

CONVERGENCE OF RESTARTED KRYLOV SUBSPACES TO INVARIANT SUBSPACES

CHRISTOPHER BEATTIE*, MARK EMBREE†, AND JOHN ROSSI*

Abstract. The performance of Krylov subspace eigenvalue algorithms for large matrices can be measured by the angle between a desired invariant subspace and the Krylov subspace. We develop general bounds for this convergence that include the effects of polynomial restarting and impose no restrictions concerning the diagonalizability of the matrix or its degree of non-normality. Associated with a desired set of eigenvalues is a maximum “reachable invariant subspace” that can be developed from the given starting vector. Convergence for this distinguished subspace is bounded in terms involving a polynomial approximation problem. Elementary results from potential theory lead to convergence rate estimates and suggest restarting strategies based on optimal approximation points (e.g., Leja or Chebyshev points); exact shifts are evaluated within this framework. Computational examples illustrate the utility of these results. Origins of superlinear effects are also described.

Key words. Krylov subspace iteration, Arnoldi methods, restarts, perturbation theory, eigenvalues, pseudospectra, potential theory, invariant subspaces.

AMS subject classification. 15A18, 15A42, 31A15, 41A25, 65F15

1. Setting. Let \mathbf{A} be an $n \times n$ complex matrix with $N \leq n$ distinct eigenvalues $\{\lambda_j\}_{j=1}^N$ with corresponding eigenvectors $\{\mathbf{u}_j\}_{j=1}^N$. (We do not label multiple eigenvalues separately and make no assertion regarding the uniqueness of the \mathbf{u}_j .) Each distinct eigenvalue λ_j has geometric multiplicity n_j and algebraic multiplicity m_j (so that $1 \leq n_j \leq m_j$ and $\sum_{j=1}^N m_j = n$). We aim to compute an invariant subspace associated with L of these eigenvalues, which for brevity we call the *good* eigenvalues, labeled $\{\lambda_1, \lambda_2, \dots, \lambda_L\}$. We intend to use a Krylov subspace algorithm to approximate this invariant subspace, possibly with the aid of restarts as described below. The remaining $N - L$ eigenvalues, the *bad* eigenvalues, are not of interest and we wish to avoid excessive expense involved in inadvertently calculating the subspaces associated with them.

The class of algorithms considered here draw eigenvector approximations from Krylov subspaces generated by the starting vector $\mathbf{v}_1 \in \mathbb{C}^n$,

$$\mathcal{K}_\ell(\mathbf{A}, \mathbf{v}_1) = \text{span}\{\mathbf{v}_1, \mathbf{A}\mathbf{v}_1, \dots, \mathbf{A}^{\ell-1}\mathbf{v}_1\}.$$

Such algorithms, including the Arnoldi and bi-orthogonal Lanczos methods reviewed in §1.1, differ in their mechanisms for generating a basis for $\mathcal{K}_\ell(\mathbf{A}, \mathbf{v}_1)$ and selecting approximate eigenvectors from this Krylov subspace. Though these approximate eigenvectors may appear to be obvious objects of study, their convergence can be greatly complicated by eigenvalue multiplicity and defectiveness; see [21]. The bounds developed in the following sections avoid these difficulties by instead studying convergence of the Krylov subspace to an invariant subspace associated with the good eigenvalues as the dimension of the Krylov subspace is increased.

*Department of Mathematics, Virginia Polytechnic Institute and State University, Blacksburg, VA 24061-0123, USA (beattie@math.vt.edu, rossi@math.vt.edu),

†Oxford University Computing Laboratory, Wolfson Building, Parks Road, Oxford OX1 3QD, United Kingdom (mark.embree@comlab.ox.ac.uk). Supported in part by UK Engineering and Physical Sciences Research Council Grant GR/M12414.

Given two subspaces, \mathcal{W} and \mathcal{V} of \mathbb{C}^n , the proximity of one to the other is measured by the *containment gap* (or just *gap*), defined as

$$\delta(\mathcal{W}, \mathcal{V}) = \sup_{\mathbf{x} \in \mathcal{W}} \inf_{\mathbf{y} \in \mathcal{V}} \frac{\|\mathbf{y} - \mathbf{x}\|}{\|\mathbf{x}\|} = \sin(\vartheta_{\max}).$$

Here ϑ_{\max} is the largest canonical angle between \mathcal{W} and a “closest” subspace $\widehat{\mathcal{V}}$ of \mathcal{V} having dimension equal to $\dim \mathcal{W}$. (Throughout, $\|\cdot\|$ denotes the vector 2-norm and the matrix norm it induces.) Notice that if $\dim \mathcal{V} < \dim \mathcal{W}$ then $\delta(\mathcal{W}, \mathcal{V}) = 1$, while $\delta(\mathcal{W}, \mathcal{V}) = 0$ if and only if $\mathcal{W} \subseteq \mathcal{V}$. The gap can be expressed directly as the norm of a composition of projections: If $\mathbf{\Pi}_{\mathcal{W}}$ and $\mathbf{\Pi}_{\mathcal{V}}$ denote orthogonal projections onto \mathcal{W} and \mathcal{V} , respectively, then $\delta(\mathcal{W}, \mathcal{V}) = \|(\mathbf{I} - \mathbf{\Pi}_{\mathcal{V}})\mathbf{\Pi}_{\mathcal{W}}\|$ (see, e.g., Chatelin [7, §1.4]).

The objective of this paper then is to measure the gap between Krylov subspaces and an m -dimensional invariant subspace \mathcal{U} of \mathbf{A} associated with the good eigenvalues. We explore how quickly $\delta(\mathcal{U}, \mathcal{K}_\ell(\mathbf{A}, \mathbf{v}_1))$ can be driven to zero, reflecting the speed of convergence, and how this behavior is influenced by the distribution of eigenvalues and non-normality of \mathbf{A} . Note that $\delta(\mathcal{U}, \mathcal{K}_\ell(\mathbf{A}, \mathbf{v}_1)) = 1$ when $\ell < m$. For $\ell \geq m$, our bounds ultimately take the form

$$(1.1) \quad \delta(\mathcal{U}, \mathcal{K}_\ell(\mathbf{A}, \mathbf{v}_1)) \leq C_0 C_1 C_2 \min_{\phi \in \mathcal{P}_{\ell-m}} \frac{\max\{|\phi(z)| : z \in \Omega_{\text{bad}}\}}{\min\{|\phi(z)| : z \in \Omega_{\text{good}}\}},$$

where \mathcal{P}_ℓ is the set of degree- ℓ polynomials, and Ω_{good} and Ω_{bad} are disjoint compact subsets of \mathbb{C} containing the good and bad eigenvalues, respectively. Here C_0 and C_2 are constants depending on the non-normality of \mathbf{A} , while the constant C_1 also incorporates starting vector biases. In §2 we identify the subspace \mathcal{U} , which in common situations will be the entire invariant subspace of \mathbf{A} associated with the good eigenvalues, but will be smaller when \mathbf{A} is derogatory or the starting vector \mathbf{v}_1 is deficient. The basic bound (1.1) is derived in §3. Section 4 addresses the polynomial approximation problem embedded in (1.1), describing those factors that determine linear convergence rates and can lead to superlinear effects. Section 5 analyzes the constants C_1 and C_2 , and §6 provides computational examples illustrating the bounds.

Since it becomes prohibitively expensive to construct and store a good basis for $\mathcal{K}_\ell(\mathbf{A}, \mathbf{v}_1)$ when the dimension of \mathbf{A} is large, practical algorithms typically limit the maximum dimension of the Krylov subspace to some $p \ll n$. If satisfactory eigenvector estimates cannot be extracted from $\mathcal{K}_p(\mathbf{A}, \mathbf{v}_1)$, then the algorithm is *restarted* by replacing \mathbf{v}_1 with some new $\mathbf{v} \in \mathcal{K}_p(\mathbf{A}, \mathbf{v}_1)$ that is, hopefully, enriched in the component lying in the subspace \mathcal{U} . Since this \mathbf{v} is chosen from the Krylov subspace, we can write $\mathbf{v} = \psi(\mathbf{A})\mathbf{v}_1$ for some polynomial ψ with $\deg(\psi) < p$. Our bounds also apply to this situation, and ideas from potential theory, outlined in §4, motivate particular choices for the polynomial ψ .

The results presented here complement and extend earlier convergence theory. Saad bounded the gap between a single eigenvector and the Krylov subspace for a matrix with simple eigenvalues [34]. Jia generalized this result to invariant subspaces associated with a single eigenvalue for a defective matrix, but these bounds involve the Jordan form of \mathbf{A} and derivatives of approximating polynomials [20]. Simoncini uses pseudospectra to describe block Arnoldi convergence for defective matrices [39]. Interpreting restarted algorithms in terms of subspace iteration, Lehoucq developed

an invariant subspace convergence theory for matrices with simple eigenvalues [26]. Calvetti, Reichel, and Sorensen have studied single eigenvector convergence for Hermitian matrices using elements of potential theory [6]. A key feature of our approach is its applicability to general invariant subspaces, which may be better conditioned than individual eigenvectors (see, e.g., [42, Ch. V]). Notably, we estimate convergence rates for defective matrices without introducing any special choice of basis and without requiring knowledge of the Jordan form or any related similarity transformation.

Finally, we note that other measures of convergence may be more appealing in certain situations. Alternatives include Ritz values [20, 24], although convergence behavior might be obscure for matrices that are defective (or nearly so). The subspace residual is computationally attractive because it doesn't require *a priori* knowledge of the good invariant subspace. This measure can be related to gap convergence [17, 41].

1.1. Arnoldi and Bi-orthogonal Lanczos Algorithms. Suppose \mathbf{V} is an $n \times n$ unitary matrix that reduces \mathbf{A} to upper Hessenberg form; i.e., $\mathbf{V}^* \mathbf{A} \mathbf{V} = \mathbf{H}$ for some upper Hessenberg matrix, \mathbf{H} . For any index $1 \leq \ell \leq n$, let \mathbf{H}_ℓ denote the ℓ th principal submatrix of \mathbf{H} :

$$\mathbf{H}_\ell = \begin{bmatrix} h_{11} & h_{12} & \cdots & h_{1\ell} \\ \beta_2 & h_{22} & \cdots & h_{2\ell} \\ & \ddots & \ddots & \vdots \\ & & \beta_\ell & h_{\ell\ell} \end{bmatrix} \in \mathbb{C}^{\ell \times \ell}.$$

The Arnoldi method [2, 34] builds up the matrices \mathbf{H} and \mathbf{V} one column at a time starting with the unit vector $\mathbf{v}_1 \in \mathbb{C}^n$, although the process is typically stopped well before completion, with $\ell \ll n$. The algorithm only accesses \mathbf{A} through matrix-vector products, making this approach attractive when \mathbf{A} is large and sparse.

