

On the use of harmonic Ritz pairs in approximating internal eigenpairs*

Gerard L.G. Sleijpen[†] Jasper van den Eshof[‡]

April, 2001 (revised August, 2001)

Abstract. The goal of this paper is to increase our understanding of harmonic Rayleigh–Ritz for real symmetric matrices. We do this by discussing different, though related topics: a priori error analysis, a posteriori error analysis, a comparison with refined Rayleigh–Ritz and the selection of a suitable harmonic Ritz vector.

Key words. Rayleigh–Ritz; harmonic Rayleigh–Ritz; refined Ritz; Lehmann intervals; optimal eigenvalue bounds; interior eigenvalues

AMS subject classification. 65F15, 65F50

1 Introduction

In some applications it is necessary to compute a few eigenvectors corresponding to eigenvalues in the interior of the spectrum of a real symmetric matrix A . Or, in other words, to find a pair (λ, x) (with $x \neq 0$) that satisfies

$$Ax = \lambda x.$$

Iterative methods are often the only option for computing λ and x when the matrix A is very large and sparse. Well-known examples of such methods include the Lanczos method [16, Chapter 13], the Davidson method [4] and Jacobi–Davidson [19], to mention only a few.

There are two distinct aspects of most methods. The first is the step-by-step construction of a good subspace. The second is the extraction of a good eigenvector approximation from that subspace by a projection technique. Both aspects are important. In this paper we focus on the extraction phase.

The subspace projection is often seen as a way to accelerate the convergence of a simple iteration in a similar fashion as, for example, GMRES for linear systems of equations can be viewed as an accelerated version of Richardson iteration. However, the situation for eigenvalue methods is often more delicate because here frequently an approximate eigenpair from the subspace is used in the computation of a vector to expand the subspace or for restarts. Then the success of the eigenvalue method crucially depends on the success of the subspace projection in constructing a good eigenvector approximation.

The best-known method for forming approximations from a subspace is Rayleigh–Ritz. This method has some optimal properties for exterior eigenvalues, see for example Section 11.4

*Revised version August, 2001

[†]Department of Mathematics, Utrecht University, P.O. Box 80.010, NL-3508 TA Utrecht, The Netherlands.
E-mail: sleijpen@math.uu.nl, eshof@math.uu.nl.

in [16], and is therefore well-suited for the computation of these eigenvalues and associated eigenvectors. Unfortunately, the situation is less favorable when searching for eigenvectors with eigenvalues in the interior of the spectrum [17, 8, 13].

There are various efforts to overcome the problems for interior eigenpairs. For example, Scott [17] argues that working with a shifted and inverted operator in Rayleigh–Ritz is preferable. Morgan recognized and proposed in [13] that the required inversion of the operator can be handled implicitly with a particular choice for the subspace. The resulting method has been given the name *harmonic* Rayleigh–Ritz in [15]. The eigenvalue approximations corresponding to this method (harmonic Ritz values) had already received considerable attention due to their connection with the polynomials of iterative minimal residual methods for linear systems (Kernel polynomials), see [5, 12, 15] for some recent work, and have also been studied in the context of Lehmann’s optimal inclusion intervals for eigenvalues [10, 11, 16, 1].

We treat the following aspects of harmonic Rayleigh–Ritz. The aspects are different but related and in the analysis we use similar arguments.

The subject of Section 5 is a priori error bounds for the harmonic Ritz pairs. We generalize well-known error bounds for Rayleigh–Ritz to the harmonic Rayleigh–Ritz context and discuss some of their limitations.

A posteriori error bounds for the harmonic Ritz values are discussed in Section 6. By changing the shift in harmonic Rayleigh–Ritz different intervals can be obtained. All intervals contain at least one eigenvalue. We give a condition for a posteriori choosing a new shift that results in a smaller inclusion interval. Repeatedly relocating the shift using this condition will ultimately result in an, evidently appealing, optimal interval with respect to the given information. This interval can be used as an a posteriori error estimator.

For each shift, the harmonic Rayleigh–Ritz method produces a set of harmonic Ritz vectors. In practice, the eigenvector is unknown, and it is not obvious how to tell which vector from this set forms the best approximation to the target-eigenvector. The problem of selecting a well-suited harmonic Ritz vector for a given shift is treated in Section 7.

We begin, in the next section, by giving a definition of harmonic Rayleigh–Ritz. Then we summarize some useful properties in Section 3.

Subsequently, in Section 4, we compare harmonic Rayleigh–Ritz to *refined Rayleigh–Ritz*. Refined Rayleigh–Ritz, popularized by Jia [6], is a different method to compute approximations from a subspace specially for eigenvectors with eigenvalues in the interior. Although the relation between these approaches is of interest on its own account, it turns out to be also useful in the rest of this paper.

By varying the shift the angle between the best harmonic Ritz vector for that shift and the target-eigenvector changes. As an application of the relation between harmonic and refined Rayleigh–Ritz we discuss also in Section 4 the question what shift minimizes this angle.

Although some of the results in this paper have practical applications, the purpose of this paper is to provide insight rather than algorithms.

2 Harmonic Rayleigh–Ritz

The matrix A is n by n and real symmetric and we assume that the eigenpairs of interest have eigenvalues close to some target value. The eigenpairs (λ_i, x_i) of A are numbered such that

$$\lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_n.$$

Let $V \in \mathbb{R}^{n \times k}$ be an orthonormal matrix, whose columns span a k dimensional subspace \mathcal{V} . We are interested in techniques that compute approximations to (interior) eigenpairs, only using information about \mathcal{V} and $A\mathcal{V}$. The most important method in this class is *Rayleigh–Ritz*.

The Rayleigh–Ritz approach gives k approximate eigenpairs (θ_i, u_i) , the so-called *Ritz pairs*, by imposing the Ritz-Galerkin condition

$$Au_i - \theta_i u_i \perp \mathcal{V} \quad \text{with} \quad u_i \in \mathcal{V} \setminus \{0\},$$

or equivalently,

$$(1) \quad V^T AV z_i - \theta_i z_i = 0 \quad \text{with} \quad u_i \equiv V z_i \neq 0.$$

The vector u_i is called a *Ritz vector*. From (1) it follows that the corresponding eigenvalue approximation θ_i (*Ritz value*) is the so-called *Rayleigh quotient* $\rho(u_i)$ of the vector u_i

$$\theta_i = \rho(u_i) \quad \text{where} \quad \rho(v) \equiv \frac{v^T Av}{v^T v}.$$

It follows from (1) that the Ritz values are real. We assume that they are increasingly ordered and that $\|u_i\|_2 = 1$.

Rayleigh–Ritz is not well-suited for eigenpairs (λ, x) with eigenvalue λ in the interior of the spectrum, see [17], [16, Section 11.6], [8]. If the eigenvector x makes a small angle with \mathcal{V} then there is a Ritz value close to the eigenvalue λ associated with x . However, for interior eigenvalues the corresponding Ritz vector may be a combination of eigenvectors that have little to do with the eigenvector x . It is known [8] that this can only happen if there is another Ritz value close to λ and its Ritz vector can also be a poor approximation to x . See also Example 1 in [13]. Due to the Interlace Property of the Ritz values (Theorem 10.1.1 in [16]) we know that there cannot be two Ritz values arbitrary close to λ_1 (or for that matter λ_n) and, therefore, Rayleigh–Ritz is robust for eigenvalues at the boundary of the spectrum.

A simple strategy to make Rayleigh–Ritz work for interior eigenpairs is to apply a spectral transformation such that the interesting eigenvalues are mapped to the boundary, for example apply Rayleigh–Ritz to $(A - \xi I)^2$ if the interesting eigenvalues are close to some target ξ , e.g. [13]. This gives

$$(A - \xi I)^2 \hat{u}_i - \hat{\theta}_i \hat{u}_i \perp \mathcal{V} \quad \text{with} \quad \hat{u}_i \in \mathcal{V} \setminus \{0\},$$

or,

$$(2) \quad V^T (A - \xi I)^2 V \hat{z}_i - \hat{\theta}_i \hat{z}_i = 0 \quad \text{with} \quad \hat{u}_i \equiv V \hat{z}_i \neq 0.$$

This approach plays an important role in the rest of this paper.

An alternative is to apply Rayleigh–Ritz to $(A - \sigma I)^{-1}$ for some shift σ in the neighborhood of the eigenvalue of interest. Morgan proposed to use the subspace $(A - \sigma I)\mathcal{V}$ to prevent the explicit inversion of the matrix $A - \sigma I$, which results in *harmonic Rayleigh–Ritz*. For details see [13]. We use the following equivalent definition [19, Theorem 5.1] which does not require the existence of the inverse of $A - \sigma I$. The harmonic Ritz pairs $(\tilde{\theta}_i, \tilde{u}_i)$ w.r.t. a shift σ are given by imposing the Petrov-Galerkin condition

$$(A - \tilde{\theta}_i I) \tilde{u}_i \perp (A - \sigma I)\mathcal{V} \quad \text{with} \quad \tilde{u}_i \in \mathcal{V} \setminus \{0\},$$

or as the generalized eigenvalue problem,

$$(3) \quad V^T (A - \sigma I)^2 V \tilde{z}_i - (\tilde{\theta}_i - \sigma) V^T (A - \sigma I) V \tilde{z}_i = 0 \quad \text{with} \quad \tilde{u}_i \equiv V \tilde{z}_i \neq 0.$$

Just as Ritz values are Rayleigh quotients of Ritz vectors, harmonic Ritz values $\tilde{\theta}_i$ are *harmonic Rayleigh quotients* $\tilde{\rho}_\sigma(\tilde{u}_i)$ of harmonic Ritz vectors \tilde{u}_i

$$(4) \quad \tilde{\theta}_i = \tilde{\rho}_\sigma(\tilde{u}_i) \text{ where } \tilde{\rho}_\sigma(v) \equiv \sigma + \frac{v^T(A - \sigma I)^2 v}{v^T(A - \sigma I)v}.$$

The harmonic Rayleigh quotient is sometimes also called Temple quotient, cf [3, Equation (8.31)]. In principle it can happen that $\tilde{u}^T(A - \sigma I)\tilde{u} = 0$ in (4), in this case we will write $\tilde{\theta} = \infty$ and $(\tilde{\theta} - \sigma)^{-1} = 0$. Furthermore, the vectors \tilde{u}_i are normalized and we index the harmonic Ritz values as follows

$$\tilde{\theta}_{-l} \leq \dots \leq \tilde{\theta}_{-1} < \sigma < \tilde{\theta}_1 \leq \dots \leq \tilde{\theta}_{k-l}.$$

Ideally we should write something like $\tilde{\theta}_i^\sigma$ and $\tilde{\theta}_i^\xi$ to express the non trivial dependence of these values on the shifts σ and ξ . However, we will drop the superscripts in order not to clutter the notation too much.

