

THE ARNOLDI EIGENVALUE ITERATION WITH EXACT SHIFTS CAN FAIL*

MARK EMBREE†

Abstract. The restarted Arnoldi algorithm, implemented in the ARPACK software library and MATLAB’s `eigs` command, is among the most common means of computing select eigenvalues and eigenvectors of a large, sparse matrix. To assist convergence, a starting vector is repeatedly refined via the application of automatically constructed polynomial filters whose roots are known as “exact shifts.” Though Sorensen proved the success of this procedure under mild hypotheses for Hermitian matrices, a convergence proof for the non-Hermitian case has remained elusive. The present note describes a class of examples for which the algorithm fails in the strongest possible sense; that is, the polynomial filter used to restart the iteration deflates the eigenspace one is attempting to compute.

Key words. implicitly restarted Arnoldi algorithm, Krylov–Schur algorithm, eigenvalues, exact shifts, ARPACK, `eigs`

AMS subject classifications. 15A18, 15A60, 65F15

DOI. 10.1137/060669097

1. Setting. Large-scale matrix eigenvalue problems typically derive from applications in which one seeks only some small subset of the spectrum, such as the largest magnitude or rightmost eigenvalues for stability analysis of dynamical systems. Given a matrix $\mathbf{A} \in \mathbb{C}^{n \times n}$, one can approximate such eigenvalues by projecting \mathbf{A} onto an appropriate low-dimensional subspace and then solving a small eigenvalue problem with a dense method such as the QR algorithm. The Arnoldi method [1, 13] orthogonally projects \mathbf{A} onto the Krylov subspace

$$\mathcal{K}_k(\mathbf{A}, \mathbf{v}) := \text{span}\{\mathbf{v}, \mathbf{A}\mathbf{v}, \dots, \mathbf{A}^{k-1}\mathbf{v}\}$$

generated by the starting vector $\mathbf{v} \in \mathbb{C}^n$. This subspace—the span of iterates of the power method—often produces eigenvalue estimates that fall on the “periphery” of the spectrum [16], though the quality of such approximations depends on the distribution of the eigenvalues of \mathbf{A} , the angles between associated eigenvectors, and the starting vector \mathbf{v} . Provided that $\dim \mathcal{K}_{k+1}(\mathbf{A}, \mathbf{v}) = k + 1$, by the end of the k th iteration the Arnoldi algorithm has produced orthonormal vectors $\mathbf{v}_1, \dots, \mathbf{v}_{k+1}$ such that

$$\text{span}\{\mathbf{v}_1, \dots, \mathbf{v}_j\} = \mathcal{K}_j(\mathbf{A}, \mathbf{v})$$

for all $j = 1, \dots, k + 1$. Accumulating $\mathbf{v}_1, \dots, \mathbf{v}_k$ into the columns of $\mathbf{V}_k \in \mathbb{C}^{n \times k}$, the algebra affecting the construction of the basis can be summarized in the Arnoldi factorization

$$(1) \quad \mathbf{A}\mathbf{V}_k = \mathbf{V}_k\mathbf{H}_k + h_{k+1,k}\mathbf{v}_{k+1}\mathbf{e}_k^*,$$

where \mathbf{e}_k denotes the last column of the $k \times k$ identity matrix and $\mathbf{H}_k = \mathbf{V}_k^*\mathbf{A}\mathbf{V}_k$ is an upper Hessenberg matrix. The eigenvalues of \mathbf{H}_k , the *Ritz values*, approximate

*Received by the editors September 5, 2006; accepted for publication by F. Dopico November 2, 2006; published electronically February 27, 2009. Supported by Department of Energy grant DE-FG03-02ER25531 and National Science Foundation grant DMS-CAREER-0449973.
<http://www.siam.org/journals/simax/31-1/66909.html>

†Department of Computational and Applied Mathematics, Rice University, 6100 Main Street, MS 134, Houston, TX 77005–1892 (embree@rice.edu).

eigenvalues of \mathbf{A} in the sense that if $\mathbf{H}_k \mathbf{u} = \theta \mathbf{u}$, then

$$\|\mathbf{A}(\mathbf{V}_k \mathbf{u}) - \theta(\mathbf{V}_k \mathbf{u})\| \leq |h_{k+1,k}| |\mathbf{e}_k^* \mathbf{u}|.$$

These Ritz values must fall within the *numerical range* (or *field of values*) of \mathbf{A} ,

$$(2) \quad W(\mathbf{A}) := \{z \in \mathbb{C} : z = \mathbf{x}^* \mathbf{A} \mathbf{x} \text{ for some } \mathbf{x} \in \mathbb{C}^n \text{ with } \|\mathbf{x}\| = 1\},$$

since $\theta = (\mathbf{V}_k \mathbf{u})^* \mathbf{A}(\mathbf{V}_k \mathbf{u})$. The numerical range of a normal matrix is the convex hull of its spectrum [5]; for nonnormal matrices, the numerical range may contain points far from any eigenvalue (see, e.g., [17, Ch. 17]), a fact critical to the examples we shall present. Given the Arnoldi algorithm's broad utility, its nontrivial convergence behavior has attracted considerable attention [3, 4, 6, 7, 8, 13].

