

# HOW DESCRIPTIVE ARE GMRES CONVERGENCE BOUNDS?

MARK EMBREE\*

**Abstract.** GMRES is a popular Krylov subspace method for solving linear systems of equations involving a general non-Hermitian coefficient matrix. The conventional bounds on GMRES convergence involve polynomial approximation problems in the complex plane. Three popular approaches pose this approximation problem on the spectrum, the field of values, or pseudospectra of the coefficient matrix. We analyze and compare these bounds, illustrating with six examples the success and failure of each. When the matrix departs from normality due only to a low-dimensional invariant subspace, we discuss how these bounds can be adapted to exploit this structure. Since the Arnoldi process that underpins GMRES provides approximations to the pseudospectra, one can estimate the GMRES convergence bounds as an iteration proceeds.

**Key words.** Krylov subspace methods, GMRES convergence, nonnormal matrices, pseudospectra, field of values

**1. Introduction.** Many algorithms for solving large, sparse systems of linear equations construct iterates that attempt to minimize the residual norm over all candidates in an affine Krylov subspace whose dimension grows at each step. For non-Hermitian matrices, the GMRES algorithm of Saad and Schultz [45] generates such optimal iterates. This method, which uses the Arnoldi process to generate an orthonormal basis for the Krylov subspace, becomes intractable for problems that converge slowly. Practical algorithms, such as restarted GMRES, BiCGSTAB, QMR (see, e.g., [24, Ch. 5], [43, Ch. 7]), reduce this computational expense but compromise the optimality, and it is tough to characterize the convergence that results. (For an indication of the complexity of restarted GMRES, see [15].) The residual norms cannot be smaller than those produced by GMRES, since the algorithms choose iterates from the same Krylov subspace; sometimes they can be related to the GMRES residual norm, as in the case of QMR [21]. Understanding GMRES convergence, facilitated by its optimality property, can thus be a step toward convergence analysis for other algorithms. Insight about GMRES convergence can also inform the construction and evaluation of preconditioners for non-Hermitian matrices.

Given a system of linear equations  $\mathbf{Ax} = \mathbf{b}$ , with  $\mathbf{A} \in \mathbb{C}^{n \times n}$  and  $\mathbf{x}, \mathbf{b} \in \mathbb{C}^n$ , the GMRES algorithm [45] iteratively generates solution estimates  $\mathbf{x}_k \approx \mathbf{x}$  based on an initial guess  $\mathbf{x}_0$ . The residuals induced by these iterates,  $\mathbf{r}_k := \mathbf{b} - \mathbf{Ax}_k$ , satisfy the minimum residual property,

$$(1.1) \quad \|\mathbf{r}_k\|_2 = \min_{\substack{p \in \mathcal{P}_k \\ p(0)=1}} \|p(\mathbf{A})\mathbf{r}_0\|_2,$$

where  $\mathcal{P}_k$  denotes the set of polynomials of degree  $k$  or less.

What properties of the coefficient matrix  $\mathbf{A}$  govern convergence? In this work, we examine three proposed answers to this question (see [24]): eigenvalues with eigenvector condition number, the field of values, and pseudospectra. When  $\mathbf{A}$  is normal (i.e., it has an orthogonal basis of eigenvectors or, equivalently, it commutes with its adjoint), convergence can be accurately bounded using the eigenvalues alone. This is not the case for nonnormal matrices, as the construction of Greenbaum, Pták, and

---

\*Department of Mathematics, Virginia Tech, Blacksburg, VA 24061 ([embree@vt.edu](mailto:embree@vt.edu)). The original version of this work appeared as Oxford University Computing Laboratory Technical Report TR 99/08 in June 1999. This updated version corrects some errors, improves the presentation, and includes additional illustrations. (August 2022)

Strakoš vividly illustrates [25]: given any set of eigenvalues  $\sigma(\mathbf{A})$ , one can construct  $\mathbf{A}$  and  $\mathbf{r}_0$  to produce any (monotonically decreasing) GMRES residual norms. When  $\mathbf{A}$  is far from normal, the residual norms  $\{\|\mathbf{r}_k\|_2\}$  often exhibits a period of initial stagnation before converging at a quicker asymptotic rate. The bounds we study here essentially differ in two ways: how they account for this delay due to nonnormality, and the sets upon which they base the asymptotic rate convergence.

Section 2 describes three standard GMRES convergence bounds. These characterizations can significantly overestimate the norm of the residual when nonnormality is only associated with a few eigenvalues (e.g., several nearly aligned eigenvectors orthogonal to all other eigenvectors). To circumvent this shortcoming, we apply spectral projectors to modify the traditional formulations. This strategy bounds GMRES convergence using the condition numbers of individual eigenvalues, and also leads to flexible generalizations of the field of values and pseudospectra bounds. In Section 3, we present six examples to illustrate that the three standard bounds can each give significant overestimates, but each can also be rather descriptive. The bounds are also compared via the relationships between the eigenvectors, field of values, and pseudospectra. These examples highlight a strength of pseudospectral bounds (elaborating an observation of Driscoll, Toh, and Trefethen [9, p. 564]): by sampling the bound based on the  $\varepsilon$ -pseudospectrum over a range of  $\varepsilon$  values, one can potentially capture *different phases of convergence* via the envelope of these bounds. While the cost of computing pseudospectra can deter the use of this bound for large-scale problems, in Section 4 we suggest an approach for obtaining GMRES convergence *estimates* at a lower computational expense based on approximate pseudospectra taken from the Hessenberg matrix generated by the Arnoldi process within the standard GMRES implementation.

Though we are concerned with GMRES convergence for a linear system with a specific initial residual, all the analysis described here first employs the inequality

$$(1.2) \quad \|\mathbf{r}_k\|_2 \leq \min_{\substack{p \in \mathcal{P}_k \\ p(0)=1}} \|p(\mathbf{A})\|_2 \|\mathbf{r}_0\|_2,$$

and then studies  $\|p(\mathbf{A})\|_2$  independent from  $\mathbf{r}_0$ . This approach leads to upper bounds for worst case GMRES convergence. With a carefully crafted example, Toh proved that this inequality can be arbitrarily misleading for some nonnormal matrices [51]. There may be *no* vector  $\mathbf{r}_0 \in \mathbb{C}^n$  for which  $\|p(\mathbf{A})\|_2$  equals  $\|\mathbf{r}_k\|_2/\|\mathbf{r}_0\|_2$  at iteration  $k$ . Examples of this extreme behavior are thought to be rare in practice [50, §3.6] and thus we are typically content to make the inequality (1.2) and proceed with the analysis of  $\|p(\mathbf{A})\|_2$  that follows from it.

**2. Three Convergence Bounds and Variations.** In this section, we describe three common convergence bounds for GMRES based on eigenvalues with the eigenvector condition number, the field of values, and pseudospectra. These bounds all fail to describe convergence accurately when nonnormality is primarily associated with a few eigenvalues. One can use spectral projectors to decouple sets of eigenvalues, leading to localized versions of these bounds that can be sharper than the conventional versions.

**2.1. Eigenvalues with Eigenvector Conditioning.** The first convergence bound suggested for GMRES predicts convergence at a rate determined by the set of eigenvalues of  $\mathbf{A}$ , denoted  $\sigma(\mathbf{A})$ . If  $\mathbf{A}$  is normal,  $\sigma(\mathbf{A})$  determines convergence, since

for any polynomial  $p$

$$\|p(\mathbf{A})\|_2 = \max_{\lambda \in \sigma(\mathbf{A})} |p(\lambda)|.$$

Nonnormality can impede the onset of convergence at this spectral rate; to account for such delay, this bound scales the spectral convergence behavior by the condition number of the matrix having the eigenvectors of  $\mathbf{A}$  as its columns [14, 45]. Provided that  $\mathbf{A}$  is diagonalizable,  $\mathbf{A} = \mathbf{V}\mathbf{\Lambda}\mathbf{V}^{-1}$ ,

$$\begin{aligned} \|\mathbf{r}_k\|_2 &= \min_{\substack{p \in \mathcal{P}_k \\ p(0)=1}} \|p(\mathbf{A})\mathbf{r}_0\|_2 \leq \|\mathbf{V}p(\mathbf{\Lambda})\mathbf{V}^{-1}\|_2 \|\mathbf{r}_0\|_2 \\ &\leq \|\mathbf{V}\|_2 \|\mathbf{V}^{-1}\|_2 \|p(\mathbf{\Lambda})\|_2 \|\mathbf{r}_0\|_2, \end{aligned}$$

implying the bound

$$(EV) \quad \frac{\|\mathbf{r}_k\|_2}{\|\mathbf{r}_0\|_2} \leq \kappa(\mathbf{V}) \min_{\substack{p \in \mathcal{P}_k \\ p(0)=1}} \max_{\lambda \in \sigma(\mathbf{A})} |p(\lambda)|.$$

Here,  $\kappa(\mathbf{V}) := \|\mathbf{V}\|_2 \|\mathbf{V}^{-1}\|_2$  is the 2-norm condition number of the eigenvector matrix  $\mathbf{V}$ . If  $\mathbf{A}$  is normal, then  $\kappa(\mathbf{V}) = 1$ ; if, in addition, the eigenvalues are real, then (EV) reduces to the standard convergence bound for MINRES [18]. If  $\mathbf{A}$  is nonnormal, then  $\kappa(\mathbf{V}) > 1$  and determining the optimal value of  $\kappa(\mathbf{V})$  can be a challenge [26]; this task is further complicated if  $\mathbf{A}$  has repeated eigenvalues. Throughout this work, we scale the columns of  $\mathbf{V}$  to have unit 2-norm; provided each eigenvalue of  $\mathbf{A}$  is simple, this scaling ensures that  $\kappa(\mathbf{V})$  is no more than  $\sqrt{n}$  times its optimal value, where  $n$  is the matrix dimension [60].

Since  $\sigma(\mathbf{A})$  is a discrete point set for finite-dimensional matrices, the polynomial approximation problem in (EV) will be zero when  $k$  reaches the matrix dimension,  $k = n$ : *thus the bound (EV) captures the finite termination of GMRES*. However (despite the small size of most of our examples), we want to apply GMRES for large  $n$ , in the hope of obtaining convergence after  $k \ll n$  iterations. Thus the bound (EV) is often employed with  $\sigma(\mathbf{A})$  replaced by some compact set  $\Omega \supset \sigma(\mathbf{A})$ . For example, if all eigenvalues of  $\mathbf{A}$  are real and positive, one might take  $\Omega$  to be the real interval connecting the extreme eigenvalues, as is common in analysis of the conjugate gradient method; see, e.g., [58, chap. 38]. (Better bounds result from including outlying eigenvalues as singletons, and bounding the rest of the spectrum in aggregate [61, 62].)

The constant  $\kappa(\mathbf{V})$  in (EV) reflects the departure of  $\mathbf{A}$  from nonnormality. The normalized residual norms  $\|\mathbf{r}_k\|_2 / \|\mathbf{r}_0\|_2$  form a nonincreasing sequence starting with the value 1 at  $k = 0$ . Thus if  $\kappa(\mathbf{V})$  is large, the bound (EV) cannot describe convergence at least until the iteration  $k$  at which the polynomial minimization term is as small as  $1/\kappa(\mathbf{V})$ . Even then, (EV) can be grossly inaccurate. For example,  $\kappa(\mathbf{V})$  can be large because all the eigenvectors are ill-conditioned, or if only two eigenvectors are nearly aligned. In the latter case, the bound usually fails to predict convergence, while it may be more appropriate in the former case. These situations are illustrated in Examples B and E of Section 3. (For a discussion about the shortcomings of scalar measures of nonnormality see [59, chap. 48].)

In an effort to avoid this difficulty, the bound (EV) can be adapted by considering the conditioning of individual eigenvalues. Suppose  $\lambda \in \sigma(\mathbf{A})$  is simple, having left and right eigenvectors  $\hat{\mathbf{v}}$  and  $\mathbf{v}$ . Then the condition number of  $\lambda$  [63, §2.8] is

$$\kappa(\lambda) := \frac{\|\hat{\mathbf{v}}\|_2 \|\mathbf{v}\|_2}{|\hat{\mathbf{v}}^* \mathbf{v}|} = \frac{1}{\cos \angle(\hat{\mathbf{v}}, \mathbf{v})}.$$

Using these condition numbers leads to a bound that can be much sharper than (EV). This result can be seen as a special case of a theorem of Joubert [34, Thm. 3.2(4)], who presents it in the context of the Jordan canonical form.

**THEOREM 2.1.** *Suppose every eigenvalue  $\lambda_j$  of  $\mathbf{A}$  is simple. Then for any  $p \in \mathcal{P}_k$ ,*

$$(2.1) \quad \|p(\mathbf{A})\|_2 \leq \sum_{j=1}^n \kappa(\lambda_j) |p(\lambda_j)|.$$

*Proof.* Since  $\mathbf{A}$  has simple eigenvalues, it is diagonalizable,  $\mathbf{A} = \mathbf{V}\mathbf{\Lambda}\mathbf{V}^{-1}$ . Let  $\{\widehat{\mathbf{v}}_j\}_{j=1}^n$  be the left eigenvectors ( $\widehat{\mathbf{v}}_j^*$  is the  $j$ th row of  $\mathbf{V}^{-1}$ ) and  $\{\mathbf{v}_j\}_{j=1}^n$  the corresponding right eigenvectors (columns of  $\mathbf{V}$ ), with  $\mathbf{\Lambda}_{jj} = \lambda_j$ . Then

$$\|p(\mathbf{A})\|_2 = \|\mathbf{V}p(\mathbf{\Lambda})\mathbf{V}^{-1}\|_2 = \left\| \sum_{j=1}^n p(\lambda_j) \mathbf{v}_j \widehat{\mathbf{v}}_j^* \right\|_2 \leq \sum_{j=1}^n |p(\lambda_j)| \|\mathbf{v}_j \widehat{\mathbf{v}}_j^*\|_2.$$

The result follows from noting that since  $\widehat{\mathbf{v}}_j^* \mathbf{v}_j = 1$  by construction,  $\|\mathbf{v}_j \widehat{\mathbf{v}}_j^*\|_2 = \|\widehat{\mathbf{v}}_j\|_2 \|\mathbf{v}_j\|_2 = \|\widehat{\mathbf{v}}_j\|_2 \|\mathbf{v}_j\|_2 / |\widehat{\mathbf{v}}_j^* \mathbf{v}_j| = \kappa(\lambda_j)$ .  $\square$

The quantity on the right of equation (2.1) is the 1-norm of  $p(\mathbf{\Lambda})\mathbf{c}$  where  $c_j = \kappa(\lambda_j)$ . Using a norm equivalence, we can obtain a conventional GMRES problem involving a normal matrix, but with a special right hand side.

**COROLLARY 2.2.** *Define the components of  $\mathbf{c} \in \mathbb{C}^n$  to be  $c_j = \kappa(\lambda_j)$ . Then*

$$(EV') \quad \frac{\|\mathbf{r}_k\|_2}{\|\mathbf{r}_0\|_2} \leq \sqrt{n} \min_{\substack{p \in \mathcal{P}_k \\ p(0)=1}} \|p(\mathbf{\Lambda})\mathbf{c}\|_2.$$

Note that  $\|\mathbf{c}\|_2 \geq \sqrt{n}$ , since each  $c_j \geq 1$ . If  $c_j$  is large, the corresponding eigenvalue  $\lambda_j$  is ill-conditioned. If only a few eigenvalues of  $\mathbf{A}$  are ill-conditioned, then the GMRES problem on the right-hand side of (EV') can exhibit an initial phase of rapid convergence (as one could choose a polynomial  $p \in \mathcal{P}_k$  with roots at the ill-conditioned eigenvalues, eliminating those components from  $\mathbf{c}$ ), leading to more rapid convergence than one would expect for a typical initial residual of similar magnitude. Figure 4.5 shows the ability of (EV') to describe convergence for a highly nonnormal matrix from a convection-diffusion problem.

When  $\kappa(\mathbf{V})$  is large (or  $\mathbf{A}$  is nondiagonalizable, which we regard as  $\kappa(\mathbf{V}) = \infty$ ), it may be appealing to obtain GMRES bounds with smaller leading constants, at the expense of polynomial approximation problems on larger sets in the complex plane.