3 Some useful properties of harmonic Rayleigh–Ritz

For the convenience of the reader we summarize some properties of harmonic Rayleigh–Ritz with shift σ that turn out to be useful in the rest of this paper but are also of interest on their own account.

Lemma 3.1. *Assume that $(A - \sigma I)V$ has full rank. Let $k = \dim(\mathcal{V})$ and let $l, m, k - l - m$ be the number of Ritz values of A w.r.t. \mathcal{V} less than σ , equal to σ , greater than σ , respectively. Then there exist k reciprocals of shifted harmonic Ritz values $(\tilde{\theta}_i - \sigma)^{-1}$ (see (3)), of which l are negative, m equal zero and $k - l - m$ are positive. There are k linear independent \tilde{u}_i . More precisely,*

$$\tilde{u}_i^T(A - \sigma I)\tilde{u}_j = 0, \quad \tilde{u}_i^T(A - \sigma I)^2\tilde{u}_j = 0, \quad \text{if } i \neq j.$$

Proof. The matrix $V^T(A - \sigma I)^2V$ is symmetric. Since it also has full rank it is strictly positive definite and the Cholesky decomposition $LL^T = V^T(A - \sigma I)^2V$ exists. Then the \tilde{z}_i from (3) equals $y_i = L^T\tilde{z}_i$ where $((\tilde{\theta}_i - \sigma)^{-1}, y_i)$ is an eigenpair of $B \equiv L^{-1}V^T(A - \sigma I)VL^{-T}$. The $(\tilde{\theta}_i - \sigma)^{-1}$ are real, possibly zero, because of the symmetry of this operator. Sylvester’s law of inertia [16, Fact 1.6] shows that the number of positive, negative and zero eigenvalues of $V^T(A - \sigma I)V$ equals these numbers for B . Finally, the $(A - \sigma I)$ - and $(A - \sigma I)^2$ -orthogonality follow easily from the orthogonality and B -orthogonality of the y_i . \square

This lemma shows that the number of Ritz values equal to the shift σ equals the number of infinite harmonic Ritz values.

3.1 A minmax characterization for harmonic Ritz values

A useful characterization of the harmonic Ritz values is the following formulation of the minmax property, see also [10].

Lemma 3.2. *Assume that $(A - \sigma I)V$ has full rank. Let $k = \dim(\mathcal{V})$ and let $l, m, k - l - m$ be the number of Ritz values less than σ , equal to σ , greater than σ , respectively. In this case*

$$\begin{aligned} \frac{1}{\tilde{\theta}_j - \sigma} &= \max_{\mathcal{S} \subset \mathcal{V}, \dim(\mathcal{S})=j} \min_{u \in \mathcal{S}, u \neq 0} \frac{1}{\tilde{\rho}_\sigma(u) - \sigma} \quad \text{for } j \in \{1, \dots, k - l - m\}, \\ \frac{1}{\tilde{\theta}_{-j} - \sigma} &= \min_{\mathcal{S} \subset \mathcal{V}, \dim(\mathcal{S})=j} \max_{u \in \mathcal{S}, u \neq 0} \frac{1}{\tilde{\rho}_\sigma(u) - \sigma} \quad \text{for } j \in \{1, \dots, l\}. \end{aligned}$$

Proof. Using the matrix B defined in the proof of Lemma 3.1, the standard minmax characterization ([16, Theorem 10.2.1]) for the eigenvalues of B yields for $j > 1$

$$\begin{aligned} \frac{1}{\tilde{\theta}_j - \sigma} &= \max_{\mathcal{S} \subset \mathbb{R}^k, \dim(\mathcal{S})=j} \min_{y \in \mathcal{S}, y \neq 0} \frac{y^T B y}{y^T y} = \max_{\mathcal{S} \subset \mathbb{R}^k, \dim(\mathcal{S})=j} \min_{\tilde{z} \in \mathcal{S}, \tilde{z} \neq 0} \frac{\tilde{z}^T V^T (A - \sigma I) V \tilde{z}}{\tilde{z}^T V^T (A - \sigma I)^2 V \tilde{z}} \\ &= \max_{\mathcal{S} \subset \mathcal{V}, \dim(\mathcal{S})=j} \min_{\tilde{u} \in \mathcal{S}, \tilde{u} \neq 0} \frac{\tilde{u}^T (A - \sigma I) \tilde{u}}{\tilde{u}^T (A - \sigma I)^2 \tilde{u}}. \end{aligned}$$

A similar argument can be used for the harmonic Ritz values with a negative index. \square

Note that due to the way we index the harmonic Ritz values it is necessary to split the minmax property in two parts.

In case the subspace \mathcal{V} is a Krylov space for A it is known that the harmonic Ritz values interlace the Ritz values and the shift σ [15, Section 7]. The following interesting corollary can be interpreted as generalization of this interlace property for more general subspaces. It is an application of the minmax property (in Lemma 3.2) for harmonic Ritz values in combination with the statement from Lemma 3.1. The corollary can also be viewed as a generalization of Theorem 2.1 in [1] for indefinite matrices.

Corollary 3.1. *Assume that $(A - \sigma I)V$ has full rank. Let $k = \dim(\mathcal{V})$ and let $l, k - l$ be the number of Ritz values less than σ , greater than σ , respectively. Then*

$$\begin{aligned} \sigma &< \theta_{l+j} \leq \tilde{\theta}_j \quad \text{for } j \in \{1, \dots, k - l\}, \\ \sigma &> \theta_{l+1-j} \geq \tilde{\theta}_{-j} \quad \text{for } j \in \{1, \dots, l\}. \end{aligned}$$

Proof. We prove the first statement.

With $T \equiv V^T (A - \sigma I) V$ and $R \equiv (A - \sigma I) V - V T$ we have that $(A - \sigma I) V = V T + R$ and $V^T R = 0$. Hence, $x^T V^T (A - \sigma I)^2 V x = x^T T^2 x + x^T R^T R x \geq x^T T^2 x$ for all x . We know from Lemma 3.1 that for $j \geq 1$, $\theta_{l+j} > \sigma$ and $\tilde{\theta}_j > \sigma$. Since $T z_i = (\theta_i - \sigma) z_i$ we have that $T z_i = (\theta_i - \sigma)^{-1} T^2 z_i$. Because T^2 is positive definite (there are no Ritz values equal to σ) we can use the minmax property for generalized eigenvalue problems and for harmonic Ritz values to get

$$\begin{aligned} \frac{1}{\theta_{l+j} - \sigma} &= \max_{\mathcal{S} \subset \mathbb{R}^k, \dim(\mathcal{S})=j} \min_{y \in \mathcal{S}, y \neq 0} \frac{y^T T y}{y^T T^2 y} \\ &\geq \max_{\mathcal{S} \subset \mathbb{R}^k, \dim(\mathcal{S})=j} \min_{y \in \mathcal{S}, y \neq 0} \frac{y^T T y}{y^T (T^2 + R^T R) y} = \frac{1}{\tilde{\theta}_j - \sigma}. \end{aligned}$$

\square

3.2 Optimal inclusion intervals for eigenvalues

Just like Ritz values, harmonic Ritz values provide information about the eigenvalues that is optimal in some sense. Paige, Parlett and Van der Vorst [15] pointed out an important relation between Lehmann intervals and harmonic Rayleigh–Ritz. They showed that the harmonic Ritz values with respect to the shift σ give Lehmann’s optimal inclusion intervals for eigenvalues as described in the following proposition.

Proposition 3.1 (Lehmann [10]). *Let $k = \dim(\mathcal{V})$ and let $l, k - l$ be the number of Ritz values less than σ , greater than σ , respectively.*

For each $i = 1, \dots, l$, the interval $[\theta_{-i}, \sigma]$ contains at least i eigenvalues of A . For each $i = 1, \dots, k - l$, the interval $[\sigma, \tilde{\theta}_i]$ contains at least i of A ’s eigenvalues. Moreover, in the absence of extra information no smaller intervals have this property.

This shows that the harmonic Ritz values provide outer bounds for the eigenvalues. Another consequence is that if $\lambda_{p-1} < \sigma < \lambda_p < \lambda_{p+1}$ there can at most be one harmonic Ritz value in the interval $[\lambda_p, \lambda_{p+1})$ and there is no $\tilde{\theta}$ such that $\sigma \leq \tilde{\theta} < \lambda_p$.

Kahan derived an explicit matrix \hat{A} with $V^T \hat{A}^2 V = V^T A^2 V$ and $V^T \hat{A} V = V^T A V$, such that the eigenvalues of \hat{A} are at the end points of the intervals. This matrix can be used to compute the harmonic Ritz values, which can offer computational advantages, for example, when \mathcal{V} is a Krylov subspace. See [15, Section 7], [16, Section 10.5] for details.

3.3 The concept of ρ -values

We see from the harmonic Rayleigh quotient (4) that if the shift σ is close to the eigenvalue λ_p , then a harmonic Ritz vector \tilde{u} must make a relatively small angle with x_p to have the harmonic Ritz value $\tilde{\theta}$ close to λ_p . This problem becomes more severe when σ is chosen closer to λ_p . Hence, the harmonic Ritz values do not always provide good approximations to the eigenvalues.

Several authors note (i.e., [13, 20, 14]) that “better” eigenvalue approximations are given by the so-called ρ -values defined as (recall that the \tilde{u}_i are normalized)

$$\rho_i \equiv \rho(\tilde{u}_i) = \tilde{u}_i^T A \tilde{u}_i.$$

In practical implementations these ρ -values are cheap to compute, i.e., they do not require an additional matrix-vector product.