The computational cost of enlarging the Krylov subspace in the Arnoldi algorithm grows with the subspace dimension, k , and for large problems storage of the basis vectors alone becomes burdensome. Unfortunately, in many cases practical values of k fail to produce sufficiently accurate eigenvalue estimates. A simple solution is to *restart* the iteration, using information culled from the factorization (1) to refine the starting vector \mathbf{v} in a manner that enriches components in the direction of desired eigenvectors while damping unwanted eigenvectors. This process is most commonly implemented through a *polynomial filter*: the starting vector for the new Arnoldi factorization takes the form

$$\mathbf{v}_+ = \psi(\mathbf{A})\mathbf{v}$$

for some polynomial ψ , generally of low degree. Strategies for constructing ψ vary. Saad advocated a choice designed to produce \mathbf{v}_+ as a linear combination of the desired Ritz vectors [13]. Later methods chose ψ to be small over an adaptively identified region of the complex plane containing only the unwanted eigenvalues, but such procedures proved difficult to reliably automate. Sorensen proposed a less elaborate strategy: take ψ to be a polynomial whose roots match the undesired Ritz values [14]. More precisely, suppose we seek m eigenvalues and have already built an Arnoldi factorization of dimension $k = m + p$ for some $p > 0$. Sort the eigenvalues of \mathbf{H}_k into the m possessing the desired trait (e.g., those rightmost or largest in magnitude), labeled $\theta_1, \dots, \theta_m$, and the remaining p values $\theta_{m+1}, \dots, \theta_{m+p}$, which we implicitly assume are coarse approximations to eigenvalues we do not seek. Set $\psi(z) = \prod_{j=1}^p (z - \theta_{m+j})$, so that

$$\mathbf{v}_+ = \prod_{j=1}^p (\mathbf{A} - \theta_{m+j} \mathbf{I}) \mathbf{v}.$$

These roots of ψ , the undesired Ritz values, are called *exact shifts*. Such shifts form an essential part of the ARPACK library [9], and their efficacy is responsible in no small part for the popularity enjoyed by that software and MATLAB's subordinate `eigs` command. (In the same paper that advocated exact shifts, Sorensen also proposed a robust method of *implicitly* restarting the Arnoldi algorithm using the roots of an arbitrary filter polynomial [14]. An alternative implementation, Stewart's Krylov-Schur algorithm, requires the use of exact shifts [15].)

Numerous computational examples demonstrate the success of exact shifts; see, e.g., [3, 4, 9]. A convergence proof would require that one make precise the notion that the p Ritz values used as shifts approximate the unwanted eigenvalues. For Hermitian

\mathbf{A} an appeal to the interlacing theorem [11] suffices. Labeling the eigenvalues of \mathbf{A} as $\lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_n$, there can be at most one Ritz value in the interval (λ_1, λ_2) , two Ritz values in (λ_2, λ_3) , and so on. This interlacing forms the basis of Sorensen’s proof that for a Hermitian matrix the restarted Arnoldi algorithm converges to extreme eigenvalues.¹ In the years since the introduction of exact shifts, a convergence proof for non-Hermitian matrices has remained elusive. Several results ensure convergence provided the exact shifts satisfy an appropriate distribution [3, 4, 8], but conditions guaranteeing that exact shifts exhibit such behavior have not yet been established. Indeed, few fine results about the Ritz values of non-Hermitian matrices are known.

The purpose of this note is to provide a counterexample to the conjecture that the restarted Arnoldi algorithm with exact shifts must converge under hypotheses resembling those sufficient for the Hermitian case. Our examples illustrate the failure of the algorithm in the strongest sense, in that the filter polynomial ψ exactly deflates a sought-after (perfectly conditioned) eigenvalue.

Throughout we assume that all computations are performed in exact arithmetic and in particular that eigenvalues of the upper Hessenberg matrix \mathbf{H}_k are determined exactly.

