

HARMONIC AND REFINED RAYLEIGH–RITZ FOR THE POLYNOMIAL EIGENVALUE PROBLEM

MICHEL E. HOCHSTENBACH* AND GERARD L. G. SLEIJPEN†

Abstract. After reviewing the harmonic Rayleigh–Ritz approach for the standard and generalized eigenvalue problem, we discuss different extraction processes for subspace methods for the polynomial eigenvalue problem. We generalize the harmonic and refined Rayleigh–Ritz approach which leads to novel approaches to extract promising approximate eigenpairs from a search space. We give theoretical as well as numerical results of the methods. In addition, we take a new look at the Jacobi–Davidson method for polynomial eigenvalue problems.

AMS subject classifications. 65F15, 65F50.

Key words. Polynomial eigenvalue problem, harmonic Rayleigh–Ritz, refined Rayleigh–Ritz, interior eigenvalues, Jacobi–Davidson, subspace extraction, subspace method, subspace expansion, Rayleigh–Ritz.

1. Introduction: Rayleigh–Ritz for the standard eigenvalue problem. Let A be a (large sparse) $n \times n$ matrix, and let \mathcal{U} be a k -dimensional search space for a sought eigenvalue λ corresponding to an eigenvalue λ of A . The search space could for instance be generated by the Arnoldi method (see, e.g., [1, Ch. VII]) or Jacobi–Davidson [2] method; typically $k \ll n$. Let the columns of the $n \times k$ search matrix U form an orthonormal basis for \mathcal{U} . The Rayleigh–Ritz extraction determines approximate eigenpairs (θ, u) with $u \in \mathcal{U}$; therefore, we can write $u = Uc$ for a small k -vector c . The Galerkin condition on the residual

$$r := Au - \theta u \perp \mathcal{U}, \quad u = Uc \in \mathcal{U},$$

leads to the projected eigenproblem

$$U^*AUc = \theta c.$$

The Ritz value θ is the Rayleigh quotient of u :

$$\theta(u) := \frac{u^*Au}{u^*u}.$$

Here c is called a primitive Ritz vector (in [3]), and $u = Uc$ is a Ritz vector.

We continue to show that the Rayleigh–Ritz method works well for the extraction of well-separated exterior eigenvalues of Hermitian matrices, but may fail for clustered or interior eigenvalues, or for non-Hermitian matrices. The following claim is well known.

CLAIM 1.1. *Rayleigh–Ritz works well for the extraction of well-separated exterior eigenvalues of Hermitian matrices.*

By this we mean that if (θ, u) is a Ritz pair of a Hermitian matrix, and if $\theta \approx \lambda$, where λ is either the leftmost or the rightmost well-separated eigenvalue, then necessarily $u \approx x$. The reason for this is the following. Let $\lambda_1 < \dots < \lambda_n$ be the eigenvalues of a Hermitian matrix A , with corresponding eigenvectors x_1, \dots, x_n . Suppose that $u \approx x$ is an approximate eigenvector. Decomposing

$$u = \sum_{i=1}^n \gamma_i x_i,$$

*Version December 2005. Department of Mathematics, Case Western Reserve University, 10900 Euclid Avenue, Cleveland, OH 44106-7058, USA, hochsten@case.edu, www.case.edu/artsci/math/hochstenbach, phone +1-216-368-5015, fax +1-216-368-5163. The research of this author was supported in part by NSF grant DMS-0405387.

†Mathematical Institute, Utrecht University, P.O. Box 80.010, NL-3508 TA Utrecht, The Netherlands, sleijpen@math.uu.nl, www.math.uu.nl/people/sleijpen.

we have $\theta = \sum_{i=1}^n |\gamma_i|^2 \lambda_i$. If $\theta \approx \lambda_1$ or $\theta \approx \lambda_n$, then we must have $\gamma_1 \approx 1$ or $\gamma_n \approx 1$ (when the eigenvectors are scaled appropriately), from which we conclude $u \approx x_1$ or $u \approx x_n$, respectively. This justifies Claim 1.1.

When the exterior eigenvalues of a Hermitian matrix are *clustered*, there is no guarantee that u is close to any particular eigenvector, but (using a similar argument as before) we do know that u is almost in the span of the relevant eigenvectors. In this case, block methods may be advantageous.

Rayleigh–Ritz does not need to work well for the extraction of *interior eigenvalues of Hermitian matrices*, as the example in [3, p. 282] shows. Claim 1.1 can also fail for *exterior eigenvalues of non-Hermitian A* , as the following example, inspired by [4], shows. Let for small $\varepsilon, \varphi > 0$ and $\alpha, \beta > 0$

$$A = \text{diag}(0, -\beta, -\varepsilon + \alpha i, -\varepsilon - \alpha i), \quad U = [\cos(\varphi)e_1 + \sin(\varphi)e_2 \quad (e_3 + e_4)/\sqrt{2}],$$

where e_i denotes the i th canonical basis vector. Then, although $\angle(U, e_1) = \varphi$, that is, the search space makes a small angle with the eigenvector of an exterior (the rightmost) eigenvalue, the Rayleigh–Ritz method will abusively select the Ritz pair $(-\varepsilon, (e_3 + e_4)/\sqrt{2})$ as an approximation to the rightmost eigenpair, as long as $-\beta \sin^2(\varphi) < -\varepsilon$.

On the other hand, Van den Eshof [5, Th. 2.2.3] has shown that if an eigenvalue is outside the field of values generated by the space orthogonal to the corresponding eigenvector—as is for instance the case in the last example—then the Rayleigh–Ritz approach can *asymptotically* (that is, close enough to the solution) be used safely. For the example discussed, this means that for φ small enough the corresponding Ritz vector is a good approximation to the true eigenvector. Also, when the field of values condition fails to hold, Rayleigh–Ritz in practice often performs much better for (relatively well-separated) exterior eigenvalues than for interior eigenvalues of non-Hermitian matrices.

In an attempt to better approximate interior eigenpairs, the harmonic Rayleigh–Ritz approach was introduced in [6]. The rest of this paper has been organized as follows. In Section 2, we review the harmonic Rayleigh–Ritz approach for the standard and generalized eigenvalue problem, adding a few novel viewpoints. Sections 3, 4, and 5 generalize standard Rayleigh–Ritz, harmonic Rayleigh–Ritz, and refined Rayleigh–Ritz for the polynomial eigenvalue problem, where our goal is to find promising approximate eigenpairs corresponding to (interior) eigenvalues from a search space. In Section 6, we revisit the Jacobi–Davidson subspace expansion. Finally, we give numerical experiments and a conclusion in Sections 7 and 8.