We note, using an application of Cauchy-Schwarz and (4), that

$$0 \leq (\rho_i - \sigma)^2 = (\tilde{u}_i^T (A - \sigma I) \tilde{u}_i)^2 \leq \tilde{u}_i^T (A - \sigma I)^2 \tilde{u}_i = (\tilde{\theta}_i - \sigma)(\rho_i - \sigma),$$

where the second inequality is strict if and only if $(A - \sigma I)\tilde{u}_i$ is not a multiple of \tilde{u}_i , i.e., \tilde{u}_i is not an eigenvector. We see that $|\tilde{\theta}_i - \sigma| \geq |\rho_i - \sigma|$ and from $(\tilde{\theta}_i - \sigma)(\rho_i - \sigma) \geq 0$ it follows that the ρ -value ρ_i always lies between the shift σ and $\tilde{\theta}_i$. This was also observed by Morgan and Zeng, who derived this as a corollary of the following lemma. This lemma from [14] allows for an inexpensive calculation of the norm of the residual of a harmonic Ritz vector with ρ_i as approximate eigenvalue. We give a different proof here.

Lemma 3.3 (Morgan, Zeng [14, Theorem 2.1]). *Let $r_i \equiv A\tilde{u}_i - \rho_i\tilde{u}_i$, then*

$$(5) \quad \|r_i\|_2^2 = (\rho_i - \sigma)(\tilde{\theta}_i - \rho_i).$$

Proof. Using the fact that $(A - \tilde{\theta}_i I)\tilde{u}_i \perp (A - \sigma I)\tilde{u}_i$ and $(A - \rho_i I)\tilde{u}_i \perp \tilde{u}_i$ leads to

$$\begin{aligned} \|r_i\|_2^2 &= (A\tilde{u}_i - \rho_i\tilde{u}_i)^T (A\tilde{u}_i - \rho_i\tilde{u}_i) = (A\tilde{u}_i - \tilde{\theta}_i\tilde{u}_i)^T (A\tilde{u}_i - \rho_i\tilde{u}_i) \\ &= (A\tilde{u}_i - \tilde{\theta}_i\tilde{u}_i)^T (\sigma\tilde{u}_i - \rho_i\tilde{u}_i) = (\rho_i\tilde{u}_i - \tilde{\theta}_i\tilde{u}_i)^T (\sigma\tilde{u}_i - \rho_i\tilde{u}_i) = (\rho_i - \sigma)(\tilde{\theta}_i - \rho_i). \end{aligned}$$

□

For readers familiar with Temple's bound (cf [2, Lemma 1.27],[3, pp. 116]) we mention that the expression in (5) is very natural given the optimality of the inclusion intervals $[\sigma, \tilde{\theta}_1]$ and $[\tilde{\theta}_{-1}, \sigma]$ from Proposition 3.1.

4 A comparison with refined Rayleigh–Ritz

A different approach for finding approximations from a subspace to eigenvectors with an eigenvalue in the interior is *refined Rayleigh–Ritz* by Jia [6]. In the real symmetric case this method is closely related to applying Rayleigh–Ritz to $(A - \xi I)^2$ as in Section 2. It is of theoretical and practical interest to have some insight on how this method is related to harmonic Rayleigh–Ritz and which method works best in which situation. Some numerical experiments were done in [13] for specific choices of the shifts in both methods, but no explanation or insight is offered there.

In this section we want to give some heuristics on this subject and we relate the quality of the approximation by refined Rayleigh–Ritz to the quality of the harmonic Ritz vector \tilde{u}_1 (or \tilde{u}_{-1}). In Section 4.1 we give a definition of the refined Ritz vector and prove a relation between harmonic and refined Rayleigh–Ritz. This relation also plays a crucial role in the rest of this paper. In Sections 4.2 and 4.3 we study the optimal tuning of the parameters for the two approaches. This is necessary in order to make a fair comparison between both methods which is done in Section 4.4.

4.1 Refined Rayleigh–Ritz and its relation to harmonic Rayleigh–Ritz

We give a formal definition of refined Rayleigh–Ritz [6]. Let ξ be a given approximation for the eigenvalue of which we want to approximate the corresponding eigenvector. Define

$$(6) \quad \nu_\xi \equiv \min_{u \in \mathcal{V}, \|u\|_2=1} \|Au - \xi u\|_2 \quad \text{and} \quad \hat{u} \equiv \operatorname{minarg}_{u \in \mathcal{V}, \|u\|_2=1} \|Au - \xi u\|_2.$$

Then \hat{u} is equal to the Ritz vector with smallest Ritz value of $(A - \xi I)^2$ with respect to \mathcal{V} , see (2). The vector \hat{u} is called the *refined Ritz vector* of A with respect to the approximate eigenvalue ξ and the search subspace \mathcal{V} . In refined Rayleigh–Ritz as in [6], it is proposed to pick ξ as the currently best known approximation for the wanted eigenvalue, for example a Ritz value. Here we allow that ξ is more general.

There is an interesting relation between harmonic Ritz vectors and the refined Ritz vectors. It is based on the following observation. If the approximate eigenvalue ξ is in the middle of the interval with endpoints σ and $\tilde{\theta}$ (either $\tilde{\theta}_1$ or $\tilde{\theta}_{-1}$), then the refined Ritz vector \hat{u} and the harmonic Ritz vector \tilde{u} (with harmonic Ritz value $\tilde{\theta}$) coincide. We will show this in the theorem below. Moreover, the radius of the interval is exactly equal to the residual norm ν_ξ of the refined Ritz vector: Fig. 1 illustrates this situation. This latter observation was also used by Lehmann (in a more general form) in the construction of his optimal intervals [10,

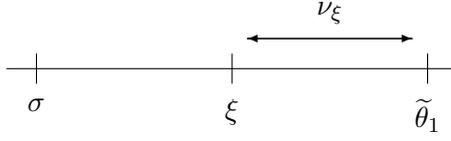


Figure 1: The location of the shift σ and the approximate eigenvalue ξ for which the harmonic Ritz vector and the refined Ritz vector coincide (see Theorem 4.1).

pp. 258] (see also [16, pp. 219]). The situation of Fig. 1 can be enforced for a given shift σ by picking ξ as the average of σ and $\tilde{\theta}$, and, for given ξ , by picking σ at distance ν_ξ from ξ .

Theorem 4.1. *If σ is given, then select $\xi = \frac{1}{2}(\sigma + \tilde{\theta}_{\pm 1})$. If ξ is given, then select $\sigma = \xi \mp \nu_\xi$. In both cases, we have that $\tilde{\theta}_{\pm 1} = \xi \pm \nu_\xi$, $\nu_\xi = |\frac{1}{2}(\sigma - \tilde{\theta}_{\pm 1})|$ and if $\tilde{u}_{\pm 1}$ and \hat{u} are unique (up to a sign), then \hat{u} equals $\tilde{u}_{\pm 1}$ up to a sign.*

Proof. We work out the details for $\tilde{\theta} = \tilde{\theta}_1$. Consider a $v \in \mathcal{V}$ and put $\tilde{\xi} \equiv \frac{1}{2}(\sigma + \tilde{\theta})$ and $\gamma \equiv \frac{1}{2}(\tilde{\theta} - \sigma)$. Then $v = \tilde{u}$ if and only if $Av - \tilde{\theta}v \perp (A - \sigma I)\mathcal{V}$ and $\|v\| = 1$, or equivalently,

$$(7) \quad (A - \sigma I)(A - \tilde{\theta} I)v = (A - \tilde{\xi} I)^2 v - \gamma^2 v \perp \mathcal{V}.$$

Note that $\gamma^2 = (\tilde{\theta} - \sigma)^2/4 = \|(A - \tilde{\xi} I)v\|_2^2$. Hence, \tilde{u}_1 satisfies (2). It must also correspond to the smallest Ritz value, otherwise this would contradict the optimality of the Lehmann interval from Proposition 3.1.

Furthermore, $v = \hat{u}$ if and only if

$$(8) \quad (A - \xi I)^2 v - \delta v \perp \mathcal{V}$$

for the pair (δ, v) with the smallest value of δ . Note that (8) implies that $\delta = \|(A - \xi I)v\|_2^2 = \|(A - \xi I)\hat{u}\|_2^2 = \nu_\xi^2$.

Now, using the fact that (7) and (8) are equivalent for appropriate scalars $\xi, \tilde{\xi}, \gamma$ and δ , the theorem follows. \square

As a side remark we note that this theorem has a nice consequence if \mathcal{V} is a Krylov space for A . It gives a way to compute the roots of the refined Ritz polynomial (see [7]): select $\sigma = \xi - \nu_\xi$. Then the roots are given by the harmonic Ritz values, excluding the one at $\sigma + 2\nu_\xi$. This is an alternative for the construction in Theorem 3.1 in [7] for symmetric matrices.

4.2 The optimal value of ξ in refined Rayleigh–Ritz

In this section we investigate for which ξ the angle between the refined Ritz vector and the unknown eigenvector x_p is as small as possible. We call this shift ξ^* . It is difficult to make general statements about the value of ξ^* because it depends not only on eigenvalue distribution but also on the structure of the space \mathcal{V} . However, some general observations can be made. For example, if ξ is closest to the eigenvalue λ_p then, for $\sin^2 \angle(x_p, \mathcal{V})$ small enough, we have the following proposition with an error bound for the refined Ritz vector. This bound is simply Theorem 3.1 from [18] with A replaced by $(A - \xi I)^2$.

Proposition 4.1. *Let*

$$\xi \in \left(\frac{1}{2}(\lambda_{p-1} + \lambda_p), \frac{1}{2}(\lambda_p + \lambda_{p+1}) \right)$$

and $q \equiv \operatorname{argmin}_{i \neq p} |\lambda_i - \xi|$, $r \equiv \operatorname{argmax}_i |\lambda_i - \xi|$ and $\epsilon \equiv \sin^2 \angle(\mathcal{V}, x_p)$.