2. Simple example. We begin with a small example that clearly demonstrates how restarting with exact shifts can produce catastrophic results. Suppose we seek the largest magnitude (and rightmost) eigenvalue $\lambda = 1$ and associated eigenvector $\mathbf{u}_1 = [1 \ 0 \ 0 \ 0]^T$ of the matrix

$$\mathbf{A} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 6 & -2 \\ 0 & 0 & 0 & 2 \\ 0 & 0 & 0 & 0 \end{bmatrix}.$$

We denote the bottom-right 3×3 submatrix of \mathbf{A} by

$$\mathbf{D} = \begin{bmatrix} 0 & 6 & -2 \\ 0 & 0 & 2 \\ 0 & 0 & 0 \end{bmatrix}.$$

The desired eigenvalue $\lambda = 1$ of \mathbf{A} is far from the rest of the spectrum, $\sigma(\mathbf{D}) = \{0\}$, and is also fairly well separated [16, p. 46]:

$$\text{sep}(1, \mathbf{D}) = \|(\mathbf{I} - \mathbf{D})^{-1}\|^{-1} = 0.0837\dots$$

In the language of pseudospectra [17], this implies $1 \notin \sigma_\varepsilon(\mathbf{D})$ for all $\varepsilon < 0.0837\dots$, where

$$(3) \quad \begin{aligned} \sigma_\varepsilon(\mathbf{D}) &= \{z \in \mathbb{C} : z \in \sigma(\mathbf{D} + \mathbf{E}) \text{ for some } \mathbf{E} \in \mathbb{C}^{n \times n} \text{ with } \|\mathbf{E}\| < \varepsilon\} \\ &= \{z \in \mathbb{C} : \|(z\mathbf{I} - \mathbf{D})^{-1}\| > \varepsilon^{-1}\}. \end{aligned}$$

Furthermore, the eigenvector \mathbf{u}_1 is orthogonal to the complementary invariant subspace associated with the zero eigenvalue; that is, $\lambda = 1$ is perfectly conditioned [2].

¹Sorensen’s theorem [14, Thm. 5.9] imposes several reasonable caveats: the starting vector cannot be deficient in any of the sought-after eigenvectors, and the iteration must not come arbitrarily close to “lucky breakdown”; i.e., all the subdiagonal entries in \mathbf{H}_k must be bounded away from zero, so none of the intermediate Krylov subspaces fall too close to an invariant subspace.

Apply the restarted Arnoldi method to \mathbf{A} with the starting vector $\mathbf{v} = [1 \ 1 \ 1 \ 1]^T$. We seek the $m = 1$ eigenvalue and will build the Krylov subspace out to dimension $k = m + p = 2$ before restarting. Two steps of the Arnoldi algorithm produce the orthonormal basis vectors

$$\mathbf{v}_1 = \frac{1}{2} \begin{bmatrix} 1 \\ 1 \\ 1 \\ 1 \end{bmatrix}, \quad \mathbf{v}_2 = \frac{\sqrt{35}}{70} \begin{bmatrix} -3 \\ 9 \\ 1 \\ -7 \end{bmatrix}$$

and upper Hessenberg matrix

$$\mathbf{H}_2 = [\mathbf{v}_1 \ \mathbf{v}_2]^* \mathbf{A} [\mathbf{v}_1 \ \mathbf{v}_2] = \begin{bmatrix} 7/4 & 3\sqrt{35}/140 \\ \sqrt{35}/4 & 5/4 \end{bmatrix}.$$

The characteristic polynomial for \mathbf{H}_2 is

$$\begin{aligned} \det(\lambda \mathbf{I} - \mathbf{H}_2) &= (\lambda - 7/4)(\lambda - 5/4) - \frac{3 \cdot 35}{4 \cdot 140} \\ &= (\lambda - 1)(\lambda - 2), \end{aligned}$$

and hence the Ritz values are $\theta_1 = 2$ and $\theta_2 = 1$ (ordered by decreasing magnitude). The smaller magnitude (and leftmost) of these is chosen as the exact shift; hence the filter polynomial $\psi(z) = z - 1$ produces the new starting vector

$$\mathbf{v}_+ = \psi(\mathbf{A})\mathbf{v} = \begin{bmatrix} 0 \\ 3 \\ 1 \\ -1 \end{bmatrix}.$$

Note that \mathbf{v}_+ has no component in the desired eigenvector \mathbf{u}_1 —that is, \mathbf{v}_+ is contained in the invariant subspace associated with the zero eigenvalue, and hence the eigenvalue $\lambda = 1$ will exert no influence upon further Arnoldi iterations.