Different choices for \mathbf{v}_1 produce distinct outcomes for \mathbf{H}_ℓ . The defining recurrence may be derived from the fundamental relation

$$\mathbf{A} \mathbf{V}_\ell = \mathbf{V}_\ell \mathbf{H}_\ell + \beta_{\ell+1} \mathbf{v}_{\ell+1} \mathbf{e}_\ell^*,$$

where \mathbf{e}_ℓ is the ℓ th column of the $\ell \times \ell$ identity matrix. The ℓ th column of \mathbf{H}_ℓ is determined so as to force $\mathbf{v}_{\ell+1}$ to be orthogonal to the columns of \mathbf{V}_ℓ , and $\beta_{\ell+1}$ then is determined so that $\|\mathbf{v}_{\ell+1}\| = 1$. After ℓ steps, the columns of \mathbf{V}_ℓ constitute an orthonormal basis for the order- ℓ Krylov subspace $\mathcal{K}_\ell(\mathbf{A}, \mathbf{v}_1) = \text{span}\{\mathbf{v}_1, \mathbf{A}\mathbf{v}_1, \dots, \mathbf{A}^{\ell-1}\mathbf{v}_1\}$. Since $\mathbf{V}_\ell^* \mathbf{A} \mathbf{V}_\ell = \mathbf{H}_\ell$, the matrix \mathbf{H}_ℓ is a Ritz–Galerkin approximation of \mathbf{A} on this subspace, as described by Saad [35]. The eigenvalues of \mathbf{H}_ℓ are called *Ritz values* and will, in many circumstances, be reasonable approximations to some of the eigenvalues of \mathbf{A} . An eigenvector of \mathbf{H}_ℓ associated with a given Ritz value θ_j can be used to construct an eigenvector approximation for \mathbf{A} . Indeed, if $\mathbf{H}_\ell \mathbf{y}_j = \theta_j \mathbf{y}_j$, then the *Ritz vector* $\hat{\mathbf{u}}_j = \mathbf{V}_\ell \mathbf{y}_j$ yields the residual

$$\|\mathbf{A} \hat{\mathbf{u}}_j - \theta_j \hat{\mathbf{u}}_j\| = |\beta_{\ell+1}| |\mathbf{e}_\ell^* \mathbf{y}_j|.$$

When $|\beta_{\ell+1}| \ll 1$, the columns of \mathbf{V}_ℓ nearly span an invariant subspace of \mathbf{A} . Small residuals more often arise from negligible trailing entries of the vector \mathbf{y}_j , indicating the most recent Krylov direction contributed negligibly to the Ritz vector $\hat{\mathbf{u}}_j$.

Alternatively, suppose $\widehat{\mathbf{V}}$ is an $n \times n$ invertible matrix that transforms \mathbf{A} via similarity to tridiagonal form; i.e., $\widehat{\mathbf{V}}^{-1}\mathbf{A}\widehat{\mathbf{V}} = \mathbf{T}$ for some tridiagonal matrix \mathbf{T} . For any index $1 \leq \ell \leq n$, let \mathbf{T}_ℓ denote the ℓ th principal submatrix of \mathbf{T} :

$$\mathbf{T}_\ell = \begin{bmatrix} \alpha_1 & \gamma_2 & & & \\ \delta_2 & \alpha_2 & \ddots & & \\ & \ddots & \ddots & \gamma_\ell & \\ & & \delta_\ell & \alpha_\ell & \\ & & & & \alpha_\ell \end{bmatrix} \in \mathbb{C}^{\ell \times \ell}.$$

Let $\widehat{\mathbf{V}}_\ell = [\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_\ell]$ denote a matrix containing the first ℓ columns of $\widehat{\mathbf{V}}$; and for $\widehat{\mathbf{W}}^* = \widehat{\mathbf{V}}^{-1}$, let $\widehat{\mathbf{W}}_\ell = [\mathbf{w}_1, \mathbf{w}_2, \dots, \mathbf{w}_\ell]$ denote a matrix containing the first ℓ columns of $\widehat{\mathbf{W}}$.

The bi-orthogonal Lanczos algorithm [25] builds up the matrices \mathbf{T} , $\widehat{\mathbf{V}}$, and $\widehat{\mathbf{W}}$ one column at a time starting with the unit vectors \mathbf{v}_1 and \mathbf{w}_1 , accessing \mathbf{A} and \mathbf{A}^* only through matrix-vector product operations. Different choices for \mathbf{v}_1 and \mathbf{w}_1 produce distinct outcomes for \mathbf{T} , if all goes well. Recovering from situations where not all goes well is a fundamental aspect of later refinements of the algorithm; two such approaches are discussed in [16] and [30].

At the ℓ th step, the basic recursion is

$$\begin{aligned} \mathbf{A}\widehat{\mathbf{V}}_\ell &= \widehat{\mathbf{V}}_\ell \mathbf{T}_\ell + \delta_{\ell+1} \mathbf{v}_{\ell+1} \mathbf{e}_\ell^*, \\ \mathbf{A}^* \widehat{\mathbf{W}}_\ell &= \widehat{\mathbf{W}}_\ell \mathbf{T}_\ell^* + \gamma_{\ell+1} \mathbf{w}_{\ell+1} \mathbf{e}_\ell^*. \end{aligned}$$

Typically, normalization is determined so that $|\gamma_j| = |\delta_j|$ and $\mathbf{v}_j^* \mathbf{w}_j = 1$. With exact arithmetic, the first $\ell - 1$ steps yield matrices $\widehat{\mathbf{V}}_\ell$ and $\widehat{\mathbf{W}}_\ell$ that satisfy

- $\widehat{\mathbf{W}}_\ell^* \widehat{\mathbf{V}}_\ell = \mathbf{I}$,
- $\text{Ran}(\widehat{\mathbf{V}}_\ell) = \text{span}\{\mathbf{v}_1, \mathbf{A}\mathbf{v}_1, \dots, \mathbf{A}^{\ell-1}\mathbf{v}_1\} = \mathcal{K}_\ell(\mathbf{A}, \mathbf{v}_1)$, and
- $\text{Ran}(\widehat{\mathbf{W}}_\ell) = \text{span}\{\mathbf{w}_1, \mathbf{A}^* \mathbf{w}_1, \dots, (\mathbf{A}^*)^{\ell-1} \mathbf{w}_1\} = \mathcal{K}_\ell(\mathbf{A}^*, \mathbf{w}_1)$.

Notice that $\mathbf{T}_\ell = \widehat{\mathbf{W}}_\ell^* \mathbf{A} \widehat{\mathbf{V}}_\ell$, and so the eigenvalues for \mathbf{T}_ℓ (also called Ritz values) are Petrov–Galerkin approximations to the eigenvalues of \mathbf{A} [35]. In both the Arnoldi and bi-orthogonal Lanczos methods, approximations to the (right) invariant subspaces of \mathbf{A} are drawn from the *same* subspace $\mathcal{K}_\ell(\mathbf{A}, \mathbf{v}_1)$, though with respect to different bases: \mathbf{V}_ℓ vs. $\widehat{\mathbf{V}}_\ell$. However, bi-orthogonal Lanczos's \mathbf{T}_ℓ is *not* in general similar to Arnoldi's \mathbf{H}_ℓ ; indeed, the bi-orthogonal Lanczos eigenvalue approximations typically differ considerably from those provided by Arnoldi.

Our focus here will remain fixed on how well a good invariant subspace \mathcal{U} is captured by $\mathcal{K}_\ell(\mathbf{A}, \mathbf{v}_1)$ without regard to how a basis for $\mathcal{K}_\ell(\mathbf{A}, \mathbf{v}_1)$ has been generated. Further algorithmic details of particular methods can be found in [4].

1.2. Polynomial Restarts. The cost of proceeding through p steps of the Arnoldi recurrence is roughly the cost of p matrix-vector products of the form $\mathbf{A}\mathbf{v}_k$ on top of the net $2np^2$ floating point operations necessary for orthogonalization. Bi-orthogonal Lanczos requires $2p$ matrix-vector products (of the form $\mathbf{A}\mathbf{v}_k$ and $\mathbf{A}^* \mathbf{w}_k$) and also on the order of np^2 floating point operations to enforce bi-orthogonality. If n is very large and \mathbf{A} is very sparse (say, with a maximum number of nonzero entries per row very much smaller than n), then the cost of orthogonalization can quickly

dominate as p grows. One general approach for alleviating the growing cost of orthogonalization is known as *polynomial restarting*. At the end of $p + 1$ steps of the recurrence, one selects some “best” vector $\mathbf{v}_1^+ \in \mathcal{K}_{p+1}(\mathbf{A}, \mathbf{v}_1)$ and restarts the recurrence from the beginning using \mathbf{v}_1^+ . Different restart strategies differ essentially in how they attempt to condense progress made in the last $p + 1$ steps into the vector \mathbf{v}_1^+ . Since any vector in $\mathcal{K}_{p+1}(\mathbf{A}, \mathbf{v}_1)$ can be represented as $\psi_p(\mathbf{A})\mathbf{v}_1$ for some polynomial ψ_p of degree p or less, a restart of this type can be expressed as

$$(1.2) \quad \mathbf{v}_1^+ \leftarrow \psi_p(\mathbf{A})\mathbf{v}_1.$$

If subsequent restarts occur (relabeling \mathbf{v}_1^+ as $\mathbf{v}_1^{(1)}$) then

$$\begin{aligned} \mathbf{v}_1^{(1)} &\leftarrow \psi_p^{[1]}(\mathbf{A})\mathbf{v}_1 && \text{(first restart)} \\ \mathbf{v}_1^{(2)} &\leftarrow \psi_p^{[2]}(\mathbf{A})\mathbf{v}_1^{(1)} && \text{(second restart)} \\ &\vdots \\ \mathbf{v}_1^{(\nu)} &\leftarrow \psi_p^{[\nu]}(\mathbf{A})\mathbf{v}_1^{(\nu-1)} && (\nu\text{th restart}), \end{aligned}$$

and we can collect the effect of the restarts into a single aggregate polynomial of degree νp :

$$(1.3) \quad \mathbf{v}_1^{(\nu)} \leftarrow \Psi_{\nu p}(\mathbf{A})\mathbf{v}_1,$$

where $\Psi_{\nu p}(\lambda) = \prod_{k=1}^{\nu} \psi_p^{[k]}(\lambda)$.

Evidently, the restart vectors should retain and amplify components of the good invariant subspace while damping and eventually purging components of the bad invariant subspace. One obvious way of encouraging such a trend is to choose the polynomial $\Psi_{\nu p}(\lambda)$ to be as large as possible when evaluated on the good eigenvalues while being as small as possible on the bad eigenvalues. If the bad eigenvalues are situated within a known compact set Ω_{bad} (not containing any good eigenvalues), Chebyshev polynomials associated with Ω_{bad} are often a reasonable choice. When integrated with the Arnoldi algorithm, this results in the Arnoldi–Chebyshev method [36] (cf. [18]).