2. Harmonic Rayleigh–Ritz. Suppose we are interested in eigenvalues near a target τ , for instance in the interior of the spectrum. We assume that τ is not an eigenvalue of A (otherwise we are fortunate). Denote the identity matrix by I . Since

$$(A - \tau I)^{-1}x = (\lambda - \tau)^{-1}x,$$

we see that eigenvalues of A near τ are exterior eigenvalues of $(A - \tau I)^{-1}$, a desirable property in view of the previous section. Consider the (Petrov–)Galerkin condition on the residual of the shift-and-inverse matrix

$$(2.1) \quad (A - \tau I)^{-1}\tilde{u} - (\tilde{\theta} - \tau)^{-1}\tilde{u} \perp \mathcal{V}, \quad \tilde{u} = U\tilde{c},$$

where \mathcal{U} is the search space and \mathcal{V} is the *test space*. For large sparse matrices, working with the inverse $(A - \tau I)^{-1}$ is often highly undesirable. To eliminate the inverse in the Galerkin condition, we want to choose \mathcal{V} in a suitable way. While $\mathcal{V} = (A - \tau I)^*\mathcal{U}$ gives rise to the standard Rayleigh–Ritz extraction $Au - \theta u \perp \mathcal{U}$, the choice $\mathcal{V} = (A - \tau I)^*(A - \tau I)\mathcal{U}$ leads to the following equivalent conditions:

- (i) $(A - \tau I)^{-1}\tilde{u} - (\tilde{\theta} - \tau)^{-1}\tilde{u} \perp (A - \tau I)^*(A - \tau I)\mathcal{U}$,
- (ii) $(A - \tau I)^{-1}\tilde{u} - (\tilde{\theta} - \tau)^{-1}\tilde{u} \perp_{(A - \tau I)^*(A - \tau I)} \mathcal{U}$,
- (iii) $(A - \tilde{\theta} I)\tilde{u} \perp (A - \tau I)\mathcal{U}$,
- (iv) $U^*(A - \tau I)^*(A - \tau I)U\tilde{c} = (\tilde{\theta} - \tau)U^*(A - \tau I)^*U\tilde{c}$;

this is called the *harmonic Rayleigh–Ritz* approach. In analogy to the Rayleigh–Ritz terminology, $\tilde{\theta}$ is a *harmonic Ritz value*, the *harmonic Rayleigh quotient* or the *Temple quotient* of a *harmonic Ritz vector* \tilde{u} :

$$(2.2) \quad \tilde{\theta} = \tilde{\theta}_\tau(\tilde{u}) := \tau + \frac{\|(A - \tau I)\tilde{u}\|^2}{\tilde{u}^*(A - \tau I)^*\tilde{u}}.$$

(In this paper, $\|\cdot\|$ stands for the Euclidean norm.) Characterization (ii) gives a Galerkin condition on the residual of the shift-and-inverse matrix using the inner product

$$(a, b)_{(A - \tau I)^*(A - \tau I)} := b^*(A - \tau I)^*(A - \tau I)a,$$

where the situation that $(a, b)_{(A - \tau I)^*(A - \tau I)} = 0$ is denoted by $a \perp_{(A - \tau I)^*(A - \tau I)} b$. Item (iii) states that the (ordinary) residual should be orthogonal to a test space different from the search space. Item (iv) expresses that a harmonic Ritz pair $(\tilde{\theta}, \tilde{u})$ is closely connected to an eigenpair of a projected generalized eigenvalue problem. It is natural to call \tilde{c} a *primitive harmonic Ritz vector*.

Denote $\xi = \tilde{\theta} - \tau$. As noted in [3, p. 293], by left-multiplying

$$(2.3) \quad U^*(A - \tau I)^*(A - \tau I)U\tilde{c} = \xi U^*(A - \tau I)^*U\tilde{c}$$

by \tilde{c}^* and using Cauchy–Schwarz, we see that a solution (ξ, \tilde{c}) of this generalized eigenproblem satisfies (using $\tilde{u} = U\tilde{c}$)

$$(2.4) \quad \|(A - \tau I)\tilde{u}\| \leq |\xi|.$$

This makes clear that if $|\xi|$ is small, i.e., if the harmonic Ritz value $\tilde{\theta}$ is close to τ , then the pair (τ, \tilde{u}) has a small residual norm and therefore forms a good approximation to an eigenpair of A , as the following well-known results show. Part (a) is Bauer–Fike. For (b), we recall (see [3, p. 44]) that if λ is a simple eigenvalue, we have a spectral decomposition $A = \lambda xy^* + X_1LY_1^*$, where $\|x\| = 1$, Y_1 is orthonormal, and $Y_1^*x = 0$.

THEOREM 2.1. (See e.g., [7, Th. 3.6], [3, p. 288])

(a) If $A = X\Lambda X^{-1}$ is diagonalizable, there exists an eigenvalue λ such that

$$|\lambda - \tau| \leq \kappa(X) \|(A - \tau I)\tilde{u}\|,$$

where $\kappa(X) = \|X\| \|X^{-1}\|$ is the condition number of X .

(b) If λ is a simple eigenvalue and $A = \lambda xy^* + X_1LY_1^*$, where $\|x\| = 1$, Y_1 is orthonormal, and $Y_1^*x = 0$, then

$$\sin(\tilde{u}, x) \leq \frac{\|(A - \tau I)\tilde{u}\|}{\sigma_{\min}(L - \tau I)} \leq \frac{\|(A - \tau I)\tilde{u}\|}{\sigma_{\min}(L - \lambda I) - |\lambda - \tau|}.$$

We remark that the quantity $\sigma_{\min}(L - \tau I)$ is also known in literature as $\text{sep}(\tau, L)$; it is nonzero since τ is not an eigenvalue of A . Since we assumed that λ is simple, also $\sigma_{\min}(L - \lambda I) > 0$, and therefore $\tau \rightarrow \lambda$ and $\|(A - \tau I)\tilde{u}\| \rightarrow 0$ are sufficient to ensure that $\tilde{u} \rightarrow x$.

The conclusion is that harmonic Rayleigh–Ritz avoids so-called “ghost eigenvalues”, also named “imposters” or “impersonators” in [3]: if there is a harmonic Ritz pair $(\tilde{\theta}, \tilde{u})$ with $\tilde{\theta}$

close to τ , the corresponding residual norm $\|(A - \tau I) \tilde{u}\|$ (and to a lesser extent $\|(A - \tilde{\theta} I) \tilde{u}\|$) is necessarily small, and hence \tilde{u} should be a good approximate eigenvector, and not some linear combination of irrelevant eigenvectors as in the example in Section 1.

The next important observation is that the harmonic Rayleigh–Ritz extraction is primarily meant to get a decent approximate eigenvector. It turns out that the harmonic Ritz value can be quite inaccurate (see for instance [4]). Indeed, given the approximate eigenvector \tilde{u} , the Rayleigh quotient $\rho(\tilde{u})$ is the best approximation to the eigenvalue, in the sense that it minimizes the residual norm:

$$\rho(\tilde{u}) = \underset{\theta}{\operatorname{argmin}} \|A\tilde{u} - \theta\tilde{u}\|.$$

Therefore, it is often advisable to take the Rayleigh quotient of the harmonic Ritz vector as an approximate eigenvalue, instead of the harmonic Ritz value.