One might suspect that this starting vector has been specially engineered to yield this behavior. While this is true of the last three components of \mathbf{v} , it is curious that the phenomenon persists even when \mathbf{v} is arbitrarily enriched in the eigenvector we wish to compute. If we set $\mathbf{v} = [\alpha \ 1 \ 1 \ 1]^T$ for any $\alpha \in \mathbb{R}$, then two steps of the Arnoldi algorithm yield the upper Hessenberg matrix

$$\mathbf{H}_2 = \begin{bmatrix} \frac{6 + \alpha^2}{3 + \alpha^2} & \frac{3\alpha^3}{(3 + \alpha^2)(24 + 11\alpha^2)^{1/2}} \\ \frac{(24 + 11\alpha^2)^{1/2}}{3 + \alpha^2} & \frac{3 + 2\alpha^2}{3 + \alpha^2} \end{bmatrix}$$

with characteristic polynomial

$$\det(\lambda \mathbf{I} - \mathbf{H}_2) = \lambda^2 - 3\lambda + 2 = (\lambda - 1)(\lambda - 2).$$

The Ritz values $\theta_1 = 2$ and $\theta_2 = 1$ are independent of the bias of the starting vector \mathbf{v} toward the desired eigenvector. As demonstrated in the next section, this behavior is an instance of a broader phenomenon that will facilitate the design of larger examples.

3. General construction. While the orthogonal basis for the Krylov subspace and the resulting Arnoldi factorization (1) have computational advantages, other bases for $\mathcal{K}_k(\mathbf{A}, \mathbf{v})$ are often better suited to analysis. Among the choices available (see [15, Thm. 2.2]), we shall use a decomposition introduced by Ruhe [12, sect. 2]. Suppose $\dim \mathcal{K}_k(\mathbf{A}, \mathbf{v}) = k$. Arrange the Krylov basis vectors $\mathbf{v}, \mathbf{A}\mathbf{v}, \dots, \mathbf{A}^{k-1}\mathbf{v}$ into the columns of

$$\mathbf{K}_k = [\mathbf{v} \quad \mathbf{A}\mathbf{v} \quad \dots \quad \mathbf{A}^{k-1}\mathbf{v}],$$

and define the companion matrix

$$\mathbf{C}_k = \begin{bmatrix} & & c_1 \\ 1 & & c_2 \\ & \ddots & \vdots \\ & & 1 & c_k \end{bmatrix} \in \mathbb{C}^{k \times k}$$

for constants c_1, \dots, c_k ; unspecified entries equal zero. A direct calculation reveals

$$(4) \quad \mathbf{A}\mathbf{K}_k - \mathbf{K}_k\mathbf{C}_k = \mathbf{r}\mathbf{e}_k^*$$

for the vector

$$\mathbf{r} := \mathbf{A}^k\mathbf{v} - \sum_{j=1}^k c_j \mathbf{A}^{j-1}\mathbf{v} \in \mathcal{K}_{k+1}(\mathbf{A}, \mathbf{v}).$$

Now since

$$\text{Ran } \mathbf{K}_j = \text{Ran } \mathbf{V}_j = \mathcal{K}_j(\mathbf{A}, \mathbf{v}), \quad j = 1, \dots, k,$$

there exists some invertible $\mathbf{R}_k \in \mathbb{C}^{k \times k}$ such that

$$\mathbf{K}_k = \mathbf{V}_k \mathbf{R}_k.$$

Recalling that $\mathbf{H}_k = \mathbf{V}_k^* \mathbf{A} \mathbf{V}_k$, pre- and postmultiply (4) by \mathbf{V}_k^* and \mathbf{R}_k^{-1} to obtain

$$\mathbf{H}_k - \mathbf{R}_k \mathbf{C}_k \mathbf{R}_k^{-1} = \mathbf{V}_k^* \mathbf{r} \mathbf{e}_k^* \mathbf{R}_k^{-1}.$$

If c_1, \dots, c_k are chosen such that $\mathbf{r} \perp \mathcal{K}_k(\mathbf{A}, \mathbf{v}) = \text{Ran } \mathbf{V}_k$, we conclude that

$$\mathbf{H}_k = \mathbf{R}_k \mathbf{C}_k \mathbf{R}_k^{-1};$$

i.e., the companion matrix \mathbf{C}_k and the upper Hessenberg matrix \mathbf{H}_k are similar and thus have the same eigenvalues. This development facilitates the following result.

THEOREM 3.1. *Let $\theta_1, \dots, \theta_k$ denote the Ritz values produced by $k \leq \ell$ steps of the Arnoldi algorithm applied to the matrix $\mathbf{D} \in \mathbb{C}^{\ell \times \ell}$ with starting vector $\mathbf{w} \in \mathbb{C}^\ell$, where $\dim \mathcal{K}_k(\mathbf{D}, \mathbf{w}) = k$. For any $\mathbf{T} \in \mathbb{C}^{m \times m}$, $m \leq k$, with spectrum $\sigma(\mathbf{T}) \subseteq \{\theta_1, \dots, \theta_k\}$ (respecting multiplicity) and any $\mathbf{x} \in \mathbb{C}^m$, k steps of the Arnoldi process applied to the matrix and starting vector*

$$\mathbf{A} = \begin{bmatrix} \mathbf{T} & \mathbf{0} \\ \mathbf{0} & \mathbf{D} \end{bmatrix}, \quad \mathbf{v} = \begin{bmatrix} \mathbf{x} \\ \mathbf{w} \end{bmatrix}$$

produce the identical Ritz values $\theta_1, \dots, \theta_k$.