Define

$$(9) \quad \kappa_\xi \equiv C_\xi + C_\xi^{-1} - 2 \quad \text{where} \quad C_\xi \equiv \frac{(\lambda_r - \xi)^2 - (\lambda_p - \xi)^2}{(\lambda_q - \xi)^2 - (\lambda_p - \xi)^2}.$$

If

$$(10) \quad \epsilon < C_\xi^{-1},$$

then we have for all $k \in \{2, \dots, n-1\}$:

$$(11) \quad \sin^2 \angle(\hat{u}, x_p) \leq \frac{1}{2}(1 + \epsilon) - \frac{1}{2}\sqrt{(1 - \epsilon)^2 - \kappa_\xi \epsilon},$$

Furthermore, if ξ is independent of \mathcal{V} then bound (11) is sharp.

The constant C_ξ can be interpreted as a condition number of the eigenvector x_p of the matrix $(A - \xi I)^2$ (see [2, Section 2.3]). From the fact that $C_\xi \geq 1$ and $\kappa_\xi \equiv C_\xi + C_\xi^{-1} - 2$, it follows that without additional information the shift ξ that minimizes C_ξ , results in the smallest attainable upper bound and gives this bound the largest area of application.

A simple analysis shows that the shift

$$(12) \quad \frac{\lambda_{p+1} + \lambda_{p-1}}{2}$$

minimizes C_ξ and, therefore, without further information, is the best shift.

In practice one often picks ξ as close as possible to the eigenvalue. This makes sense for general non-normal problems, but in our context the singular- and eigen-vectors coincide and selecting $\xi \approx \lambda_p$ is only expected to be optimal if the eigenvalue distribution is uniform around λ_p .

4.3 The optimal value of σ in harmonic Rayleigh–Ritz

Now we look at the optimal position of σ , denoted by σ^* , for harmonic Rayleigh–Ritz. We expect there to be two (locally) optimal positions, one left of λ_p in which case \tilde{u}_1 is offering the approximation and, of course, one right of λ_p where \tilde{u}_{-1} is the vector of interest.

Given the statement of Theorem 4.1 it is straightforward to give an expression for the location of the optimal shifts in terms of the quantities ξ^* and ν_{ξ^*} : if ξ^* is the best shift for the refined Ritz approach then the harmonic Ritz vector $\tilde{u}_{\pm 1}$ is the best possible approximation for x_p if and only if $\sigma^* = \xi^* \mp \nu_{\xi^*}$. Indeed we see that, because $\nu_{\xi^*} \geq |\lambda_p - \xi^*|$, there are two equally good and optimal positions for σ (left and right of λ_p).

Again let $\epsilon \equiv \sin^2 \angle(\mathcal{V}, x_p)$ and note that ν_{ξ^*} and $|\lambda_p - \xi^*|$ approach the same value if $\epsilon \rightarrow 0$ (recall that ξ^* depends on the space \mathcal{V}). Hence, we expect the optimal shifts σ^* to be arbitrarily close to λ_p and $2\xi^* - \lambda_p$ for the angle between \mathcal{V} and x_p small enough. More precisely, if ξ^* is a bounded number then $\nu_{\xi^*} = |\lambda_p - \xi^*| + O(\sqrt{\epsilon})$ and the optimal shifts lie at $\lambda_p + O(\sqrt{\epsilon})$ and $2\xi^* - \lambda_p + O(\sqrt{\epsilon})$. So, if no additional information is at hand it follows using (12) that the optimal shifts are at

$$(13) \quad \lambda_{p-1} + \lambda_{p+1} - \lambda_p + O(\sqrt{\epsilon}) \quad \text{and} \quad \lambda_p + O(\sqrt{\epsilon}) \quad (\epsilon \rightarrow 0).$$

4.4 Discussion

Theorem 4.1 allows us to interpret the harmonic Ritz vector $\tilde{u}_{\pm 1}$ as a refined Ritz vector with shift $\tilde{\xi} \equiv (\sigma + \tilde{\theta}_{\pm 1})/2$ and using this interpretation we can try to explain some often observed differences in the behavior of both methods. It is clear from Theorem 4.1 that we cannot say that one approach is best for general choices of ξ and σ . But some heuristic can be given for more practical situations in which σ and ξ are chosen fixed and close to λ_p , as in the experiments in [13].

Assume for the moment that $\lambda_{p+1} > \sigma > \lambda_p$ such that we can restrict our attention to $(\tilde{\theta}, \tilde{u}) = (\tilde{\theta}_{-1}, \tilde{u}_{-1})$. If $\tilde{\theta}$ is far from λ_p , hence $\angle(\mathcal{V}, x_p)$ is large compared to the distance between σ and λ_p (see also Section 3.3), then the shift $\tilde{\xi} \equiv (\sigma + \tilde{\theta})/2$ can be at some distance from the optimal ξ^* . This is not necessarily a problem but can cause the refined Ritz approach to pick up information about x_p a little easier if \mathcal{V} contains a poor approximation to x_p . This can be viewed as an analogy of the difference between inverse iteration and Rayleigh quotient iteration for a bad approximating vector.

For the asymptotic situation ($\angle(\mathcal{V}, x_p) \rightarrow 0$) and $\sigma \neq \lambda_p$, comparing the quality of the refined Ritz vector \hat{u} and the harmonic Ritz vector \tilde{u} amounts to comparing the shifts ξ and $\tilde{\xi} = (\sigma + \lambda_p)/2$. From the previous two sections we expect that the method that constructs the best approximation asymptotically depends on the distribution of the eigenvalues of A and one method is not best in general.

We give a simple illustration with the matrix

$$(14) \quad A = \text{diag}(1, 2, \dots, 100, 110, 114, \dots, 200)$$

and $\lambda_p = 110$. As subspace \mathcal{V} we take a Krylov space for A with a starting vector with all components equal and of dimension $k = 10, 20, \dots, 100$. We illustrate for different choices of σ (for the harmonic case) and ξ (for refined) the quality of the approximation in Fig. 2.

For this matrix we expect from (12) that $\xi = 107$ is a good choice whereas the real minimum value, ξ^* , is close to this value. Note that in this situation the position of the optimal ξ does not depend much on $\angle(\mathcal{V}, x_p)$. This is in contrast to the harmonic case where the optimal positions of σ vary with \mathcal{V} . The two pictures on the right in Fig. 2 show this. For every \mathcal{V} there are two optimal values that for larger dimensional \mathcal{V} seem to settle close to fixed positions as expected. From (13) these positions are expected in the neighborhood of $\lambda_p = 110$ and $\lambda_{p-1} + \lambda_{p+1} - \lambda_p = 104$. This seems in accordance with Fig. 2. Although one optimal shift becomes closer and closer to $\lambda_p = 110$, the shift $\sigma = 110$ is not the best candidate.

Using our new insight we can try to explain some observed differences between harmonic Rayleigh–Ritz and Rayleigh–Ritz applied to $(A - \xi I)^2$ for Example 2, page 299 in [13], where there is a matrix with $\lambda_{p-1} = 9$, $\lambda_p = 10$, $\lambda_{p+1} = 11$ and $\lambda_{p+2} \approx 11.03488$ and some subspaces \mathcal{V} . Furthermore, $\sigma = \xi = 10.1$. It is observed that the quality of \tilde{u}_{-1} and \hat{u}_1 as approximations to x_p is competitive. But the vector \tilde{u}_1 is a better approximation to x_{p+1} than \hat{u}_2 from (2).

The optimal shift for refined Ritz for λ_p is expected with (12) to be at 10 and for λ_{p+1} at approximately 10.51. In case the harmonic Ritz values $\tilde{\theta}_{-1}$ and $\tilde{\theta}_1$ are sufficiently close to λ_p and λ_{p+1} respectively, the harmonic Ritz vectors \tilde{u}_{-1} and \tilde{u}_1 are equivalent to refined Ritz with shifts 10.05 and 10.55. Notice that these shifts are close to the predicted optimal shifts for refined Rayleigh–Ritz and we can expect good approximations for both vectors with harmonic Rayleigh–Ritz with $\sigma = 10.1$. If $\xi = 10.1$, the eigenvalue λ_{p+1} becomes almost a double eigenvalue after shifting and squaring which explains the poor approximation. In

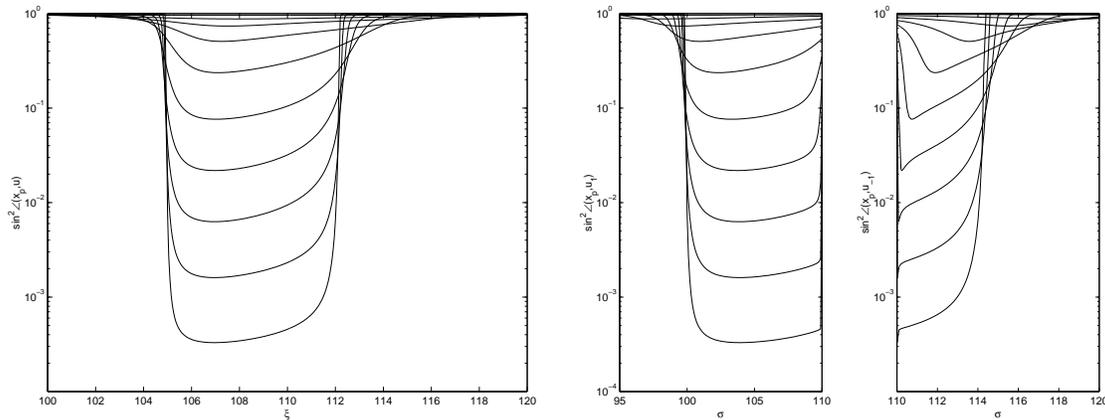


Figure 2: Left picture: $\sin^2 \angle(\hat{u}, x_p)$ on log-scale for different choices of ξ . Middle picture: $\sin^2 \angle(\tilde{u}_1, x_p)$ on log-scale for $\sigma \leq 110$. Right picture: $\sin^2 \angle(\tilde{u}_{-1}, x_p)$ on log-scale for $\sigma \geq 110$.

general, for refined Ritz or Rayleigh–Ritz applied to $(A - \xi I)^2$, a good shift for λ_p , as in (12), introduces a double eigenvalue and therefore potentially a poor approximation for x_{p-1} or x_{p+1} . This seems less likely for harmonic Rayleigh–Ritz and according to [13] somehow easier to detect.