For the generalized eigenproblem $Ax = \lambda Bx$, we can use a similar harmonic Rayleigh–Ritz approach. Based on $(A - \tau B)^{-1}Bx = (\lambda - \tau)^{-1}x$, we impose the Galerkin condition (cf. (2.1))

$$(A - \tau B)^{-1}B\tilde{u} - (\tilde{\theta} - \tau)^{-1}\tilde{u} \perp \mathcal{V}, \quad \tilde{u} = \tilde{U}\tilde{c}.$$

With $\mathcal{V} = (A - \tau B)^*(A - \tau B)\mathcal{U}$, the following are equivalent:

- (i) $(A - \tau B)^{-1}B\tilde{u} - (\tilde{\theta} - \tau)^{-1}\tilde{u} \perp (A - \tau B)^*(A - \tau B)\mathcal{U}$,
- (ii) $(A - \tau B)^{-1}B\tilde{u} - (\tilde{\theta} - \tau)^{-1}\tilde{u} \perp_{(A - \tau B)^*(A - \tau B)} \mathcal{U}$,
- (iii) $(A - \tilde{\theta} B)\tilde{u} \perp (A - \tau B)\mathcal{U}$,
- (iv) $U^*(A - \tau B)^*(A - \tau B)U\tilde{c} = (\tilde{\theta} - \tau)U^*(A - \tau B)^*BU\tilde{c}$;

Therefore, we are interested in the eigenpair(s) (ξ, \tilde{c}) of the projected generalized eigenproblem

$$(2.5) \quad U^*(A - \tau B)^*(A - \tau B)U\tilde{c} = \xi U^*(A - \tau B)^*BU\tilde{c},$$

for which ξ , playing the role of $\tilde{\theta} - \tau$, should be of minimal magnitude. Similar to (2.4), this gives

$$\|(A - \tau B)\tilde{u}\| \leq |\xi| \|B\tilde{u}\| \leq |\xi| \|BU\|,$$

which, again, “discourages imposters” [3, p. 296].

A justification of the harmonic approach for both the standard and generalized eigenvalue problem is that eigenvectors that are present in the search space are detected indeed, see the following proposition, which is a straightforward generalization of [3, p. 293].

PROPOSITION 2.2. *Let (λ, x) be an eigenpair of $Ax = \lambda Bx$, then it is also a harmonic Ritz pair with respect to the search space $\operatorname{span}(x)$.*

Proof. This follows for instance from substituting (λ, x) in characterization (iii). \square

We note that if x is part of a larger subspace \mathcal{U} the harmonic approach may still have difficulties selecting the right vector if there are (nearly) equal harmonic Ritz values, see [8]. This, however, rarely forms a problem in practical processes, since we can just continue the subspace method by expanding the search space and performing a new extraction process.

3. The polynomial eigenvalue problem. Now we switch our attention to the polynomial eigenvalue problem

$$(3.1) \quad p(\lambda)x = (\lambda^l A_l + \lambda^{l-1} A_{l-1} + \cdots + \lambda A_1 + A_0)x = 0,$$

where all matrices have size $n \times n$. In this section, we summarize some known extraction techniques: (standard) Rayleigh–Ritz and two-sided Rayleigh–Ritz.

3.1. Rayleigh–Ritz for the polynomial eigenvalue problem. Given an approximate eigenvalue $\theta \approx \lambda$ and an approximate eigenvector $u \approx x$, the *residual* r is defined by

$$r(\theta, u) := p(\theta) u = (\theta^l A_l + \theta^{l-1} A_{l-1} + \cdots + A_0) u.$$

Then the Rayleigh–Ritz approach for the polynomial eigenvalue problem imposes the Galerkin condition

$$p(\theta) u \perp \mathcal{U},$$

with $u = Uc$ leading to the projected polynomial eigenproblem

$$(\theta^l U^* A_l U + \theta^{l-1} U^* A_{l-1} U + \cdots + U^* A_0 U) c = 0,$$

which, as a low-dimensional problem, may for instance be solved using techniques employing a companion matrix. For alternative extraction methods, we consider a Petrov–Galerkin condition

$$p(\theta) u \perp \mathcal{V},$$

where $\mathcal{V} \neq \mathcal{U}$ is a suitably chosen test space. For every *test vector* $v \in \mathcal{V}$ this implies $p(\theta) u \perp v$. We will examine one option in the next subsection, and come back to the subject in Sections 4 and 6.

3.2. Two-sided Rayleigh–Ritz. We now argue that a search space for the left eigenvector is the best possible test space \mathcal{V} in the following sense. Given any nonzero $v \in \mathcal{V}$, the equation $v^* p(\theta) u = 0$ defines $\theta = \theta(u, v)$ implicitly as a function of u and v . When $v^* p'(\lambda) x \neq 0$, then by the Implicit Function Theorem we get [9]

$$(3.2) \quad \frac{\partial \theta}{\partial u}(x, v) = -\frac{v^* p(\lambda)}{v^* p'(\lambda) x}.$$

So when we choose $v = y$, the left eigenvector of (3.1) satisfying $y^* p(\lambda) = 0$, and if $y^* p'(\lambda) x \neq 0$, then θ is stationary in the eigenvector point (x, y) with respect to changes in x (and y).

In practice we will not have the exact left eigenvector to our disposal, but we may try to approximate both the left and the right eigenvector in the process, this is the essence of two-sided subspace methods, as for instance two-sided Lanczos [10], [1, Ch. VIII] or two-sided Jacobi–Davidson [11, 12, 9]. Equation (3.2) implies that if u is a first order accurate approximation to the right eigenvector (that is, $u = x + d$, with $\|d\|$ small), and v is a first order accurate approximation to the left eigenvector ($v = y + e$, with $\|e\|$ small), then $\theta(u, v)$ is a second order accurate approximation to the eigenvalue ($|\theta - \lambda| = \mathcal{O}(\|d\| \|e\|)$). However, in the rest of the paper we will not consider two-sided methods and assume that we have no information about the left eigenvector at hand.

4. Harmonic Rayleigh–Ritz for the polynomial eigenvalue problem. Assume that we are interested in (interior) eigenpairs (λ, x) of (3.1), where λ is close to a known target τ , and where $p(\tau)$ is nonsingular (otherwise τ is an eigenvalue). The aim of this section is to generalize the harmonic Rayleigh–Ritz approach to the polynomial eigenproblem. First, we use a simple Taylor expansion which we will write out for completeness: we have

$$\begin{aligned} p(\tau) x &= (p(\tau) - p(\lambda)) x \\ &= \left(\int_0^1 \frac{d}{d\alpha} p(\alpha(\tau - \lambda) + \lambda) d\alpha \right) x \\ &= (\tau - \lambda) \left(\int_0^1 p'(\alpha(\tau - \lambda) + \lambda) d\alpha \right) x. \end{aligned}$$

So

$$(4.1) \quad \|p(\tau)x\| \leq \mu_1 |\lambda - \tau|,$$

where

$$(4.2) \quad \mu_1 = \max_{t \in [\tau, \lambda]} \|p'(t)\| := \max_{t \in [0, 1]} \|p'(t(\lambda - \tau) + \tau)\|.$$

Moreover,

$$\begin{aligned} (p(\tau) - (\tau - \lambda)p'(\tau))x &= (\tau - \lambda) \left(\int_0^1 (p'(\alpha(\tau - \lambda) + \lambda) - p'(\tau)) d\alpha \right) x \\ &= (\tau - \lambda) \left(\int_0^1 \int_0^1 \frac{d}{d\beta} p'(\beta(\alpha(\tau - \lambda) + \lambda - \tau) + \tau) d\beta d\alpha \right) x \\ &= (\tau - \lambda)^2 \left(\int_0^1 (\alpha - 1) \int_0^1 p''(\beta(\alpha - 1)(\tau - \lambda) + \tau) d\beta d\alpha \right) x. \end{aligned}$$