Saad introduced a different restart strategy that may appear more direct and involves less *a priori* information about bad eigenvalue locations [34]. In this strategy, one computes the eigenvalues of \mathbf{H}_ℓ and sorts the resulting $\ell = k + p$ Ritz values into disjoint sets S_{good} and S_{bad} . The k Ritz values in S_{good} are regarded as approximations to the good eigenvalues of \mathbf{A} , and the restart vector is defined as a linear combination of good Ritz vectors,

$$(1.4) \quad \mathbf{v}_1^+ \leftarrow \sum_{j=1}^k \gamma_j \hat{\mathbf{u}}_j.$$

Saad proposed using either a selected Ritz vector by itself, or forming a linear combination that is weighted to favor Ritz vectors that yield large residuals, so as to provide balanced convergence to the good invariant subspace [34]. Since each Ritz vector $\hat{\mathbf{u}}_j$ is in $\mathcal{K}_\ell(\mathbf{A}, \mathbf{v}_1)$, one may interpret either of these strategies as polynomial restarting, as in (1.2) above.

Sorensen introduced a related strategy, called *exact shifts* [40], that has proved extremely successful in practice. As before, one computes the eigenvalues of \mathbf{H}_ℓ and sorts the resulting $\ell = k + p$ Ritz values into two disjoint sets S_{good} and S_{bad} . But instead of using any explicitly determined linear combination of Ritz vectors associated with the k good Ritz values, the remaining p Ritz values in the set S_{bad} are used to define the restart polynomial $\psi_p(\lambda) = \prod_{j=k+1}^{k+p} (\lambda - \theta_j)$. Morgan discovered a remarkable consequence of this restart strategy: The updated Krylov subspace, $\mathcal{K}_\ell(\mathbf{A}, \mathbf{v}_1^+)$ generated by the new starting vector \mathbf{v}_1^+ in (1.2) using exact shifts satisfies $\mathcal{K}_\ell(\mathbf{A}, \mathbf{v}_1^+) = \text{span}\{\hat{\mathbf{u}}_1, \hat{\mathbf{u}}_2, \dots, \hat{\mathbf{u}}_k, \mathbf{A}\hat{\mathbf{u}}_j, \mathbf{A}^2\hat{\mathbf{u}}_j, \dots, \mathbf{A}^p\hat{\mathbf{u}}_j\}$ for each index $j = 1, 2, \dots, k$ [28]. Thus, Sorensen's exact shifts will provide, in the stage following a restart, a subspace containing every possible Krylov subspace of dimension p that Saad's explicit Ritz vector restart could generate. Furthermore, Sorensen showed how to apply shifts implicitly, regenerating the Krylov subspace $\mathcal{K}_\ell(\mathbf{A}, \mathbf{v}_1^+)$ with only p matrix-vector products in a numerically stable way. Analogous features can be verified for the restarted bi-orthogonal Lanczos method using bad Lanczos values as polynomial roots. Such a strategy has been explored in [16, 9].

In all that follows we assume the Arnoldi or bi-orthogonal Lanczos process has proceeded ℓ steps past the last of ν restarts, each of which (for the sake of simplicity) has the same order p . For the j th restart, with $j = 1, 2, \dots, \nu$, we use a set of shifts $\{\mu_{jk}\}_{k=1}^p$. Define

$$\Psi_{\nu p}(\lambda) = \prod_{j=1}^{\nu} \prod_{k=1}^p (\lambda - \mu_{jk})$$

to be the aggregate restart polynomial after ν restarts. An iteration without restarts will have $p = \nu = 0$ and $\Psi_{\nu p}(\lambda) = 1$.

Let $\mathcal{K}_\tau(\mathbf{A}, \mathbf{v}_1^{(\nu)})$ denote the Krylov subspace of order τ generated by the starting vector $\mathbf{v}_1^{(\nu)}$ that is obtained after ν restarts. The following basic result follows immediately from the observation that $\mathbf{v}_1^{(\nu)} = \Psi_{\nu p}(\mathbf{A})\mathbf{v}_1$.

LEMMA 1.1. *For all $\tau \geq 0$, $\mathcal{K}_\tau(\mathbf{A}, \mathbf{v}_1^{(\nu)}) = \Psi_{\nu p}(\mathbf{A})\mathcal{K}_\tau(\mathbf{A}, \mathbf{v}_1)$.*

2. Reachable Invariant Subspaces. If the good eigenvalues are all simple, then the associated invariant subspace is uniquely determined as the span of all the good eigenvectors. However, if some of these eigenvalues are multiple, there could be a variety of invariant subspaces associated with them. Nonetheless, single-vector Krylov eigenvalue algorithms with polynomial restarts are only capable of revealing one of the many possible invariant subspaces for any given initial vector. Before developing convergence bounds, our first task is to characterize this distinguished invariant subspace precisely.

Let \mathcal{M} be the cyclic subspace generated by the initial starting vector \mathbf{v}_1 ,

$$\mathcal{M} = \text{span}\{\mathbf{v}_1, \mathbf{A}\mathbf{v}_1, \mathbf{A}^2\mathbf{v}_1, \dots\}.$$

\mathcal{M} is evidently an invariant subspace of \mathbf{A} and $s \equiv \dim(\mathcal{M}) \leq n$. Since any invariant subspace of \mathbf{A} that contains \mathbf{v}_1 must also contain $\mathbf{A}^\tau \mathbf{v}_1$ for each integer $\tau > 0$, \mathcal{M} is the *smallest* invariant subspace of \mathbf{A} that contains \mathbf{v}_1 . The first s vectors of the Krylov sequence $\{\mathbf{v}_1, \mathbf{A}\mathbf{v}_1, \mathbf{A}^2\mathbf{v}_1, \dots, \mathbf{A}^{s-1}\mathbf{v}_1\}$ are linearly independent, and so must constitute a basis for \mathcal{M} .

Recall that a linear transformation is *non-derogatory* if each eigenvalue has geometric multiplicity equal to one, i.e., each distinct eigenvalue has precisely one eigenvector associated with it, determined up to scaling.

Define $\mathbf{A}|_{\mathcal{M}}$ to be the restriction of \mathbf{A} to \mathcal{M} .

LEMMA 2.1. $\mathbf{A}|_{\mathcal{M}}$ is non-derogatory and $\mathcal{K}_\tau(\mathbf{A}, \mathbf{v}_1^{(\nu)}) = \mathcal{K}_\tau(\mathbf{A}|_{\mathcal{M}}, \mathbf{v}_1^{(\nu)}) \subset \mathcal{M}$.

Proof. Consider the matrix representation of $\mathbf{A}|_{\mathcal{M}}$ with respect to the basis $\{\mathbf{v}_1, \mathbf{A}\mathbf{v}_1, \mathbf{A}^2\mathbf{v}_1, \dots, \mathbf{A}^{s-1}\mathbf{v}_1\}$. Since \mathcal{M} is s -dimensional,

$$\begin{aligned} \mathbf{A}|_{\mathcal{M}}\mathbf{A}^{s-1}\mathbf{v}_1 &= \mathbf{A}^s\mathbf{v}_1 \\ &= \sum_{j=0}^{s-1} c_j \mathbf{A}^j \mathbf{v}_1 \end{aligned}$$

for some constants c_0, \dots, c_{s-1} . Thus,

$$\mathbf{A}|_{\mathcal{M}} \sim \begin{bmatrix} 0 & \cdots & 0 & c_0 \\ 1 & \ddots & \vdots & \vdots \\ & \ddots & 0 & c_{s-2} \\ & & 1 & c_{s-1} \end{bmatrix} = \widehat{\mathbf{H}}.$$

That is, $\mathbf{A}|_{\mathcal{M}}$ is similar to an irreducible upper Hessenberg matrix, which is necessarily non-derogatory since $\text{rank}(\widehat{\mathbf{H}} - \lambda) \geq s - 1$ for all λ . The second assertion follows immediately from $\mathbf{v}_1^{(\nu)} \in \mathcal{M}$. \square

Define α_j to be the *ascent* (or *index*) of the eigenvalue λ_j , i.e., the minimum positive integer α such that $\text{Ker}(\mathbf{A} - \lambda_j)^\alpha = \text{Ker}(\mathbf{A} - \lambda_j)^{\alpha+1}$. This α_j is the maximum dimension of the n_j different Jordan blocks associated with λ_j and $\text{Ker}(\mathbf{A} - \lambda_j)^{\alpha_j}$ then is the span of all generalized eigenvectors associated with λ_j .

The spectral projection onto each subspace $\text{Ker}(\mathbf{A} - \lambda_j)^{\alpha_j}$ can be constructed in the following coordinate-free manner; see, e.g., [23, §I.5.3]. For each eigenvalue λ_j , $j = 1, \dots, N$, let Γ_j be some positively-oriented Jordan curve in \mathbb{C} containing λ_j in its interior and all other eigenvalues $\lambda_k \neq \lambda_j$ in its exterior. Then the spectral projection is defined as

$$\mathbf{P}_j \equiv \frac{1}{2\pi i} \int_{\Gamma_j} (z - \mathbf{A})^{-1} dz.$$

\mathbf{P}_j is a projection onto the span of all generalized eigenvectors associated with λ_j . In particular, $\mathbf{P}_j\mathbf{v}_1$ will be a generalized eigenvector associated with λ_j and will generate a cyclic subspace $\mathcal{K}_{\alpha_j}(\mathbf{A}, \mathbf{P}_j\mathbf{v}_1) \subseteq \text{Ker}(\mathbf{A} - \lambda_j)^{\alpha_j}$. Let $\widehat{\alpha}_j$ be the minimum index $\widehat{\alpha}$ so that $\mathcal{K}_{\widehat{\alpha}}(\mathbf{A}, \mathbf{P}_j\mathbf{v}_1) = \mathcal{K}_{\widehat{\alpha}+1}(\mathbf{A}, \mathbf{P}_j\mathbf{v}_1)$. This $\widehat{\alpha}_j$ is called the *ascent with respect to \mathbf{v}_1* of the eigenvalue λ_j . Notice that $1 \leq \widehat{\alpha}_j \leq \alpha_j$ and $\mathcal{K}_{\widehat{\alpha}_j}(\mathbf{A}, \mathbf{P}_j\mathbf{v}_1)$ is the *smallest* invariant subspace of \mathbf{A} that contains $\mathbf{P}_j\mathbf{v}_1$. Furthermore, $\mathbf{P}_j\mathbf{v}_1$ is a generalized eigenvector of grade $\widehat{\alpha}_j$ associated with λ_j and $\widehat{\alpha}_j < \alpha_j$ only if \mathbf{v}_1 is deficient in all generalized eigenvectors of maximal grade α_j associated with λ_j .