5 A priori error estimation

For Rayleigh–Ritz the *a priori* error bounds

$$(15) \quad \theta_1 - \lambda_1 \leq (\lambda_n - \lambda_1) \sin^2 \angle(\mathcal{V}, x_1) \quad \text{and} \quad \sin^2 \angle(u_1, x_1) \leq \frac{\theta_1 - \lambda_1}{\lambda_2 - \lambda_1}$$

are standard bounds for the smallest Ritz value and the corresponding Ritz vector [9]. For $\lambda_1 \leq \theta_1 < \lambda_2$ we rewrite the bound for the Ritz vector in (15) to

$$(16) \quad \tan^2 \angle(u_1, x_1) \leq \frac{\theta_1 - \lambda_1}{\lambda_2 - \theta_1}.$$

We translate (15) and (16) to harmonic Rayleigh–Ritz. Similar statements can be found in [13, Theorem 3].

We assume in this section that the eigenvalue of interest is λ_p with p such that $\lambda_p > \sigma > \lambda_{p-1}$ and we concentrate on the pair $(\tilde{\theta}_1, \tilde{u}_1)$. Statements for \tilde{u}_{-1} can be obtained by replacing A by $-A$.

We first give a variant of (16) for the harmonic Ritz vector \tilde{u}_1 .

Theorem 5.1. *Let $q \equiv \operatorname{argmin}_{i \neq p} |\lambda_i - \frac{1}{2}(\tilde{\theta}_1 + \sigma)|$. If $\tilde{\theta}_1$ exists and $\lambda_p \leq \tilde{\theta}_1 < \lambda_{p+1}$, then*

$$\tan^2 \angle(\tilde{u}_1, x_p) \leq \left(\frac{\lambda_p - \sigma}{\lambda_q - \sigma} \right) \left(\frac{\tilde{\theta}_1 - \lambda_p}{\lambda_q - \tilde{\theta}_1} \right).$$

Proof. According to Theorem 4.1, the pair $(\frac{1}{4}(\tilde{\theta}_1 - \sigma)^2, \tilde{u}_1)$ is the Ritz pair of $(A - \frac{1}{2}(\tilde{\theta}_1 + \sigma)I)^2$ with respect to \mathcal{V} with the smallest Ritz value. Now, apply (16). \square

Besides the factor $|(\lambda_p - \sigma)/(\lambda_q - \sigma)|$ (which can be less than one), the expression in this theorem is similar to (16). But it is well known that the distance between λ_p and $\tilde{\theta}_1$ can be large even if $\angle(\mathcal{V}, x_p)$ is small, see also Section 3.3. In fact, it can take a very small angle $\angle(\mathcal{V}, x_p)$ for a harmonic Ritz value larger than σ to exist. The following simple example illustrates this.

We consider the one dimensional subspace $\mathcal{V} \equiv \text{span}(u)$ with $u \equiv \sqrt{1 - \epsilon}x_p + \sqrt{\epsilon}x_1$. From Lemma 3.1 we know that there can only be a harmonic Ritz value larger than σ if there is a positive Ritz value for $A - \sigma I$. The only Ritz value is given by $(\lambda_p - \sigma)(1 - \epsilon) + (\lambda_1 - \sigma)\epsilon$ which is positive if $\epsilon/(1 - \epsilon) < (\sigma - \lambda_p)/(\lambda_1 - \sigma)$. Apparently, we may not expect that there are harmonic Ritz values larger than σ if $\tan^2 \angle(x_p, \mathcal{V}) = \epsilon/(1 - \epsilon) > (\sigma - \lambda_p)/(\lambda_1 - \sigma)$.

The following theorem shows that for smaller angles, there is at least one harmonic Ritz value larger than σ and it gives an a priori error bound for $\tilde{\theta}_1$.

Theorem 5.2. *Let $\tau \equiv \tan^2 \angle(\mathcal{V}, x_p) < (\sigma - \lambda_p)/(\lambda_1 - \sigma)$. Then there exists a harmonic Ritz value $\tilde{\theta}_1$ of A with respect to shift σ for which*

$$(17) \quad 0 \leq \tilde{\theta}_1 - \lambda_p \leq \tau \max_i \frac{(\lambda_i - \sigma)(\lambda_i - \lambda_p)}{(\lambda_p - \sigma) + \tau(\lambda_i - \sigma)}.$$

Proof. For ease of notation, first assume that $\sigma = 0$.

Let $x_{\mathcal{V}}$ be the normalized projection of x_p on \mathcal{V} . Decompose $x_{\mathcal{V}} = \sqrt{1 - \epsilon}x_p + \sqrt{\epsilon}e$ with $e \perp x_p$ and $\|e\|_2 = 1$. Then the Rayleigh quotient $\rho(x_{\mathcal{V}}) > 0$ and $\lambda_p + \tau\lambda_1 > 0$. Hence, because $\tau = \epsilon/(1 - \epsilon)$, we have

$$\tilde{\rho}_0 \equiv \tilde{\rho}_0(x_{\mathcal{V}}) = \frac{(1 - \epsilon)\lambda_p^2 + \epsilon e^T A^2 e}{(1 - \epsilon)\lambda_p + \epsilon e^T A e} = \frac{\lambda_p^2 + \tau e^T A^2 e}{\lambda_p + \tau e^T A e}.$$

Therefore, $\lambda_p(\tilde{\rho}_0 - \lambda_p) = \tau e^T A(A - \tilde{\rho}_0 I)e \leq \tau \lambda_i(\lambda_i - \tilde{\rho}_0)$ with $i = 1$ (if $\tilde{\rho}_0 > \lambda_1 + \lambda_n$) or $i = n$ (otherwise), which implies $\tilde{\rho}_0 - \lambda_p \leq \lambda_i(\lambda_i - \lambda_p)/(\lambda + \tau\lambda_i)$. An application of Lemma 3.2 concludes the proof for $\sigma = 0$. The more general statement follows by noting that the harmonic Ritz values w.r.t. shift σ are the harmonic Ritz values minus σ of $A - \sigma I$ w.r.t. the shift 0. \square

True a priori bounds for small enough τ ($\equiv \tan^2 \angle(\mathcal{V}, x_p)$) can be obtained by substituting the result of Theorem 5.2 in the bound of Theorem 5.1. This shows that if $\lambda_p \neq \sigma$, $\tan^2 \angle(\tilde{u}_1, x_p) = O(\tau)$ for $\tau \rightarrow 0$. Unfortunately, this bound becomes useless when λ_p lies too close to σ . Sharper asymptotic a priori bounds can, for example, be obtained, by combining Theorem 4.1 with Proposition 4.1 or using a technique as used in [18] for Rayleigh–Ritz. However, this does not remove the problem of the small applicability of these bounds when λ_p close to σ .

The question is if this means, in case λ_p is close to σ , that all harmonic Ritz vectors can be poor approximations to x_p in some instances. If $\sigma = \lambda_p$, then all the eigenvectors of the pencil $(A - \sigma I, (A - \sigma I)^2)$ can have arbitrary components in the direction of x_p . This non-uniqueness seems to cause in (3) that many harmonic Ritz vectors can point in a direction close to x_p . Also if $\angle(\mathcal{V}, x_p)$ is large compared to $\lambda_p - \sigma$ this behavior is often observed, see [13].

It would be interesting to have an error bound for the harmonic Ritz vectors in case $\lambda_p = \sigma$ to better understand this behavior. Numerical experiments suggest that this upper bound

only depends on the dimension of \mathcal{V} (k) and not on some measure of the gap as is the case for $\lambda_p \neq \sigma$. In fact, we expect that for $\sigma = \lambda_p$:

$$\tan^2 \angle(\tilde{u}_j, x_p) \leq k \tan^2 \angle(\mathcal{V}, x_p) \text{ for some } j.$$

However, we have no proof for this in general.

6 A posteriori error estimation

We assume that we apply harmonic Rayleigh–Ritz with some shift σ that lies left of the interesting eigenvalue λ_p such that we can restrict our attention to $(\tilde{\theta}_1, \tilde{u}_1)$. In Section 3.3 we observed that, in general, ρ_1 provides a better approximation to λ_p than $\tilde{\theta}_1$. We, furthermore, noted that ρ_1 lies in the Lehmann interval $[\sigma, \tilde{\theta}_1]$ that contains at least one eigenvalue of A (Proposition 3.1). Hence, this interval can serve as an *a posteriori* error estimator for the approximate eigenvalue ρ_1 . However, this is not expected to be a very effective error estimator. In the first place, the size of the interval is bounded from below by $|\lambda_p - \sigma|$ and is not going to zero if \tilde{u}_1 approaches an eigenvector. Secondly, we have seen in Section 5 that if λ_p is close to σ , $\tilde{\theta}_1$ can approach λ_p more slowly for $\angle(\mathcal{V}, x) \rightarrow 0$ than might be expected from the quality of the subspace \mathcal{V} .

In this section we are interested in how to choose the shift σ such that its distance to either $\tilde{\theta}_1$ or $\tilde{\theta}_{-1}$ is as small as possible. This means that we then have located one eigenvalue with maximal precision in the absence of additional information, see Proposition 3.1.

6.1 A condition for the minimizing shift

We denote the shift σ that results in the smallest possible Lehmann interval $[\sigma, \tilde{\theta}_1]$ or $[\tilde{\theta}_{-1}, \sigma]$ with σ^+ . This shift can be easily characterized.

Theorem 6.1. *The shift σ^+ results in the smallest possible Lehmann interval if and only if*

$$(18) \quad \sigma^+ = u^{+T} A u^+ \pm \|A u^+ - \rho(u^+) u^+\|_2$$

and

$$(19) \quad u^+ = \operatorname{argmin}_{u \in \mathcal{V}, \|u\|_2=1} \|A u - \rho(u) u\|_2.$$

Proof. According to Theorem 4.1 minimizing $|\tilde{\theta}_{\pm 1} - \sigma|$ with respect to σ is equivalent to minimizing ν_ξ from (6) with respect to ξ . Since, $\|A u - \rho(u) u\|_2 \leq \|A u - \xi u\|_2$ for all u and ξ , we see that the claim follows. \square

We give a simple example. A similar illustration can be found in [16, Section 10.5].