From this it follows that

$$\|p(\tau)x - (\tau - \lambda)p'(\tau)x\| \leq \frac{1}{2} \mu_2 |\lambda - \tau|^2,$$

where

$$\mu_2 = \max_{t \in [\tau, \lambda]} \|p''(t)\| := \max_{t \in [0, 1]} \|p''(t(\lambda - \tau) + \tau)\|.$$

Summarizing, we have

$$(4.3) \quad p(\tau)x = (\tau - \lambda)p'(\tau)x + Fx, \quad \|F\| = \mathcal{O}(|\tau - \lambda|^2).$$

Neglecting the $\mathcal{O}(|\tau - \lambda|^2)$ term for the moment, this yields

$$p(\tau)^{-1}p'(\tau)x \approx (\tau - \lambda)^{-1}x.$$

Therefore, when we look for eigenvalues close to τ , we can also search for eigenvalues with large magnitude of $p(\tau)^{-1}p'(\tau)$. In the (Petrov-)Galerkin equation

$$p(\tau)^{-1}p'(\tau)\tilde{u} - (\tau - \tilde{\theta})^{-1}\tilde{u} \perp \mathcal{V}, \quad \tilde{u} = U\tilde{c},$$

we want to choose \mathcal{V} , such that it is unnecessary to work with the inverse of $p(\tau)$. With the choice $\mathcal{V} = p(\tau)\mathcal{U}$ we get back (when including again the second order term Fu) the standard Rayleigh–Ritz condition $p(\theta)u \perp \mathcal{U}$.

Therefore, also inspired by Section 2, we take $\mathcal{V} = p(\tau)^*p(\tau)\mathcal{U}$. This leads to the following four equivalent conditions:

- (i) $p^{-1}(\tau)p'(\tau)\tilde{u} - (\tau - \tilde{\theta})^{-1}\tilde{u} \perp p(\tau)^*p(\tau)\mathcal{U},$
- (ii) $p^{-1}(\tau)p'(\tau)\tilde{u} - (\tau - \tilde{\theta})^{-1}\tilde{u} \perp_{p(\tau)^*p(\tau)} \mathcal{U},$
- (iii) $p(\tau)\tilde{u} - (\tau - \tilde{\theta})p'(\tau)\tilde{u} \perp p(\tau)\mathcal{U},$
- (iv) $U^*p(\tau)^*p(\tau)U\tilde{c} = (\tau - \tilde{\theta})U^*p(\tau)^*p'(\tau)U\tilde{c}.$

Any pair (ξ, \tilde{c}) solving

$$(4.4) \quad U^*p(\tau)^*p(\tau)U\tilde{c} = \xi U^*p(\tau)^*p'(\tau)U\tilde{c}$$

satisfies (left-multiply by \tilde{c}^* and use Cauchy–Schwarz)

$$(4.5) \quad \|p(\tau)\tilde{u}\| \leq |\xi| \|p'(\tau)\tilde{u}\| \leq |\xi| \|p'(\tau)U\|.$$

Here, ξ plays the role of $\tau - \tilde{\theta}$. Therefore, to ensure a minimal residual norm, we are interested in the eigenpair of (4.4) with the smallest $|\xi|$; we call this approach the *linearized harmonic Rayleigh–Ritz method for polynomial eigenproblems*. The word “linearized” reflects the fact that we take the linear part of the Taylor series in (4.3). The vector $\tilde{u} = U\tilde{c}$ is called a *linearized harmonic Ritz vector*. Note that the *linearized harmonic Ritz value* $\tilde{\theta}$ satisfies (cf. (2.2))

$$\tilde{\theta} = \tau - \frac{\|p(\tau)\tilde{u}\|^2}{\tilde{u}^*p(\tau)^*p'(\tau)\tilde{u}}.$$

For the practical computation of the linearized harmonic Ritz pair, we proceed as follows. During the process, we compute the QR decomposition $p(\tau)U = QR$ incrementally (that is, in every step add an extra column to Q and R) and hence cheaply. Then condition (iv) is equivalent to the low-dimensional projected generalized eigenproblem

$$Q^*p'(\tau)U\tilde{c} = \xi^{-1}R\tilde{c}$$

where we are interested in the maximal $|\xi^{-1}|$; see [3, p. 294] for a similar formulation for the standard eigenvalue problem. As the harmonic Ritz value $\tilde{\theta}$ may be a dissatisfactory approximation to the sought eigenvalue (see [4] for the standard eigenvalue problem), we may subsequently determine a new approximate eigenvalue from the derived linearized harmonic Ritz vector; see [13] for a number of possible alternatives to do this.

However, in the derivation of (4.4) we have neglected the second-order term involving F . Reintroducing this term in characterization (iii) above, we get $p(\tau)\tilde{u} - (\tau - \tilde{\theta})p'(\tau)\tilde{u} - F\tilde{u} \perp p(\tau)\mathcal{U}$, and this is exactly the Galerkin condition

$$(4.6) \quad p(\tilde{\theta})\tilde{u} \perp p(\tau)\mathcal{U},$$

(cf. the two characterizations (iii) in Section 2) where we are interested in the $\tilde{\theta}$ closest to τ . We will call this the *harmonic Rayleigh–Ritz method for polynomial eigenproblems*.

Both the linearized harmonic method (4.4) and the harmonic method (4.6) reduce to (2.3) and (2.5) in case of the standard and generalized eigenvalue problem, respectively. However, the following result indicates that the non-linearized version may be more promising than the linearized one.

PROPOSITION 4.1. *If (λ, x) is an eigenpair of $p(\lambda)x = 0$ then it is also a harmonic Ritz pair with respect to the search space $\text{span}(x)$.*

Proof. This follows directly from substitution in (4.6). \square

Therefore, if an eigenvector is in the search space, the harmonic Rayleigh–Ritz method will find it—unless we are in an unlucky (but not serious) circumstance, see the remarks after Proposition 2.2. The linearized harmonic Rayleigh–Ritz method, however, will in general miss the present exact eigenvector in the search space, as may easily be checked.

Of course, we are interested in the quality of the approximate eigenpair. Denote a perturbation of the polynomial $p(\theta)$ by $\Delta p(\theta) := \theta^l \Delta A_l + \theta^{l-1} \Delta A_{l-1} + \dots + \Delta A_0$. The backward error $\zeta(\theta, u)$ of an approximate eigenpair (θ, u) can be defined as (see Tisseur [14])

$$\zeta(\theta, u) := \min\{\varepsilon : (p(\theta) + \Delta p(\theta))u = 0, \|\Delta A_j\| \leq \varepsilon \|E_j\|, j = 1, \dots, l\},$$

where the E_j are error matrices; in practice one often takes either $E_j = I$ (corresponding to absolute perturbations) or $E_j = A_j$ (corresponding to relative perturbations). The backward error $\zeta(\theta)$ of an approximate eigenvalue θ can be defined as

$$\zeta(\theta) := \min_{\|u\|=1} \zeta(\theta, u).$$

Then we have [14]

$$\zeta(\theta, u) = \|r(\theta, u)\| \left/ \sum_{j=1}^l \theta^j \|E_j\| \right.,$$

and

$$\zeta(\theta) = \sigma_{\min}(p(\theta)) \left/ \sum_{j=1}^l \theta^j \|E_j\| \right..$$

The pseudospectrum for the PEP can be defined as in Tisseur and Higham [15]

$$\Lambda_\varepsilon(p) := \{z \in \mathbb{C} \mid \zeta(z) \leq \varepsilon\}.$$

Then in terms of the pseudospectrum $\Lambda_\varepsilon(p)$ we have

$$\theta \in \Lambda_{\zeta(\theta)}(p).$$

So from (4.5) we get that the pair (τ, \tilde{u}) , where \tilde{u} is the harmonic Ritz vector, has a backward error with respect to absolute perturbations of at most

$$|\xi| \|p'(\tau) \tilde{u}\| \left/ (1 + |\tau| \cdots + |\tau|^l) \right. \leq |\xi| \|p'(\tau)U\| \left/ (1 + |\tau| \cdots + |\tau|^l) \right..$$

Also, the approximate eigenvalue τ has a backward error with respect to absolute perturbations of at most $|\xi| \sigma_{\min}(p'(\tau)U) \left/ (1 + |\tau| \cdots + |\tau|^l) \right.$.