Proof. To write a Krylov factorization for $\mathcal{K}_k(\mathbf{D}, \mathbf{w})$, we define

$$\widehat{\mathbf{K}}_k = [\mathbf{w} \quad \mathbf{D}\mathbf{w} \quad \cdots \quad \mathbf{D}^{k-1}\mathbf{w}], \quad \mathbf{C}_k = \begin{bmatrix} & & & c_1 \\ 1 & & & c_2 \\ & \ddots & & \vdots \\ & & 1 & c_k \end{bmatrix} \in \mathbb{C}^{k \times k},$$

with coefficients c_1, \dots, c_k chosen, e.g., via the Gram–Schmidt process, to ensure

$$\widehat{\mathbf{r}} := \mathbf{D}^k \mathbf{w} - \sum_{j=1}^k c_j \mathbf{D}^{j-1} \mathbf{w}$$

is orthogonal to $\mathcal{K}_k(\mathbf{D}, \mathbf{w})$. This choice gives $\sigma(\mathbf{C}_k) = \{\theta_1, \dots, \theta_k\}$. Now with

$$\mathbf{K}_k = [\mathbf{v} \quad \mathbf{A}\mathbf{v} \quad \cdots \quad \mathbf{A}^{k-1}\mathbf{v}],$$

observe that

$$\mathbf{A}\mathbf{K}_k - \mathbf{K}_k \mathbf{C}_k = \mathbf{r} \mathbf{e}_k^*,$$

where

$$\mathbf{r} := \mathbf{A}^k \mathbf{v} - \sum_{j=1}^k c_j \mathbf{A}^{j-1} \mathbf{v} = \begin{bmatrix} \mathbf{T}^k \mathbf{x} - \sum_{j=1}^k c_j \mathbf{T}^{j-1} \mathbf{x} \\ \mathbf{D}^k \mathbf{w} - \sum_{j=1}^k c_j \mathbf{D}^{j-1} \mathbf{w} \end{bmatrix}.$$

Since $\sigma(\mathbf{T}) \subseteq \{\theta_1, \dots, \theta_k\}$ (respecting multiplicity), the polynomial

$$z^k - \sum_{j=1}^k c_j z^{j-1} = \prod_{j=1}^k (z - \theta_j)$$

annihilates \mathbf{T} , leaving

$$\mathbf{r} = \begin{bmatrix} \mathbf{0} \\ \widehat{\mathbf{r}} \end{bmatrix}.$$

It follows that \mathbf{r} is orthogonal to $\mathcal{K}_k(\mathbf{A}, \mathbf{v})$, and hence the Ritz values for \mathbf{A} and \mathbf{v} drawn from this subspace must be $\theta_1, \dots, \theta_k$. \square

As Ritz values are invariant to unitary similarity transformations, the matrix \mathbf{A} and \mathbf{v} in the theorem could be replaced by $\mathbf{Q}^* \mathbf{A} \mathbf{Q}$ and $\mathbf{Q}^* \mathbf{v}$ for any unitary matrix $\mathbf{Q} \in \mathbb{C}^{n \times n}$. Of fundamental importance is the fact that the invariant subspace associated with $\sigma(\mathbf{T})$ be orthogonal to the one associated with $\sigma(\mathbf{D})$. The matrix \mathbf{T} need not be normal nor even diagonalizable.

This result (which relates to the example of maximal Krylov subspaces for derogatory matrices given in [3, p. 1081]) suggests a procedure for manufacturing examples like the one presented in section 2. The following approach presumes that one seeks the largest magnitude eigenvalues of a matrix (as is typical, e.g., after performing a shift-invert or Cayley transformation of the original matrix [10]).

1. Find a matrix \mathbf{D} and starting vector \mathbf{w} such that the Ritz values produced by $k \geq 2m$ steps of the Arnoldi process, when ordered by decreasing magnitude, satisfy

$$(5) \quad |\theta_m| > |\theta_{m+1}| \geq \cdots \geq |\theta_{2m}| > \max_{\lambda \in \sigma(\mathbf{D})} |\lambda|.$$

2. Construct a matrix $\mathbf{T} \in \mathbb{C}^{m \times m}$ with eigenvalues $\theta_{m+1}, \dots, \theta_{2m}$; take $\mathbf{x} \in \mathbb{C}^m$ to be any vector.
3. Build the matrix and starting vector

$$\mathbf{A} = \begin{bmatrix} \mathbf{T} & \mathbf{0} \\ \mathbf{0} & \mathbf{D} \end{bmatrix}, \quad \mathbf{v} = \begin{bmatrix} \mathbf{x} \\ \mathbf{w} \end{bmatrix}.$$

We wish to compute the largest magnitude eigenvalues of \mathbf{A} , which by (5) are the eigenvalues $\theta_{m+1}, \dots, \theta_{2m}$ of \mathbf{T} . Apply k steps of the Arnoldi process, resulting in the same Ritz values obtained in step 1.