Example 6.1. Let the matrix A be 3 dimensional and diagonal with $A_{ii} = i$ and let $\mathcal{V} = \operatorname{span}(u)$ with $u = (1, 1/2, 1/2)^T$. The one harmonic Ritz value in this case is given by the harmonic Rayleigh quotient

$$\tilde{\theta} = \tilde{\rho}_\sigma(u) = \sigma + \frac{u^T (A - \sigma I)^2 u}{u^T (A - \sigma I) u}$$

and the only Ritz value is $\theta = \rho = 1.5$. Fig. 3 shows the size of the Lehmann interval $|\tilde{\theta} - \sigma|$ for some values of $\sigma < \theta$. A simple computation shows that the smallest interval is attained for $\sigma^+ = \theta - \|A u - \theta u\|_2 \approx 0.74$ for which $|\tilde{\theta} - \sigma| = 2 \|A u - \theta u\|_2 \approx 1.53$. Note that there is of course a second minimal value at $\theta + \|A u - \theta u\|_2$.

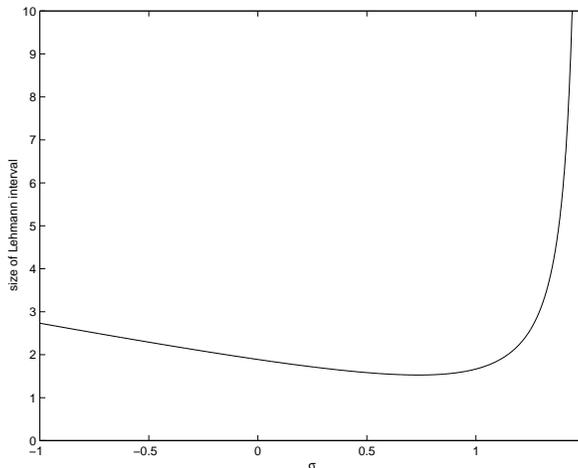


Figure 3: Size of Lehmann interval $|\tilde{\theta} - \sigma|$ as a function of σ (see Example 6.1).

In [16, pp. 218] it is remarked that no explicit expression is known for the best shift: considering the nonlinearity of (19) it is not likely that such an expression can exist. Equations (18) and (19) do, however, give a computable expression for the best shift σ^+ . Note that the computation based on (19) only requires availability of the low dimensional matrices $V^T A^2 V$ and $V^T A V$ and can be done with a suitable iterative method. The method we will propose in the next subsection is similar. It is iterative and only requires $V^T A V$ and $V^T A^2 V$ but it computes a sequence of shifts $\sigma^0, \sigma^1, \dots$ for harmonic Rayleigh–Ritz where for the computation of σ^{m+1} only the ρ -values and harmonic Ritz values are used from the previous shift σ^m .

6.2 A condition for the shift σ^{m+1}

Let $\tilde{\theta}_i^m$ be the harmonic Ritz values w.r.t. shift σ^m . We will also denote the corresponding ρ -values and residuals for this shift with a superscript m .

Given some harmonic Ritz value $\tilde{\theta}_i^m$ w.r.t. shift σ^m we want an expression for σ^{m+1} such that

$$(20) \quad |\tilde{\theta}_1^{m+1} - \sigma^{m+1}| < |\tilde{\theta}_i^m - \sigma^m|.$$

We remark that $|\tilde{\theta}_{-1}^{m+1} - \sigma^{m+1}| < |\tilde{\theta}_i^m - \sigma^m|$ is also possible but we do not consider it here.

We consider here i to be more general than plus or minus one. It can be useful to start for σ^0 with another harmonic Ritz value than $\tilde{\theta}_1$ or $\tilde{\theta}_{-1}$, we return to this in Section 6.3. If in the next steps we take i to be constantly 1, we get a sequence of shifts $\sigma^0, \sigma^1, \sigma^2, \dots$ for which the size of the Lehmann interval $[\sigma^m, \tilde{\theta}_1^m]$ becomes smaller and smaller and will ultimately result in an evidently appealing minimal interval. The idea we use to get an expression for a new shift σ^{m+1} , given σ^m and the harmonic Ritz values and ρ -values for this shift, is based on the following observation.

Every Lehmann interval with boundaries $\tilde{\theta}_i$ and σ contains ρ_i and at least one eigenvalue. An alternative interval with this inclusion property can be given using the Bauer–Fike theorem [16, Theorem 4.5.1]. This theorem implies that there is at least one eigenvalue of A in the interval $[\rho_i - \|r_i\|_2, \rho_i + \|r_i\|_2]$ where r_i is the residual $r_i \equiv A\tilde{u}_i - \rho_i\tilde{u}_i$ of the approximate eigenpair (ρ_i, \tilde{u}_i) . It turns out that the size of this interval is always smaller than the size of the corresponding Lehmann interval as follows from the following lemma.

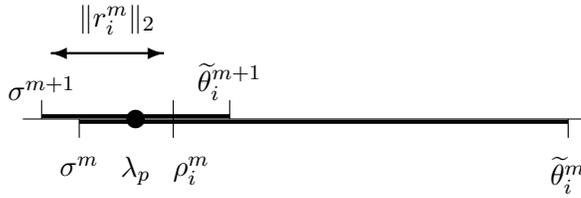


Figure 4: Relocation of the shift σ^m can lead to Lehmann intervals of smaller size, but the shorter interval can not be a contained in the taller one: see Section 6.2.

Lemma 6.1. *Let $(\tilde{\theta}_i, \tilde{u}_i)$ be a harmonic Ritz pair and ρ_i the corresponding ρ -value and define $r_i \equiv A\tilde{u}_i - \rho_i\tilde{u}_i$. Then*

$$2\|r_i\|_2 \leq |\tilde{\theta}_i - \sigma|$$

with equality if and only if $\rho_i = (\tilde{\theta}_i + \sigma)/2$.

Proof. Apply Lemma 3.3 and note that the expression on the right is maximal for $\rho_i = \frac{1}{2}(\tilde{\theta}_i + \sigma)$. \square

The idea now is as follows. Although the harmonic Ritz values w.r.t. shift σ^m provide the best upper bound on the eigenvalues closest to σ^m , they provide no better error estimates for the computed ρ -value, ρ_i^m , than a simple application of the Bauer–Fike theorem, see Lemma 6.1. Because in the Bauer–Fike interval only information is used about the vector \tilde{u}_i^m it is suggested that we can find a smaller interval with the same inclusion property by choosing σ^{m+1} at the boundary of the Bauer–Fike interval, for example $\sigma^{m+1} = \rho_i^m - \|r_i^m\|_2$. We use Lemma 3.3 to get a cheaply computable expression for $\|r_i^m\|_2$ and select

$$(21) \quad \sigma^{m+1} \equiv \rho_i^m - \sqrt{(\tilde{\theta}_i^m - \rho_i^m)(\rho_i^m - \sigma_i^m)}.$$

Fig. 4 illustrates the idea for $\tilde{\theta}_i^m$. Note that the intersection of the Bauer–Fike interval and the Lehmann interval cannot be empty because ρ_i is in both intervals. On the other hand the Bauer–Fike interval cannot be strictly contained in the Lehmann interval, this would contradict the optimality from Proposition 3.1.

Indeed, the property in (20) can be proven for the shift σ^{m+1} given by (21).

Lemma 6.2. *Take σ^{m+1} as in (21). Then*

$$|\tilde{\theta}_1^{m+1} - \sigma^{m+1}| \leq |\tilde{\theta}_i^m - \sigma^m|$$

with equality if and only if $\rho_i^m = (\tilde{\theta}_i^m + \sigma^m)/2$.

Proof. If ρ_i^m is not precisely in the middle of the Lehmann interval then

$$(\tilde{\theta}_1^{m+1} - \sigma^{m+1})^2 \leq 4(\tilde{\theta}_i^m - \rho_i^m)(\rho_i^m - \sigma^m) < (\tilde{\theta}_i^m - \sigma^m)^2.$$

The first inequality follows from Theorem 4.1 applied to the one dimensional subspace \tilde{u}_i^m followed by the minmax property of harmonic Ritz values (Lemma 3.2).

If ρ_i^m is precisely in the middle of the interval $\sigma^{m+1} = \sigma^m$ which concludes the proof. \square

We give a simple numerical illustration for the matrix defined in (14) and \mathcal{V} is a Krylov subspace for the matrix A of dimension 150. Again, $\lambda_p = 110$. We started with two different initial shifts: $\sigma^0 = 110$ and $\sigma^0 = 109.75$. In all applications of (21) we choose $i = 1$ except for σ^0 : because we are interested in a small interval in the neighborhood of σ^0 it is wise to select a harmonic Ritz vector that makes a small angle with the eigenvector x_p and this is not necessarily \tilde{u}_1 . So, for σ^0 we have selected a suitable harmonic Ritz vector using a new selection strategy that is discussed in Section 7. The numbers are in Table 1. For both initial shifts σ^0 applying (21) a few times results in almost the same interval.

m	σ^m	Lehmann interval		Bauer–Fike interval	
		interval	size	interval	size
0	110.0000	[110.0000, 518.0980]	408.0980	[109.9964, 110.0036]	0.007171
1	109.9964	[109.9964, 110.0012]	0.004779	[109.9979, 110.0021]	0.004136
2	109.9979	[109.9979, 110.0021]	0.004136	[109.9979, 110.0021]	0.004136
0	109.7500	[109.7500, 110.0000]	0.250020	[109.9979, 110.0021]	0.004136
1	109.9979	[109.9979, 110.0021]	0.004136	[109.9979, 110.0021]	0.004136
2	109.9979	[109.9979, 110.0021]	0.004136	[109.9979, 110.0021]	0.004136

Table 1: The effect of applying (21) for $\sigma^0 = 110$ and $\sigma^0 = 109.75$.