5. Refined Rayleigh–Ritz. Another idea is to generalize *refined Rayleigh–Ritz*, where, for a τ , one minimizes

$$\|A\hat{u} - \tau\hat{u}\|, \quad \hat{u} \in \mathcal{U}, \quad \|\hat{u}\| = 1.$$

This τ can be either be a fixed target or the computed Ritz value in each step; the latter was advocated by Jia [16] in the context of the Arnoldi method. Generalizing this idea to the polynomial eigenproblem, we minimize

$$(5.1) \quad \|p(\tau)\hat{u}\|, \quad \hat{u} \in \mathcal{U}, \quad \|\hat{u}\| = 1,$$

where again τ may be a fixed target or a Ritz value. Computationally, this approach amounts to determining the smallest right singular vector \hat{c} of the tall and thin $n \times k$ matrix $p(\tau)U$. In line with refined Rayleigh–Ritz terminology, we call the minimizing vector $\hat{u} = U\hat{c}$ a *refined Ritz vector*. Having this vector, we can determine an approximate eigenvalue along the lines in [13]—as for the harmonic approach.

Taking the target in (5.1) instead of the Rayleigh quotient is often more sensible in the beginning of the process, and has the additional advantage that the QR decomposition of $p(\tau)U$ can be computed incrementally; since

$$\|p(\tau)U\hat{c}\| = \|QR\hat{c}\| = \|R\hat{c}\|,$$

in this case we only have to determine the smallest right singular vector of the $k \times k$ upper triangular matrix R in the k th step of the process.

This refined approach has the following equivalent characterizations:

- (i) $\hat{u} = \operatorname{argmin}_{u \in \mathcal{U}, \|u\|=1} \|p(\tau)u\|$ with $\hat{\theta}$ as minimum,
- (ii) $(p(\tau)p(\tau)^*)^{-1} \hat{V}\hat{c} - \hat{\theta}^{-2} \hat{V}\hat{c} \perp \hat{V} := p(\tau)U$ where $\hat{\theta}^{-2}$ is maximal,
- (iii) $U^*p(\tau)^*p(\tau)U\hat{c} = \hat{\theta}^2\hat{c}$ where $\hat{\theta}^2$ is minimal.

The next theorem is a generalization of a part of [3, p. 290].

THEOREM 5.1. *For the residual of the refined Ritz vector we have*

$$\|p(\tau)\hat{u}\| \leq \frac{\mu_1 |\tau - \lambda| + \|p(\tau)\| \sin \angle(\mathcal{U}, x)}{\sqrt{1 - \sin^2 \angle(\mathcal{U}, x)}},$$

where μ_1 is as in (4.2).

Proof. Decompose $x = \gamma_U x_U + \sigma_U e_U$, where $x_U := UU^*x/\|UU^*x\|$ is the orthogonal projection of x onto \mathcal{U} , $\|x_U\| = \|e_U\| = 1$, $\gamma_U = \cos \angle(\mathcal{U}, x)$, and $\sigma_U = \sin \angle(\mathcal{U}, x)$. Since $p(\tau)x_U = (p(\tau)x - \sigma_U p(\tau)e_U)/\gamma_U$, we have by the definition of a refined Ritz vector and (4.1)

$$\|p(\tau)\hat{u}\| \leq \|p(\tau)x_U\| \leq (\mu_1 |\tau - \lambda| + \sigma_U \|p(\tau)\|) / \gamma_U.$$

□

This theorem indicates a disadvantage of the refined approach: in order to have $\|p(\tau)\hat{u}\| \rightarrow 0$, we not only need $\sin \angle(\mathcal{U}, x) \rightarrow 0$ (the search space \mathcal{U} should contain x), but also $\tau \rightarrow \lambda$ (the target must converge to the wanted eigenvalue). Indeed, as can be checked, if $x \in \mathcal{U}$ but $\tau \neq \lambda$, then the refined extraction will in general not select x as best approximate eigenvector. The practical difficulty of taking a varying $\tau \rightarrow \lambda$ is that it is computationally more expensive: we have to recompute $p(\tau)U$ (and possibly its QR decomposition) whenever τ changes. Therefore, it could be a good idea to change the target τ only every now and then, for instance, change it to the Rayleigh quotient at a *restart* of the method if we are close enough to the eigenvalue (indicated by a small $\|r\|$). At a restart, we start over with $\mathcal{U}_{\text{new}} \subset \mathcal{U}_{\text{old}}$ with a basis of the form $U_{\text{new}} := U_{\text{old}}C$, where U_{old} has size $n \times \text{maxdim}$ and C has orthonormal columns with size $\text{maxdim} \times \text{mindim}$ (see also Table 7.2). In this case, we would have to recompute a QR decomposition of $p(\tau)UC = QRC$ in any case, so this is a relatively good point to completely recompute $p(\tau)U$ and its QR decomposition for a different τ .

As an alternative, we can use the refined extraction in the beginning of the process and change to the harmonic extraction when converging; see also the numerical experiments in Section 7.

We remark that the difficulty of the refined approach sketched above is also present in the refined approach for the standard and generalized eigenproblem: asymptotically, one has to change the target to ensure convergence to the wanted eigenpair.

6. Subspace expansion for the polynomial eigenvalue problem. We take a new look at the Jacobi–Davidson subspace expansion for the polynomial eigenvalue problem (see [11] and [9] for previous work). Assume that we have an approximate eigenpair (θ, u) , where the residual $r = p(\theta)u$ is orthogonal to a certain test vector v . We are interested in an update $s \perp u$ such that

$$p(\lambda)(u + s) = 0,$$

that is, $u + s$ is a multiple of the true eigenvector. Rewriting this equation in terms of the known approximate eigenvalue θ gives

$$(6.1) \quad p(\theta)s = -p(\theta)u + (p(\theta) - p(\lambda))u + (p(\theta) - p(\lambda))s.$$

We will neglect the third term on the right-hand side of (6.1), which is—under a mild assumption—of second order ($\mathcal{O}(\|s\|^2)$) as is shown in the next lemma.

LEMMA 6.1. *Let v be a test vector with $r \perp v$. If $v^*p'(\lambda)x \neq 0$ then $|\lambda - \theta| = \mathcal{O}(\|s\|)$.*

Proof. This follows directly from $v^*p(\lambda)(u + s) = 0$, $v^*p(\theta)u = 0$, $v^*p'(\lambda)x \neq 0$, and (3.2).

