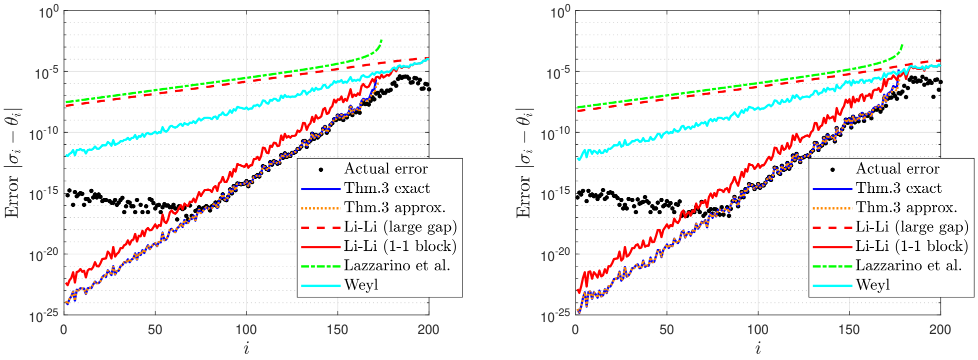

(a) Petrov-Galerkin approximation (b) Randomized SVD

Fig. 5: Error $|\sigma_i - \theta_i|$ and bounds for geometrically distributed singular values where the trial subspaces were found with double power iteration. Left: estimation with Petrov-Galerkin approximation ; right: estimation with randomized SVD. In some cases the bounds lie below the actual error; this is due to roundoff errors (which our bounds do not account for). Indeed, the effect of round-off errors was not accounted for in the perturbation matrix $\bar{A}$ (from (1.2)) considered in our analysis. The operations (e.g. orthogonal multiplication) involved in both methods above are backward stable, so in finite precision arithmetic, we can consider that σ_i is actually the exact singular value of $\bar{A} + E_u$ with $\|E_u\|_2 = \mathcal{O}(u\|A\|_2)$. Therefore, Weyl's theorem implies that the contribution of round-off errors in $|\sigma_i - \theta_i|$ is fortunately bounded by $\mathcal{O}(u\|A\|_2)$. That is, even in finite precision arithmetic, the bounds can be trusted up to working precision.

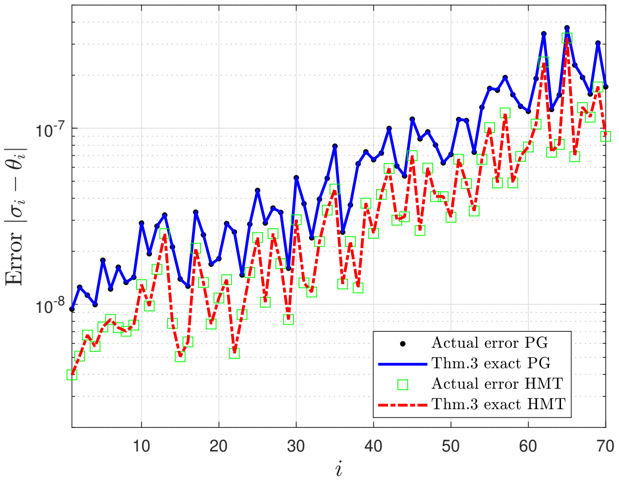

Fig. 6: Error $|\theta_i - \sigma_i|$ and theoretical bound Theorem 3 for the Petrov-Galerkin and randomized SVD methods, computed based on the same left trial subspace Q_1 (with Q_1 obtained from a single power iteration).

Also as mentioned for the asymptotic case at the end of Section 4, the bound (1.6) (which is Li-Li (large gap)) is sharper than (1.7). However, one must keep in mind that our bounds use stronger assumptions and information: while the bounds (1.6) and (1.7) only use the spectral norm of the full error blocks E, F, we assume that all the norms of each column ($\|E_i\|_2, \|F_i\|_2$) are known. We repeat that these are indeed typically available or computable.

Finally, we briefly compare the results from HMT and PG. In theory, for the same left trial subspace Q_1, HMT gives more accurate singular values than PG, as HMT is based on an orthogonal projection. Experimentally we obtained nearly the same values of the residual norms $\|F_i\|_2$ with both methods (especially for small i), so our Theorem 3 suggests that the error bound for HMT is better than the bound for PG by about a factor 2. This is confirmed in Figure 6. In all cases our bounds reflect the actual errors, and the PG estimates give accuracy comparable to that of HMT.