Define spectral projections \mathbf{P}_{good} and \mathbf{P}_{bad} having ranges that are the maximal invariant subspaces associated with the good and bad eigenvalues, respectively:

$$\mathbf{P}_{\text{good}} = \sum_{j=1}^L \mathbf{P}_j \quad \text{and} \quad \mathbf{P}_{\text{bad}} = \sum_{j=L+1}^N \mathbf{P}_j.$$

Note that $\mathbf{P}_{\text{good}} + \mathbf{P}_{\text{bad}} = \mathbf{I}$.

The following result characterizes \mathcal{M} . The first statement, included for comparison, is well-known; the second is also understood, though we are unaware of its explicit appearance in the literature. Related issues are discussed in [1], [13, Ch. VII].

LEMMA 2.2. $\mathbb{C}^n = \bigoplus_{j=1}^N \text{Ker}(\mathbf{A} - \lambda_j)^{\alpha_j}$ with $\sum_{j=1}^N \alpha_j \leq n$, and

$$\mathcal{M} = \bigoplus_{j=1}^N \mathcal{K}_{\hat{\alpha}_j}(\mathbf{A}, \mathbf{P}_j \mathbf{v}_1) \text{ with } \sum_{j=1}^N \hat{\alpha}_j = \dim \mathcal{M}.$$

Proof. Since $\sum_{j=1}^N \mathbf{P}_j = \mathbf{I}$, any $\mathbf{x} \in \mathbb{C}^n$ can be written as $\mathbf{x} = \mathbf{I}\mathbf{x} = \sum_{j=1}^N \mathbf{P}_j \mathbf{x}$, which shows that $\mathbb{C}^n \subseteq \bigoplus_{j=1}^N \text{Ker}(\mathbf{A} - \lambda_j)^{\alpha_j}$. The reverse inclusion is trivial and yields the first statement.

For the second statement, use $\sum_{j=1}^N \mathbf{P}_j = \mathbf{I}$ to get, for any integer $\tau > 0$,

$$\mathbf{v}_1 = \sum_{j=1}^N \mathbf{P}_j \mathbf{v}_1, \quad \mathbf{A} \mathbf{v}_1 = \sum_{j=1}^N \mathbf{A} \mathbf{P}_j \mathbf{v}_1, \quad \dots, \quad \mathbf{A}^\tau \mathbf{v}_1 = \sum_{j=1}^N \mathbf{A}^\tau \mathbf{P}_j \mathbf{v}_1.$$

Thus, for each integer $\tau > 0$, $\mathcal{K}_\tau(\mathbf{A}, \mathbf{v}_1) \subseteq \bigoplus_{j=1}^N \mathcal{K}_{\hat{\alpha}_j}(\mathbf{A}, \mathbf{P}_j \mathbf{v}_1)$, and, in particular, for τ sufficiently large this yields $\mathcal{M} \subseteq \bigoplus_{j=1}^N \mathcal{K}_{\hat{\alpha}_j}(\mathbf{A}, \mathbf{P}_j \mathbf{v}_1)$.

To show the reverse inclusion, note that for every $j = 1, \dots, N$, one can construct a polynomial p_j such that $p_j(\mathbf{A}) = \mathbf{P}_j$. (This polynomial interpolates zero at eigenvalues $\lambda_k \neq \lambda_j$, one at λ_j , and has $\alpha_j - 1$ zero derivatives at λ_j ; see, e.g., [19, §6.1] for related information.) Thus for any $\mathbf{x} \in \sum_{j=1}^N \mathcal{K}_{\hat{\alpha}_j}(\mathbf{A}, \mathbf{P}_j \mathbf{v}_1)$, one can write

$$\mathbf{x} = \sum_{j=1}^N g_j(\mathbf{A}) \mathbf{P}_j \mathbf{v}_1 = \sum_{j=1}^N g_j(\mathbf{A}) p_j(\mathbf{A}) \mathbf{v}_1 \in \mathcal{M}$$

for polynomials g_j with degree not exceeding $\hat{\alpha}_j - 1$. Thus $\bigoplus_{j=1}^N \mathcal{K}_{\hat{\alpha}_j}(\mathbf{A}, \mathbf{P}_j \mathbf{v}_1) \subseteq \mathcal{M}$, and so $\mathcal{M} = \bigoplus_{j=1}^N \mathcal{K}_{\hat{\alpha}_j}(\mathbf{A}, \mathbf{P}_j \mathbf{v}_1)$. \square

Let $\mathcal{X}_{\text{good}}$ and \mathcal{X}_{bad} be the invariant subspaces of \mathbf{A} associated with the good and bad eigenvalues, respectively. Then define $\mathcal{U}_{\text{good}} \equiv \mathcal{M} \cap \mathcal{X}_{\text{good}}$ and $\mathcal{U}_{\text{bad}} \equiv \mathcal{M} \cap \mathcal{X}_{\text{bad}}$. The following lemma develops a representation for $\mathcal{U}_{\text{good}}$ and \mathcal{U}_{bad} ; it shows that $\mathcal{U}_{\text{good}}$ is the *maximum reachable invariant subspace* associated with the good eigenvalues that can be obtained from a Krylov subspace algorithm started with \mathbf{v}_1 . “Maximum reachable invariant subspace” means that any invariant subspace \mathcal{U} associated with the good eigenvalues and strictly larger than $\mathcal{U}_{\text{good}}$ is *unreachable*: The angle between \mathcal{U} and any computable subspace generated from \mathbf{v}_1 is bounded away from zero independent of ℓ , p , ν , and choice of filter shifts $\{\mu_{jk}\}$.

LEMMA 2.3.

$$\begin{aligned} \mathcal{U}_{\text{good}} &= \bigoplus_{j=1}^L \mathcal{K}_{\hat{\alpha}_j}(\mathbf{A}, \mathbf{P}_j \mathbf{v}_1), & \mathcal{U}_{\text{bad}} &= \bigoplus_{j=L+1}^N \mathcal{K}_{\hat{\alpha}_j}(\mathbf{A}, \mathbf{P}_j \mathbf{v}_1), \\ \dim \mathcal{U}_{\text{good}} &= \sum_{j=1}^L \hat{\alpha}_j \equiv m, & \text{and} & \\ \dim \mathcal{U}_{\text{bad}} &= \sum_{j=L+1}^N \hat{\alpha}_j = s - m. \end{aligned}$$

Furthermore, for any subspace \mathcal{U} of $\mathcal{X}_{\text{good}}$ that properly contains $\mathcal{U}_{\text{good}}$,

$$\mathcal{U}_{\text{good}} \subset \mathcal{U} \subseteq \mathcal{X}_{\text{good}},$$

convergence in gap cannot occur: For all integers $\ell \geq 1$,

$$\delta(\mathcal{U}, \mathcal{K}_\ell(\mathbf{A}, \mathbf{v}_1^{(\nu)})) \geq \frac{1}{\|\mathbf{P}_{\text{good}}\|} > 0.$$

Proof. Since $\mathcal{K}_{\hat{\alpha}_j}(\mathbf{A}, \mathbf{P}_j \mathbf{v}_1) \subseteq \text{Ker}(\mathbf{A} - \lambda_j)^{\alpha_j}$, Lemma 2.2 leads to $\mathcal{M} \cap \mathcal{X}_{\text{good}} = \bigoplus_{j=1}^L \mathcal{K}_{\hat{\alpha}_j}(\mathbf{A}, \mathbf{P}_j \mathbf{v}_1)$. Furthermore, $\dim \mathcal{K}_{\hat{\alpha}_j}(\mathbf{A}, \mathbf{P}_j \mathbf{v}_1) = \hat{\alpha}_j$ implies that $\dim \mathcal{U}_{\text{good}} = m$ as defined above. The analogous results for \mathcal{U}_{bad} follow similarly.

Note that $\mathcal{X}_{\text{bad}} = \bigoplus_{j=L+1}^N \text{Ker}(\mathbf{A} - \lambda_j)^{\alpha_j}$ so, for all $\ell \geq 0$,

$$\mathcal{K}_\ell(\mathbf{A}, \mathbf{v}_1^{(\nu)}) \subseteq \mathcal{M} \subseteq \mathcal{U}_{\text{good}} \oplus \mathcal{X}_{\text{bad}}.$$

Thus any $\mathbf{v} \in \mathcal{K}_\ell(\mathbf{A}, \mathbf{v}_1^{(\nu)})$ can be decomposed as $\mathbf{v} = \mathbf{w}_1 + \mathbf{w}_2$ for some $\mathbf{w}_1 \in \mathcal{U}_{\text{good}}$ and $\mathbf{w}_2 \in \mathcal{X}_{\text{bad}}$. When $\mathcal{U}_{\text{good}}$ is a proper subspace of \mathcal{U} , there exists an $\hat{\mathbf{x}} \in \mathcal{U}$ so that $\hat{\mathbf{x}} \perp \mathcal{U}_{\text{good}}$ and $\|\hat{\mathbf{x}}\| = 1$. Note that $\|\hat{\mathbf{x}} - \mathbf{w}_1\| \geq \|\hat{\mathbf{x}}\| = 1$. Now,

$$\begin{aligned} \min_{\mathbf{v} \in \mathcal{K}_\ell(\mathbf{A}, \mathbf{v}_1^{(\nu)})} \|\mathbf{v} - \hat{\mathbf{x}}\| &\geq \min_{\substack{\mathbf{w}_1 \in \mathcal{U}_{\text{good}} \\ \mathbf{w}_2 \in \mathcal{X}_{\text{bad}}}} \|\mathbf{w}_1 + \mathbf{w}_2 - \hat{\mathbf{x}}\| \\ &\geq \min_{\substack{\mathbf{w}_1 \in \mathcal{U}_{\text{good}} \\ \mathbf{w}_2 \in \mathcal{X}_{\text{bad}}}} \frac{\|\mathbf{w}_2 - (\hat{\mathbf{x}} - \mathbf{w}_1)\|}{\|\hat{\mathbf{x}} - \mathbf{w}_1\|} \\ &\geq \min_{\substack{\mathbf{y} \in \mathcal{X}_{\text{good}} \\ \mathbf{w}_2 \in \mathcal{X}_{\text{bad}}}} \frac{\|\mathbf{w}_2 - \mathbf{y}\|}{\|\mathbf{y}\|} \\ &\geq \left(\max_{\substack{\mathbf{y} \in \mathcal{X}_{\text{good}} \\ \mathbf{w}_2 \in \mathcal{X}_{\text{bad}}}} \frac{\|\mathbf{P}_{\text{good}}(\mathbf{w}_2 - \mathbf{y})\|}{\|\mathbf{w}_2 - \mathbf{y}\|} \right)^{-1} = \frac{1}{\|\mathbf{P}_{\text{good}}\|}. \end{aligned}$$

Thus,

$$\begin{aligned} \delta(\mathcal{U}, \mathcal{K}_\ell(\mathbf{A}, \mathbf{v}_1^{(\nu)})) &= \max_{\mathbf{x} \in \mathcal{U}} \min_{\mathbf{v} \in \mathcal{K}_\ell(\mathbf{A}, \mathbf{v}_1^{(\nu)})} \frac{\|\mathbf{v} - \mathbf{x}\|}{\|\mathbf{x}\|} \\ &\geq \min_{\mathbf{v} \in \mathcal{K}_\ell(\mathbf{A}, \mathbf{v}_1^{(\nu)})} \|\mathbf{v} - \hat{\mathbf{x}}\| \geq \frac{1}{\|\mathbf{P}_{\text{good}}\|}. \quad \square \end{aligned}$$

This means that we have no hope of capturing any invariant subspace that contains a (generalized) eigenspace associated with multiple Jordan blocks — at least when using a single vector iteration in exact arithmetic. On the other hand, convergence can occur to the good invariant subspace $\mathcal{U}_{\text{good}}$, with a rate that depends on properties of \mathbf{A} , \mathbf{v}_1 , and the choice of filter shifts $\{\mu_{jk}\}$, as we shall see.