□

We note that the condition $v^*p'(\lambda)x \neq 0$ is not unnatural: if we take for v the left eigenvector y , then $y^*p'(\lambda)x = 0$ implies that the condition number of the eigenvalue λ is infinite, see [14, Th. 5].

In (6.1), we now would like to project out the term $(p(\theta) - p(\lambda))u$, rather than disregarding it, but the difficulty is that $p(\lambda)u$ is unknown. Therefore, we approximate the term by a multiple of a known quantity

$$(p(\theta) - p(\lambda))u \approx (\theta - \lambda)p'(\theta)u,$$

and we ensure that we project out the factor $p'(\theta)u$. Actually, by doing this we disregard only second order terms, which forms the basis for a quadratically convergent method, see Proposition 6.2. We also want to keep the information contained in the residual $r = p(\theta)u$. Since we assumed that $r \perp v$, the projection $P = I - \frac{p'(\theta)uv^*}{v^*p'(\theta)u}$ satisfies our wishes; a possible correction equation is then

$$(6.2) \quad \left(I - \frac{p'(\theta)uv^*}{v^*p'(\theta)u} \right) p(\theta)(I - uu^*)s = -p(\theta)u, \quad s \perp u.$$

We may take the Rayleigh quotient or one of the other alternatives mentioned in [13] for θ . The vector v may depend on this choice. For instance, if we take the ordinary Rayleigh quotient θ (the most appropriate root of $u^*p(\theta)u = 0$), then the residual r is orthogonal to u and we may take $v = u$.

It may be advantageous to replace the θ on the left-hand side of (6.2) by τ for two different reasons:

- in particular in the beginning of the process, the target τ will often be more accurate than the Rayleigh quotient or any other of the approximations in [13];
- the computation of $p'(\tau)u$ for a fixed τ may be a bit cheaper than that of $p'(\theta)u$ for a varying θ , in particular in combination with the linearized harmonic approach of Section 4, where we already have to compute $p'(\tau)U$.

To determine what “the beginning of the process” is, we can monitor the residual norm. As long as $\|r\|$ is larger than a certain parameter `fix`, we use the target τ in (6.2), after which we switch to a moving approximation θ ; see also Algorithm 1 in the next section.

The following proposition states conditions under which the Jacobi–Davidson method, where we take $v = u$, has asymptotically quadratic convergence.

PROPOSITION 6.2. *Suppose the Jacobi–Davidson method for the polynomial eigenvalue problem, where we solve the correction equation (6.2) exactly, converges to an eigenvalue λ with left and right eigenvector y and x , respectively. If λ is simple, $x^*p'(\lambda)x \neq 0$, and $y^*p'(\lambda)x \neq 0$, the method converges asymptotically quadratically.*

Proof. Denote $P = I - \frac{p'(\theta)uu^*}{u^*p'(\theta)u}$. While the true update $s \perp u$ satisfies (cf. (6.1))

$$(6.3) \quad Pp(\theta)Ps = -r + P(p(\theta) - p(\lambda))u + P(p(\theta) - p(\lambda))s,$$

for the computed update $\tilde{s} \perp u$ we have

$$Pp(\theta)P\tilde{s} = -r.$$

(Note that compared with (6.2), we have replaced the projection $I - uu^*$ by P , but since $s \perp u$ and $\tilde{s} \perp u$, this is equivalent.) Therefore,

$$Pp(\theta)P(s - \tilde{s}) = P(p(\theta) - p(\lambda))u + P(p(\theta) - p(\lambda))s.$$

Because of the assumption and Lemma 6.1, the last term is $\mathcal{O}(\|s\|^2)$. For the second term we have similar to (4.3) and using Lemma 6.1

$$\|P(p(\theta) - p(\lambda))u\| = \|PFu\| = \mathcal{O}(|\theta - \lambda|^2) = \mathcal{O}(\|s\|^2).$$

To make sure that $\|s - \tilde{s}\| = \mathcal{O}(\|s\|^2)$, the last part of the proof is to show that the operator

$$(6.4) \quad \left(I - \frac{p'(\lambda) x x^*}{x^* p'(\lambda) x} \right) p(\lambda) \left(I - \frac{p'(\lambda) x x^*}{x^* p'(\lambda) x} \right)$$

is an invertible operator from x^\perp to x^\perp . Suppose therefore that this operator maps $z \perp x$ to 0. Then it follows that $p(\lambda)z = \alpha p'(\lambda)x$ for an $\alpha \in \mathbb{C}$. Left multiplying by the left eigenvector gives $0 = y^* p(\lambda)z = \alpha y^* p'(\lambda)x$. Because of the assumption, we must have $\alpha = 0$. So $p(\lambda)z = 0$ and since $z \perp x$ and the assumption that λ is a simple eigenvalue, we have $z = 0$. Therefore, we have proved the injectivity of (6.4); the bijectivity follows from comparing dimensions. \square

We remark that for a two-sided variant of Jacobi–Davidson (that is, the test vector v converges to the left eigenvector y), the condition $x^* p'(\lambda)x \neq 0$ is no longer necessary.

The following proposition proves what is often experienced in practice: inexact Jacobi–Davidson typically converges asymptotically linearly. By “inexact” we mean that the correction equation is solved until the residual norm is reduced by a fixed factor.

PROPOSITION 6.3. *We assume the same hypotheses as in Proposition 6.2, only now we solve the correction equation (6.2) inexactly, such that $\tilde{s} \perp u$ satisfies*

$$\|Pp(\theta)P\tilde{s} + r\| \leq \eta \|r\|,$$

for $\eta < 1$. If $\eta < \kappa^{-1}$, where κ is the condition number of the operator in (6.4), seen as operator from x^\perp to x^\perp , the asymptotic convergence of the Jacobi–Davidson method is linear.

Proof. In this case, we have

$$Pp(\theta)P(s - \tilde{s}) = \eta_1 \|r\| f + P(p(\theta) - p(\lambda))u + P(p(\theta) - p(\lambda))s,$$

for a certain vector $f \perp u$, $\|f\| = 1$, and $0 \leq \eta_1 \leq \eta$. In the previous proposition we saw that the last two terms on the right-hand side are $\mathcal{O}(\|s\|^2)$, and that the operator $Pp(\theta)P$ is asymptotically invertible. Moreover, from (6.3) it is clear that, neglecting higher order terms, $\|r\| \leq \|Pp(\lambda)P\| \|s\|$. Therefore, $\|s - \tilde{s}\| \leq \eta \kappa \|s\| +$ higher order terms, where κ denotes the condition number of (6.4). \square

Rather than a fixed residual reduction, we will take a fixed number of inner iterations to solve the correction equation (6.2) in the numerical experiments in the following section.

7. Numerical experiments. The following experiments were carried out in MATLAB. For all experiments we first put the “seed” to 0 by the command “`rand('seed', 0)`” so that the results are reproducible.

EXPERIMENT 7.1. To get some feeling for the different extractions, we first take 100×100 matrices A_2, A_1 , and A_0 with elements taken randomly from a normal distribution over $[0, 1]$. Figure 7.1 shows the 200 eigenvalues. Such an eigenvalue distribution (one, or a few outliers) is characteristic for random matrices and an interesting phenomenon in its own right (for work on the eigenvalues of standard eigenvalue problems, see Edelman [17]). Also indicated in Figure 7.1 are the targets $\tau = -5$ and $\tau = -2.5$.