**2.2. Field of Values.** The *field of values* (or *numerical range*),

$$(2.2) \quad W(\mathbf{A}) := \left\{ \frac{\mathbf{x}^* \mathbf{A} \mathbf{x}}{\mathbf{x}^* \mathbf{x}} : \mathbf{x} \in \mathbb{C}^n, \mathbf{x} \neq 0 \right\},$$

is a popular alternative to eigenvalues for understanding the behavior of functions of nonnormal matrices. The field of values is always a closed, convex set that contains  $\sigma(\mathbf{A})$ , and so it is possible that  $0 \in W(\mathbf{A})$  even when  $\mathbf{A}$  is nonsingular. When  $\mathbf{A}$  is far from normal,  $W(\mathbf{A})$  can contain points far beyond the convex hull of  $\sigma(\mathbf{A})$ . While the eigenvalues of  $\mathbf{A}$  can be sensitive to perturbations, the field of values is robust:

$$W(\mathbf{A} + \mathbf{E}) \subseteq W(\mathbf{A}) + \{z \in \mathbb{C} : |z| \leq \|\mathbf{E}\|_2\}.$$

Moreover, the extreme eigenvalues of the Hermitian  $((\mathbf{A} + \mathbf{A}^*)/2)$  and skew-Hermitian  $(\mathbf{A} - \mathbf{A}^*)/2$  parts of  $\mathbf{A}$  give tight bounds on the real and imaginary extent of  $W(\mathbf{A})$ ; see [32, chap. 2] for the basic algorithm for computing  $W(\mathbf{A})$ , and Bracconier and Higham [3] for an approach for large-scale problems. For many of the examples that follow, we use Higham's `fv` code [30]. For additional properties of  $W(\mathbf{A})$ , see [28], [32, chap. 1], and [59, chap. 17].

To develop a GMRES bound based on the field of values, we seek an alternative to (EV) that replaces maximization over  $\sigma(\mathbf{A})$  by maximization over  $W(\mathbf{A})$ . Enlarging this maximizing set generally increases the polynomial approximation term, which will hopefully be counterbalanced by a leading constant that is smaller than  $\kappa(\mathbf{V})$  in (EV).

How does  $\|p(\mathbf{A})\|$  relate to  $|p(z)|$  for  $z \in W(\mathbf{A})$ ? This question dates back at least to the 1960s, but major progress has been made in recent years by Crouzeix and collaborators. Indeed, Crouzeix has shown that for any  $p \in \mathcal{P}_k$ , there exists a constant  $C_{\text{FOV}}$  (independent of  $\mathbf{A}$ , the degree  $k$ , and the dimension  $n$ ) such that

$$(2.3) \quad \|p(\mathbf{A})\| \leq C_{\text{FOV}} \max_{z \in W(\mathbf{A})} |p(z)|.$$

How large is  $C_{\text{FOV}}$ ? Crouzeix [4] proved that  $2 \leq C_{\text{FOV}} \leq 11.08$  and conjectured that  $C_{\text{FOV}} = 2$ . More recently, Crouzeix and Palencia [6] showed that  $2 \leq C_{\text{FOV}} \leq 1 + \sqrt{2}$ . We thus have

$$(FOV) \quad \frac{\|\mathbf{r}_k\|_2}{\|\mathbf{r}_0\|_2} \leq C_{\text{FOV}} \min_{\substack{p \in \mathcal{P}_k \\ p(0)=1}} \max_{z \in W(\mathbf{A})} |p(z)|,$$

with  $C_{\text{FOV}} \leq 1 + \sqrt{2}$ . We note that this bound complements important earlier work by Eiermann.<sup>1</sup> While Crouzeix's conjecture is known to hold for certain classes of matrices, in all the illustrations that follow we use  $C_{\text{FOV}} = 1 + \sqrt{2} = 2.4142\dots$

The bound (FOV) has a small leading constant and is appealingly simple, but it suffers from several notable limitations.

1. It is possible that  $0 \in W(\mathbf{A})$  even when  $\mathbf{A}$  is nonsingular ( $0 \notin \sigma(\mathbf{A})$ ). Then

$$\min_{\substack{p \in \mathcal{P}_k \\ p(0)=1}} \max_{z \in W(\mathbf{A})} |p(z)| \geq \min_{\substack{p \in \mathcal{P}_k \\ p(0)=1}} |p(0)| = 1,$$

and so the bound (FOV) fails to describe any convergence, despite the fact that the full GMRES algorithm must converge. This limitation makes (FOV) unsuitable for Hermitian indefinite  $\mathbf{A}$ , not to mention nonnormal problems for which 0 is embedded within  $W(\mathbf{A})$  despite the spectrum being relatively well separated from the origin.

2. Since  $W(\mathbf{A})$  is a convex set in  $\mathbb{C}$ , it hides information about the *distribution of the eigenvalues* within  $W(\mathbf{A})$ . Suppose  $\mathbf{A}$  is Hermitian with  $W(\mathbf{A}) = [\alpha, \beta] \subset \mathbb{R}$  for  $0 < \alpha < \beta$ . GMRES will converge very differently if  $\sigma(\mathbf{A})$  is uniformly distributed throughout  $[\alpha, \beta]$ , or if  $\sigma(\mathbf{A})$  consists of a cluster near  $\alpha$  and a single eigenvalue near  $\beta$ , yet both scenarios give the same  $W(\mathbf{A})$ . The bound (FOV) cannot capture so-called *superlinear* convergence effects associated with isolated outlying eigenvalues [9, 62].

---

<sup>1</sup>To the best of our knowledge, Eiermann was the first to use the field of values to bound the convergence of iterative linear solvers [11, 12]. Working before Crouzeix's analysis, he bounded  $W(\mathbf{A})$  by an ellipse in  $\mathbb{C}$ , then used Chebyshev polynomials to bound the optimal polynomial on this ellipse. Eiermann and Ernst proposed a different field of values bound involving  $W(\mathbf{A}^{-1})$  in [13, sect. 6].

3. Similarly, suppose  $W(\mathbf{A})$  contains points far from  $\sigma(\mathbf{A})$  due to a large departure from normality associated with a low-degree invariant subspace of  $\mathbf{A}$ . In the GMRES approximation problem, the optimizing polynomial  $p$  could target these ill-conditioned eigenvalues, effectively eliminating the nonnormality from the problem at an early iteration: later iterations can focus on the rest of the spectrum. In contrast, the bound (FOV) must continue to optimize  $p$  over all  $W(\mathbf{A})$  at every iteration.

4. As the constant term  $C_{\text{FOV}}$  in (FOV) is small and the approximation problem on a convex region predicts asymptotic linear convergence (see equation (2.11)), the bound (FOV) cannot be entirely descriptive for iterations that initially stagnate before converging at a more rapid asymptotic rate. This behavior is observed by Higham and Trefethen for matrix powers [29], and identified by Ernst in the context of GMRES applied to convection-diffusion problems [17], as we see in Section 4.1.

We can remedy some limitations of (FOV) by working with projectors onto invariant subspaces of  $\mathbf{A}$ . Partition the spectrum of  $\mathbf{A}$  into disjoint sets  $\Lambda_j$ , such that  $\sigma(\mathbf{A}) = \cup_{j=1}^m \Lambda_j$ . Define the spectral projector

$$\mathbf{P}_j := \frac{1}{2\pi i} \int_{\Gamma_j} (z\mathbf{I} - \mathbf{A})^{-1} dz,$$

where  $\Gamma_j$  is the union of Jordan curves containing the eigenvalues  $\Lambda_j$  in their collective interior, but not enclosing any other eigenvalues. Then  $\mathbf{P}_j$  is a projector onto the invariant subspace of  $\mathbf{A}$  associated with the eigenvalues  $\Lambda_j$  (see, e.g., [35, sect. I.5.3]).

**THEOREM 2.3.** *Let  $\{\Lambda_j\}_{j=1}^m$  be a partition of  $\sigma(\mathbf{A})$  into  $m$  disjoint sets. For each  $1 \leq j \leq m$ , let  $\mathbf{P}_j$  be the spectral projector onto the invariant subspace associated with  $\Lambda_j$ , and let the columns of  $\mathbf{U}_j^{n \times \text{rank}(\mathbf{P}_j)}$  be an orthonormal basis for  $\text{Ran}(\mathbf{P}_j)$ . Then for any polynomial  $p \in \mathcal{P}_k$ ,*

$$\|p(\mathbf{A})\|_2 \leq \sum_{j=1}^m \|\mathbf{P}_j\|_2 \|p(\mathbf{U}_j^* \mathbf{A} \mathbf{U}_j)\|_2.$$

*Proof.* Note that  $\mathbf{\Pi}_j := \mathbf{U}_j \mathbf{U}_j^*$  is the orthogonal projector onto  $\text{Ran}(\mathbf{P}_j)$ , an invariant subspace of  $\mathbf{A}$ . (If  $\mathbf{A}$  is normal, then  $\mathbf{\Pi}_j = \mathbf{P}_j$ ; for nonnormal  $\mathbf{A}$ , we generally have  $\mathbf{\Pi}_j \neq \mathbf{P}_j$  though both  $\mathbf{\Pi}_j$  and  $\mathbf{P}_j$  are projectors onto the same subspace.) We apply several important identities for spectral projectors:  $\sum_{j=1}^m \mathbf{P}_j = \mathbf{I}$ ,  $\mathbf{A} \mathbf{P}_j = \mathbf{A} \mathbf{\Pi}_j \mathbf{P}_j$ , and  $\mathbf{A} \mathbf{\Pi}_j = \mathbf{\Pi}_j \mathbf{A} \mathbf{\Pi}_j$  (see, e.g., [35, §I.5.3], [44, § 3.1]). Substituting the first identity into  $\|p(\mathbf{A})\|_2$  yields

$$\begin{aligned} \|p(\mathbf{A})\|_2 &= \left\| p(\mathbf{A}) \sum_{j=1}^m \mathbf{P}_j \right\|_2 \leq \sum_{j=1}^m \|p(\mathbf{A}) \mathbf{P}_j\|_2 = \sum_{j=1}^m \|\mathbf{\Pi}_j p(\mathbf{A}) \mathbf{\Pi}_j \mathbf{P}_j\|_2 \\ (2.4) \qquad &\leq \sum_{j=1}^m \|\mathbf{\Pi}_j p(\mathbf{A}) \mathbf{\Pi}_j\|_2 \|\mathbf{P}_j\|_2. \end{aligned}$$

Notice that for each  $j$ ,  $\|\mathbf{\Pi}_j p(\mathbf{A}) \mathbf{\Pi}_j\|_2 = \|p(\mathbf{\Pi}_j \mathbf{A} \mathbf{\Pi}_j)\|_2 = \|\mathbf{U}_j p(\mathbf{U}_j^* \mathbf{A} \mathbf{U}_j) \mathbf{U}_j^*\|_2 \leq \|p(\mathbf{U}_j^* \mathbf{A} \mathbf{U}_j)\|_2$ . Using this bound in (2.4) completes the proof.  $\square$

Joubert [34, Thm. 3.2(4)] presents a similar result that is presented using the language of the Jordan canonical form rather than spectral projectors.

Theorem 2.3 provides a natural tool for transitioning between global statements like (EV) and localized statements like (EV'). In the former case, Theorem 2.3 is vacuous since the spectrum is partitioned into a single set; in the latter case, Theorem 2.3 reduces to Theorem 2.1: each set  $\Lambda_j$  is a single eigenvalue, which implies  $\|\mathbf{P}_j\|_2 = \kappa(\lambda_j)$ . In summary, *one can eliminate the nonnormal coupling between sets of eigenvalues at the cost of scaling by the norm of the associated spectral projectors.*

Theorem 2.3 can be combined with the analysis leading to the bound (FOV) to provide a field of values analogue to (EV').

**COROLLARY 2.4.** *Partition the eigenvalues  $\sigma(\mathbf{A})$  into disjoint sets,  $\{\Lambda_j\}_{j=1}^m$ , with the columns of each  $\mathbf{U}_j$  giving an orthonormal basis for the invariant subspace of  $\mathbf{A}$  associated with  $\Lambda_j$ . Then*

$$(FOV') \quad \frac{\|\mathbf{r}_k\|_2}{\|\mathbf{r}_0\|_2} \leq (1 + \sqrt{2}) \min_{\substack{p \in \mathcal{P}_k \\ p(0)=1}} \sum_{j=1}^m \left( \|\mathbf{P}_j\|_2 \max_{z \in W(\mathbf{U}_j^* \mathbf{A} \mathbf{U}_j)} |p(z)| \right).$$

This localization procedure enables one to work around numerous scenarios in which  $0 \in W(\mathbf{A})$ , but some matrices remain out of reach. For example, if  $\mathbf{A}$  is a matrix with a single eigenvalue, Corollary 2.4 does not permit any splitting of the invariant subspaces (though similar ideas could be applied if a nontrivial block diagonalization is possible [34, Thm. 3.2(4)]). If additionally  $0 \in W(\mathbf{A})$ , analysis based on the field of values cannot predict any convergence.

Recent analysis by Crouzeix and Greenbaum [5] provides a different way to handle cases where  $0 \in W(\mathbf{A})$ . Let  $\mu(\mathbf{A}) := \max_{z \in W(\mathbf{A})} |z|$  denote the *numerical radius* of  $\mathbf{A}$ . Suppose  $0 \in W(\mathbf{A})$  and let us denote by  $\Omega_{CG}$  the region formed by the intersection of  $W(\mathbf{A})$  with the exterior of the disk centered at the origin with radius  $1/\mu(\mathbf{A}^{-1})$ :

$$(2.5) \quad \Omega_{CG} := W(\mathbf{A}) \cap \{z \in \mathbb{C} : |z| \geq 1/\mu(\mathbf{A}^{-1})\}.$$

This set effectively results from “carving out” from  $W(\mathbf{A})$  the disk of radius  $1/\mu(\mathbf{A}^{-1})$  centered at the origin. Crouzeix and Greenbaum [5, p. 1098] show that  $\Omega_{CG}$  is a  $(3 + 2\sqrt{3})$ -spectral set, giving the GMRES bound

$$(CG) \quad \frac{\|\mathbf{r}_k\|_2}{\|\mathbf{r}_0\|_2} \leq (3 + 2\sqrt{3}) \min_{\substack{p \in \mathcal{P}_k \\ p(0)=1}} \max_{z \in \Omega_{CG}} |p(z)|.$$

As noted in [5], when  $\Omega_{CG}$  surrounds the origin the polynomial approximation problem in (CG) must equal 1 (due to the maximum-modulus theorem); see the left plot in Figure 3.11 for an example. If  $\Omega_{CG}$  does not surround the origin – as can occur even in some cases where  $\mathbf{A}$  has just one eigenvalue, as in the right plot in Figure 3.11 – the bound (CG) gives convergence at an asymptotic rate determined by  $\Omega_{CG}$ . However, since  $3 + 2\sqrt{3} \approx 6.464$  is still a small constant, like (FOV) the bound (CG) will not describe iterations that initially converge slowly, before accelerating at later iterations. The bound (CG) seems to be especially useful when  $0 \in W(\mathbf{A})$  because of the requisite convexity of  $W(\mathbf{A})$ , rather than nonnormality. For example, suppose

$$(2.6) \quad \mathbf{A} = \mathbf{A}_{-1} \oplus \mathbf{A}_{+1} := \begin{bmatrix} -1 & 1/4 & & & \\ & -1 & \ddots & & \\ & & \ddots & 1/4 & \\ & & & -1 & \\ & & & & -1 \end{bmatrix} \oplus \begin{bmatrix} +1 & 1/2 & & & \\ & +1 & \ddots & & \\ & & \ddots & 1/2 & \\ & & & -1 & \\ & & & & +1 \end{bmatrix},$$




FIG. 2.1. For the matrix  $\mathbf{A} = \mathbf{A}_{-1} \oplus \mathbf{A}_{+1}$  in (2.6): on the left,  $\sigma(\mathbf{A}) = \{\pm 1\}$  (dots),  $W(\mathbf{A})$  (gray region), the boundaries of  $W(\mathbf{A}_{-1})$  and  $W(\mathbf{A}_{+1})$  (black lines), and the circle of radius  $1/\mu(\mathbf{A}^{-1})$  (red line); on the right, the Crouzeix–Greenbaum region  $\Omega_{CG}$ . Note that  $0 \in W(\mathbf{A})$  but  $0 \notin \Omega_{CG}$ .

where  $\mathbf{A}_{-1}$  and  $\mathbf{A}_{+1}$  are bidiagonal matrices of dimension  $n/2 = 32$ . Though the matrices in the direct sum are both far from normal,  $0 \notin W(\mathbf{A}_{\pm 1})$ . However, since  $W(\mathbf{A})$  is convex,  $0 \in W(\mathbf{A})$ ; see Figure 2.1. The set  $\Omega_{CG}$  remedies this “incidental” inclusion of the origin in  $W(\mathbf{A})$ , and (CG) gives a convergent bound. (Figure 2.1 shows several interesting features: the left lobe of  $\Omega_{CG}$  is larger than  $W(\mathbf{A}_{-1})$ , and moving the eigenvalue  $-1$  further to the left would enlarge this component; the right lobe of  $\Omega_{CG}$  omits some points in  $W(\mathbf{A}_{+1})$ .) In contrast, for this example the bound (FOV’) would give

$$\frac{\|\mathbf{r}_k\|_2}{\|\mathbf{r}_0\|_2} \leq (1 + \sqrt{2}) \min_{\substack{p \in \mathcal{P}_k \\ p(0)=1}} \left( \left( \max_{z \in W(\mathbf{A}_{-1})} |p(z)| \right) + \left( \max_{z \in W(\mathbf{A}_{+1})} |p(z)| \right) \right),$$

since in this case the spectral projectors for the two distinct eigenvalues are orthogonal.