Almost every vector in an invariant subspace is a generalized eigenvector of maximal grade and so almost every starting vector will capture maximally defective Jordan blocks. While easily acknowledged, this fact can have perplexing consequences for the casual Arnoldi or bi-orthogonal Lanczos user, since eigenvectors of other Jordan blocks may be unexpectedly “washed out.”

Suppose \mathbf{A} is defined as

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

\mathbf{A} is in Jordan canonical form and has eigenvectors \mathbf{e}_2 and \mathbf{e}_5 associated with the eigenvalue 1. (\mathbf{e}_j denotes the j th column of the 5×5 identity matrix.) The eigenvalue 1 also has generalized eigenvectors of grade 1 given by \mathbf{e}_1 and \mathbf{e}_4 associated with the 2×2 and 3×3 Jordan blocks, respectively, and a generalized eigenvector of grade 2 given by \mathbf{e}_3 associated with the 3×3 Jordan block.

Let $\beta \in \mathbb{C}$ be arbitrary and note that the vector $\mathbf{v}_1 = [1 \ \beta \ 1 \ 1 \ 1]^T$ generates a cyclic subspace spanned by the first three vectors in the Krylov sequence,

$$\mathbf{v}_1 = \begin{bmatrix} 1 \\ \beta \\ 1 \\ 1 \\ 1 \end{bmatrix}, \quad \mathbf{v}_2 = \mathbf{A}\mathbf{v}_1 = \begin{bmatrix} 1 \\ 1 + \beta \\ 1 \\ 2 \\ 2 \end{bmatrix}, \quad \text{and} \quad \mathbf{v}_3 = \mathbf{A}^2\mathbf{v}_1 = \begin{bmatrix} 1 \\ 2 + \beta \\ 1 \\ 3 \\ 4 \end{bmatrix}.$$

By choosing $|\beta|$ to be large, the starting vector \mathbf{v}_1 can be made to have an arbitrarily large component in the direction of \mathbf{e}_2 , the eigenvector associated with the 2×2 Jordan block.

Defining $\mathbf{M} = [\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3]$ and $\hat{\mathbf{H}} = \begin{bmatrix} 0 & 0 & 1 \\ 1 & 0 & -3 \\ 0 & 1 & 3 \end{bmatrix}$, a simple calculation reveals

$$\mathbf{A}\mathbf{M} = \mathbf{M}\hat{\mathbf{H}}.$$

The Jordan form of $\hat{\mathbf{H}}$ is easy to calculate:

$$(2.1) \quad \mathbf{R}^{-1}\hat{\mathbf{H}}\mathbf{R} = \begin{bmatrix} 1 & 0 & 0 \\ 1 & 1 & 0 \\ 0 & 1 & 1 \end{bmatrix}, \quad \text{where} \quad \mathbf{R} = \begin{bmatrix} 1 & -1 & 1 \\ 0 & 1 & -2 \\ 0 & 0 & 1 \end{bmatrix}.$$

The cyclic subspace generated by the single vector \mathbf{v}_1 has captured a *three*-dimensional invariant subspace, associated with the maximally defective 3×3 Jordan block. But this subspace is not the expected $\text{span}\{\mathbf{e}_3, \mathbf{e}_4, \mathbf{e}_5\}$. Using the change of basis defined by \mathbf{R} in (2.1), one may calculate $\mathbf{A}(\mathbf{M}\mathbf{R}) = (\mathbf{M}\mathbf{R})(\mathbf{R}^{-1}\hat{\mathbf{H}}\mathbf{R})$, which is

$$\begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 & 0 \\ 0 & 0 & 0 & 1 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 \\ \beta & 1 & 0 \\ 1 & 0 & 0 \\ 1 & 1 & 0 \\ 1 & 1 & 1 \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 \\ \beta & 1 & 0 \\ 1 & 0 & 0 \\ 1 & 1 & 0 \\ 1 & 1 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 \\ 1 & 1 & 0 \\ 0 & 1 & 1 \end{bmatrix}.$$

Note that \mathbf{e}_5 alone is revealed as the eigenvector associated with the eigenvalue 1; \mathbf{e}_2 has been washed out in spite of \mathbf{v}_1 having an arbitrarily large component in that

direction. Indeed the eigenvector \mathbf{e}_2 (and so any subspace containing it) is unreachable from *any* starting vector \mathbf{v}_1 for which $\mathbf{e}_3^* \mathbf{v}_1 \neq 0$. In this example, \mathbf{v}_1 itself emerges as a generalized eigenvector of grade 2. Note that *every* vector \mathbf{v} in \mathbb{C}^5 with $\mathbf{e}_3^* \mathbf{v} \neq 0$ is a generalized eigenvector of grade 2 associated with the eigenvalue 1.

We close this section with a computational example that both confirms the gap stagnation lower bound for derogatory matrices given in Lemma 2.3 and illustrates other convergence properties explored in future sections. Consider two matrices \mathbf{A}_1 and \mathbf{A}_2 , each of dimension $n = 150$ with eigenvalues spaced uniformly in the interval $[0, 1]$. In both cases, all the eigenvalues are simple except for the single good eigenvalue $\lambda = 1$, which has algebraic multiplicity five. In the first case, the geometric multiplicity also equals five, so the matrix is diagonalizable but derogatory. In the second case, there is only one eigenvector associated with $\lambda = 1$, so it is defective but not derogatory. Both matrices are constructed so that $\|\mathbf{P}_{\text{good}}\| \approx 10^4$. Figure 2.1 illustrates the gap convergence for the Krylov subspace to the invariant subspace $\mathcal{X}_{\text{good}}$ associated with $\lambda = 1$. The starting vector \mathbf{v}_1 has $1/\sqrt{n}$ in each component; no restarting is used here. Convergence cannot begin until the fifth iteration, when the Krylov subspace dimension matches the dimension of $\mathcal{X}_{\text{good}}$. This initial period of stagnation is followed by a sublinear phase of convergence leading to a second stagnation period. This is the end of the story for the derogatory case, but for the nonderogatory case, the second stagnation period is transient and the convergence rate eventually settles towards a nearly linear rate. In fact, this rate improves slightly over the final iterations shown here, yielding so-called “superlinear” convergence, the subject of §4.3. These convergence phases resemble those observed for the GMRES iteration, as described by Nevanlinna [29].

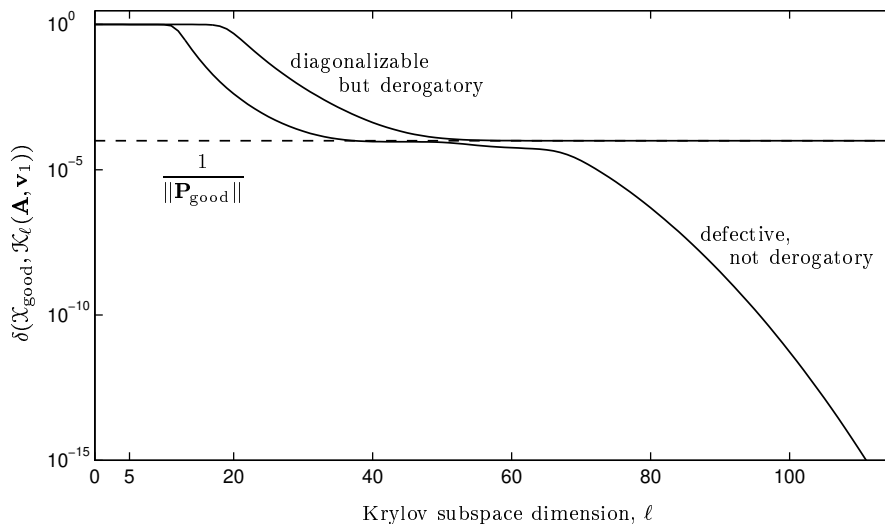


FIG. 2.1. The Krylov subspace can never capture $\mathcal{X}_{\text{good}}$ when this subspace is associated with a derogatory eigenvalue; convergence is possible, however, when the associated eigenvalues are defective but not derogatory, as described by Lemma 2.3.

3. Basic Estimates. Since all reachable subspaces are contained in \mathcal{M} and $\mathbf{A}|_{\mathcal{M}}$ is non-derogatory, henceforth we assume without loss of generality that \mathbf{A} itself is non-derogatory so that $n = \dim \mathcal{M}$, and \mathbf{v}_1 is not deficient in any generalized eigenvector of maximal grade. To summarize the current situation: \mathbf{A} is an $n \times n$ matrix with $N \leq n$ distinct eigenvalues, $\{\lambda_j\}_{j=1}^N$, each having geometric multiplicity 1 and algebraic multiplicity m_j , so that $\sum_{j=1}^N m_j = n$. We seek L ($1 \leq L < N$) of these eigenvalues $\{\lambda_1, \lambda_2, \dots, \lambda_L\}$ (the “good” eigenvalues) together with the corresponding (maximal) invariant subspace $\mathcal{U}_{\text{good}}$ of dimension $m = \sum_{j=1}^L m_j$, which is now the net algebraic multiplicity of good eigenvalues since \mathbf{A} is non-derogatory.

We begin by establishing two lemmas that are used to develop a bound for the gap in terms of a polynomial approximation problem in the subsequent theorems.