We build up a four-dimensional search space \mathcal{U} to find the eigenpair corresponding to the eigenvalue $\lambda = -2.69 \dots$, which is the eigenvalue closest to both targets. The first basis vector of \mathcal{U} is taken to be $u_1 = x + \varepsilon w$, where x is the corresponding eigenvector and w is a random vector. We complement \mathcal{U} by three random vectors. For ε we take the values 10^{-j} for $j = 1, \dots, 5$, meaning the search space ranges from bad to good quality. Furthermore, the two different targets $\tau = -5$ and $\tau = -2.5$ represent a not very accurate and a more accurate target, respectively.




FIG. 7.1. The eigenvalues of the 100×100 quadratic eigenvalue problem and targets $\tau = -5$ and $\tau = -2.5$.

TABLE 7.1

Extraction results for 100×100 random matrices, targets $\tau = -5$ and $\tau = -2.5$, and search spaces of different quality.

target	$\angle(\mathcal{U}, x)$	$3.2e-5$	$3.2e-4$	$3.2e-3$	$3.2e-2$	$3.1e-1$
	standard	$1.6e-4$	$1.6e-3$	$1.5e-2$	$1.0e-1$	$4.4e-1$
-5	lin-harmonic	$6.1e-3$	$6.1e-3$	$6.4e-3$	$3.3e-2$	$4.2e-1$
	harmonic	$3.3e-5$	$3.3e-4$	$3.3e-3$	$3.3e-2$	$3.4e-1$
	refined	$4.8e-2$	$4.8e-2$	$4.8e-2$	$5.8e-2$	$3.2e-1$
-2.5	lin-harmonic	$1.7e-4$	$3.5e-4$	$3.3e-3$	$3.4e-2$	$1.4e-0$
	harmonic	$3.3e-5$	$3.3e-4$	$3.3e-3$	$3.4e-2$	$1.4e-0$
	refined	$6.7e-3$	$6.7e-3$	$7.5e-3$	$3.4e-2$	$3.2e-1$

Table 7.1 gives for $\tau = -5$ and $\tau = -2.5$ the results of the four different extraction processes. We display $\angle(\mathcal{U}, x)$ (this gives the best possible approximate vector in \mathcal{U}) and $\angle(u, x)$, the angle of the approximate eigenvector with the exact eigenvector.

It turns out that the results of the standard extraction are independent of the target τ ; the other extraction methods do depend on τ . For an inaccurate search space, say $\angle(\mathcal{U}, x) \approx 10^{-1}$, the refined approach is the best, having a near-optimal extraction, which means $\angle(u, x) \approx \angle(\mathcal{U}, x)$. For better search spaces, the refined approach does not improve any longer, as already expected from the discussion after Theorem 5.1, and the harmonic and linearized harmonic approach do the best job. “Asymptotically”, that is, for quite accurate search spaces, the harmonic approach is the best with an almost optimal extraction, with the standard extraction following in second place but with a clear difference.

These results suggest to combine the strengths of the different extraction methods, giving the JDPOLY algorithm as outlined in Algorithm 1.

The “rgs” in Line 3 stands for repeated Gram–Schmidt or any other numerical stable way to form an orthonormal basis for the search space. In Line 11/12, we have chosen $v = u$ in the correction equation (6.2), corresponding to looking for an orthogonal updates to an approximate eigenvector. Table 7.2 gives an oversight of some of the available parameters and the default value.

EXPERIMENT 7.2. We test the JDPOLY algorithm for the above problem with target $\tau = -5$ and the four different extraction methods. An inexact LU preconditioner for $\tau^2 A_2 + \tau A_1 + A_0$ with drop tolerance 0.01 is used. Inspired by the previous experiment, we switch from the linearized harmonic extraction and the refined extraction to the harmonic extraction if $\|r\| \leq \text{threshold}$; we test various values for this parameter, see Table 7.3.

Algorithm 1 Jacobi–Davidson for the polynomial eigenvalue problem

Input: A device to compute Ax for arbitrary x , a starting vector u_1 , a target τ , a tolerance ε , and parameters **threshold** and **fix**

Output: An approximate eigenpair (θ, u) with $\theta \approx \tau$

```

1:  $s = u_1$ 
2: for  $k = 1, 2, \dots$  do
3:    $\text{rgs}(U_{k-1}, s) \rightarrow U_k$ 
4:   Extract a Ritz pair  $(\theta, c)$  close to  $\tau$  using:
5:     harmonic (or standard) extraction  $(\|r\| \leq \text{threshold})$ 
6:     harmonic (or refined, linearized harmonic, standard) extraction  $(\|r\| > \text{threshold})$ 
7:    $u = U_k c$ 
8:    $r = p(\theta) u$ 
9:   Stop if  $\|r\| \leq \varepsilon$ 
10:  Solve (approximately) an  $s \perp u$  from:
11:     $(I - \frac{zu^*}{u^*z}) p(\theta)(I - uu^*) s = -r, \quad z = p'(\theta) u, \quad (\|r\| \leq \text{fix})$ 
12:     $(I - \frac{zu^*}{u^*z}) p(\tau)(I - uu^*) s = -r, \quad z = p'(\tau) u, \quad (\|r\| > \text{fix})$ 
13: end for

```

TABLE 7.2

Parameters of the Jacobi–Davidson method for the polynomial eigenvalue problem with their default value.

Parameter	Meaning	Default value
tol	(absolute) tolerance for the outer iteration	10^{-6}
maxit	maximum number of outer iterations	1000
extraction	extraction process	harmonic
target	complex number of interest	largest $ \lambda $
mindim	dimension search spaces after restart	10
maxdim	maximum dimension search spaces	20
maxit_inner	maximum number of inner iterations	10
inner_tol	tolerance inner iteration	0
fix	fix target until $\ r\ \leq \text{fix}$	0.01
u1	initial search space	random vector
M1, M2	preconditioner $M = M_1 M_2$ (possibly in factored form)	$M = I$
Mtype	left or right preconditioning	left
verbosity	output desired	no

After extraction of an approximate eigenpair, all extraction methods except the standard approach employ the “two-dimensional Galerkin method” [13] to determine a more accurate eigenvalue (for the standard extraction it turns out in practice that this extra step is generally unfavorable since a good quality vector is required).

TABLE 7.3

Number of iterations for a random 100×100 matrix with target $\tau = -5$.

Extraction	threshold	Iterations
Standard	—	47
Lin-harmonic	—	—
	1	20
Harmonic	—	20
Refined	—	—
	1	62
	5	20

The harmonic approach takes much fewer iterations (20) than the standard extraction (47). The linearized harmonic and refined extraction in themselves do not converge in 100 outer iterations; indeed, this is what we expect for a fixed target as was predicted in Sections 4 and 5: using a fixed target they generally do not recognize an eigenvector that is present in the search space.

However, if we combine them with the harmonic extraction, that is, switch to the harmonic extraction as soon as $\|r\| < \text{threshold}$, they converge much faster. Especially for the refined extraction we have to select a reasonably large `threshold`, which means that we switch to the harmonic extraction quite early in the process; not surprisingly, the combined approaches behave much like the harmonic extraction. As can be seen, the new (combination) approaches converge much faster than the standard extraction.