4. The exact shift procedure will select Ritz values $\theta_{m+1}, \dots, \theta_k$ as the roots of the filter polynomial. These shifts include the eigenvalues of \mathbf{A} we seek, which are thus deflated from \mathbf{v} and hence cannot be recovered by the restarted iteration.

Examples of this sort rely on Ritz values that fall well beyond all the eigenvalues of \mathbf{D} . If that matrix is normal, then all its Ritz values would be contained in the convex hull of its spectrum. Hence no normal \mathbf{D} will give Ritz values larger in magnitude than (or further to the right of) its eigenvalues, and consequently no such matrix is suitable for use in the above construction. Nonnormality thus plays a central role in examples of this form. Both the example of section 2 and the one we shall next present have been designed so that the eigenvalues of interest are not unduly influenced by this nonnormality and hence would likely be meaningful in applications (e.g., \mathbf{D} may be far from normal, but the desired eigenvalues of \mathbf{A} are well separated from \mathbf{D}). As in [4, section 4.4], nonnormality associated with undesired eigenvalues complicates convergence to desired, ideally conditioned eigenvalues.

4. Larger example. Next we follow the procedure just described to produce a scenario that perhaps appears less contrived than the 4×4 example of section 2. To begin, consider the upper triangular matrix $\mathbf{D} = \mathbf{A} + \beta \mathbf{S}^\gamma$, where

$$\mathbf{A} = \text{diag}(\lambda_1, \dots, \lambda_n), \quad \mathbf{S} = \begin{bmatrix} 0 & s_1 & & & \\ & 0 & \ddots & & \\ & & \ddots & s_{n-1} & \\ & & & & 0 \end{bmatrix},$$

with

$$\lambda_j = -\alpha + \frac{2\alpha(j-1)}{n-1}, \quad s_j = \frac{j}{n-1}$$

for parameters $\alpha, \beta, \gamma \geq 0$. The spectrum of this matrix is uniformly distributed over the interval $[-\alpha, \alpha]$. For $\alpha > 0$ the matrix is diagonalizable, with β and γ controlling the conditioning of the eigenvalues. Qualitatively, the growth of s_j with j causes eigenvalues on the right end of the spectrum to exhibit greater sensitivity than those on the left. Increasing β magnifies the ill-conditioning throughout the spectrum, while increasing γ improves the conditioning, most acutely for the leftmost eigenvalues.

In the example that follows, \mathbf{D} has dimension $n = 100$ with $\alpha = 1/2$ and $\beta = \gamma = 4$. Figure 1(a) shows the spectrum, numerical range (2), and ε -pseudospectra (3) of \mathbf{D} . The matrix exhibits a moderate departure from normality, with the right half of the spectrum especially sensitive.

Figure 1(b) shows $k = 10$ Ritz values for \mathbf{D} with the starting vector $\mathbf{w} = [1, \dots, 1]^T$. If we seek $m = 5$ eigenvalues, then the exact shift strategy would choose