LEMMA 3.1. *Given $\mathcal{U}, \mathcal{V} \subseteq \mathbb{C}^n$, suppose $\hat{\mathbf{u}} \in \mathcal{U}$ ($\|\hat{\mathbf{u}}\| = 1$) and $\hat{\mathbf{v}} \in \mathcal{V}$ satisfy*

$$\delta(\mathcal{U}, \mathcal{V}) = \max_{\mathbf{u} \in \mathcal{U}} \min_{\mathbf{v} \in \mathcal{V}} \frac{\|\mathbf{u} - \mathbf{v}\|}{\|\mathbf{u}\|} = \|\hat{\mathbf{u}} - \hat{\mathbf{v}}\|.$$

Then $\hat{\mathbf{u}} - \hat{\mathbf{v}} \perp \mathcal{V}$ and $\hat{\mathbf{u}} - \hat{\mathbf{v}} - \delta(\mathcal{U}, \mathcal{V})^2 \hat{\mathbf{u}} \perp \mathcal{U}$.

Proof. The first assertion is a fundamental property of least squares approximation. To show the second, consider an arbitrary unit vector $\mathbf{u} \in \mathcal{U}$ and take $\varepsilon > 0$. Letting $\mathbf{\Pi}_{\mathcal{V}}$ denote the orthogonal projection onto \mathcal{V} , the optimality of $\hat{\mathbf{u}}$ and $\hat{\mathbf{v}}$ implies

$$\|\hat{\mathbf{u}} - \hat{\mathbf{v}}\|^2 \geq \frac{\|(\mathbf{I} - \mathbf{\Pi}_{\mathcal{V}})(\hat{\mathbf{u}} + \varepsilon \mathbf{u})\|^2}{\|\hat{\mathbf{u}} + \varepsilon \mathbf{u}\|^2}.$$

Expanding this inequality, noting $\hat{\mathbf{v}} = \mathbf{\Pi}_{\mathcal{V}} \hat{\mathbf{u}}$, and using the first assertion gives

$$\delta(\mathcal{U}, \mathcal{V})^2 (1 + 2\varepsilon \operatorname{Re}(\hat{\mathbf{u}}^* \mathbf{u}) + \varepsilon^2) \geq \delta(\mathcal{U}, \mathcal{V})^2 + 2\varepsilon \operatorname{Re}((\hat{\mathbf{u}} - \hat{\mathbf{v}})^* \mathbf{u}) + \varepsilon^2 \|(\mathbf{I} - \mathbf{\Pi}_{\mathcal{V}}) \mathbf{u}\|^2.$$

Collecting terms quadratic in ε on the left hand side,

$$\varepsilon^2 (\delta(\mathcal{U}, \mathcal{V})^2 - \|(\mathbf{I} - \mathbf{\Pi}_{\mathcal{V}}) \mathbf{u}\|^2) \geq 2\varepsilon \operatorname{Re}((\hat{\mathbf{u}} - \hat{\mathbf{v}} - \delta(\mathcal{U}, \mathcal{V})^2 \hat{\mathbf{u}})^* \mathbf{u}).$$

Note that the left hand side must be non-negative. Repeating the above argument with \mathbf{u} multiplied by a complex scalar of unit modulus, we can replace the right hand side with $2\varepsilon |(\hat{\mathbf{u}} - \hat{\mathbf{v}} - \delta(\mathcal{U}, \mathcal{V})^2 \hat{\mathbf{u}})^* \mathbf{u}|$. Thus for any unit vector $\hat{\mathbf{u}} \in \mathcal{U}$,

$$\varepsilon (\delta(\mathcal{U}, \mathcal{V})^2 - \|(\mathbf{I} - \mathbf{\Pi}_{\mathcal{V}}) \mathbf{u}\|^2) \geq 2 |(\hat{\mathbf{u}} - \hat{\mathbf{v}} - \delta(\mathcal{U}, \mathcal{V})^2 \hat{\mathbf{u}})^* \mathbf{u}| \geq 0.$$

Taking $\varepsilon \rightarrow 0$, we conclude that $\hat{\mathbf{u}} - \hat{\mathbf{v}} - \delta(\mathcal{U}, \mathcal{V})^2 \hat{\mathbf{u}}$ is orthogonal to every $\mathbf{u} \in \mathcal{U}$. \square

As the gap between subspaces closes ($\delta(\mathcal{U}, \mathcal{V}) \rightarrow 0$), notice that $\hat{\mathbf{u}} - \hat{\mathbf{v}}$ becomes “almost” orthogonal to \mathcal{U} in the sense that the projection of $\hat{\mathbf{u}} - \hat{\mathbf{v}}$ onto \mathcal{U} has norm $\delta(\mathcal{U}, \mathcal{V})^2$.

LEMMA 3.2. *Let \mathcal{P}_{m-1} denote the space of polynomials of degree $m-1$ or less. The mapping $\iota: \mathcal{P}_{m-1} \rightarrow \mathcal{U}_{\text{good}}$ defined by*

$$(3.1) \quad \iota(\psi) = \psi(\mathbf{A}) \mathbf{P}_{\text{good}} \mathbf{v}_1$$

is an isomorphism between \mathcal{P}_{m-1} and $\mathcal{U}_{\text{good}}$. Furthermore, there exist positive constants c_1 and c_2 so that

$$(3.2) \quad c_1 \|\psi\|_{\mathcal{P}_{m-1}} \leq \|\psi(\mathbf{A}) \mathbf{P}_{\text{good}} \mathbf{v}_1\| \leq c_2 \|\psi\|_{\mathcal{P}_{m-1}},$$

uniformly for all $\psi \in \mathcal{P}_{m-1}$ for any fixed norm $\|\cdot\|_{\mathcal{P}_{m-1}}$ defined on the space \mathcal{P}_{m-1} .

Proof. ι is clearly linear. To see that ι maps \mathcal{P}_{m-1} onto $\mathcal{U}_{\text{good}}$, observe that for any given $\mathbf{y} \in \mathcal{U}_{\text{good}}$, there exist polynomials $\{g_j(\lambda)\}_{j=1}^L$ with $\deg(g_j) \leq m_j - 1$ such that

$$\mathbf{y} = \sum_{j=1}^L g_j(\mathbf{A}) \mathbf{P}_j \mathbf{v}_1.$$

The L polynomials $\{g_j\}_{j=1}^L$ provide L separate ‘‘slices’’ of a single polynomial that can be recovered by (generalized) Hermite interpolation. Let ψ be a polynomial interpolant that interpolates g_j and its derivatives at λ_j :

$$\psi^{(k)}(\lambda_j) = g_j^{(k)}(\lambda_j)$$

for $k = 0, 1, \dots, m_j - 1$ and $j = 1, 2, \dots, L$. Theorem VIII.3.16 of [11] leads us first to observe that $\psi(\mathbf{A}) \mathbf{P}_j = g_j(\mathbf{A}) \mathbf{P}_j$ for each $j = 1, \dots, L$. Then since $\deg(\psi) \leq \sum_{j=1}^L m_j - 1 = m - 1$, we have from (3.1) that

$$\mathbf{y} = \sum_{j=1}^L \psi(\mathbf{A}) \mathbf{P}_j \mathbf{v}_1 = \psi(\mathbf{A}) \mathbf{P}_{\text{good}} \mathbf{v}_1 = \iota(\psi).$$

Since $\dim(\mathcal{P}_{m-1}) = \dim(\mathcal{U}_{\text{good}})$, $\text{nullity}(\iota) = 0$ and ι is bijective from \mathcal{P}_{m-1} to $\mathcal{U}_{\text{good}}$. The last statement is an immediate consequence of the fact that linear bijections are bounded linear transformations with bounded inverses. \square

THEOREM 3.3. *Suppose that \mathbf{A} and \mathbf{v}_1 satisfy the assumptions of this section, and that none of the filter shifts $\{\mu_{jk}\}$ coincides with any of the good eigenvalues $\{\lambda_j\}_{j=1}^L$. For all indices $\ell \geq m$, the gap between the good invariant subspace, $\mathcal{U}_{\text{good}}$, and the Krylov subspace of order ℓ , $\mathcal{K}_\ell(\mathbf{A}, \mathbf{v}_1^{(\nu)})$, generated from the ν -fold restarted vector, $\mathbf{v}_1^{(\nu)}$ satisfies*

$$\delta(\mathcal{U}_{\text{good}}, \mathcal{K}_\ell(\mathbf{A}, \mathbf{v}_1^{(\nu)})) \leq C_0 \max_{\psi \in \mathcal{P}_{m-1}} \min_{\phi \in \mathcal{P}_{\ell-m}} \frac{\|\phi(\mathbf{A}) \psi(\mathbf{A}) \Psi_{\nu p}(\mathbf{A}) \mathbf{P}_{\text{bad}} \mathbf{v}_1\|}{\|\phi(\mathbf{A}) \psi(\mathbf{A}) \Psi_{\nu p}(\mathbf{A}) \mathbf{P}_{\text{good}} \mathbf{v}_1\|},$$

where $C_0 \equiv 1$ if $\mathcal{U}_{\text{good}} \perp \mathcal{U}_{\text{bad}}$ and $C_0 \equiv \sqrt{2}$ otherwise.

Proof. First, suppose $\mathcal{U}_{\text{good}} \perp \mathcal{U}_{\text{bad}}$. This implies that \mathbf{P}_{good} and \mathbf{P}_{bad} are orthogonal projections, $\mathcal{U}_{\text{good}}$ is an invariant subspace for both $\Psi_{\nu p}(\mathbf{A})$ and $[\Psi_{\nu p}(\mathbf{A})]^*$, and, as we will see, that $\delta(\mathcal{U}_{\text{good}}, \mathcal{K}_\ell(\mathbf{A}, \mathbf{v}_1^{(\nu)})) < 1$. Indeed, suppose instead that $\delta(\mathcal{U}_{\text{good}}, \mathcal{K}_\ell(\mathbf{A}, \mathbf{v}_1^{(\nu)})) = 1$. Then there is a vector $\hat{\mathbf{u}} \in \mathcal{U}_{\text{good}}$ with $\|\hat{\mathbf{u}}\| = 1$ such that $\hat{\mathbf{u}} \perp \mathcal{K}_\ell(\mathbf{A}, \mathbf{v}_1^{(\nu)})$. Define $\hat{\mathbf{y}} \equiv [\Psi_{\nu p}(\mathbf{A})]^* \hat{\mathbf{u}} \in \mathcal{U}_{\text{good}}$, and note that by Lemma 3.2, there exists a polynomial $\hat{\psi} \in \mathcal{P}_{m-1}$ such that $\hat{\mathbf{y}} = \hat{\psi}(\mathbf{A}) \mathbf{P}_{\text{good}} \mathbf{v}_1$. Now, for each $j = 1, 2, \dots, \ell$ we have

$$\begin{aligned} 0 &= \langle \hat{\mathbf{u}}, \mathbf{A}^{j-1} \mathbf{v}_1^{(\nu)} \rangle = \langle \hat{\mathbf{u}}, \mathbf{A}^{j-1} \Psi_{\nu p}(\mathbf{A}) \mathbf{v}_1 \rangle \\ &= \langle \hat{\mathbf{y}}, \mathbf{A}^{j-1} \mathbf{P}_{\text{good}} \mathbf{v}_1 \rangle \\ &= \langle \hat{\psi}(\mathbf{A}) \mathbf{P}_{\text{good}} \mathbf{v}_1, \mathbf{A}^{j-1} \mathbf{P}_{\text{good}} \mathbf{v}_1 \rangle. \end{aligned}$$

Since $\ell \geq m$, this implies first that $\|\widehat{\psi}(\mathbf{A})\mathbf{P}_{\text{good}}\mathbf{v}_1\| = 0$ and then $\widehat{\mathbf{u}} = \mathbf{0}$. (Recall that $[\Psi_{\nu p}(\mathbf{A})]^*$ is bijective on $\mathcal{U}_{\text{good}}$ since $\Psi_{\nu p}$ has no roots in common with good eigenvalues.) But $\widehat{\mathbf{u}}$ was given to be a unit vector, so it must be that $\delta(\mathcal{U}_{\text{good}}, \mathcal{K}_\ell(\mathbf{A}, \mathbf{v}_1^{(\nu)})) < 1$.