EXPERIMENT 7.3. Using the command `sprand`, we create random 1000×1000 matrices A_2 , A_1 , and A_0 with density $\approx 10\%$ and condition numbers $\kappa = \mathcal{O}(10^5)$. With the `threshold` value 1, target $\tau = 0$ and an ILU preconditioner with drop tolerance $\delta = 0.001$, JDPOLY with harmonic extraction converges in 11 outer iterations versus the standard extraction in 12 iterations. However, if we solve the correction equation less accurately, so that a good extraction process may become more important, the difference becomes larger: for 5 inner iterations instead of 10, the number of outer iterations is 11 (harmonic) and 14 (standard); if we expand the search space by the preconditioned residual (no inner iterations), the number of outer iterations is 17 (harmonic) versus 27 (standard). This experiment indicates that the harmonic extraction may be an important advantage if the quality of the search space expansion is modest, for instance because a (good) preconditioner is unavailable.

EXPERIMENT 7.4. Finally, we consider an example arising from a damped gyroscopic system as in [18]: define $m \times m$ tridiagonal matrices

$$\begin{aligned} B_2 &= \text{tridiag}(1, 4, 1)/6, & B_1 &= \text{tridiag}(1, 0, -1), \\ B_0 &= \text{tridiag}(1, -2, 1), & C_1 &= \text{tridiag}(1, 2, 1) \end{aligned}$$

where $m = 90$ and the 8100×8100 matrices

$$\begin{aligned} A_2 &= c_{11}I_m \otimes B_2 - c_{12}B_2 \otimes I_m, \\ A_1 &= c_{21}I_m \otimes B_1 - c_{22}B_1 \otimes I_m + 10^{-3}(c_{41}I_m \otimes C_1 - c_{42}C_1 \otimes I_m), \\ A_0 &= c_{31}I_m \otimes B_0 - c_{32}B_0 \otimes I_m, \end{aligned}$$

where $c_{11} = 1$, $c_{12} = 1.3$, $c_{21} = 0.1$, $c_{22} = 1.1$, $c_{31} = 1$, $c_{32} = 1.2$, $c_{41} = 1.05$, and $c_{42} = 0.9$. With target $\tau = 0$, `tol` = 10^{-8} , and an ILU preconditioner ($\delta = 0.001$), the standard extraction did not converge to the required tolerance in 100 iterations; the harmonic, linearized harmonic, and refined extractions (the last two with `threshold=1`) took 48, 50, and 60 steps, respectively. From this experiment it seems that the harmonic extraction may asymptotically also be more accurate than the standard extraction, as was also suggested by Experiment 7.1.

8. Conclusions. We reviewed the harmonic Rayleigh–Ritz extraction for the standard and generalized eigenvalue problem, and discussed the problem of subspace extraction for the polynomial eigenvalue problem. While this paper focuses on how to extract an approximate vector from a search space, we refer to [13] for methods to derive an approximate eigenvalue from the approximate eigenvector.

We introduced harmonic Rayleigh–Ritz (two variants) and refined Rayleigh–Ritz for the polynomial eigenvalue problem. The harmonic approaches are slightly more expensive than the standard approach, but for interior eigenpairs in particular the (non-linearized) harmonic approach is theoretically more sound and numerically gives a better subspace extraction than the standard Rayleigh–Ritz extraction results in almost all experiments. The refined approach

may be worthwhile for a search space of poor quality, as is typical for the beginning of the process. To retain accuracy for the refined extraction towards the end of the process, we may have to let the target converge to the desired eigenvalue, for instance by setting the target equal to the Rayleigh quotient when restarting, or switch to the harmonic approach.

The linearized harmonic extraction has the advantage that we have a bound (4.5) on the residual norm; on the other hand, an exact eigenpair does not satisfy (4.4) in general. For the non-linearized extraction the situation is exactly opposite: there is no direct upper bound for the residual, but an exact eigenpair satisfies (4.6). This last property may be considered important, and numerical experiments confirm that the (non-linearized) harmonic approach is often the most reliable.

We embedded the subspace extraction techniques in a Jacobi–Davidson method, with resulting quadratic (exact) or linear (inexact) convergence.

REFERENCES

- [1] Van der Vorst HA. Computational methods for large eigenvalue problems. In *Handbook of numerical analysis, Vol. VIII*, pages 3–179. North-Holland, Amsterdam, 2002.
- [2] Sleijpen GLG and Van der Vorst HA. A Jacobi–Davidson iteration method for linear eigenvalue problems. *SIAM J. Matrix Anal. Appl.*, 17(2):401–425, 1996.
- [3] Stewart GW. *Matrix algorithms. Vol. II*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 2001.
- [4] Sleijpen GLG, Van der Vorst HA, and Meijerink E. Efficient expansion of subspaces in the Jacobi–Davidson method for standard and generalized eigenproblems. *Electron. Trans. Numer. Anal.*, 7:75–89, 1998.
- [5] Van den Eshof J. *Nested Iteration Methods for Nonlinear Matrix Problems*. PhD thesis, Utrecht University, 2003.
- [6] Paige CC, Parlett BN, and Van der Vorst HA. Approximate solutions and eigenvalue bounds from Krylov subspaces. *Numer. Lin. Alg. Appl.*, 2(2):115–133, 1995.
- [7] Saad Y. *Numerical Methods for Large Eigenvalue Problems*. Manchester University Press, Manchester, UK, 1992.
- [8] Jia Z. The convergence of harmonic Ritz values, harmonic Ritz vectors, and refined harmonic Ritz vectors. *Math. Comp.*, 74(251):1441–1456, 2005.
- [9] Hochstenbach ME and Sleijpen GLG. Two-sided and alternating Jacobi–Davidson. *Linear Algebra Appl.*, 358(1-3):145–172, 2003.
- [10] Lanczos C. An iteration method for the solution of the eigenvalue problem of linear differential and integral operators. *J. Res. Nat. Bur. Stand.*, 45(4):255–282, 1950.
- [11] Sleijpen GLG, Booten AGL, Fokkema DR, and Van der Vorst HA. Jacobi–Davidson type methods for generalized eigenproblems and polynomial eigenproblems. *BIT*, 36(3):595–633, 1996.
- [12] Stathopoulos A. A case for a biorthogonal Jacobi–Davidson method: restarting and correction equation. *SIAM J. Matrix Anal. Appl.*, 24(1):238–259, 2002.
- [13] Hochstenbach ME and Van der Vorst HA. Alternatives to the Rayleigh quotient for the quadratic eigenvalue problem. *SIAM J. Sci. Comput.*, 25(2):591–603, 2003.
- [14] Tisseur F. Backward error and condition of polynomial eigenvalue problems. *Linear Algebra Appl.*, 309(1-3):339–361, 2000.
- [15] Tisseur F and Higham NJ. Structured pseudospectra for polynomial eigenvalue problems, with applications. *SIAM J. Matrix Anal. Appl.*, 23(1):187–208, 2001.
- [16] Jia Z. Refined iterative algorithms based on Arnoldi’s process for large unsymmetric eigenproblems. *Linear Algebra Appl.*, 259:1–23, 1997.
- [17] Edelman A. Eigenvalues and condition numbers of random matrices. *SIAM J. Matrix Anal. Appl.*, 9(4):543–560, 1988.
- [18] Hwang T-M, Lin W-W, and Mehrmann V. Numerical solution of quadratic eigenvalue problems with structure-preserving methods. *SIAM J. Sci. Comput.*, 24(4):1283–1302, 2003.