**2.3. Pseudospectra.** Pseudospectra provide another generalization of  $\sigma(\mathbf{A})$  upon which to base GMRES convergence bounds. The  $\varepsilon$ -pseudospectrum [56, 59]  $\sigma_\varepsilon(\mathbf{A})$  has as its boundary the  $\varepsilon$ -level set of the norm of the resolvent:

$$\sigma_\varepsilon(\mathbf{A}) := \{z \in \mathbb{C} : \|(z\mathbf{I} - \mathbf{A})^{-1}\|_2 > \varepsilon^{-1}\}.$$

The  $\varepsilon$ -pseudospectrum can equivalently be defined in terms of eigenvalues of perturbations of  $\mathbf{A}$ :  $\sigma_\varepsilon(\mathbf{A}) = \{z \in \mathbb{C} : z \in \sigma(\mathbf{A} + \mathbf{E}), \|\mathbf{E}\|_2 < \varepsilon\}$ . Like the field of values,  $\sigma_\varepsilon(\mathbf{A})$  is robust to perturbations, in the sense that  $\sigma_\varepsilon(\mathbf{A} + \mathbf{E}) \subseteq \sigma_{\varepsilon + \|\mathbf{E}\|}(\mathbf{A})$ . Indeed, the pseudospectra can be regarded as a bridge between the eigenvalues and the field of values [29]. Pseudospectra are more expensive to compute than  $\sigma(\mathbf{A})$  and  $W(\mathbf{A})$ , but robust algorithms exist for their computation and approximation in the large-scale case; see [57, 64, 66]. For many of the examples below, we use Trefethen’s code in [57] or Wright’s EigTool software [65].

In an early application of pseudospectral theory, Trefethen [55] developed GMRES bounds by working from the Dunford–Taylor integral [35, §I.5.6]

$$(2.7) \quad p(\mathbf{A}) = \frac{1}{2\pi i} \int_{\Gamma} p(z)(z\mathbf{I} - \mathbf{A})^{-1} dz$$

for  $p \in \mathcal{P}_k$ , where  $\Gamma$  is any union of Jordan curves containing  $\sigma(\mathbf{A})$  in its collective interior. For a fixed  $\varepsilon > 0$ , we can take the contour  $\Gamma_\varepsilon$  to be the boundary of  $\sigma_\varepsilon(\mathbf{A})$ . (If this is not the union of Jordan curves, take  $\Gamma_\varepsilon$  to be slightly exterior.) Coarsely

approximating the resolvent norm about  $\Gamma_\varepsilon$  yields

$$\|p(\mathbf{A})\|_2 \leq \frac{1}{2\pi} \int_{\Gamma_\varepsilon} |p(z)| \|(z\mathbf{I} - \mathbf{A})^{-1}\|_2 \, d|z| \leq \frac{\mathcal{L}(\Gamma_\varepsilon)}{2\pi\varepsilon} \max_{z \in \sigma_\varepsilon(\mathbf{A})} |p(z)|,$$

where  $\mathcal{L}(\Gamma_\varepsilon)$  is the contour length of  $\Gamma_\varepsilon$ . When applied to the GMRES problem, this inequality gives

$$(PSA) \quad \frac{\|\mathbf{r}_k\|_2}{\|\mathbf{r}_0\|_2} \leq \frac{\mathcal{L}(\Gamma_\varepsilon)}{2\pi\varepsilon} \min_{\substack{p \in \mathcal{P}_k \\ p(0)=1}} \max_{z \in \sigma_\varepsilon(\mathbf{A})} |p(z)|.$$

What value  $\varepsilon$  should one use in (PSA) to get a concrete bound for a given problem? Notice that (PSA) is actually a *family of bounds* when sampled over all  $\varepsilon > 0$  [9, p. 564]). This crucial aspect of (PSA) is often overlooked. Larger values of  $\varepsilon$  tend to give smaller leading constants  $\mathcal{L}(\Gamma_\varepsilon)/(2\pi\varepsilon)$  but larger sets  $\sigma_\varepsilon(\mathbf{A})$  over which to maximize  $|p(z)|$ : such bounds can potentially capture the slow initial convergence of GMRES commonly observed for nonnormal  $\mathbf{A}$ . Smaller values of  $\varepsilon$  tend to give larger values of  $\mathcal{L}(\Gamma_\varepsilon)/(2\pi\varepsilon)$  but small  $\sigma_\varepsilon(\mathbf{A})$ : this regime can describe the faster later stage of GMRES convergence, with the large constant corresponding to the delayed onset of the fast convergence phase. To see an illustration of this phenomenon, look ahead to Figure 3.10.

Several different strategies lead to localized versions of (PSA). One can use Theorem 2.3 to decompose  $\mathbf{A}$  into invariant subspaces (at the cost of multiplication by the norms of the spectral projectors), or one can simply replace  $\sigma_\varepsilon(\mathbf{A})$  with a conglomeration of disjoint components of pseudospectra using several different values of  $\varepsilon$ .

**COROLLARY 2.5.** *Partition the eigenvalues  $\sigma(\mathbf{A})$  into disjoint sets,  $\{\Lambda_j\}_{j=1}^m$ , with the columns of each  $\mathbf{U}_j$  giving an orthonormal basis for the invariant subspace of  $\mathbf{A}$  associated with  $\Lambda_j$ . Then for any  $\varepsilon > 0$ ,*

$$(PSA') \quad \frac{\|\mathbf{r}_k\|_2}{\|\mathbf{r}_0\|_2} \leq \min_{\substack{p \in \mathcal{P}_k \\ p(0)=1}} \sum_{j=1}^m \left( \|\mathbf{P}_j\|_2 \frac{\mathcal{L}(\Gamma_\varepsilon(\mathbf{U}_j^* \mathbf{A} \mathbf{U}_j))}{2\pi\varepsilon} \max_{z \in \sigma_\varepsilon(\mathbf{U}_j^* \mathbf{A} \mathbf{U}_j)} |p(z)| \right),$$

where  $\Gamma_\varepsilon(\mathbf{U}_j^* \mathbf{A} \mathbf{U}_j)$  is a Jordan curve enclosing  $\sigma_\varepsilon(\mathbf{U}_j^* \mathbf{A} \mathbf{U}_j)$ .

**COROLLARY 2.6.** *Partition the eigenvalues  $\sigma(\mathbf{A})$  into disjoint sets  $\{\Lambda_j\}_{j=1}^m$ , and let  $\{\Gamma_j\}_{j=1}^m$  be a set of non-intersecting Jordan curves and  $\{\varepsilon_j\}_{j=1}^m$  positive constants such that for each  $j$ :*

- *The interior of  $\Gamma_j$  contains all eigenvalues in  $\Lambda_j$ , and all other eigenvalues of  $\mathbf{A}$  are exterior to  $\Gamma_j$ ;*
- *$\|(z\mathbf{I} - \mathbf{A})^{-1}\| \leq 1/\varepsilon_j$  for all  $z \in \Gamma_j$ .*

Then

$$(PSA'') \quad \frac{\|\mathbf{r}_k\|_2}{\|\mathbf{r}_0\|_2} \leq \min_{\substack{p \in \mathcal{P}_k \\ p(0)=1}} \sum_{j=1}^m \left( \frac{\mathcal{L}(\Gamma_j)}{2\pi\varepsilon_j} \max_{z \in \Gamma_j} |p(z)| \right).$$

This bound follows from choosing  $\Gamma = \cup_{j=1}^m \Gamma_j$  in the integral (2.7). Notice that when  $\varepsilon_1 = \dots = \varepsilon_m$ , this bound reduces to (PSA). However, in some situations it might be advantageous to choose, for example, a very small  $\varepsilon_j$  for a few eigenvalues near the origin, but larger  $\varepsilon_j$  for the remaining eigenvalues farther from the origin.

**2.4. Computing the Convergence Bounds.** With each of the bounds (EV), (FOV), and (PSA) is associated a constant, defined as

$$(2.8) \quad C_{\text{EV}} := \kappa(\mathbf{V}), \quad C_{\text{FOV}} := 1 + \sqrt{2}, \quad \text{and} \quad C_{\text{PSA}}(\varepsilon) := \frac{\mathcal{L}(\Gamma_\varepsilon)}{2\pi\varepsilon}.$$

The asymptotic behavior of each bound is determined by the associated complex approximation problem over  $\sigma(\mathbf{A})$ ,  $W(\mathbf{A})$ , or  $\sigma_\varepsilon(\mathbf{A})$ .

Let  $\Omega \subset \mathbb{C}$  be a compact domain (without isolated points) that tightly bounds  $W(\mathbf{A})$ ,  $\sigma_\varepsilon(\mathbf{A})$ , or the clustered eigenvalues of  $\sigma(\mathbf{A})$ .<sup>2</sup> Provided  $0 \notin \Omega$ , the error of the approximation problem

$$\min_{\substack{p \in \mathcal{P}_k \\ p(0)=1}} \max_{z \in \Omega} |p(z)|$$

decreases at an asymptotically linear rate in the degree  $k$  (see, e.g., [31, Ch. 16]):

$$(2.9) \quad \limsup_{k \rightarrow \infty} \left( \min_{\substack{p \in \mathcal{P}_k \\ p(0)=1}} \max_{z \in \Omega} |p(z)| \right)^{1/k} = \rho.$$

We call  $\rho \in (0, 1)$  the *asymptotic (linear) convergence rate* for  $\Omega$ . Driscoll, Toh, and Trefethen demonstrate how this constant can be computed via conformal mapping [9]. When the set is a line segment or a disk, the rate is simple to compute. For arbitrary polygons, the rate can be computed in MATLAB using Driscoll's Schwarz–Christoffel Toolbox for numerical conformal mapping [7]; recent versions Chebfun [8, 23] also include conformal mapping capabilities. One might prefer to take  $\Omega$  to be a disconnected set if  $\sigma(\mathbf{A})$  has outliers or for  $\sigma_\varepsilon(\mathbf{A})$  with sufficiently small  $\varepsilon$ . If each connected component of a disconnected set is a polygon on the real axis and is symmetric about the real axis, the rate can still be computed [16]. More general sets present greater difficulty, and it may in practice be necessary to bound  $\sigma(\mathbf{A})$  or  $\sigma_\varepsilon(\mathbf{A})$  with a single over-sized polygon.

The asymptotic characterization (2.9) does not directly provide a convergence bound at a specific iteration number,  $k$ : effectively, the  $k$ th root in (2.9) obscures a leading constant relating the minimization problem to  $\rho^k$ . We are guaranteed that

$$(2.10) \quad \min_{\substack{p \in \mathcal{P}_k \\ p(0)=1}} \max_{z \in \Omega} |p(z)| \geq \rho^k$$

(see, e.g., [9]), and if  $\Omega$  is a disk, equation (2.10) holds with equality. When  $\Omega$  is a segment  $[a, b]$  of a line passing through the origin, shifted and scaled Chebyshev polynomials are optimal. In this case, the minimax error is bounded above by  $2\rho^k$  and known explicitly for each  $k$  (see, e.g., [43, §6.11]). If  $\Omega$  is convex, Eiermann [10, 12] uses Faber polynomial analysis based on an approximation theorem of Kövari and Pommerenke [36] to show that

$$(2.11) \quad \min_{\substack{p \in \mathcal{P}_k \\ p(0)=1}} \max_{z \in \Omega} |p(z)| \leq \frac{2\rho^k}{1 - \rho^k}.$$

<sup>2</sup>If  $\mathbf{A}$  has finite dimension,  $\sigma(\mathbf{A})$  is a discrete point set with no finite asymptotic convergence rate. If the eigenvalues are clustered, the asymptotic convergence rate of the bounding set  $\Omega$  typically describes convergence. Outlying eigenvalues do not affect this convergence *asymptotic* rate; see [9] for details.

In particular, this bound always holds when  $\Omega = W(\mathbf{A})$  [12]. In other circumstances, one can construct some polynomial  $\phi \in \mathcal{P}_k$  (e.g., by interpolating at well-chosen points on  $\Omega$ , or by constructing the Faber polynomials associated with  $\Omega$  from the conformal map that determines  $\rho$ ) to obtain an upper bound:

$$\min_{\substack{p \in \mathcal{P}_k \\ p(0)=1}} \max_{z \in \Omega} |p(z)| \leq \max_{z \in \Omega} |\phi(z)|.$$

To unify notation, we label the rate associated with each of the three sets  $\sigma(\mathbf{A})$ ,  $W(\mathbf{A})$ , and  $\sigma_\varepsilon(\mathbf{A})$  as  $\rho_{\text{EV}}$ ,  $\rho_{\text{FOV}}$ , and  $\rho_{\text{PSA}}(\varepsilon)$ . In most situations we will base these constants on compact domains  $\Omega$  that contain  $\sigma(\mathbf{A})$ ,  $W(\mathbf{A})$ , and  $\sigma_\varepsilon(\mathbf{A})$ . (For a few examples we will handle outlying eigenvalues separately.)

**3. Which Bounds are Most Useful?** The previous section described the standard bounds (EV), (FOV), and (PSA), along with “localized” variants. How do these bounds compare? Are they redundant, or can each provide specific insight? We begin by exploring the relationships between the sets  $\sigma(\mathbf{A})$ ,  $W(\mathbf{A})$ , and  $\sigma_\varepsilon(\mathbf{A})$ , and then turn to concrete examples illustrating the relative merits of the three standard bounds. We denote the spectral radius as

$$\rho(\mathbf{A}) := \max_{z \in \sigma(\mathbf{A})} |z|$$

and the numerical radius as

$$\mu(\mathbf{A}) := \max_{z \in W(\mathbf{A})} |z|.$$

Let  $\Delta_r := \{z \in \mathbb{C} : |z| < r\}$  and  $\bar{\Delta}_r := \{z \in \mathbb{C} : |z| \leq r\}$  denote the open and closed disks of radius  $r$ .

If for a small  $\varepsilon > 0$  the  $\varepsilon$ -pseudospectrum contains points far from any eigenvalue, then the field of values must also contain points far from  $\sigma(\mathbf{A})$ , and the condition number  $\kappa(\mathbf{V})$  of the eigenvector matrix must also be large, as made precise by the following theorems. The first, a version of the Bauer–Fike theorem [1],[59, Thm. 3.2], bounds the  $\varepsilon$ -pseudospectrum by the union of balls of radius  $\varepsilon\kappa(\mathbf{V})$  centered at the eigenvalues. The second (known already to Stone in 1932 [49, Thm. 4.20]; see also [28, §4.6], [59, Thm. 17.2]) bounds the  $\varepsilon$ -pseudospectrum in terms of the field of values. The third relates the field of values to the eigenvector condition number, a consequence of the basic inequality  $\mu(\mathbf{A}) \leq \|\mathbf{A}\|_2$ .

**THEOREM 3.1 (Bauer–Fike).** *Let  $\mathbf{A}$  be diagonalizable,  $\mathbf{A} = \mathbf{V}\mathbf{\Lambda}\mathbf{V}^{-1}$ . Then for any  $\varepsilon > 0$ ,  $\sigma_\varepsilon(\mathbf{A}) \subseteq \sigma(\mathbf{A}) + \Delta_{\varepsilon\kappa(\mathbf{V})}$ .*

**THEOREM 3.2.** *For any  $\varepsilon > 0$ ,  $\sigma_\varepsilon(\mathbf{A}) \subseteq W(\mathbf{A}) + \Delta_\varepsilon$ .*

**THEOREM 3.3.** *Let  $\mathbf{A}$  be diagonalizable,  $\mathbf{A} = \mathbf{V}\mathbf{\Lambda}\mathbf{V}^{-1}$ . Then  $\mu(\mathbf{A}) \leq \kappa(\mathbf{V})\rho(\mathbf{A})$ .*

Theorems 3.1 and 3.3 hold with equality when  $\mathbf{A}$  is normal. Theorem 3.2 is sharp if  $\mathbf{A}$  is a multiple of the identity, or if  $\mathbf{A}$  is a Jordan block in the limit  $n = \infty$  (see Example D). For nonnormal matrices, all three bounds can be far from equality.