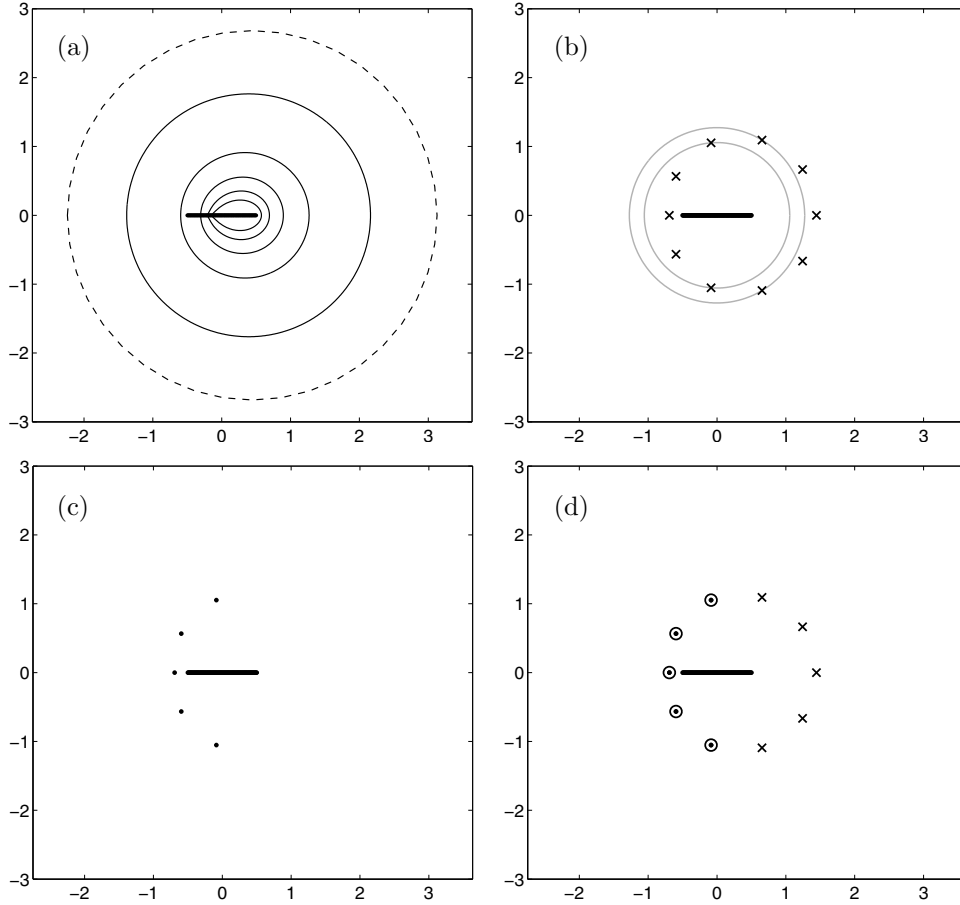


FIG. 1. Illustration of the example discussed in section 4. (a) The eigenvalues (dots covering the segment $[-1/2, 1/2]$), ε -pseudospectra (boundaries shown as solid lines for $\varepsilon = 10^{-1}, 10^{-3}, \dots, 10^{-9}$), and the numerical range (boundary shown as a dashed line) for \mathbf{D} . (b) The eigenvalues of \mathbf{D} with the $k = 10$ Ritz values for \mathbf{D} and \mathbf{w} (crosses); the gray circles indicate the magnitude of the fifth and sixth largest Ritz values. (c) The eigenvalues of \mathbf{A} . (d) The eigenvalues of \mathbf{A} (dots) with $k = 10$ Ritz values for \mathbf{A} and \mathbf{v} (circles and crosses); the smallest five in magnitude (circles) match the largest magnitude eigenvalues of \mathbf{A} and will be used as exact shifts.

the $k - m = 5$ smallest magnitude Ritz values as roots of the filter polynomial; these Ritz values are the leftmost ones shown in the figure. Comparing this with Figure 1(a), notice that these Ritz values are outside the ε -pseudospectra of \mathbf{A} for all $\varepsilon \leq 10^{-3}$; equivalently, the “sep” of all these Ritz values from \mathbf{D} is larger than 10^{-3} .

We now follow the recipe outlined above to obtain an \mathbf{A} and \mathbf{v} for which these shifts would be catastrophic. Let $\mathbf{T} \in \mathbb{C}^{5 \times 5}$ be a diagonal matrix whose eigenvalues equal those five smallest magnitude Ritz values, and set \mathbf{A} to be the 105×105 matrix $\mathbf{A} = \text{diag}(\mathbf{T}, \mathbf{D})$; the eigenvalues of \mathbf{A} are shown in Figure 1(c). By design, the eigenvalues of \mathbf{T} are the largest magnitude eigenvalues of \mathbf{A} , and they all fall beyond the $\varepsilon = 10^{-3}$ pseudospectrum of \mathbf{D} .

Now compute $k = 10$ Ritz values for \mathbf{A} with $\mathbf{v} = [\mathbf{x}^T \mathbf{w}^T]^T$ for any choice of \mathbf{x} , shown in Figure 1(d). (The computation to produce this illustration used $\mathbf{x} = [1, \dots, 1]^T$, but that choice has no influence on the figure.) As ensured by the theorem,

these Ritz values are identical to those obtained from \mathbf{D} and \mathbf{w} . To compute the largest magnitude eigenvalues of \mathbf{A} , the exact shift strategy picks as shifts the five smallest Ritz values—which coincide with the five largest magnitude eigenvalues of \mathbf{A} . Again, the exact shift procedure will deflate precisely those eigenvalues we wish to compute.

5. Discussion. Our aim here has been to address a theoretical question concerning the convergence of the restarted Arnoldi algorithm with exact shifts. In no way do we suggest that the behavior our examples exhibit is commonplace. The exact shift procedure remains the most robust general-purpose method for restarting the Arnoldi method, endorsed by years of successful computation. Our constructions rely on special choices for the starting vector (the components in \mathbf{w}) and a fixed number of steps; by changing \mathbf{w} , or increasing or decreasing k , one may well converge to the correct eigenvalues without incident. Even with problematic \mathbf{w} and k , the rounding errors that occur in practical computations can allow the desired eigenvectors to emerge after numerous additional iterations. (Indeed, when applied to the example in section 4, `eigs` eventually finds the five desired eigenvalues.)