6. Discussion and future directions.

One might wonder if the graded structure of the residuals could be exploited further in the derivation of our bounds. We discuss this below by looking at two cases.

To analyze the behavior of our bound when the residuals have a graded structure we look closer at the denominator of our bound, especially at the term $s_i := \sum_{\substack{1\leq j\leq k\\ j\neq i}} \|E_i\|_2^2/\delta_{i,j}$. We discuss the symmetric case for simplicity but the ideas also apply to the singular value case. The idea is to examine which conditions s_i is small, as it corresponds to getting close to the asymptotic case. Take i such that $1 < i < k$ so that θ_i is somewhere in the middle of the approximated eigenvalues. Then in terms of s_i we consider a j that is far from i and distinguish roughly two cases :

(i) if $j > i$, then $|\theta_i - \theta_j|$ is large and $\|E_j\|_2 \gg \|E_i\|_2$, therefore the fact that the denominator $|\theta_i - \theta_j|$ is large compensates for the large residual ;

(ii) if $j < i$, then $|\theta_i - \theta_j|$ is large and $\|E_j\|_2 \ll \|E_i\|_2$, which is the best case as it gives a small contribution to s_i.

Therefore for small i most j are larger so we are in the less favorable situation (i), but then it is compensated by the fact that the numerator $\|E_i\|_2^2$ of the bound is small. Conversely as $i \to k$ the numerator gets larger but the terms in s_i become smaller so d_i decreases.

On the other hand, taking the θ_i to be uniformly distributed and the structure not graded we can assume $\|E_i\|_2 \approx \epsilon = \mathit{constant}, \forall i$, which leads to $s_i \approx \epsilon^2 \sum_{j \neq i}(|\theta_i - \theta_j| - \epsilon)^{-1}$ being almost independent of i. This implies that the bound is virtually only determined by $\mathrm{Gap}_i = \min_j |\theta_i - \lambda_j(A_2)|$ and the residual $\|E_i\|_2^2$ (weak dependence on other terms $j \neq i$).

One way of looking at this is that our bounds have enough *parameters* to account for a more structured error matrix than the previous bounds from the literature. Note that it also takes into account the other Ritz values θ_j, although their influence is quite limited, so all the available information is used in the bound.

One might suspect that there is room for improvement in finding a more precise definition of δ_i, δ_{ij} in order to make the bound sharper. However in practice the gap Gap_i is the dominant term in the denominator of our bound: in our experiments from Subsection 5.1 $\mathrm{Gap}_i = O(1)$ while $\|E_i\|_2 \approx 10^{-5}$ and $s_i \lesssim 10^{-6}$. Removing these terms and keeping only Gap_i in the denominator kept our bound valid and did not change them much.

Hence our theoretical work implies that, if the gap Gap_i is known to be very large compared to the residual norms, one could even use the bound (1.11) with $c = 1$ as an approximation to the error. Our result echoes the improved bounds for the accuracy of Ritz vectors that was derived in [16], as it also improves a classical bound by showing that a bigger gap in the denominator can govern the perturbation of eigenvectors.

We believe that the most interesting extension of this work would be to derive a similar bound that is still applicable when the error matrix also has diagonal terms, or to find a trick that would put the error matrix in the right form in order to apply our bound. This could notably allow us to derive sharp error bounds for the approximate singular values computed from the (Generalized) Nyström method, which is often seen to be more accurate than Petrov-Galerkin and the randomized SVD methods. A somewhat related problem is to derive error bounds for eigenvalues computed by the *sketched* Rayleigh-Ritz method [18], which is also not based on orthogonal projection. Another idea could be finding block-wise bounds inspired by our Thm.1 that would be adapted to specific eigensolvers. One can notably think of Krylov methods with restarting (see for example the recent bounds from [25]), which could lead to a more complex structured perturbation matrix, but this idea is out of the scope of this paper.

Acknowledgment. We would like to thank the reviewers for their perceptive comments and suggestions, Andrew Knyazev for bringing several papers to our attention, and Ren-Cang Li for a discussion about how the results compare with [12]. This project was supported by the EPSRC grants EP/Y010086/1 and EP/Y030990/1. For the purpose of open access, the author has applied a CC BY public copyright licence to any author accepted manuscript arising from this submission.

References.

[1] Z. Bai, J. Demmel, J. Dongarra, A. Ruhe, and H. van der Vorst, *Templates for the Solution of Algebraic Eigenvalue Problems: A Practical Guide*, SIAM, Philadelphia, 2000.