Theorems 3.1 and 3.2 shed light on the constant  $C_{\text{PSA}}(\varepsilon) = \mathcal{L}(\Gamma_\varepsilon)/(2\pi\varepsilon)$  defined in (2.8). More precisely, they permit insight about  $C_{\text{PSA}}(\varepsilon)$  when  $\Gamma_\varepsilon$  is taken to be a convenient curve that encloses  $\sigma_\varepsilon(\mathbf{A})$  in its interior, rather than the generally more

complicated boundary of  $\sigma_\varepsilon(\mathbf{A})$ . (The bound (PSA) holds when  $\sigma_\varepsilon(\mathbf{A})$  is replaced by such a larger set:  $C_{\text{PSA}}(\varepsilon)$  is now defined using  $\Gamma_\varepsilon$ , and the polynomial approximation problem associated with  $\rho_{\text{PSA}}(\varepsilon)$  in (PSA) is now posed over the interior of  $\Gamma_\varepsilon$ , rather than  $\sigma_\varepsilon(\mathbf{A})$  itself.)

By Theorem 3.1,  $\sigma_\varepsilon(\mathbf{A})$  is bounded by the union of  $n$  disks each with radius  $\varepsilon\kappa(\mathbf{V})$ . Taking  $\Gamma_\varepsilon$  to be the boundary of this union,  $\mathcal{L}(\Gamma_\varepsilon)$  can be no larger than  $2\pi n\varepsilon\kappa(\mathbf{V})$ , so  $C_{\text{PSA}}(\varepsilon) \leq n\kappa(\mathbf{V})$  for this  $\Gamma_\varepsilon$ . But since  $\sigma(\mathbf{A}) \subset \sigma_\varepsilon(\mathbf{A}) \subseteq \text{interior}(\Gamma_\varepsilon)$ , we must have  $\rho_{\text{EV}} < \rho_{\text{PSA}}(\varepsilon)$  for all  $\varepsilon > 0$ , and thus (PSA) is generally only useful for those  $\varepsilon$  for which  $C_{\text{PSA}}(\varepsilon) \ll n\kappa(\mathbf{V})$ .

By Theorem 3.2, if one takes  $\Gamma_\varepsilon$  to be the boundary of the disk centered at the origin having radius  $\mu(\mathbf{A}) + \varepsilon$ , then  $C_{\text{PSA}}(\varepsilon) = 1 + \mu(\mathbf{A})/\varepsilon$ . Notice then that for such  $\Gamma_\varepsilon$  we have  $C_{\text{PSA}}(\varepsilon) \rightarrow 1$  as  $\varepsilon \rightarrow \infty$ . (Of course,  $\rho_{\text{PSA}}(\varepsilon) = 1$  for such sets.)

When the containment  $\sigma_\varepsilon(\mathbf{A}) \subset W(\mathbf{A}) + \Delta_\varepsilon$  is nearly equality even for small values of  $\varepsilon$ , the bound (FOV) can be slightly sharper than (PSA), as is seen in Figures 3.2 and 3.5. In cases where the bound in Theorem 3.2 is far from equality (i.e.,  $\sigma_\varepsilon(\mathbf{A})$  does not contain points near the boundary of  $W(\mathbf{A}) + \Delta_\varepsilon$  for moderate values of  $\varepsilon$ ), one often finds that (FOV) suggests slow, consistent convergence, while (PSA) reflects convergence that eventually accelerates to a more rapid rate, as in Figures 3.8, 3.10, and 4.3.

**3.1. The Examples.** The bounds (EV), (FOV), and (PSA) are descriptive in different situations. We demonstrate with six examples where the bounds succeed together, fail together, and, in turn, fail and succeed alone. These examples are summarized in Table 3.1. We only discuss the standard bounds, though in some instances a localized version would fix the flaw that causes the corresponding standard bound to fail. It is difficult to show the failure of (PSA) with the simultaneous success of (EV) or (FOV). Example C, showing success of (EV) with pessimistic (PSA) bounds, is the least satisfying of our six examples. We also discuss why (FOV) cannot significantly outperform (PSA).

TABLE 3.1  
Predicted iterations for the six main examples in Section 3.1.

<i>example</i>	(EV)	(FOV)	(PSA)	<i>true iterations</i>
A all descriptive	1	1	1	1
B none descriptive	$\infty$	$\infty$	$\infty$	<i>see note 1</i>
C (EV) wins	<i>see note 2</i>			
D (EV) loses	$\infty$	1	1	1
(FOV) wins	<i>see note 3</i>			
E (FOV) loses	2	$\infty$	2	2
F (PSA) wins	$\infty$	$\infty$	2	2
(PSA) loses	<i>see note 3</i>			

1.  $\lceil 2 + \log(\text{TOL})/\log(\rho) \rceil$  iterations, for parameters  $\text{TOL} < 1$  and  $\rho \in (0, 1)$ .
2. “(EV) wins” is more involved; details are given in the text.
3. (FOV) cannot significantly beat (PSA); see the “No Example” sections.

In the illustrations that follow, the Ideal GMRES value

$$\min_{\substack{p \in \mathcal{P}_k \\ p(0)=1}} \|p(\mathbf{A})\|_2$$

is drawn as a solid black line with dots superimposed at each iteration  $k$ . The bound (EV) is drawn as a solid red line, (FOV) with a broken blue line, and (PSA) with gray lines for various values of  $\varepsilon$ .

• **Example A: All descriptive.** All bounds accurately describe GMRES convergence for a scalar multiple of the identity,

$$\mathbf{A} = \alpha \mathbf{I}, \quad \alpha \in \mathbb{C} \setminus \{0\}.$$

Since  $\mathbf{A}$  is normal with a single eigenvalue,  $\sigma(\mathbf{A}) = W(\mathbf{A}) = \{\alpha\}$  and  $\sigma_\varepsilon(\mathbf{A}) = \{\alpha\} + \Delta_\varepsilon$ . The approximation problems in (EV) and (FOV) are on singleton sets, and thus these bounds ensure convergence in one iteration; the associated constants are  $C_{\text{EV}} = 1$  and  $C_{\text{FOV}} = 1 + \sqrt{2}$ . The pseudospectral bound also gives convergence in a single iteration, but the explanation is a bit more elaborate. The constant term in (PSA) is  $C_{\text{PSA}} = 1$  for all  $\varepsilon$ ; with its approximation problem on the disk  $\alpha + \Delta_\varepsilon$ , (PSA) gives

$$\min_{\substack{p \in \mathcal{P}_k \\ p(0)=1}} \|p(\mathbf{A})\|_2 \leq (\varepsilon/\alpha)^k, \quad \text{for all } \varepsilon \in (0, \alpha),$$

implying, as  $\varepsilon \rightarrow 0$ , convergence to any given tolerance in a single iteration.

All three bounds also perform well if the single eigenvalue is expanded to a positive real interval  $[a, b] \not\ni 0$ , provided  $\mathbf{A}$  remains normal. To get  $\sigma(\mathbf{A}) = [a, b]$  requires an

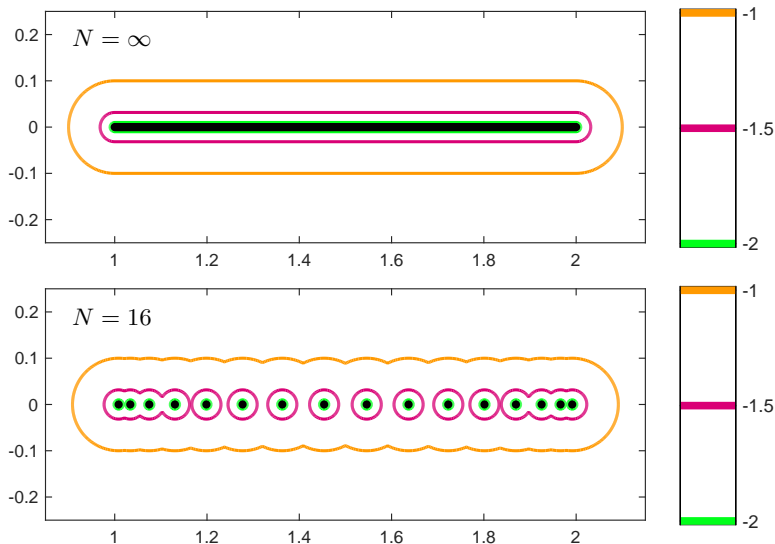


FIG. 3.1. Spectra (black dots in  $[1, 2]$ ) and  $\varepsilon$ -pseudospectra ( $\varepsilon = 10^{-1}, 10^{-1.5}, 10^{-2}$ ) for the operator (3.1) (top) with  $[a, b] = [1, 2]$ , and its finite section of dimension  $N = 16$  (bottom). (In this and all following plots of  $\sigma_\varepsilon(\mathbf{A})$ , the color bar denotes  $\log_{10}(\varepsilon)$ , and the corresponding colored lines show the boundary of  $\sigma_\varepsilon(\mathbf{A})$ .)



ple B illustrates this shortcoming:<sup>3</sup>

$$(3.2) \quad \mathbf{A} = \begin{bmatrix} 1 & \alpha \\ 0 & 1 \end{bmatrix} \oplus \widehat{\mathbf{\Lambda}}, \quad \alpha \gg 2,$$

where  $\widehat{\mathbf{\Lambda}}$  is a diagonal matrix with uniformly distributed entries in the positive real interval  $[1, b]$  for  $b > 1$ . When  $\alpha$  is large, the  $2 \times 2$  Jordan block will dominate  $W(\mathbf{A})$  and  $\sigma_\varepsilon(\mathbf{A})$  even for some fairly small values of  $\varepsilon$ . (Figure 3.3 shows an example.) A polynomial  $p(z)$  with two roots at  $z = 1$  would annihilate this Jordan block, leaving a normal matrix with eigenvalues in the interval  $[1, b]$  to handle at later iterations. This structure allows one to bound the norm of the GMRES residual *independent of*  $\alpha$ . To see this, replace the optimal GMRES polynomial  $p(z)$  with  $(1 - z)^2 q(z)$  at iteration  $k \geq 2$  to get the upper bound

$$(3.3) \quad \begin{aligned} \min_{\substack{p \in \mathcal{P}_k \\ p(0)=1}} \|p(\mathbf{A})\|_2 &\leq \min_{\substack{q \in \mathcal{P}_{k-2} \\ q(0)=1}} \|(\mathbf{I} - \mathbf{A})^2 q(\mathbf{A})\|_2 = \min_{\substack{q \in \mathcal{P}_{k-2} \\ q(0)=1}} \|\mathbf{0} \oplus (\mathbf{I} - \widehat{\mathbf{\Lambda}})^2 q(\widehat{\mathbf{\Lambda}})\|_2 \\ &\leq \min_{\substack{q \in \mathcal{P}_{k-2} \\ q(0)=1}} \max_{z \in [1, b]} |1 - z|^2 |q(z)| \\ &\leq |1 - b|^2 \min_{\substack{q \in \mathcal{P}_{k-2} \\ q(0)=1}} \max_{z \in [1, b]} |q(z)| \leq 2|1 - b|^2 \left( \frac{\sqrt{b} - 1}{\sqrt{b} + 1} \right)^{k-2}. \end{aligned}$$

This last step just uses Chebyshev approximation on the interval  $[1, b]$ . Now if  $b \leq 1 + 1/\sqrt{2}$  (so  $2|1 - b|^2 \leq 1$ ),

$$\min_{\substack{p \in \mathcal{P}_k \\ p(0)=1}} \|p(\mathbf{A})\|_2 \leq \left( \frac{\sqrt{b} - 1}{\sqrt{b} + 1} \right)^{k-2}.$$

This bound ensures convergence to the tolerance TOL in  $\lceil 2 + \log(\text{TOL})/\log(\rho) \rceil$  iterations, where  $\rho = (\sqrt{b} - 1)/(\sqrt{b} + 1)$ , as given in Table 3.1.

Since  $\mathbf{A}$  is nondiagonalizable,  $C_{\text{EV}} = \infty$  and (EV) does not apply. The field of values of  $\mathbf{A}$  grows ever larger with  $\alpha$ . Note that

$$W\left(\begin{bmatrix} 1 & \alpha \\ 0 & 1 \end{bmatrix}\right) = 1 + \overline{\Delta}_{\alpha/2}.$$

Thus if  $b \leq \alpha/2$ , the normal eigenvalues on the diagonal of  $\widehat{\mathbf{\Lambda}}$  are all embedded within the field of values of the  $2 \times 2$  Jordan block, and hence

$$W(\mathbf{A}) = 1 + \overline{\Delta}_{\alpha/2}, \quad b \in (1, \alpha/2].$$

(The left plot of Figure 3.3 shows this scenario.) In any case,  $1 + \overline{\Delta}_{\alpha/2} \subseteq W(\mathbf{A})$ , and so if  $\alpha \geq 2$ , then  $0 \in W(\mathbf{A})$  and (FOV) cannot give convergence.

Analysis of the pseudospectral bound is more involved. In Example F we shall see that (PSA) accurately predicts convergence for a Jordan block by taking  $\varepsilon$  sufficiently small. The  $\widehat{\mathbf{\Lambda}}$  component in the present example (3.2) adds the crucial complication:

<sup>3</sup>This example is essentially an extreme version of the diagonalizable example constructed by Greenbaum and Strakoš [26] to demonstrate the failure of the pseudospectral bound (PSA).




FIG. 3.3. Field of values (gray disk),  $\varepsilon$ -pseudospectra ( $\varepsilon = 10^{-5}, \dots, 10^0$ ) and eigenvalues (small dots) for Example B with  $\alpha = 10$ ,  $b = 3/2$ , and  $N = 102$ . The  $2 \times 2$  Jordan block determines  $W(\mathbf{A})$  and dominates  $\sigma_\varepsilon(\mathbf{A})$  for the  $\varepsilon$  values shown, exerting strong influence on the associated convergence bounds. Nevertheless, this Jordan block does not much delay GMRES convergence.

as seen in Example A, for Hermitian matrices the bound (PSA) is best when  $\varepsilon$  is relatively large ( $10^{-1}$  and  $10^{-2}$  in Figure 3.2). By explicitly computing the norm of the resolvent, one can show that

$$(3.4) \quad 1 + \Delta_{\sqrt{\alpha\varepsilon + \varepsilon^2}} = \sigma_\varepsilon \left( \begin{bmatrix} 1 & \alpha \\ 0 & 1 \end{bmatrix} \right) \subseteq \sigma_\varepsilon(\mathbf{A}).$$

Since  $\sigma_\varepsilon(\mathbf{A})$  must contain the disk  $1 + \Delta_{\sqrt{\alpha\varepsilon + \varepsilon^2}}$ , the boundary of  $\sigma_\varepsilon(\mathbf{A})$  must encircle this set, ensuring that, for any  $\varepsilon > 0$ ,

$$C_{\text{PSA}}(\varepsilon) \geq \frac{2\pi\sqrt{\alpha\varepsilon + \varepsilon^2}}{2\pi\varepsilon} \geq \sqrt{\alpha/\varepsilon}.$$

Moreover, equation (3.4) implies that

$$0 \in \sigma_\varepsilon(\mathbf{A}) \quad \text{for all } \varepsilon > \frac{1}{2}(\sqrt{\alpha^2 + 4} - \alpha),$$

i.e.,  $\varepsilon > \frac{1}{2}(\sqrt{\alpha^2 + 4} - \alpha) \approx 1/\alpha$  for large  $\alpha$ : thus ruling out the large values of  $\varepsilon$  that could control  $C_{\text{PSA}}(\varepsilon)$ . By increasing  $\alpha$ , we can make  $C_{\text{PSA}}(\varepsilon)$  arbitrarily large, even though (3.3) bounds GMRES convergence independent of  $\alpha$ . On the other hand, the  $\widehat{\mathbf{A}}$  block gives a nonzero asymptotic rate of convergence for any  $\varepsilon > 0$  (unlike Example F, where a single eigenvalue will permit arbitrarily fast convergence rates as  $\varepsilon \rightarrow 0$ ). Thus in Table 3.1 we say that (PSA) predicts infinitely many iterations.

• **Example C: Only (EV) descriptive.** Aside from trivial cases, the field of values and pseudospectral bounds both involve approximation problems on regions in the complex plane that give nonzero asymptotic convergence rates  $\rho_{\text{FOV}}$  and  $\rho_{\text{PSA}}(\varepsilon)$ . However, these sets could be unduly influenced by parts of the spectrum that could be effectively eliminated at an early stage of the GMRES iteration. When a few eigenvalues are far from the rest of the spectrum, (EV) can be more descriptive than (PSA) and (FOV), because the approximation problem in (EV) is posed on a discrete point set, and isolated outliers do not influence  $\rho_{\text{EV}}$ . Define

$$\mathbf{A} = \delta \oplus \widehat{\mathbf{A}}, \quad \delta \in \mathbb{R},$$

where  $\widehat{\mathbf{A}}$  is a diagonal matrix with entries uniformly distributed in the real, positive interval  $[a, b]$ , and  $0 < |\delta| \ll a < b$ .