Now, there are optimal vectors $\widehat{\mathbf{v}} \in \mathcal{K}_\ell(\mathbf{A}, \mathbf{v}_1^{(\nu)})$ and $\widehat{\mathbf{x}} \in \mathcal{U}_{\text{good}}$ with $\|\widehat{\mathbf{x}}\| = 1$ that satisfy

$$(3.3) \quad \delta(\mathcal{U}_{\text{good}}, \mathcal{K}_\ell(\mathbf{A}, \mathbf{v}_1^{(\nu)})) = \max_{\mathbf{x} \in \mathcal{U}_{\text{good}}} \min_{\mathbf{v} \in \mathcal{K}_\ell(\mathbf{A}, \mathbf{v}_1^{(\nu)})} \frac{\|\mathbf{v} - \mathbf{x}\|}{\|\mathbf{x}\|} = \|\widehat{\mathbf{v}} - \widehat{\mathbf{x}}\|.$$

Since $\delta(\mathcal{U}_{\text{good}}, \mathcal{K}_\ell(\mathbf{A}, \mathbf{v}_1^{(\nu)})) < 1$, it must be that $\widehat{\mathbf{v}} \neq \mathbf{0}$. Furthermore, optimality for $\widehat{\mathbf{v}}$ means $\widehat{\mathbf{v}} - \widehat{\mathbf{x}} \perp \mathcal{K}_\ell(\mathbf{A}, \mathbf{v}_1^{(\nu)})$ (viz., Lemma 3.1) and in particular, $\widehat{\mathbf{v}}^*(\widehat{\mathbf{v}} - \widehat{\mathbf{x}}) = 0$. So, $\widehat{\mathbf{v}} \neq \mathbf{0}$ implies $\widehat{\mathbf{v}} \notin \mathcal{U}_{\text{bad}}$. There is a polynomial $\pi_{\ell-1} \in \mathcal{P}_{\ell-1}$ such that

$$\widehat{\mathbf{v}} = \pi_{\ell-1}(\mathbf{A})\mathbf{v}_1^{(\nu)} = \pi_{\ell-1}(\mathbf{A})\Psi_{\nu p}(\mathbf{A})\mathbf{v}_1.$$

Define $\mathcal{Q} = \mathcal{U}_{\text{good}} \cap \text{Ker}(\pi_{\ell-1}(\mathbf{A}))$ and let \widehat{q} be the minimum (monic) annihilating polynomial for \mathcal{Q} .¹ Evidently, $\pi_{\ell-1}$ must contain \widehat{q} as a factor.

Since $\widehat{\mathbf{v}} \notin \mathcal{U}_{\text{bad}}$, $\pi_{\ell-1}$ cannot be an annihilating polynomial for $\mathcal{U}_{\text{good}}$, so $\mathcal{Q} \neq \mathcal{U}_{\text{good}}$ and $\deg(\widehat{q}) \leq m - 1$. One may factor $\pi_{\ell-1}$ as the product of a polynomial, ϕ , of degree $\ell - m$ and a polynomial, q , of degree $m - 1$ containing \widehat{q} as a factor,

$$\pi_{\ell-1}(\lambda) = \phi(\lambda)q(\lambda).$$

Observing that $\mathcal{U}_{\text{good}}$ is invariant for both $\phi(\mathbf{A})$ and $\phi(\mathbf{A})^*$, we may decompose $\widehat{\mathbf{x}}$ as $\widehat{\mathbf{x}} = \phi(\mathbf{A})\widehat{\mathbf{y}} + \mathbf{n}$ for some $\widehat{\mathbf{y}} \in \mathcal{U}_{\text{good}}$ and some $\mathbf{n} \in \text{Ker}(\phi(\mathbf{A})^*) \cap \mathcal{U}_{\text{good}}$. Notice that $\widehat{\mathbf{v}}^*\phi(\mathbf{A})\widehat{\mathbf{y}} = \widehat{\mathbf{v}}^*\widehat{\mathbf{x}} = \widehat{\mathbf{v}}^*\widehat{\mathbf{v}} > 0$, so $\phi(\mathbf{A})\widehat{\mathbf{y}} \neq \mathbf{0}$. However, we'll see that it must happen that $\mathbf{n} = \mathbf{0}$. Indeed, Lemma 3.1 shows that if $\mathbf{z} \in \mathcal{U}_{\text{good}}$ is orthogonal to $\widehat{\mathbf{x}}$, $\widehat{\mathbf{x}}^*\mathbf{z} = 0$, then $\widehat{\mathbf{v}}^*\mathbf{z} = 0$ as well. In particular, for $\mathbf{z} = \|\mathbf{n}\|^2\phi(\mathbf{A})\widehat{\mathbf{y}} - \|\phi(\mathbf{A})\widehat{\mathbf{y}}\|^2\mathbf{n}$ we have $\widehat{\mathbf{x}}^*\mathbf{z} = 0$. Since $\text{Ker}(\phi(\mathbf{A})^*) = \text{Ran}(\phi(\mathbf{A}))^\perp$ implies $\widehat{\mathbf{v}}^*\mathbf{n} = 0$, we have

$$0 = \widehat{\mathbf{v}}^*\mathbf{z} = \|\mathbf{n}\|^2\widehat{\mathbf{v}}^*\phi(\mathbf{A})\widehat{\mathbf{y}}.$$

We have already seen that $\widehat{\mathbf{v}}^*\phi(\mathbf{A})\widehat{\mathbf{y}} > 0$, and so $\mathbf{n} = \mathbf{0}$. Thus we can safely exclude from the maximization in (3.3) all $\mathbf{x} \in \mathcal{U}_{\text{good}}$ except for those vectors having the special form $\mathbf{x} = \phi(\mathbf{A})\mathbf{y}$ for $\mathbf{y} \in \mathcal{U}_{\text{good}}$ and ϕ as defined above.

We can now begin our process of bounding the gap. Note that

$$(3.4) \quad \begin{aligned} \delta(\mathcal{U}_{\text{good}}, \mathcal{K}_\ell(\mathbf{A}, \mathbf{v}_1^{(\nu)})) &= \max_{\mathbf{x} \in \mathcal{U}_{\text{good}}} \min_{\mathbf{v} \in \mathcal{K}_\ell(\mathbf{A}, \mathbf{v}_1^{(\nu)})} \frac{\|\mathbf{v} - \mathbf{x}\|}{\|\mathbf{x}\|} \\ &= \max_{\mathbf{x} \in \mathcal{U}_{\text{good}}} \min_{\phi \in \mathcal{P}_{\ell-m}} \min_{q \in \mathcal{P}_{m-1}} \frac{\|\Psi_{\nu p}(\mathbf{A})\phi(\mathbf{A})q(\mathbf{A})\mathbf{v}_1 - \mathbf{x}\|}{\|\mathbf{x}\|} \\ &= \max_{\mathbf{y} \in \mathcal{U}_{\text{good}}} \min_{\phi \in \mathcal{P}_{\ell-m}} \min_{q \in \mathcal{P}_{m-1}} \frac{\|\Psi_{\nu p}(\mathbf{A})\phi(\mathbf{A})[q(\mathbf{A})\mathbf{v}_1 - \mathbf{y}]\|}{\|\Psi_{\nu p}(\mathbf{A})\phi(\mathbf{A})\mathbf{y}\|}, \end{aligned}$$

where we are able to justify the substitution $\mathbf{x} = \Psi_{\nu p}(\mathbf{A})\phi(\mathbf{A})\mathbf{y}$ since $\Psi_{\nu p}(\mathbf{A})$ is an invertible map of $\mathcal{U}_{\text{good}}$ to itself.

¹That is, \widehat{q} is the minimum degree monic polynomial such that $\widehat{q}(\mathbf{A})\mathbf{r} = \mathbf{0}$ for all $\mathbf{r} \in \mathcal{Q}$.

Now by Lemma 3.2, $\mathbf{y} \in \mathcal{U}_{\text{good}}$ can be represented as $\mathbf{y} = \psi(\mathbf{A})\mathbf{P}_{\text{good}}\mathbf{v}_1$ for some $\psi \in \mathcal{P}_{m-1}$. Since $\mathbf{I} = \mathbf{P}_{\text{bad}} + \mathbf{P}_{\text{good}}$, one finds

$$\psi(\mathbf{A})\mathbf{v}_1 - \mathbf{y} = \psi(\mathbf{A})\mathbf{P}_{\text{bad}}\mathbf{v}_1.$$

Continuing with (3.4), assign $q \equiv \psi \in \mathcal{P}_{m-1}$. Then

$$\begin{aligned} \delta(\mathcal{U}_{\text{good}}, \mathcal{K}_\ell(\mathbf{A}, \mathbf{v}_1^{(\nu)})) &\leq \max_{\substack{\mathbf{y} \in \mathcal{U}_{\text{good}} \\ (\mathbf{y} = \psi(\mathbf{A})\mathbf{P}_{\text{good}}\mathbf{v}_1)}} \min_{\phi \in \mathcal{P}_{\ell-m}} \frac{\|\Psi_{\nu p}(\mathbf{A})\phi(\mathbf{A})[\psi(\mathbf{A})\mathbf{v}_1 - \mathbf{y}]\|}{\|\Psi_{\nu p}(\mathbf{A})\phi(\mathbf{A})\mathbf{y}\|} \\ &= \max_{\psi \in \mathcal{P}_{m-1}} \min_{\phi \in \mathcal{P}_{\ell-m}} \frac{\|\Psi_{\nu p}(\mathbf{A})\phi(\mathbf{A})\psi(\mathbf{A})\mathbf{P}_{\text{bad}}\mathbf{v}_1\|}{\|\Psi_{\nu p}(\mathbf{A})\phi(\mathbf{A})\psi(\mathbf{A})\mathbf{P}_{\text{good}}\mathbf{v}_1\|}, \end{aligned}$$

as required, concluding the proof when $\mathcal{U}_{\text{good}} \perp \mathcal{U}_{\text{bad}}$.