On those occasions when ARPACK fails to converge, the culprit is likely more mundane than the extreme failure exhibited here. For example, the requested eigenvalues may form part of a tight cluster, in which case any shift procedure would struggle to develop, in a tractable number of iterations, a polynomial filter that is large on the desired eigenvalues while being small on those nearby undesired eigenvalues. Such behavior is explained by the conventional restarted Arnoldi convergence analysis cited in the introduction.

The mere existence of examples for which the restarted Arnoldi algorithm with exact shifts deflates the desired eigenvalues underscores the need for a deeper understanding of the behavior of Ritz values of non-Hermitian matrices. Though these eigenvalue estimates must fall within the numerical range, little else is known, deterministically or stochastically, about their distribution. Progress on this important problem could provide a foundation for a robust convergence theory for the restarted Arnoldi algorithm and illuminate many other corners of iterative linear algebra.

Acknowledgments. I thank Rich Lehoucq for drawing my attention to Ruhe's paper [12], and Howard Elman for his kind hospitality during a visit to the University of Maryland, where this work was begun. The pseudospectra in the figure were computed using the EigTool package of T. G. Wright [18].

REFERENCES

- [1] W. E. ARNOLDI, *The principle of minimized iterations in the solution of the matrix eigenvalue problem*, Quart. Appl. Math., 9 (1951), pp. 17–29.
- [2] Z. BAI, J. DEMMEL, AND A. MCKENNEY, *On computing condition numbers for the nonsymmetric eigenproblem*, ACM Trans. Math. Software, 19 (1993), pp. 202–223.
- [3] C. BEATTIE, M. EMBREE, AND J. ROSSI, *Convergence of restarted Krylov subspaces to invariant subspaces*, SIAM J. Matrix Anal. Appl., 25 (2004), pp. 1074–1109.
- [4] C. A. BEATTIE, M. EMBREE, AND D. C. SORENSEN, *Convergence of polynomial restart Krylov methods for eigenvalue computations*, SIAM Review, 47 (2005), pp. 492–515.
- [5] R. A. HORN AND C. R. JOHNSON, *Topics in Matrix Analysis*, Cambridge University Press, Cambridge, UK, 1991.
- [6] Z. JIA, *The convergence of generalized Lanczos methods for large unsymmetric eigenproblems*, SIAM J. Matrix Anal. Appl., 16 (1995), pp. 843–862.
- [7] L. KNIZHNERMAN, *Error bounds for the Arnoldi method: A set of extreme eigenpairs*, Linear Algebra Appl., 296 (1999), pp. 191–211.
- [8] R. B. LEHOUCQ, *Implicitly restarted Arnoldi methods and subspace iteration*, SIAM J. Matrix Anal. Appl., 23 (2001), pp. 551–562.

- [9] R. B. LEHOUCQ, D. C. SORENSEN, AND C. YANG, *ARPACK Users' Guide: Solution of Large-Scale Eigenvalue Problems with Implicitly Restarted Arnoldi Methods*, Software Environ. Tools 6, SIAM, Philadelphia, 1998.
- [10] K. MEERBERGEN, A. SPENCE, AND D. ROOSE, *Shift-invert and Cayley transforms for detection of rightmost eigenvalues of nonsymmetric matrices*, BIT, 34 (1994), pp. 409–423.
- [11] B. N. PARLETT, *The Symmetric Eigenvalue Problem*, Prentice–Hall, Englewood Cliffs, NJ, 1980.
- [12] A. RUHE, *Rational Krylov sequence methods for eigenvalue computation*, Linear Algebra Appl., 58 (1984), pp. 391–405.
- [13] Y. SAAD, *Variations on Arnoldi's method for computing eigenelements of large unsymmetric matrices*, Linear Algebra Appl., 34 (1980), pp. 269–295.
- [14] D. C. SORENSEN, *Implicit application of polynomial filters in a k-step Arnoldi method*, SIAM J. Matrix Anal. Appl., 13 (1992), pp. 357–385.
- [15] G. W. STEWART, *A Krylov–Schur algorithm for large eigenproblems*, SIAM J. Matrix Anal. Appl., 23 (2001), pp. 601–614.
- [16] G. W. STEWART, *Matrix Algorithms Volume II: Eigensystems*, SIAM, Philadelphia, 2001.
- [17] L. N. TREFETHEN AND M. EMBREE, *Spectra and Pseudospectra: The Behavior of Nonnormal Matrices and Operators*, Princeton University Press, Princeton, NJ, 2005.
- [18] T. G. WRIGHT, *EigTool*, 2002. Software available at <http://www.comlab.ox.ac.uk/pseudospectra/eigtool>.