Since  $\mathbf{A}$  is normal, convergence is determined by the spectrum. The bound (EV), with  $C_{\text{EV}} = 1$ , is exact. This convergence can be bounded using the polynomial  $p_k(z) = (1 - z/\delta)q_{k-1}(z)$ , where  $q_{k-1}$  is the optimal degree- $k - 1$  residual polynomial for the interval  $[a, b]$ . The eigenvalue  $\delta$  near the origin causes an initial stagnation [9]; the polynomial  $p_k$  suggests that this plateau will last no longer than the number of iterations it takes for  $q_{k-1}(z)$  to overcome  $(1 - z/\delta)$  for  $z \in [a, b]$ , i.e.,

$$\left\lceil 1 + \frac{\log |\delta| - \log 2(b - \delta)}{\log(\rho_{\text{EV}})} \right\rceil \text{ iterations,}$$

where

$$\rho_{\text{EV}} = \frac{\sqrt{b/a} - 1}{\sqrt{b/a} + 1}$$

is the asymptotic convergence rate associated with the interval  $[a, b]$ . Figure 3.4 shows the stagnation caused by the eigenvalue near the origin.

If  $\delta < 0 < a < b$ , then  $W(\mathbf{A}) = [\delta, b]$  contains the origin, and hence  $\rho_{\text{FOV}} = 1$ , and so (FOV) does not predict any convergence. If  $0 < \delta < a < b$ , then  $W(\mathbf{A}) = [\delta, b]$  does not contain the origin, and the asymptotic convergence rate

$$\rho_{\text{FOV}} = \frac{\sqrt{b/\delta} - 1}{\sqrt{b/\delta} + 1}$$

will be close to one when  $0 < \delta \ll a < b$ . Provided  $b$  is not too large,  $\rho_{\text{EV}} \ll \rho_{\text{FOV}}$ : (FOV) predicts slow convergence, accurately describing the initial period of stagnation but missing the transition to more rapid asymptotic convergence.

The pseudospectral bound suffers from the fact that it cannot treat  $\delta$  as a single simple eigenvalue, eliminated at an early stage of convergence. Moreover, to give con-




FIG. 3.4. Convergence bounds for Example C with  $\delta = 0.01$  and  $n = \infty$ , using underestimates of  $\sigma_\varepsilon(\mathbf{A})$  for the approximation problem in (PSA).

vergence we need  $0 \notin \sigma_\varepsilon(\mathbf{A})$ , requiring  $0 < \varepsilon < |\delta|$ . For any  $\varepsilon \in (0, (a-\delta)/2)$ , the component of  $\sigma_\varepsilon(\mathbf{A})$  about  $\lambda = \delta$  is a disk of radius  $\varepsilon$ , and this disk will always influence the asymptotic convergence rate associated with  $\sigma_\varepsilon(\mathbf{A})$ . This effect diminishes as  $\varepsilon$  decreases, but such small values of  $\varepsilon$  will give large constants ( $C_{\text{PSA}}(\varepsilon) = 2 + (b-a)/(\pi\varepsilon)$  for  $n = \infty$ ) due to the interval  $[a, b]$ .

This case is the least analytically compelling of our six examples. In particular, it is difficult to cleanly describe the convergence rates associated with (PSA). In Figure 3.4, we illustrate (PSA) for  $\delta = 0.01$ ,  $[a, b] = [1, 2]$ , and  $n = \infty$ . Since we argue that (PSA) gives poor bounds for this scenario, we do not want our portrayal of the bounds to suffer from an overestimate  $\sigma_\varepsilon(\mathbf{A})$  that would yield an easily computable asymptotic convergence rate that is not true to the bound. Instead, Figure 3.4 shows an *underestimate* of (PSA) using the slightly faster convergence rate for the union of two intervals,  $[\delta - \varepsilon, \delta + \varepsilon] \cup [a - \varepsilon, b + \varepsilon] \subset \sigma_\varepsilon(\mathbf{A})$ . The asymptotic convergence rate  $\rho(\delta, \varepsilon, a, b)$  for the union of these two real intervals can be expressed in terms of elliptic integrals, as described and implemented in MATLAB by Fischer [18]. The pseudospectral constants are  $C_{\text{PSA}}(\varepsilon) = 2 + (b-a)/(\pi\varepsilon)$ . In place of (PSA), Figure 3.4 shows the underestimates  $C_{\text{PSA}}(\varepsilon)\rho(\delta, \varepsilon, a, b)^k$ .

Of course, one could adapt (FOV) and (PSA) to handle a single outlying eigenvalue by splitting up the spectrum, as in Theorem 2.3; however, many problems present a range of outlying eigenvalues at different scales, which are more difficult to identify and handle.

• **Example D: Only (EV) not descriptive.** The bound (EV) fails for any nondiagonalizable matrix, yet defectiveness need not imply poor GMRES convergence. Consider a small perturbation to the identity matrix,

$$(3.5) \quad \mathbf{A} = \begin{bmatrix} 1 & \delta & & \\ & 1 & \ddots & \\ & & \ddots & \delta \\ & & & 1 \end{bmatrix}, \quad 0 < \delta \ll 1.$$

This matrix is completely defective for all  $\delta \neq 0$ , but small values of  $\delta$  exert only the slightest impact on convergence. Unlike the case of  $\delta = 0$ , for most initial residuals GMRES will require  $n$  iterations to converge *exactly*; however, for small  $\delta > 0$  GMRES will make excellent progress at each step. (For bounds describing how small changes to  $\mathbf{A}$  affect the convergence of GMRES for a fixed  $\mathbf{r}_0$ , see [46].) If  $\mathbf{r}_0$  is the  $n$ th column of the identity matrix, Ipsen [33] derived the exact formula

$$(3.6) \quad \frac{\|\mathbf{r}_k\|_2}{\|\mathbf{r}_0\|_2} = \delta^k \sqrt{\frac{1 - \delta^2}{1 - \delta^{2(k+1)}}}.$$

For small  $\delta$  this convergence must be close to the worst case, since

$$\min_{\substack{p \in \mathcal{P}_k \\ p(0)=1}} \|p(\mathbf{A})\|_2 \leq \|(\mathbf{I} - \mathbf{A})^k\| = \delta^k.$$

Since  $\mathbf{A}$  is nondiagonalizable, (EV) fails to predict convergence. The field of values is known explicitly for this example,  $W(\mathbf{A}) = 1 + \Delta_{\delta \cos(\pi/(n+1))}$  [28, §1.3], leading to the exact formulation of (FOV):

$$\min_{\substack{p \in \mathcal{P}_k \\ p(0)=1}} \|p(\mathbf{A})\|_2 \leq (1 + \sqrt{2}) \left( \delta \cos\left(\frac{\pi}{n+1}\right) \right)^k.$$

In the case of  $n = \infty$ , the infinite dimensional Toeplitz matrix has spectrum  $\sigma(\mathbf{A}) = 1 + \Delta_\delta$ , and the closure of the field of values is  $\overline{W(\mathbf{A})} = 1 + \Delta_\delta$ ; see [2, sect. 7.5].

The pseudospectra of  $\mathbf{A}$  are also disks [59], but the radii of these disks are not known in closed form for general  $n$ . In the limit  $n \rightarrow \infty$ , a theorem of Reichel and Trefethen [42] shows that  $\sigma_\varepsilon(\mathbf{A}) = 1 + \Delta_{\delta+\varepsilon}$ ; in this case, the bounds give:

$$(\text{FOV}) : \min_{\substack{p \in \mathcal{P}_k \\ p(0)=1}} \|p(\mathbf{A})\|_2 \leq (1 + \sqrt{2}) \delta^k \quad (\text{PSA}) : \min_{\substack{p \in \mathcal{P}_k \\ p(0)=1}} \|p(\mathbf{A})\|_2 \leq (1 + \delta/\varepsilon)(\delta + \varepsilon)^k.$$

In the limit as  $\delta \rightarrow 0$ , both these bounds predict convergence to arbitrary desired tolerance in a single iteration, though the asymptotic rate  $\rho_{\text{FOV}} = \delta^k$  is slightly sharper than the pseudospectral rate  $\rho_{\text{PSA}} = (\delta + \varepsilon)^k$ . One must balance the size of  $C_{\text{PSA}}$  against the accompanying convergence rate, just as for the normal matrix with  $\sigma(\mathbf{A}) \subseteq [a, b]$  discussed in Example A. Figure 3.5 illustrates this example for  $\delta = 1/2$  with  $n = 32$  and  $n = \infty$ . For the ‘‘exact’’ curve, we plot the lower bound given by equation (3.6).

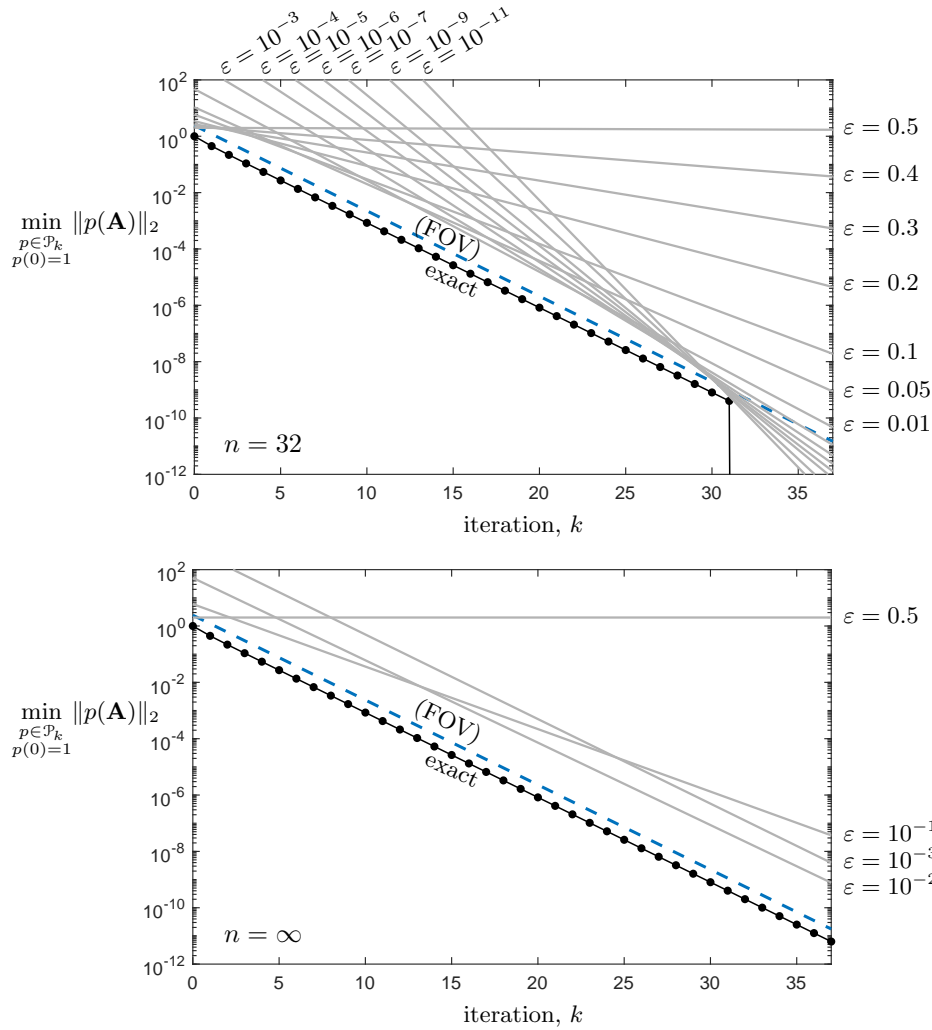


FIG. 3.5. Convergence bounds for Example D with  $\delta = 1/2$  for  $n = 32$  and  $n = \infty$ .




FIG. 3.6. Field of values  $W(\mathbf{A})$  and  $\varepsilon$ -pseudospectra  $\sigma_\varepsilon(\mathbf{A})$  for Example D with  $\delta = 1/2$ . On the left,  $n = 32$  and  $\varepsilon = 0.5, 0.4, 0.3, 0.2, 0.1, 0.05,$  and  $0.01$ ; on the right,  $n = \infty$  and  $\varepsilon = 0.5, 10^{-1}, 10^{-2},$  and  $10^{-3}$  (the last of these is barely visible).

The bound (PSA) is particularly interesting in the finite-dimensional case: For very small values of  $\varepsilon$ , (PSA) predicts convergence rates that are too quick, associated with large constants that ensure the bound does not intersect the convergence curve for  $k < n$ . Figure 3.6 shows  $W(\mathbf{A})$  and  $\sigma_\varepsilon(\mathbf{A})$  for  $n = 32$  and  $n = \infty$ .

Taking  $\mathbf{A}$  to be nondiagonalizable makes for a clean example, but it is not necessary. Perturbing the diagonal entries of  $\mathbf{A}$  from  $\lambda$  to distinct nearby values, one can obtain finite but arbitrarily large values of  $C_{\text{EV}} = \kappa(\mathbf{V})$ , which will depend on the off-diagonal  $\delta$  value; at the same time, taking all the eigenvalues close to 1 makes  $\rho_{\text{EV}}$  arbitrary close to 0, independent of  $\delta$ . Yet  $\delta$  can be seen to have a crucial role in determining the asymptotic rate of convergence.

- **No Example: Only (FOV) descriptive.** Theorem 3.2 indicates that examples where (FOV) significantly outperforms (PSA) will be difficult to find. Since  $\sigma_\varepsilon(\mathbf{A}) \subset W(\mathbf{A}) + \Delta_\varepsilon$ , the rate  $\rho_{\text{PSA}}(\varepsilon)$  can only be significantly slower than  $\rho_{\text{FOV}}$  when this containment is sharp,  $\varepsilon$  is relatively large, and the sets  $W(\mathbf{A})$  and  $\sigma_\varepsilon(\mathbf{A})$  are near the origin. In such cases, values of  $\varepsilon$  that give convergence rates  $\rho_{\text{PSA}}(\varepsilon)$  similar to  $\rho_{\text{FOV}}$  will be associated with small  $\varepsilon$  and thus large constant terms  $C_{\text{PSA}}(\varepsilon)$ . But proximity to the origin implies that  $\rho_{\text{FOV}}$  will predict slow convergence, and the pseudospectral bounds, while less sharp, should still provide a decent indication of the nature of convergence. For an example in this vein, modify Example D: take a Jordan block with eigenvalue  $\lambda$  near the origin and superdiagonal having the constant  $\delta$  for  $0 \ll \delta < |\lambda|/\cos(\pi/(n+1))$  to ensure the field of values extends near the origin but does not contain it.

- **Example E: Only (FOV) not descriptive.** The field of values bound (FOV) is not descriptive when there is initial stagnation followed by more rapid convergence. The simplest example of GMRES stagnation at the first iteration occurs for the Hermitian indefinite matrix

$$\mathbf{A} = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix},$$

which was already used by Saad and Schultz [45] to illustrate stagnation of GMRES(1). Since  $\mathbf{A} \in \mathbb{C}^{2 \times 2}$ , the second iteration gives exact convergence.

Since  $\mathbf{A}$  is normal, (EV) is exact, correctly predicting two iterations to solve the polynomial approximation problem on the discrete set of two eigenvalues. Since  $W(\mathbf{A})$

is the convex hull of  $\sigma(\mathbf{A})$ , we have  $0 \in W(\mathbf{A}) = [-1, 1]$ , and thus (FOV) predicts no convergence. The  $\varepsilon$ -pseudospectrum consists of the union of two disks of radius  $\varepsilon$ , each centered at an eigenvalue. Thus,  $C_{\text{PSA}}(\varepsilon) = 2$  independent of  $\varepsilon \in (0, 1]$ . Taking  $\varepsilon \rightarrow 0$ , (PSA) predicts convergence to arbitrary accuracy in two iterations.

The bound (FOV) can also fail when  $\mathbf{A}$  is a nonnormal matrix for which  $0 \in W(\mathbf{A})$  but  $0 \notin \sigma_\varepsilon(\mathbf{A})$  for  $\varepsilon$  sufficiently small. The following example, with uniformly ill-conditioned eigenvectors, is the only case we discuss in which  $\kappa(\mathbf{V})$  is large and yet (EV) is still descriptive. Let  $\mathbf{A}$  be diagonal with eigenvalues uniformly distributed in the interval  $[1, 2]$  and define

$$(3.7) \quad \mathbf{A} = \mathbf{V}\mathbf{\Lambda}\mathbf{V}^{-1}, \quad \text{where } \mathbf{V} = \begin{bmatrix} 1 & \sqrt{1-\delta} & \sqrt{1-\delta} & \cdots & \sqrt{1-\delta} \\ & \sqrt{\delta} & 0 & \cdots & 0 \\ & & \sqrt{\delta} & \ddots & \vdots \\ & & & \ddots & 0 \\ & & & & \sqrt{\delta} \end{bmatrix}$$

is upper triangular and  $0 < \delta \ll 1$  is a small positive parameter. (The matrix  $\mathbf{V}$  is inspired by an example of Greenbaum and Strakoš [26].)