6.3 Discussion

In the previous section we discussed the idea of relocating the shift σ in harmonic Rayleigh–Ritz to the boundary of the Bauer–Fike interval for one of the harmonic Ritz vectors with associated ρ -value. A point of concern is which i in (20) to begin with or how to pick σ^0 . We note that in fact any boundary point of any inclusion interval (e.g. Gershgorin discs, etc) can be used as a starting point for σ . But it is not necessarily true that there is enough information in \mathcal{V} and $A\mathcal{V}$ such that this interval can be reduced. However, it can be useful to start with intervals at different positions if one wants to outline the spectrum of the matrix A .

If we take as starting point a Bauer–Fike interval for one of the harmonic Ritz vectors it is true that this results in intervals that do reduce in size. However, if we are interested in some small inclusion interval in the neighborhood of σ^0 , it is not always a good idea to start with $\tilde{\theta}_1^0$ or $\tilde{\theta}_{-1}^0$ or, equivalently, pick $i = \pm 1$ in (20). This is the problem of the next section, where we want to select the harmonic Ritz vector w.r.t. shift σ that is making a small angle with the eigenvector with eigenvalue closest to σ .

7 The selection of a harmonic Ritz pair

In this section the shift σ for harmonic Ritz is fixed and we are interested in the eigenvector x_p that is associated with the eigenvalue λ_p closest to σ . We want to select an approximation to x_p from the set of all harmonic Ritz vectors \tilde{u}_i with respect to the shift σ . An important practical question is which harmonic Ritz vector is a good candidate.

Although theoretical justification is still incomplete, numerical experiments [13] suggest that there is at least one harmonic Ritz vector, say \tilde{u}_{j_p} , that offers a good approximation

to x_p , whenever the angle between x_p and \mathcal{V} is small. In this section we will propose a new strategy for selecting a \tilde{u}_j . We will argue that the new strategy is more successful in finding the desired \tilde{u}_{j_p} than the obvious strategies that select the harmonic Ritz vector with harmonic Ritz value or with ρ -value closest to σ .

7.1 The selection strategies

Ideally we want to detect the harmonic Ritz vector \tilde{u}_j that makes the smallest angle with the eigenvector x_p . The first strategy is based on the harmonic Ritz values. We take the \tilde{u}_j with $j = \operatorname{argmin}_i |\tilde{\theta}_i - \sigma|$ and call this *harmonic selection*. An alternative (ρ -selection) is to use the ρ -values as indication, i.e., select \tilde{u}_j with $j = \operatorname{argmin}_i |\rho_i - \sigma|$. The third and new alternative that we propose is a mixture of these two strategies (*product selection*): we select \tilde{u}_j with

$$(22) \quad j = \operatorname{argmin}_i \sqrt{(\rho_i - \sigma)(\tilde{\theta}_i - \sigma)}.$$

Note that

$$\sqrt{(\rho_i - \sigma)(\tilde{\theta}_i - \sigma)} = \|A\tilde{u}_i - \sigma\tilde{u}_i\|_2.$$

All three strategies can be applied equally cheaply. It is interesting to observe that because $|\rho_i - \sigma| \leq |\tilde{\theta}_i - \sigma|$

$$|\rho_i - \sigma| \leq \sqrt{(\rho_i - \sigma)(\tilde{\theta}_i - \sigma)} \leq |\tilde{\theta}_i - \sigma|.$$

In order to illustrate some characteristic properties of these three selection methods we applied harmonic Rayleigh–Ritz with shift $\sigma = 0$ to the following matrix A and subspace \mathcal{V} :

$$A = A_\mu \equiv \begin{bmatrix} \mu & 0 & 0 \\ 0 & -2 & 0 \\ 0 & 0 & 1 \end{bmatrix} \text{ and } \mathcal{V} = \mathcal{V}_{\epsilon, \phi} \equiv \operatorname{span} \left\{ \begin{bmatrix} \sqrt{1-\epsilon} \\ \sqrt{\epsilon} \sin(\phi) \\ \sqrt{\epsilon} \cos(\phi) \end{bmatrix}, \begin{bmatrix} 0 \\ -\cos(\phi) \\ \sin(\phi) \end{bmatrix} \right\}.$$

The eigenpair of interest is (μ, e_1) and note that $\sin^2 \angle(\mathcal{V}, e_1) = \epsilon$. The best possible approximation from the subspace $\mathcal{V}_{\epsilon, \phi}$ is denoted by $x_{\mathcal{V}_{\epsilon, \phi}}$. We have applied harmonic Rayleigh–Ritz with zero shift and subspace $\mathcal{V} = \mathcal{V}_{\epsilon, \phi}$, applied the three selection strategies and calculated the maximal angle over ϕ between the selected harmonic Ritz vector and $x_{\mathcal{V}_{\epsilon, \phi}}$ for a fixed ϵ . The results are in Fig. 5 for the matrix A_μ with $\mu = 0$ and $\mu = -0.5$.

For $\mu \neq \sigma = 0$ Theorem 5.1 shows that for ϵ small enough, harmonic selection selects a relevant harmonic Ritz vector. This can be seen from the right figure in Fig. 5. In case $\mu = \sigma = 0$ the graph for harmonic selection equals 1, this means that the selected harmonic Ritz vector can be perpendicular to $x_{\mathcal{V}_{\epsilon, \phi}}$ for some subspaces $\mathcal{V}_{\epsilon, \phi}$ and harmonic selection can perform very poorly.

For ρ -selection the situation is the reverse. If $\mu = \sigma = 0$ it seems that the ρ -values provide useful information about the quality of the harmonic Ritz vector. However, if $\mu \neq \sigma = 0$ an irrelevant harmonic Ritz vector can be selected. In general we expect that ρ -selection is not effective if σ is far from the eigenvalue of interest. Our new selection strategy based on the product $\sqrt{\tilde{\theta}\rho}$ gives a reasonable compromise between the two strategies.

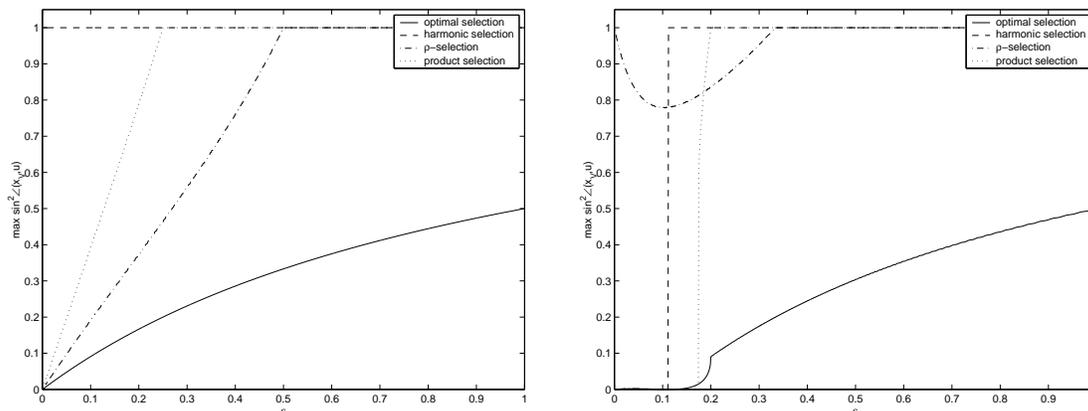


Figure 5: Numerical approximation of the maximal angle over ϕ between \tilde{u}_j and $x_{\mathcal{V}}$ for harmonic Rayleigh–Ritz and the three selection strategies. Also included is the situation of optimal selection. Left picture is for $\mu = 0$, right for $\mu = -0.5$

7.2 Numerical experiments

To illustrate the three selection strategies in the framework of the computation of an eigenvector with an iterative method, we adapt Example 5 from [13].

We search for the eigenvalue closest to 27.0 of the tridiagonal matrix A with 0.2, 0.4, \dots , 58.8, 60.0 on the diagonal and one on the sub- and super-diagonal. In every iteration step the space \mathcal{V} is expanded with a correction v given by the Davidson correction-equation

$$v = (\text{diag}(A) - \sigma I)^{-1}r, \text{ with the residual } r \equiv Au - \rho(u)u.$$

Here u is some approximation to the wanted eigenvector constructed from the subspace. For this construction we have used harmonic Rayleigh–Ritz with shift σ and selected the vector $u = \tilde{u}_j$ using our three selection methods. Furthermore, we computed the results when u was taken the refined Ritz vector \hat{u} from (6) with shift $\xi = \sigma$. Fig. 6 shows the convergence history for two values of σ .

From this picture the best strategy to use seems refined Rayleigh–Ritz due to faster convergence in the initial phase. For $\sigma = 27.05$ the figure shows a very irregular behavior for ρ -selection, the other two selection strategies perform equally well. When σ is decreased to $\sigma = 27.0001$ the convergence for ρ -selection is still irregular but in the end not slower than for selection with (22). In the setting discussed here, irregular convergence is not really a problem, but if restarts are performed then a restart at a peak in the convergence curve may be fatal. With harmonic selection the process converges to a different eigenpair, in contrast to the other strategies.