[2] K. L. Clarkson and D. P. Woodruff, *Numerical linear algebra in the streaming model*, in Proceedings of the ACM Symposium on Theory of Com-

puting (STOC), 2009, pp. 205–214.

[3] C. Davis and W. M. Kahan, *The rotation of eigenvectors by a perturbation. III*, SIAM J. Numer. Anal., 7 (1970), pp. 1–46.

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

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

[6] R. A. Horn and C. R. Johnson, *Matrix Analysis*, Cambridge University Press, second, corrected reprint ed., 2013.

[7] Z. Jia and G. W. Stewart, *An analysis of the Rayleigh-Ritz method for approximating eigenspaces*, Math. Comput., 70 (2001), pp. 637–647.

[8] A. V. Knyazev, *lobpcg.m at MATLAB Central File Exchange.* Available at http://www.mathworks.com/matlabcentral/fileexchange/48-lobpcg-m. Accessed: February 2025.

[9] A. V. Knyazev, *Toward the optimal preconditioned eigensolver: Locally Optimal Block Preconditioned Conjugate Gradient method*, SIAM J. Sci. Comp, 23 (2001), pp. 517–541.

[10] C. Lanczos, *An iteration method for the solution of the eigenvalue problem of linear differential and integral operators*, J. Res. Natl. Bur. Stand. B, 45 (1950), pp. 255–282.

[11] L. Lazzarino, H. A. Daas, and Y. Nakatsukasa, *Matrix perturbation analysis of methods for extracting singular values from approximate singular subspaces*, arXiv preprint arXiv:2409.09187, (2024), https://arxiv.org/abs/2409.09187.

[12] C.-K. Li and R.-C. Li, *A note on eigenvalues of perturbed Hermitian matrices*, Linear Algebra Appl., 395 (2005), pp. 183–190.

[13] R.-C. Li, Y. Nakatsukasa, N. Truhar, and W. Wang, *Perturbation of multiple eigenvalues of Hermitian matrices*, Linear Algebra Appl., 437 (2012), pp. 202–213.

[14] R. Mathias, *Quadratic residual bounds for the Hermitian eigenvalue problem*, SIAM J. Matrix Anal. Appl, 19 (1998), pp. 541–550.

[15] Y. Nakatsukasa, *Eigenvalue perturbation bounds for Hermitian block tridiagonal matrices*, Appl. Numer. Math., 62 (2012), pp. 67–78.

[16] Y. Nakatsukasa, *Sharp error bounds for Ritz vectors and approximate singular vectors*, Math. Comput., 89 (2018), pp. 1843–1866.

[17] Y. Nakatsukasa, *Fast and stable randomized low-rank matrix approximation*, arXiv preprint arXiv:2009.11392v1, (2020), https://arxiv.org/abs/2009.11392v1.

[18] Y. Nakatsukasa and J. A. Tropp, *Fast and accurate randomized algorithms for linear systems and eigenvalue problems*, SIAM J. Matrix Anal. Appl., 45 (2024), pp. 1183–1214.

[19] C. C. Paige, *Error analysis of the Lanczos algorithm for tridiagonalizing a symmetric matrix*, IMA J. Appl. Math., 18 (1976), pp. 341–349.

[20] B. N. Parlett, *The Symmetric Eigenvalue Problem*, SIAM, Philadelphia, 1998.

[21] A. K. Saibaba, *Randomized subspace iteration: Analysis of canonical angles and unitarily invariant norms*, SIAM J. Matrix Anal. Appl., 40 (2019), pp. 23–48.

[22] G. W. Stewart, *Matrix Algorithms*, vol. 1, SIAM, Philadelphia, 1998.

[23] G. W. Stewart and J.-G. Sun, *Matrix Perturbation Theory*, Computer Science and Scientific Computing, Academic Press, 1990.

[24] F. Woolfe, E. Liberty, V. Rokhlin, and M. Tygert, *A fast randomized algorithm for the approximation of matrices*, Appl. Comput. Harmon. Anal., 25

(2008), pp. 335–366.

[25] M. Zhou, A. Knyazev, and K. Neymeyr, *Angle-free cluster-robust Ritz value bounds for restarted block eigensolvers*, Numerical Linear Algebra with Applications, 32 (2025), p. e2607, https://doi.org/https://doi.org/10.1002/nla.2607.

[26] P. Zhu, M. E. Argentati, and A. V. Knyazev, *Bounds for the rayleigh quotient and the spectrum of self-adjoint operators*, SIAM Journal on Matrix Analysis and Applications, 34 (2013), pp. 244–256, https://doi.org/10.1137/120884468.