Clearly, the eigenvectors that form the column of  $\mathbf{V}$  are severely ill-conditioned; looking at the upper left  $2 \times 2$  block of  $\mathbf{V}$  alone shows that  $\kappa(\mathbf{V}) \geq (1 + \sqrt{\delta})/\sqrt{\delta}$ . One can show  $\kappa(\mathbf{V}) = \mathcal{O}(n/\sqrt{\delta})$  as  $n \rightarrow \infty$  and  $\delta \rightarrow 0$ . For sufficiently small  $\delta > 0$ ,  $0 \in W(\mathbf{A})$  and (FOV) gives no convergence; indeed, as seen in Figure 3.7, the field of values is quite large for small  $\delta$ . The pseudospectra are also large, yet as  $\varepsilon$  is taken small enough such that  $0 \notin \sigma_\varepsilon(\mathbf{A})$ , the bound becomes increasingly descriptive. This is another case where the inclusion  $\sigma_\varepsilon(\mathbf{A}) \subseteq W(\mathbf{A}) + \Delta_\varepsilon$  in Theorem 3.2 gives a poor upper bound on  $\sigma_\varepsilon(\mathbf{A})$  for small  $\varepsilon$ .

Figure 3.8 shows the bounds for  $\delta = 10^{-8}$  and  $n = 64$ , the same parameters used in Figure 3.7. The constant  $C_{\text{EV}} = \kappa(\mathbf{V})$  is obtained by computing the condition number in MATLAB, with the first column of  $\mathbf{V}$  multiplied by  $\sqrt{n}$  to improve conditioning. The rate  $\rho_{\text{EV}}$  is taken from the optimal polynomial on the interval  $[1, 2]$ . The pseudospectral bound was determined for each  $\varepsilon$  by bounding  $\sigma_\varepsilon(\mathbf{A})$  from the outside




FIG. 3.7. Field of values  $W(\mathbf{A})$  (left) and pseudospectra  $\sigma_\varepsilon(\mathbf{A})$  for  $\varepsilon = 10^{-5}$ ,  $5 \times 10^{-6}$ ,  $10^{-6}$  (right) for the matrix (3.7) with  $\delta = 10^{-8}$  and  $n = 64$ . Note the scale of the plot on the left.




FIG. 3.8. Convergence bounds for the matrix (3.7) with  $\delta = 10^{-8}$  and  $n = 64$ .

by an convex polygon, obtaining the convergence rate through numerical conformal mapping, and then applying the bound (2.11). The “exact” curve was computed using the semidefinite programming strategy of Toh and Trefethen [54], as implemented in the SDPT3 Toolbox [52]. For this example, random initial residuals typically do not lead to the long plateau obtained in the “exact” (Ideal GMRES) curve shown here. (Whether any initial residual attains the Ideal GMRES curve for this example is not known.) Not only does  $W(\mathbf{A})$  contain the origin for these parameters; it contains a circle of radius  $10^4$  centered at the origin. For the values of  $\varepsilon$  used in Figure 3.7, the pseudospectra  $\sigma_\varepsilon(\mathbf{A})$  omit the origin, giving the convergent bounds in Figure 3.8.

• **Example F: Only (PSA) descriptive.** For the Jordan block in Example D, the field of values bound (FOV) captured the single convergence rate perfectly. For the present example, we seek a nondiagonalizable matrix that initially stagnates but eventually converges more rapidly, making (FOV) misleading. Though it might seem like a gimmick, we could say the most extreme example of this behavior occurs for the matrix

$$\mathbf{A} = \begin{bmatrix} 1 & \alpha \\ 0 & 1 \end{bmatrix}, \quad \alpha \geq 2,$$

which featured as a submatrix in Example B. For  $\alpha \geq 2$ , there exist right hand sides for which GMRES makes no progress at the first step. Yet at the second step, there is exact convergence since  $\mathbf{A} \in \mathbb{C}^{2 \times 2}$ .

Since  $\mathbf{A}$  is nondiagonalizable,  $C_{\text{EV}} = \infty$  and (EV) does not apply. Nor does (FOV) give convergence, since  $0 \in W(\mathbf{A}) = 1 + \Delta_{\alpha/2}$ , a disk centered at 1 with radius  $\alpha/2$  [28, §1.3]. The pseudospectral bound, however, captures the exact convergence in two iterations. From (3.4) we have  $\sigma_\varepsilon(\mathbf{A}) = 1 + \Delta_{\sqrt{\alpha\varepsilon + \varepsilon^2}}$ , giving the constant  $C_{\text{PSA}}(\varepsilon) = \sqrt{\alpha/\varepsilon + 1}$ . Using the residual polynomial  $p(z) = (1 - z)^k$ , (PSA) gives the bound

$$\min_{\substack{p \in \mathcal{P}_k \\ p(0)=1}} \|p(\mathbf{A})\|_2 \leq \sqrt{\alpha/\varepsilon + 1} \left( \alpha\varepsilon + \varepsilon^2 \right)^{k/2}.$$

At step  $k = 2$ , the upper bound becomes  $\alpha^{3/2}\varepsilon^{1/2} + O(\varepsilon^{3/2})$  as  $\varepsilon \rightarrow 0$ , and hence by

taking  $\varepsilon > 0$  sufficiently small, (PSA) predicts convergence to arbitrary accuracy at the second iteration.

For a more interesting example that does not rely on the dimension  $n = 2$ , take a Jordan block and progressively scale down the off-diagonal entries: for fixed  $\beta > 0$ ,

$$(3.8) \quad \mathbf{A} = \begin{bmatrix} 1 & \beta & & & \\ & 1 & \frac{\beta}{2} & & \\ & & 1 & \ddots & \\ & & & \ddots & \frac{\beta}{(n-1)} \\ & & & & 1 \end{bmatrix}.$$

Reichel and Trefethen called a related example an “integration matrix” [42] and noted that its pseudospectra are disks. Driscoll, Toh, and Trefethen showed that Ideal GMRES exhibits an improving convergence rate for this matrix, and linked that behavior to the notable shrinking of  $\sigma_\varepsilon(\mathbf{A})$  as  $\varepsilon \rightarrow 0$  [9, p. 564].

Since  $\mathbf{A}$  is nondiagonalizable, the bound (EV) cannot be usefully applied. Both the field of values and the pseudospectra of  $\mathbf{A}$  are circular disks, as can be seen via a diagonal unitary similarity transformation (see [59, p. 268] for details). If  $n \geq 2$  and  $\beta \geq 2$ , then  $0 \in W(\mathbf{A})$ , and so for such values of  $\beta$ , (FOV) cannot give convergence. Figure 3.9 shows the field of values and pseudospectra for  $\beta = 5/2$  and  $N = 64$ . For small  $\varepsilon$ ,  $W(\mathbf{A})$  is evidently considerably larger than  $\sigma_\varepsilon(\mathbf{A})$ , and thus (PSA) predicts much faster convergence rates than (FOV) as  $\varepsilon \rightarrow 0$ .

Figure 3.10 shows the corresponding GMRES bounds. The “exact” curve was again computed using the SDPT3 Toolbox [52]. The pseudospectral bounds were obtained by numerically computing the radius of each pseudospectral boundary. Increasing the dimension  $n$  does not significantly alter the pseudospectra for the values of  $\varepsilon$  shown here. (For one thing, such an extension would be a norm  $\beta/n$  perturbation to the block diagonal matrix  $\mathbf{A} \oplus \mathbf{I}$ .)

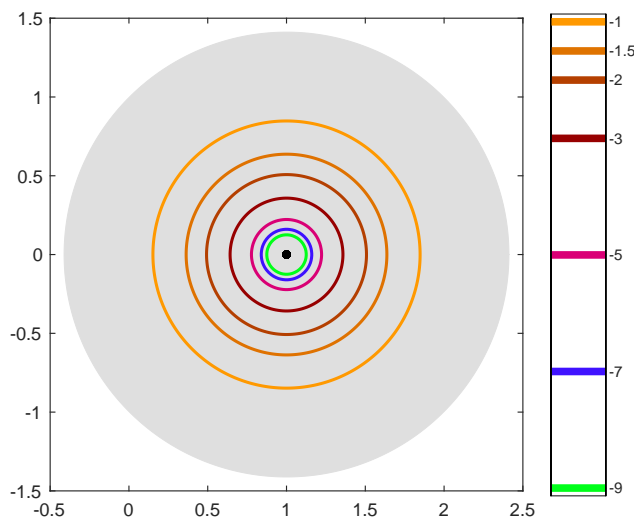


FIG. 3.9. The field of values  $W(\mathbf{A})$  (gray region) and  $\varepsilon$ -pseudospectra  $\sigma_\varepsilon(\mathbf{A})$  (for  $\varepsilon = 10^{-1}, 10^{-1.5}, 10^{-2}, 10^{-3}, 10^{-5}, 10^{-7}, 10^{-9}$ ) for the matrix (3.8) with  $\beta = 5/2$  and  $N = 64$ . Notice that  $0 \in W(\mathbf{A})$ , but  $0 \notin \sigma_\varepsilon(\mathbf{A})$  for the pseudospectra shown here.




FIG. 3.10. Convergence bounds for the integration matrix (3.8) with  $\alpha = 5/2$  and  $n = 64$ . The bounds (PSA) correspond to the same values of  $\varepsilon$  for which  $\sigma_\varepsilon(\mathbf{A})$  is shown in Figure 3.9.




FIG. 3.11. The Crouzeix–Greenbaum sets  $\Omega_{\text{CG}}$  for matrix (3.8),  $N = 64$ . For  $\beta = 5/2$  (left),  $\Omega_{\text{CG}}$  surrounds the origin (marked with +), and the bound (CG) does not describe convergence. For  $\beta = 2$  (right),  $\Omega_{\text{CG}}$  does not surround the origin, and so (CG) will give a convergent bound.

As an alternative to (FOV), one might instead consider the Crouzeix–Greenbaum bound (CG) [5]. Figure 3.11 shows the sets  $\Omega_{\text{CG}}$  defined in (2.5) for the matrix (3.8) of dimension  $N = 64$ . The plot on the left uses  $\beta = 5/2$  (as in Figures 3.9 and 3.10), giving a set  $\Omega_{\text{CG}}$  in (2.5) that does not include the origin but surrounds it, and so the bound (CG) cannot give convergence. The right plot shows  $\Omega_{\text{CG}}$  for the smaller value  $\beta = 2$ : although  $\mathbf{A}$  is a nondiagonalizable matrix with just one Jordan block, the set  $\Omega_{\text{CG}}$  excludes the origin and will yield a convergent bound with an asymptotic rate determined by  $\Omega_{\text{CG}}$ .

When faced with a nondiagonalizable matrix, one might naturally think of using the Jordan canonical form as the basis for GMRES analysis. The matrix (3.8) provides a cautionary example. Suppose we take the Jordan form  $\mathbf{A} = \mathbf{X}\mathbf{J}\mathbf{X}^{-1}$ , and then follow



• **No Example: Only (PSA) not descriptive.** This scenario would require GMRES to (eventually) converge steadily, and for this convergence to be captured by (EV) and (FOV) but not (PSA). If GMRES initially stagnates, then  $\rho_{\text{FOV}}$  must be close to 1 (since  $C_{\text{FOV}} = 1 + \sqrt{2}$  is small), and so (FOV) could only predict slow overall convergence. Thus, the example we seek could not exhibit initial stagnation. In this case, if (EV) is to be accurate, then  $C_{\text{EV}} = \kappa(\mathbf{V})$  must be small: implying that  $\mathbf{A}$  must be nearly normal. Theorems 3.1 and 3.2 then insure that  $\sigma_\varepsilon(\mathbf{A})$  cannot be much larger than  $\sigma(\mathbf{A})$  and  $W(\mathbf{A})$ , so it is impossible to get an example where (PSA) gives a significantly slower asymptotic convergence rates than (EV) and (FOV). (One could take  $\sigma(\mathbf{A})$  to contain an eigenvalue very close to the origin, requiring one to take  $\varepsilon$  very small to ensure  $0 \notin \sigma_\varepsilon(\mathbf{A})$ , but this effect is limited by the fact that  $\mathbf{A}$  must be close to normal. Moreover, if  $\mathbf{A}$  has other eigenvalues much farther from the origin, the small eigenvalue will cause GMRES to exhibit initial stagnation, and (FOV) will not capture the eventual convergence.)

**3.2. Summary of the Examples.** Let us collect some of the points highlighted in these examples.

- (EV) This bound works well for normal matrices and for (3.7), where all eigenvalues were uniformly ill-conditioned, but fails when the matrix was nondiagonalizable; it can also fail to be descriptive when  $\kappa(\mathbf{V})$  is large primarily because of a small number of ill-conditioned eigenvalues. (See the example in Section 4.1.)
- (FOV) This bound performs well when only one convergence stage was observed, as in Examples A and D. Its primary advantages over (PSA) for these examples was sharpness (Figures 3.2 and 3.5) and ease of computability. When GMRES exhibits an initial period of transient stagnation, (FOV) fails to capture the eventual convergence, as in Examples B, C, E, and F.
- (PSA) This bound inherits properties of both (EV) and (FOV), but can also capture interesting information between these extremes, as seen in Example F. The primary flaw in (PSA), exploited in Examples B and C, is its inability to recognize that the spectrum is a discrete point set, and thus it tends to overestimate the influence of outlying eigenvalues that GMRES can effectively eliminate at an early stage of convergence. The bounds (PSA') and (PSA'') suggest a way to address this shortcoming. (That said, pseudospectral techniques are not a panacea for bounding the 2-norm of matrix polynomials. For some extreme examples of their limitations, see [27, 41].)

In the next section, we illustrate how pseudospectra can yield convergence estimates during an iteration, and apply the bounds surveyed here to a matrix derived from a convection-diffusion problem.

**4. Adaptive Pseudospectral Bounds.** The pseudospectral bound (PSA) often provides a good indication of GMRES convergence, especially when considering a collection of bounds based on a wide range of  $\varepsilon$  values. Pseudospectra can be expensive to compute for large  $\mathbf{A}$ ; however, one can use elements from early GMRES iterations to *approximate* the pseudospectra of  $\mathbf{A}$ , and hence obtain an *estimate* for how the convergence will proceed. Such insight could, for example, give some indication of when to restart the GMRES algorithm. (An entirely different method for predicting future convergence based on early iterations has been suggested by Liesen [38].)

Suppose we have taken  $k$  steps of GMRES. The implementation of Saad and Schultz [45] uses the Arnoldi process to build an orthonormal basis  $\{\mathbf{v}_1, \dots, \mathbf{v}_{k+1}\}$  for the Krylov subspace  $\mathcal{K}_{k+1}(\mathbf{A}, \mathbf{r}_0) := \text{span}\{\mathbf{r}_0, \mathbf{A}\mathbf{r}_0, \dots, \mathbf{A}^k \mathbf{r}_0\}$ . Organize the basis

vectors into  $\mathbf{V}_k = [\mathbf{v}_1 \ \cdots \ \mathbf{v}_k] \in \mathbb{C}^{n \times k}$  and  $\mathbf{V}_{k+1} = [\mathbf{V}_k \ \mathbf{v}_{k+1}] \in \mathbb{C}^{n \times (k+1)}$ . The Arnoldi process gives a partial upper Hessenberg decomposition of  $\mathbf{A}$ ,

$$(4.1) \quad \mathbf{A}\mathbf{V}_k = \mathbf{V}_{k+1}\tilde{\mathbf{H}}_k \quad \text{and} \quad \mathbf{V}_k^*\mathbf{A}\mathbf{V}_k = \mathbf{H}_k,$$

where  $\tilde{\mathbf{H}}_k \in \mathbb{C}^{(k+1) \times k}$  is upper Hessenberg and  $\mathbf{H}_k \in \mathbb{C}^{k \times k}$  consists of the first  $k$  rows of  $\tilde{\mathbf{H}}_k$ ; see, e.g., [43, §6.3]. We can take the subdiagonal entries of  $\tilde{\mathbf{H}}_k$  to be nonnegative.