Now consider the case where the eigenvalue closest to the left of σ is to be computed, in other words the λ_i for which $1/(\sigma - \lambda_i)$ is maximal. It is not immediately clear how to adapt the refined Ritz method with fixed shift for this situation. Harmonic selection is changed to selecting \tilde{u}_j with $j = \text{argmax}_i 1/(\sigma - \tilde{\theta}_i)$. Similarly, for ρ -selection we pick \tilde{u}_j with $j = \text{argmax}_i 1/(\sigma - \rho_i)$. Our new strategy becomes, choosing \tilde{u}_j with

$$(23) \quad j = \text{argmax}_i \frac{\text{sign}(\sigma - \tilde{\theta}_i)}{\sqrt{(\rho_i - \sigma)(\tilde{\theta}_i - \sigma)}}.$$

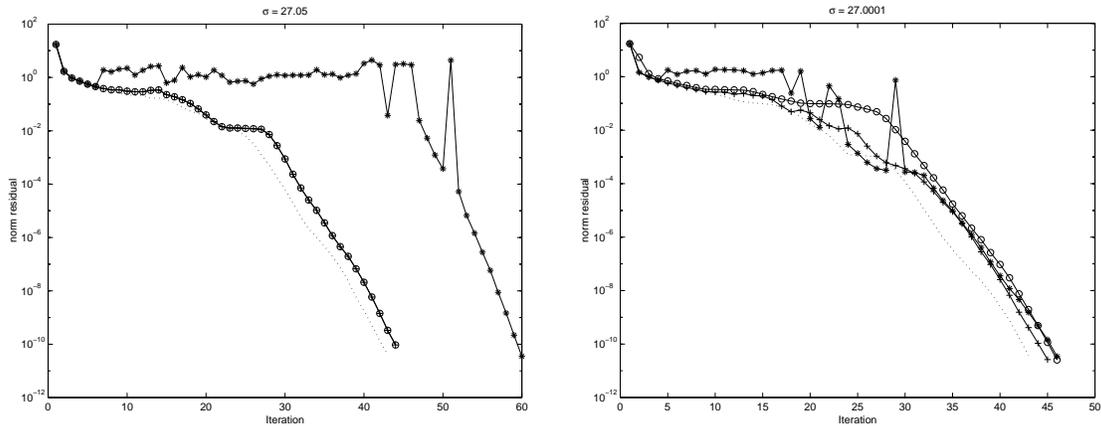


Figure 6: Finding the eigenpair with eigenvalue closest to σ with refined Ritz with shift σ (..) and harmonic Rayleigh–Ritz using harmonic selection (o), ρ -selection (*) and selection with (22) (+). Harmonic selection in the right picture finds the eigenvalue ≈ 27.2 , which is not the eigenvalue closest to σ

Fig. 7 gives the convergence history for the new situation. Again, the convergence for ρ -selection is quite irregular. Harmonic selection works in both situations, but again, if the shift is chosen any closer to 27 the method finds the eigenvalue ≈ 26.8 (not shown here). The convergence for condition (23) is again smooth and robust in this situation.

We conclude with the remark that (22) and (23) can also be used for non-normal problems.

Acknowledgments. The authors are thankful to the referee for his or her remarks that improved our presentation. The authors also thank Jan Brandts for helpful comments.

The research of the second author was financially supported by the Dutch scientific organization (NWO), under project number 613.002.035.

References

- [1] Christopher Beattie, *Harmonic Ritz and Lehmann bounds*, Electron. Trans. Numer. Anal. **7** (1998), 18–39 (electronic), Large scale eigenvalue problems (Argonne, IL, 1997). MR 1665476 [2](#), [5](#)
- [2] F. Chatelin, *Spectral approximation of linear operators*, Academic Press, New York, USA, 1983. [7](#), [9](#)
- [3] Lothar Collatz, *Funktionalanalysis und numerische Mathematik*, Springer-Verlag, Berlin, 1964. MR 29 # 2931 [4](#), [7](#)
- [4] Ernest R. Davidson, *The iterative calculation of a few of the lowest eigenvalues and corresponding eigenvectors of large real-symmetric matrices*, J. Comput. Phys. **17** (1975), 87–94. MR 52:2168 [1](#)

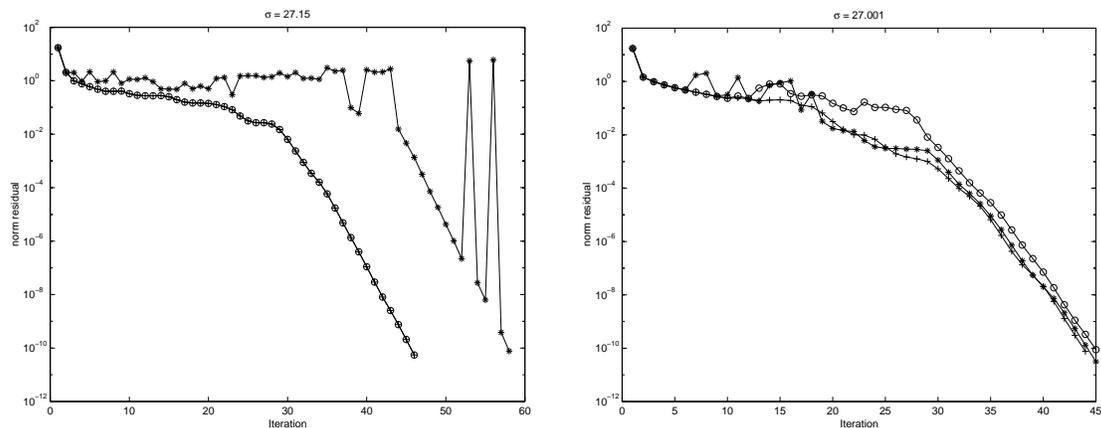


Figure 7: Finding the eigenpair with the largest eigenvalue that is less than σ with harmonic Rayleigh–Ritz using harmonic selection (o), ρ -selection (*) and selection with (23) (+). In all cases convergence is towards the desired eigenvalue.

- [5] Roland W. Freund, *Quasi-kernel polynomials and their use in non-Hermitian matrix iterations*, J. Comput. Appl. Math. **43** (1992), no. 1-2, 135–158. MR 94a:65018 [2](#)
- [6] Zhongxiao Jia, *A refined iterative algorithm based on the block Arnoldi process for large unsymmetric eigenproblems*, Linear Algebra Appl. **270** (1998), 171–189. MR 98m:65055 [2, 7](#)
- [7] ———, *Polynomial characterizations of the approximate eigenvectors by the refined Arnoldi method and an implicitly restarted refined Arnoldi algorithm*, Linear Algebra Appl. **287** (1999), no. 1-3, 191–214, Special issue celebrating the 60th birthday of Ludwig Elsner. MR 99j:65046 [8](#)
- [8] Zhongxiao Jia and G. W. Stewart, *An analysis of the Rayleigh-Ritz method for approximating eigenspaces*, Math. Comp. **70** (2001), no. 234, 637–647. MR 2001g:65040 [2, 3](#)
- [9] Shmuel Kaniel, *Estimates for some computational techniques in linear algebra*, Math. Comp. **20** (1966), 369–378. MR 38 # 2934 [11](#)
- [10] N. J. Lehmann, *Optimale Eigenwerteinschliessungen*, Numer. Math. **5** (1963), 246–272. MR 27 # 5359 [2, 4, 6, 8](#)
- [11] ———, *Zur Verwendung optimaler Eigenwerteingrenzungen bei der Lösung symmetrischer Matrizenaufgaben*, Numer. Math. **8** (1966), 42–55. MR 33 # 6820 [2](#)
- [12] Thomas A. Manteuffel and James S. Otto, *On the roots of the orthogonal polynomials and residual polynomials associated with a conjugate gradient method*, Numer. Linear Algebra Appl. **1** (1994), no. 5, 449–475. MR 95i:65051 [2](#)
- [13] Ronald B. Morgan, *Computing interior eigenvalues of large matrices*, Linear Algebra Appl. **154/156** (1991), 289–309. MR 92e:65050 [2, 3, 6, 7, 10, 11, 12, 16, 18](#)

- [14] Ronald B. Morgan and Min Zeng, *Harmonic projection methods for large non-symmetric eigenvalue problems*, Numer. Linear Algebra Appl. **5** (1998), no. 1, 33–55. MR 99g:65046 [6](#)
- [15] Chris C. Paige, Beresford N. Parlett, and Henk A. van der Vorst, *Approximate solutions and eigenvalue bounds from Krylov subspaces*, Numer. Linear Algebra Appl. **2** (1995), no. 2, 115–133. MR 96a:65054 [2](#), [5](#), [6](#)
- [16] Beresford N. Parlett, *The symmetric eigenvalue problem*, Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 1998, Corrected reprint of the 1980 original. MR 99c:65072 [1](#), [2](#), [3](#), [4](#), [5](#), [6](#), [8](#), [13](#), [14](#)
- [17] D. S. Scott, *The advantages of inverted operators in Rayleigh-Ritz approximations*, SIAM J. Sci. Statist. Comput. **3** (1982), no. 1, 68–75. [2](#), [3](#)
- [18] Gerard L. G. Sleijpen, Jasper van den Eshof, and Paul Smit, *Optimal a priori error bounds for the Rayleigh-Ritz method*, Preprint 1160, Dep. Math., University Utrecht, Utrecht, the Netherlands, September 2000. [8](#), [12](#)
- [19] Gerard L. G. Sleijpen and Henk A. van der Vorst, *A Jacobi-Davidson iteration method for linear eigenvalue problems*, SIAM J. Matrix Anal. Appl. **17** (1996), no. 2, 401–425. MR 96m:65042 [1](#), [3](#)
- [20] Gerard L. G. Sleijpen, Henk A. van der Vorst, and Ellen Meijerink, *Efficient expansion of subspaces in the Jacobi-Davidson method for standard and generalized eigenproblems*, Electron. Trans. Numer. Anal. **7** (1998), 75–89 (electronic), Large scale eigenvalue problems (Argonne, IL, 1997). MR 99i:65040 [6](#)

Contents

1	Introduction	1
2	Harmonic Rayleigh–Ritz	2
3	Some useful properties of harmonic Rayleigh–Ritz	4
3.1	A minmax characterization for harmonic Ritz values	4
3.2	Optimal inclusion intervals for eigenvalues	6
3.3	The concept of ρ -values	6
4	A comparison with refined Rayleigh–Ritz	7
4.1	Refined Rayleigh–Ritz and its relation to harmonic Rayleigh–Ritz	7
4.2	The optimal value of ξ in refined Rayleigh–Ritz	8
4.3	The optimal value of σ in harmonic Rayleigh–Ritz	9
4.4	Discussion	10
5	A priori error estimation	11
6	A posteriori error estimation	13
6.1	A condition for the minimizing shift	13
6.2	A condition for the shift σ^{m+1}	14
6.3	Discussion	16
7	The selection of a harmonic Ritz pair	16
7.1	The selection strategies	17
7.2	Numerical experiments	18