Toh and Trefethen [53] show that the pseudospectra of  $\mathbf{H}_k$  and  $\tilde{\mathbf{H}}_k$  can potentially yield good approximations to those of  $\sigma_\varepsilon(\mathbf{A})$  even when  $k \ll n$ . A point  $z \in \mathbb{C}$  is in the  $\varepsilon$ -pseudospectrum of the  $(k+1) \times k$  rectangular matrix  $\tilde{\mathbf{H}}_k$ ,  $z \in \sigma_\varepsilon(\tilde{\mathbf{H}}_k)$ , provided  $s_{\min}(z\tilde{\mathbf{I}}_k - \tilde{\mathbf{H}}_k) < \varepsilon$ , where  $\tilde{\mathbf{I}}_k$  is the  $k \times k$  identity matrix augmented by a row of zeros and  $s_{\min}(\cdot)$  denotes the smallest singular value; for details, see [53, 67]. With this definition,  $\sigma_\varepsilon(\mathbf{H}_k)$ ,  $\sigma_\varepsilon(\tilde{\mathbf{H}}_k)$ , and  $\sigma_\varepsilon(\mathbf{A})$  are related as follows.

**THEOREM 4.1.** *Suppose  $\mathbf{V}_k^*\mathbf{A}\mathbf{V}_k = \mathbf{H}_k$  and  $\mathbf{A}\mathbf{V}_k = \mathbf{V}_{k+1}\tilde{\mathbf{H}}_k$ . Then*

- (i)  $\sigma_\varepsilon(\tilde{\mathbf{H}}_1) \subseteq \sigma_\varepsilon(\tilde{\mathbf{H}}_2) \subseteq \cdots \subseteq \sigma_\varepsilon(\tilde{\mathbf{H}}_{n-1}) \subseteq \sigma_\varepsilon(\mathbf{H}_n) = \sigma_\varepsilon(\mathbf{A})$ .
- (ii)  $\sigma_\varepsilon(\mathbf{H}_k) \subseteq \sigma_{\hat{\varepsilon}}(\tilde{\mathbf{H}}_k) \subseteq \sigma_{\hat{\varepsilon}}(\mathbf{A})$ , where  $\hat{\varepsilon} := \varepsilon + h_{k+1,k}$ .

*Proof.* Toh and Trefethen proved part (i), which follows immediately from noting that  $s_{\min}(z\tilde{\mathbf{I}}_k - \tilde{\mathbf{H}}_k) \geq s_{\min}(z\tilde{\mathbf{I}}_{k+1} - \tilde{\mathbf{H}}_{k+1})$ . For part (ii), suppose that  $z \in \sigma_\varepsilon(\mathbf{H}_k)$ . Observe that  $s_{\min}(z\tilde{\mathbf{I}}_k - \tilde{\mathbf{H}}_k) \leq s_{\min}(z\mathbf{I} - \mathbf{H}_k) + h_{k+1,k} < \varepsilon + h_{k+1,k}$ , where the first inequality follows from [32, Thm. 3.3.16]. Applying part (i) to this bound completes the proof.  $\square$

Since the pseudospectra of  $\mathbf{H}_k$  and  $\tilde{\mathbf{H}}_k$  approximate those of  $\mathbf{A}$ , it is natural to approximate the bound (PSA) by replacing  $\sigma_\varepsilon(\mathbf{A})$  by  $\sigma_\varepsilon(\mathbf{H}_k)$  or  $\sigma_\varepsilon(\tilde{\mathbf{H}}_k)$ . The resulting expressions are no longer convergence bounds, but only estimates (as we indicate with the “ $\lesssim$ ” symbol):

$$(4.2) \quad \frac{\|\mathbf{r}_k\|_2}{\|\mathbf{r}_0\|_2} \lesssim \frac{\mathcal{L}(\Gamma_\varepsilon^{(k)})}{2\pi\varepsilon} \min_{\substack{p \in \mathcal{P}_k \\ p(0)=1}} \max_{z \in \sigma_\varepsilon(\mathbf{H}_k)} |p(z)|,$$

$$(4.3) \quad \frac{\|\mathbf{r}_k\|_2}{\|\mathbf{r}_0\|_2} \lesssim \frac{\mathcal{L}(\tilde{\Gamma}_\varepsilon^{(k)})}{2\pi\varepsilon} \min_{\substack{p \in \mathcal{P}_k \\ p(0)=1}} \max_{z \in \sigma_\varepsilon(\tilde{\mathbf{H}}_k)} |p(z)|,$$

where  $\Gamma_\varepsilon^{(k)}$  and  $\tilde{\Gamma}_\varepsilon^{(k)}$  denote Jordan curves enclosing  $\sigma_\varepsilon(\mathbf{H}_k)$  and  $\sigma_\varepsilon(\tilde{\mathbf{H}}_k)$ .

What value of  $\varepsilon$  is relevant at a specific iteration? The following bounds, while not necessarily sharp, suggest one way to approach this question. This proposition gives pseudospectral interpretations (cf. [48, Lemma 2.1]) of results about Ritz and Harmonic values. For the result about Ritz values, see, e.g., [37, §4.6]). The result about harmonic Ritz values follows from Simoncini and Gallopoulos [47]; see also Goossens and Roose [22].

**PROPOSITION 4.2.** *For  $k < n$ , the eigenvalues of  $\mathbf{H}_k$  (Ritz values) are contained in the  $\varepsilon$ -pseudospectrum of  $\mathbf{A}$  for  $\varepsilon = h_{k+1,k}$ . If  $\mathbf{H}_k$  is nonsingular, the roots of the GMRES residual polynomial (harmonic Ritz values [20, 22]) are contained in the  $\varepsilon$ -pseudospectrum of  $\mathbf{A}$  for  $\varepsilon = h_{k+1,k} + h_{k+1,k}^2/s_{\min}(\mathbf{H}_k)$ .*

*Proof.* To prove the first part, form the perturbation  $\mathbf{E} := -h_{k+1,k}\mathbf{v}_{k+1}\mathbf{e}_k^*\mathbf{V}_k^*$ . Then, using (4.1) and the fact that  $\mathbf{V}_k^*\mathbf{V}_k = \mathbf{I}$ ,

$$\begin{aligned} (\mathbf{A} + \mathbf{E})\mathbf{V}_k &= \mathbf{A}\mathbf{V}_k - h_{k+1,k}\mathbf{v}_{k+1}\mathbf{e}_k^*\mathbf{V}_k^*\mathbf{V}_k \\ &= (\mathbf{V}_k\mathbf{H}_k + h_{k+1,k}\mathbf{v}_{k+1}\mathbf{e}_k^*) - h_{k+1,k}\mathbf{v}_{k+1}\mathbf{e}_k^* = \mathbf{V}_k\mathbf{H}_k, \end{aligned}$$

where  $\mathbf{e}_k \in \mathbb{C}^k$  is the  $k$ th column of the  $k \times k$  identity. Since  $(\mathbf{A} + \mathbf{E})\mathbf{V}_k = \mathbf{V}_k\mathbf{H}_k$ ,  $\text{Ran}(\mathbf{V}_k)$  is an invariant subspace of  $\mathbf{A} + \mathbf{E}$ , and thus  $\sigma(\mathbf{H}_k) \subseteq \sigma(\mathbf{A} + \mathbf{E}) \subset \sigma_\varepsilon(\mathbf{A})$ , where  $\varepsilon := h_{k+1,k} = \|\mathbf{E}\|_2$ .

The harmonic Ritz values are the eigenvalues of  $(\mathbf{H}_k + h_{k+1,k}^2\mathbf{f}_k\mathbf{e}_k^*)$ , where  $\mathbf{f}_k := \mathbf{H}_k^{-*}\mathbf{e}_k$  [22, 47]. Defining the perturbation  $\mathbf{E} := (h_{k+1,k}^2\mathbf{V}_k\mathbf{f}_k\mathbf{e}_k^* - h_{k+1,k}\mathbf{v}_{k+1}\mathbf{e}_k^*)\mathbf{V}_k^*$ , from (4.1) it follows that

$$(\mathbf{A} + \mathbf{E})\mathbf{V}_k = \mathbf{V}_k(\mathbf{H}_k + h_{k+1,k}^2\mathbf{f}_k\mathbf{e}_k^*),$$

and thus  $\sigma(\mathbf{H}_k + h_{k+1,k}^2\mathbf{f}_k\mathbf{e}_k^*) \subseteq \sigma(\mathbf{A} + \mathbf{E}) \subset \sigma_\varepsilon(\mathbf{A})$ , where  $\|\mathbf{E}\|_2 \leq \varepsilon := h_{k+1,k} + h_{k+1,k}^2/s_{\min}(\mathbf{H}_k)$ .  $\square$

GMRES convergence *estimates* based on  $\sigma_\varepsilon(\mathbf{H}_k)$  have several applications. If  $k$  is small, the estimate might hint at GMRES behavior at future iterations. If  $k$  is larger (e.g., the  $k$  that satisfies the GMRES convergence criterion) and  $\sigma_\varepsilon(\mathbf{H}_k) \approx \sigma_\varepsilon(\mathbf{A})$  over a range of  $\varepsilon$  values, one could estimate the upper bound on GMRES convergence that might inform future runs of GMRES with the same  $\mathbf{A}$ . Note that  $\sigma_\varepsilon(\mathbf{H}_k)$  depends on the initial residual  $\mathbf{r}_0$ . If  $\mathbf{r}_0$  is deficient in all eigenvector directions associated with a particular eigenvalue, that eigenvalue cannot influence  $\mathbf{H}_k$  nor the GMRES estimates derived from it. If  $\mathbf{r}_0$  only has a small component in a certain eigenvector direction, that component may not exert much influence on early iterations (and  $\mathbf{H}_k$  for small  $k$ ), but become significant at later iterations.

Toh and Trefethen observe qualitative links between the pseudospectra and the GMRES iteration polynomial [50, 54]. A deeper quantitative understanding of this relationship could give insight about the ability of pseudospectral bounds to describe GMRES convergence, the value of  $\varepsilon$  for which  $\sigma_\varepsilon(\mathbf{A})$  gives the best bound at the  $k$ th iteration, and the merits of adaptive strategies, like the one described here, to capture significant features of convergence behavior.

Figure 4.1 shows adaptive convergence estimates drawn from the “integration matrix” (3.8), again with  $n = 64$  and  $\beta = 5/2$ , based on the pseudospectra of  $\mathbf{H}_k$  for  $k = 4$  and  $k = 12$  for a specific choice of initial residual, shown in Figure 4.2. In Figure 3.10, we saw that the bound (PSA) was descriptive for this example, and thus hope these Arnoldi estimates would perform similarly well. The estimates shown here are based on a random initial residual with entries drawn from the standard normal distribution. Taking estimates at iteration  $k = 4$  gives some hint of the quick convergence that follows; when  $k = 12$ , the pseudospectra of  $\mathbf{H}_k$  match those of  $\mathbf{A}$  for relevant values of  $\varepsilon$  and characterize the worst case convergence curve. The curve labeled  $\|\mathbf{r}_k\|_2/\|\mathbf{r}_0\|_2$  is the actual GMRES convergence obtained for this particular initial residual, with an asterisk marking the iteration from which the convergence estimate was drawn.

**4.1. A Practical Example.** We illustrate the use of this estimation technique, along with the bounds (EV), (EV'), (FOV), and (PSA), for a model problem from fluid dynamics. Let  $\mathbf{A}$  be the matrix generated by a streamline upwinded Petrov–Galerkin finite element discretization of the two-dimensional convection-diffusion equation,

$$-\nu\Delta u + \mathbf{w} \cdot \nabla u = f \quad \text{on } \Omega = [0, 1] \times [0, 1],$$

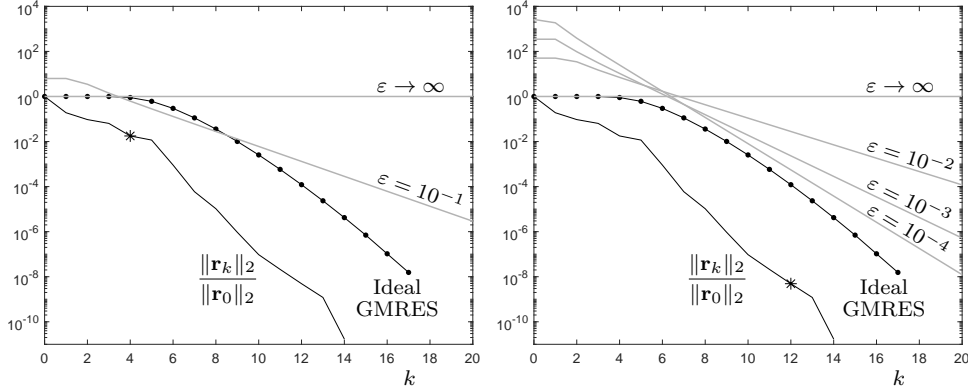


FIG. 4.1. Adaptive convergence estimates for the integration matrix (3.8) with  $\beta = 5/2$  and  $n = 64$ , generated at iteration  $k = 4$  on the left and  $k = 12$  on the right, based on pseudospectra of  $\mathbf{H}_k$  shown in Figure 4.2.




FIG. 4.2. At iterations  $k = 4$  and  $k = 12$ , the field of values  $W(\mathbf{H}_k)$  (gray region) and  $\epsilon$ -pseudospectra  $\sigma_\epsilon(\mathbf{H}_k)$  (for  $\epsilon = 10^{-1}, 10^{-1.5}, 10^{-2}, 10^{-3}, 10^{-4}$ ) for the matrix (3.8) with  $\beta = 5/2$  and  $n = 64$ , generated using the same  $\mathbf{r}_0$  whose GMRES convergence is illustrated in Figure 4.1. The color levels use the same scale as those in Figure 3.9 for the full matrix  $\mathbf{A}$ , to facilitate comparison.

with diffusion coefficient  $\nu = 0.01$ , constant advection in the vertical direction,  $\mathbf{w} = [0, 1]^T$ , and Dirichlet boundary conditions that induce an interior layer and a boundary layer. The solution is approximated using bilinear finite elements on a regular square grid with  $N$  unknowns in each coordinate direction, yielding a matrix  $\mathbf{A}$  of dimension  $n = N^2$ . This problem is discussed by Fischer et al. [19]; we apply the upwinding parameter they suggest, and focus on the case of  $N = 13$  ( $n = 169$ ). Though this upwinding parameter can give good approximate solutions to the partial differential equation, the corresponding matrix  $\mathbf{A}$  is highly nonnormal. Its eigenvalues, though very sensitive to perturbations, are known explicitly for this special wind direction [19]. These eigenvalues fall on  $N = 13$  lines in the complex plane with constant real part, with  $N = 13$  eigenvalues per line; the eigenvalues with largest real part are the most ill-conditioned. Liesen and Strakos have investigated the influence of the spectral properties of this class of discretizations on GMRES convergence [39].

Figure 4.3 shows GMRES convergence, along with the bounds (EV), (FOV), and (PSA) for this model problem. The Ideal GMRES curve was computed using the SDPT3 Toolbox [52], and we compare it to GMRES convergence for an initial residual

derived from boundary conditions that induce an internal layer and a boundary layer.

The matrix  $\mathbf{A}$  has a significant departure from normality, as illustrated in Figure 4.4. GMRES exhibits a period of slow convergence for typical initial residuals, followed by a more rapid phase of convergence. Ernst investigated the field of values bound (FOV) for this general problem [17]. Since  $\mathbf{A}$  results from a coercive finite element discretization,  $W(\mathbf{A})$  is contained in the open right-half plane and (FOV) guarantees convergence. Indeed, in this case the bound (FOV) gives  $\rho_{\text{FOV}} \approx 0.968$ . As expected given the initial period of slow convergence, this bound is descriptive at early iterations but fails to capture the faster second phase of convergence. As in Examples E and F, the bound (PSA) does better. Though this bound somewhat underestimates the convergence rate attained during the second phase of convergence (at least for the values of  $\varepsilon$  shown in Figure 4.3), it accurately captures the end of the slow first phase, a feature that eludes the bounds (EV) and (FOV).

Explicit formulas are available for the eigenvalues and eigenvectors of this matrix [19]. For our chosen parameters all the eigenvalues are distinct, and hence the matrix is diagonalizable; however, scaling each column of the eigenvector matrix to have unit 2-norm, we compute  $\kappa(\mathbf{V}) \approx 4.6 \times 10^{16}$ : the matrix is close to being non-diagonalizable. To get an upper bound on  $\rho_{\text{EV}}$ , we use the convergence rate associated with the convex hull of the true eigenvalues of  $\mathbf{A}$ ; Figure 4.3 shows that this *rate* seems to agree with the second phase of convergence (as observed in [19, p. 191]), but  $C_{\text{EV}}$  is much too large to make the bound descriptive. The pseudospectral bounds are better; we estimate the rates  $\rho_{\text{PSA}}(\varepsilon)$  by calculating the convergence rate of an approximate convex hull of  $\sigma_\varepsilon(\mathbf{A})$  and applying the convergence bound for convex sets (2.11).

The left plot in Figure 4.4 shows the eigenvalues, field of values, and pseudospectra for this example. While the eigenvalues of  $\mathbf{A}$  are ill-conditioned, the degree of ill-conditioning is not uniform across the spectrum: the eigenvalues closest to the origin are less sensitive than the rightmost eigenvalues. The right plot in Figure 4.4 plots  $\log_{10}(\kappa(\lambda_j))$  (rounded to the nearest integer) at the location of  $\lambda_j$ , for those eigenvalues on or above the real axis. Given the extreme range of  $\kappa(\lambda_j)$  values, how does the bound (EV') perform in this situation? Figure 4.5 illustrates that (EV') can handle such disparate ill-conditioning quite well. In the second phase of conver-




FIG. 4.3. Convergence bounds for the convection-diffusion problem with  $N = 13$ .




FIG. 4.4. On the left, the eigenvalues (black dots), field of values (gray region), and  $\varepsilon$ -pseudospectra ( $\varepsilon = 10^{-3}, 10^{-4}, \dots, 10^{-7}$ ) for the convection-diffusion problem with  $N = 13$ . On the right, the condition number of the eigenvalues used in (EV'), displayed as  $\log_{10}(\kappa(\lambda_j))$  rounded to the nearest integer and located at  $\lambda_j \in \mathbb{C}$ , for the eigenvalues on and above the real axis.




FIG. 4.5. Convergence bound (EV') for the convection-diffusion problem with  $N = 13$ .

gence, (EV') is more accurate than the three standard bounds. For this example, the eigenvalue condition numbers were computed from explicit formulas for left and right eigenvectors, and (EV') was calculated in quadruple precision arithmetic.

We test the pseudospectral estimates described earlier in this section for  $N = 13$  with the same initial residual described above. The adaptive estimates taken during the initial phase of slow convergence give little hint of future convergence behavior; the approximate pseudospectra do not improve much from iteration to iteration during these early steps. For  $k = 13$ , at the onset of more rapid convergence,  $\sigma_\varepsilon(\mathbf{H}_k) \approx \sigma_\varepsilon(\mathbf{A})$  for  $\varepsilon \geq 10^{-4}$  and one gets an indication of the improved convergence to come. By the time the convergence criterion is satisfied at  $k = 26$ ,  $\sigma_\varepsilon(\mathbf{H}_k) \approx \sigma_\varepsilon(\mathbf{A})$  for those values of  $\varepsilon$  relevant to the bound (PSA). Figure 4.6 shows these convergence estimates for  $k = 13$  and  $k = 26$ , based on the pseudospectra shown in Figure 4.7. Notice how the agreement between these pseudospectra of  $\mathbf{H}_k$  and those of the full matrix  $\mathbf{A}$  (shown in Figure 4.4) improve as  $k$  increases, in agreement with the observations of Toh and Trefethen [53].




FIG. 4.6. Adaptive convergence bounds for the convection-diffusion problem with  $N = 13$ , generated at iteration  $k = 13$  on the left and  $k = 26$  on the right, based on pseudospectra of  $\mathbf{H}_k$  shown in Figure 4.7.




FIG. 4.7. At iterations  $k = 13$  and  $k = 26$ , the field of values  $W(\mathbf{H}_k)$  (gray region) and  $\varepsilon$ -pseudospectra  $\sigma_\varepsilon(\mathbf{H}_k)$  (for  $\varepsilon = 10^{-3}, 10^{-4}, 10^{-5}, 10^{-6}, 10^{-7}$ ) for the convection-diffusion problem with  $N = 13$ , generated using the  $\mathbf{r}_0$  whose GMRES convergence is illustrated in Figure 4.6. The color levels use the same scale as those in Figure 4.4 for the full matrix  $\mathbf{A}$ , to facilitate comparison.

**5. Summary.** We have explored some of the relative merits of convergence bounds based on eigenvalues (with the eigenvector condition number), the field of values, and pseudospectra. In particular, these bounds have distinct weaknesses that indicate situations in which one bound may be preferred over the others. The standard bounds are global statements that can be refined; for (EV') and (FOV') this localization introduces spectral projector norms. Pseudospectra provide a convenient tool for bridging between the eigenvalues and the field of values, but they can be expensive to compute. Approximate pseudospectra drawn from the Arnoldi process yield convergence estimates at a fraction of the cost of full pseudospectral computation.

**Acknowledgements.** I thank Andy Wathen for guiding this research and suggesting numerous improvements to this presentation. I am also grateful for Nick Trefethen's many helpful comments. The title and the format of Section 3 were inspired by a paper by Nachtigal, Reddy, and Trefethen [40]. I thank Anne Greenbaum for providing a copy of reference [12], and Chris Beattie, Bernd Fischer, and Henk van der Vorst for stimulating discussions related to this research. The helpful referees who reviewed the original version of this document pointed out several references, and made other suggestions that have improved this work.

## REFERENCES

- [1] F. L. BAUER AND C. T. FIKE, *Norms and exclusion theorems*, Numer. Math., (1960), pp. 137–141.
- [2] A. BÖTTCHER AND S. M. GRUDSKY, *Spectral Properties of Banded Toeplitz Matrices*, SIAM, Philadelphia, 2005.
- [3] T. BRACONNIER AND N. J. HIGHAM, *Computing the field of values and pseudospectra using the Lanczos method with continuation*, BIT, 36 (1996), pp. 422–440.
- [4] M. CROUZEIX, *Numerical range and functional calculus in Hilbert space*, J. Functional Anal., 244 (2007), pp. 668–690.
- [5] M. CROUZEIX AND A. GREENBAUM, *Spectral sets: numerical range and beyond*, SIAM J. Matrix Anal. Appl., 40 (2019), pp. 1087–1101.
- [6] M. CROUZEIX AND C. PALENCIA, *The numerical range is a  $(1+\sqrt{2})$ -spectral set*, SIAM J. Matrix Anal. Appl., 38 (2017), pp. 649–655.
- [7] T. A. DRISCOLL, *A MATLAB toolbox for Schwarz–Christoffel mapping*, ACM Trans. Math. Software, 22 (1996), pp. 168–186. For accompanying software, see <https://tobydriscoll.net/project/sc-toolbox/>.
- [8] T. A. DRISCOLL, N. HALE, AND L. N. TREFETHEN, eds., *Chebfun Guide*, Pafnuty Publications, Oxford, 2014.
- [9] T. A. DRISCOLL, K.-C. TOH, AND L. N. TREFETHEN, *From potential theory to matrix iterations in six steps*, SIAM Review, 40 (1998), pp. 547–578.
- [10] M. EIERMANN, *On semiiterative methods generated by Faber polynomials*, Numer. Math., 56 (1989), pp. 139–156.
- [11] ———, *Fields of values and iterative methods*, Linear Algebra Appl., 180 (1993), pp. 167–197.
- [12] ———, *Field of values and iterative methods*. Slides from the Oberwolfach Conference on Iterative Methods, Oberwolfach, Germany, April 1997.
- [13] M. EIERMANN AND O. G. ERNST, *Geometric aspects in the theory of Krylov subspace methods*, Acta Numerica, 10 (2001), pp. 251–312.
- [14] S. C. EISENSTAT, H. C. ELMAN, AND M. H. SCHULTZ, *Variational iterative methods for non-symmetric systems of linear equations*, SIAM J. Numer. Anal., 20 (1983), pp. 345–357.
- [15] M. EMBREE, *The tortoise and the hare restart GMRES*, SIAM Review, 45 (2003), pp. 259–266.
- [16] M. EMBREE AND L. N. TREFETHEN, *Green’s functions for multiply connected domains via conformal mapping*, SIAM Review, 41 (1999), pp. 745–761.
- [17] O. G. ERNST, *Residual-minimizing Krylov subspace methods for stabilized discretizations of convection-diffusion equations*, SIAM J. Matrix Anal. Appl., 21 (2000), pp. 1079–1101.
- [18] B. FISCHER, *Polynomial Based Iteration Methods for Symmetric Linear Systems*, Wiley–Teubner, Chichester, 1996.
- [19] B. FISCHER, A. RAMAGE, D. J. SILVESTER, AND A. J. WATHEN, *On parameter choice and iterative convergence for stabilised discretizations of advection–diffusion problems*, Comp. Methods Appl. Mech. Eng., 179 (1999), pp. 179–195.
- [20] R. W. FREUND, *Quasi-kernel polynomials and their use in non-Hermitian matrix iterations*, J. Comp. Appl. Math., 43 (1992), pp. 135–158.
- [21] R. W. FREUND AND N. M. NACHTIGAL, *QMR: A quasi-minimal residual method for non-Hermitian linear systems*, Numer. Math., 60 (1991), pp. 315–339.
- [22] S. GOOSSENS AND D. ROOSE, *Ritz and harmonic Ritz values and the convergence of FOM and GMRES*, Numer. Linear Algebra Appl., 6 (1999), pp. 281–293.
- [23] A. GOPAL AND L. N. TREFETHEN, *Representation of conformal maps by rational functions*, Numer. Math., 142 (2019), pp. 359–382.
- [24] A. GREENBAUM, *Iterative Methods for Solving Linear Systems*, SIAM, Philadelphia, 1997.
- [25] A. GREENBAUM, V. PTÁK, AND Z. STRAKOŠ, *Any nonincreasing convergence curve is possible for GMRES*, SIAM J. Matrix Anal. Appl., 17 (1996), pp. 465–469.
- [26] A. GREENBAUM AND Z. STRAKOŠ, *Matrices that generate the same Krylov residual spaces*, in Recent Advances in Iterative Methods, G. Golub, A. Greenbaum, and M. Luskin, eds., Springer-Verlag, New York, 1994, pp. 95–118.
- [27] A. GREENBAUM AND L. N. TREFETHEN, *Do the pseudospectra of a matrix determine its behavior?*, Tech. Rep. TR 93-1371, Computer Science Department, Cornell University, August 1993.
- [28] K. E. GUSTAFSON AND D. K. M. RAO, *Numerical Range: The Field of Values of Linear Operators and Matrices*, Springer-Verlag, New York, 1997.
- [29] D. J. HIGHAM AND L. N. TREFETHEN, *Stiffness of ODEs*, BIT, 33 (1993), pp. 285–303.
- [30] N. J. HIGHAM, *The test matrix toolbox for MATLAB (version 3.0)*, Tech. Rep. Numerical Analysis Report No. 276, University of Manchester, September 1995.

- [31] E. HILLE, *Analytic Function Theory*, vol. 2, Chelsea, New York, 1962.
- [32] R. A. HORN AND C. R. JOHNSON, *Topics in Matrix Analysis*, Cambridge University Press, Cambridge, 1991.
- [33] I. C. F. IPSEN, *Expressions and bounds for the GMRES residual*, BIT, 40 (2000), pp. 524–535.
- [34] W. JOUBERT, *On the convergence behavior of the restarted GMRES algorithm for solving nonsymmetric linear systems*, Numer. Linear Algebra Appl., 1 (1994), pp. 427–447.
- [35] T. KATO, *Perturbation Theory for Linear Operators*, Springer-Verlag, Berlin, corrected second ed., 1980.
- [36] T. KÖVARI AND C. POMMERENKE, *On Faber polynomials and Faber expansions*, Math. Zeit., 99 (1967), pp. 193–206.
- [37] R. B. LEHOUCQ, D. C. SORENSEN, AND C. YANG, *ARPACK Users' Guide: Solution of Large-Scale Eigenvalue Problems with Implicitly Restarted Arnoldi Methods*, SIAM, Philadelphia, 1998.
- [38] J. LIESEN, *Computable convergence bounds for GMRES*, SIAM J. Matrix Anal. Appl., 21 (2000), pp. 882–903.
- [39] J. LIESEN AND Z. STRAKOŠ, *GMRES convergence analysis for a convection-diffusion model problem*, SIAM J. Sci. Comput., 26 (2005), pp. 1989–2009.
- [40] N. M. NACHTIGAL, S. C. REDDY, AND L. N. TREFETHEN, *How fast are nonsymmetric matrix iterations?*, SIAM J. Matrix Anal. Appl., 13 (1992), pp. 778–795.
- [41] T. RANSFORD AND J. ROSTAND, *Pseudospectra do not determine norm behavior, even for matrices with only simple eigenvalues*, Linear Algebra Appl., 435 (2011), pp. 3024–3028.
- [42] L. REICHEL AND L. N. TREFETHEN, *Eigenvalues and pseudo-eigenvalues of Toeplitz matrices*, Linear Algebra Appl., 162–164 (1992), pp. 153–185.
- [43] Y. SAAD, *Iterative Methods for Sparse Linear Systems*, SIAM, Philadelphia, second ed., 2003.
- [44] ———, *Numerical Methods for Large Eigenvalue Problems*, SIAM, Philadelphia, second ed., 2011.
- [45] Y. SAAD AND M. H. SCHULTZ, *GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems*, SIAM J. Sci. Stat. Comput., 7 (1986), pp. 856–869.
- [46] J. A. SIFUENTES, M. EMBREE, AND R. B. MORGAN, *GMRES convergence for perturbed coefficient matrices, with application to approximate deflation preconditioning*, SIAM J. Matrix Anal. Appl., 34 (2013), pp. 1066–1088.
- [47] V. SIMONCINI AND E. GALLOPOULOS, *Convergence properties of block GMRES and matrix polynomials*, Linear Algebra Appl., 247 (1996), pp. 97–119.
- [48] ———, *Transfer functions and resolvent norm approximation of large matrices*, Elec. Trans. Numer. Anal., 7 (1998), pp. 190–201.
- [49] M. H. STONE, *Linear Transformations in Hilbert Space*, American Mathematical Society, New York, 1932.
- [50] K.-C. TOH, *Matrix Approximation Problems and Nonsymmetric Iterative Methods*, PhD thesis, Cornell University, August 1996.
- [51] ———, *GMRES vs. ideal GMRES*, SIAM J. Matrix Anal. Appl., 18 (1997), pp. 30–36.
- [52] K.-C. TOH, M. J. TODD, AND R. H. TÜTÜNCÜ, *SDPT3—a Matlab software package for semidefinite programming*, Opt. Meth. Softw., 11 (1999), pp. 545–581. Software available at <https://github.com/sqlp/sdpt3>.
- [53] K.-C. TOH AND L. N. TREFETHEN, *Calculation of pseudospectra by the Arnoldi iteration*, SIAM J. Sci. Comput., 17 (1996), pp. 1–15.
- [54] ———, *The Chebyshev polynomials of a matrix*, SIAM J. Matrix Anal. Appl., 20 (1999), pp. 400–419.
- [55] L. N. TREFETHEN, *Approximation theory and numerical linear algebra*, in Algorithms for Approximation II, J. C. Mason and M. G. Cox, eds., Chapman and Hall, London, 1990.
- [56] L. N. TREFETHEN, *Pseudospectra of matrices*, in Numerical Analysis 1991, D. F. Griffiths and G. A. Watson, eds., Longman Scientific and Technical, Harlow, Essex, UK, 1992, pp. 234–266.
- [57] ———, *Computation of pseudospectra*, Acta Numerica, (1999), pp. 247–295.
- [58] L. N. TREFETHEN AND D. BAU, III, *Numerical Linear Algebra*, SIAM, Philadelphia, 1997.
- [59] L. N. TREFETHEN AND M. EMBREE, *Spectra and Pseudospectra: The Behavior of Nonnormal Matrices and Operators*, Princeton University Press, Princeton, NJ, 2005.
- [60] A. VAN DER SLUIS, *Condition numbers and equilibration of matrices*, Numer. Math., 14 (1969), pp. 14–23.
- [61] A. VAN DER SLUIS AND H. VAN DER VORST, *The rate of convergence of conjugate gradients*, Numer. Math., 48 (1986), pp. 543–560.
- [62] H. A. VAN DER VORST AND C. VUIK, *The superlinear convergence behaviour of GMRES*, J. Comp. Appl. Math., 48 (1993), pp. 327–341.

- [63] J. H. WILKINSON, *The Algebraic Eigenvalue Problem*, Oxford University Press, Oxford, 1965.
- [64] T. G. WRIGHT, *Algorithms and Software for Pseudospectra*, 2002. D.Phil. thesis, Oxford University.
- [65] ———, *EigTool*, 2002. Software available at <https://github.com/eigtool>.
- [66] T. G. WRIGHT AND L. N. TREFETHEN, *Large-scale computation of pseudospectra using ARPACK and eigs*, SIAM J. Sci. Comput., 23 (2001), pp. 591–605.
- [67] ———, *Pseudospectra of rectangular matrices*, IMA J. Numer. Anal., 22 (2002), pp. 501–519.