Now, if the two subspaces $\mathcal{U}_{\text{good}}$ and \mathcal{U}_{bad} are *not* orthogonal, introduce a new inner product on \mathbb{C}^n with respect to which they are orthogonal. For any $\mathbf{u}, \mathbf{v} \in \mathbb{C}^n$, define

$$\langle \mathbf{u}, \mathbf{v} \rangle_* \equiv \langle \mathbf{P}_{\text{good}}\mathbf{u}, \mathbf{P}_{\text{good}}\mathbf{v} \rangle + \langle \mathbf{P}_{\text{bad}}\mathbf{u}, \mathbf{P}_{\text{bad}}\mathbf{v} \rangle,$$

and define the gap with respect to the new norm $\|\cdot\|_* = \sqrt{\langle \cdot, \cdot \rangle_*}$ to be

$$\delta_*(\mathcal{W}, \mathcal{V}) = \sup_{\mathbf{x} \in \mathcal{W}} \inf_{\mathbf{y} \in \mathcal{V}} \frac{\|\mathbf{y} - \mathbf{x}\|_*}{\|\mathbf{x}\|_*}.$$

Notice that for any vector $\mathbf{w} \in \mathbb{C}^n$,

$$\begin{aligned} \|\mathbf{w}\|^2 &= \|\mathbf{P}_{\text{good}}\mathbf{w} + \mathbf{P}_{\text{bad}}\mathbf{w}\|^2 \leq 2(\|\mathbf{P}_{\text{good}}\mathbf{w}\|^2 + \|\mathbf{P}_{\text{bad}}\mathbf{w}\|^2) = 2\|\mathbf{w}\|_*^2, \\ \|\mathbf{P}_{\text{good}}\mathbf{w}\|_* &= \|\mathbf{P}_{\text{good}}\mathbf{w}\|, \quad \text{and} \quad \|\mathbf{P}_{\text{bad}}\mathbf{w}\|_* = \|\mathbf{P}_{\text{bad}}\mathbf{w}\|. \end{aligned}$$

In particular, for any $\mathbf{x} \in \mathcal{U}_{\text{good}}$ and $\mathbf{y} \in \mathbb{C}^n$ these relationships directly imply

$$\frac{\|\mathbf{y} - \mathbf{x}\|}{\|\mathbf{x}\|} \leq \sqrt{2} \frac{\|\mathbf{y} - \mathbf{x}\|_*}{\|\mathbf{x}\|_*},$$

and so $\delta(\mathcal{U}_{\text{good}}, \mathcal{K}_\ell(\mathbf{A}, \mathbf{v}_1^{(\nu)})) \leq \sqrt{2} \delta_*(\mathcal{U}_{\text{good}}, \mathcal{K}_\ell(\mathbf{A}, \mathbf{v}_1^{(\nu)}))$. Since $\mathcal{U}_{\text{good}}$ and \mathcal{U}_{bad} are orthogonal in this new inner product, we can apply the previous argument to conclude

$$\begin{aligned} \delta(\mathcal{U}_{\text{good}}, \mathcal{K}_\ell(\mathbf{A}, \mathbf{v}_1^{(\nu)})) &\leq \sqrt{2} \max_{\psi \in \mathcal{P}_{m-1}} \min_{\phi \in \mathcal{P}_{\ell-m}} \frac{\|\phi(\mathbf{A})\psi(\mathbf{A})\Psi_{\nu p}(\mathbf{A})\mathbf{P}_{\text{bad}}\mathbf{v}_1\|_*}{\|\phi(\mathbf{A})\psi(\mathbf{A})\Psi_{\nu p}(\mathbf{A})\mathbf{P}_{\text{good}}\mathbf{v}_1\|_*} \\ &= \sqrt{2} \max_{\psi \in \mathcal{P}_{m-1}} \min_{\phi \in \mathcal{P}_{\ell-m}} \frac{\|\phi(\mathbf{A})\psi(\mathbf{A})\Psi_{\nu p}(\mathbf{A})\mathbf{P}_{\text{bad}}\mathbf{v}_1\|}{\|\phi(\mathbf{A})\psi(\mathbf{A})\Psi_{\nu p}(\mathbf{A})\mathbf{P}_{\text{good}}\mathbf{v}_1\|}, \end{aligned}$$

as required.² \square

²A more precise value for C_0 can be found as

$$1 \leq C_0 = \sqrt{\frac{2\|\mathbf{I} - 2\mathbf{P}_{\text{good}}\|^2}{1 + \|\mathbf{I} - 2\mathbf{P}_{\text{good}}\|^2}} \leq \sqrt{2},$$

however the marginal improvement in the final bound would not appear to merit the substantial complexity added.

If \mathbf{N} is a square matrix with an invariant subspace \mathcal{V} , define

$$\|\mathbf{N}\|_{\mathcal{V}} \equiv \max_{\mathbf{v} \in \mathcal{V}} \frac{\|\mathbf{N}\mathbf{v}\|}{\|\mathbf{v}\|} = \|\mathbf{N}\mathbf{\Pi}_{\mathcal{V}}\|,$$

where $\mathbf{\Pi}_{\mathcal{V}}$ here denotes the orthogonal projection onto \mathcal{V} .

THEOREM 3.4. *Suppose \mathbf{A} , \mathbf{v}_1 , and the shifts $\{\mu_{jk}\}$ satisfy the conditions of Theorem 3.3. Then for $\ell \geq m$,*

$$\delta(\mathcal{U}_{\text{good}}, \mathcal{K}_{\ell}(\mathbf{A}, \mathbf{v}_1^{(\nu)})) \leq C_0 C_1 \min_{\phi \in \mathcal{P}_{\ell-m}} \|\phi(\mathbf{A})\Psi_{\nu p}(\mathbf{A})\|_{\mathcal{U}_{\text{good}}}^{-1} \|\phi(\mathbf{A})\Psi_{\nu p}(\mathbf{A})\|_{\mathcal{U}_{\text{bad}}},$$

where C_0 is as defined in Theorem 3.3 and

$$(3.5) \quad C_1 \equiv \max_{\psi \in \mathcal{P}_{m-1}} \frac{\|\psi(\mathbf{A})\mathbf{P}_{\text{bad}}\mathbf{v}_1\|}{\|\psi(\mathbf{A})\mathbf{P}_{\text{good}}\mathbf{v}_1\|}$$

is a constant independent of ℓ , ν , p , or the filter shifts $\{\mu_{jk}\}$.

Proof. Let $\mathbf{\Pi}_{\text{good}}$ and $\mathbf{\Pi}_{\text{bad}}$ denote the orthogonal projections onto $\mathcal{U}_{\text{good}}$ and \mathcal{U}_{bad} , respectively. Then

$$\begin{aligned} \|\Psi_{\nu p}(\mathbf{A})\phi(\mathbf{A})\mathbf{P}_{\text{bad}}\psi(\mathbf{A})\mathbf{v}_1\| &= \|\Psi_{\nu p}(\mathbf{A})\phi(\mathbf{A})\mathbf{\Pi}_{\text{bad}}\mathbf{P}_{\text{bad}}\psi(\mathbf{A})\mathbf{v}_1\| \\ &\leq \|\Psi_{\nu p}(\mathbf{A})\phi(\mathbf{A})\mathbf{\Pi}_{\text{bad}}\| \|\mathbf{P}_{\text{bad}}\psi(\mathbf{A})\mathbf{v}_1\| \end{aligned}$$

and, assuming for the moment that $\phi(\mathbf{A})$ is invertible,

$$\begin{aligned} \|\mathbf{P}_{\text{good}}\psi(\mathbf{A})\mathbf{v}_1\| &= \|[\Psi_{\nu p}(\mathbf{A})\phi(\mathbf{A})]^{-1}\mathbf{\Pi}_{\text{good}}\mathbf{P}_{\text{good}}\Psi_{\nu p}(\mathbf{A})\phi(\mathbf{A})\psi(\mathbf{A})\mathbf{v}_1\| \\ &\leq \|[\Psi_{\nu p}(\mathbf{A})\phi(\mathbf{A})]^{-1}\mathbf{\Pi}_{\text{good}}\| \|\mathbf{P}_{\text{good}}\Psi_{\nu p}(\mathbf{A})\phi(\mathbf{A})\psi(\mathbf{A})\mathbf{v}_1\|. \end{aligned}$$

Hence,

$$\begin{aligned} &\frac{\|\Psi_{\nu p}(\mathbf{A})\phi(\mathbf{A})\mathbf{P}_{\text{bad}}\psi(\mathbf{A})\mathbf{v}_1\|}{\|\Psi_{\nu p}(\mathbf{A})\phi(\mathbf{A})\mathbf{P}_{\text{good}}\psi(\mathbf{A})\mathbf{v}_1\|} \\ &\leq \|[\Psi_{\nu p}(\mathbf{A})\phi(\mathbf{A})]^{-1}\|_{\mathcal{U}_{\text{good}}} \|\Psi_{\nu p}(\mathbf{A})\phi(\mathbf{A})\|_{\mathcal{U}_{\text{bad}}} \frac{\|\psi(\mathbf{A})\mathbf{P}_{\text{bad}}\mathbf{v}_1\|}{\|\psi(\mathbf{A})\mathbf{P}_{\text{good}}\mathbf{v}_1\|}. \end{aligned}$$

Minimizing with respect to ϕ and maximizing with respect to ψ yields the conclusion provided the expression for C_1 is finite. This is assured since, as an immediate consequence of (3.2), $\|\psi(\mathbf{A})\mathbf{P}_{\text{good}}\mathbf{v}_1\| = 0$ can occur only when $\psi = 0$. \square

It is instructive to consider the situation where we seek only a single good eigenvalue, λ_1 , which is simple. In this case $m = \dim \mathcal{U}_{\text{good}} = 1$; the conclusion of Theorem 3.3 may be stated as

$$\delta(\mathcal{U}_{\text{good}}, \mathcal{K}_{\ell}(\mathbf{A}, \mathbf{v}_1^{(\nu)})) \leq C_0 C_1 \min_{\phi \in \mathcal{P}_{\ell-1}} \frac{\|\phi(\mathbf{A})\Psi_{\nu p}(\mathbf{A})\mathbf{w}\|}{|\phi(\lambda_1)\Psi_{\nu p}(\lambda_1)|},$$

where $\mathbf{w} = \mathbf{P}_{\text{bad}}\mathbf{v}_1/\|\mathbf{P}_{\text{bad}}\mathbf{v}_1\|$ and $C_1 = \|\mathbf{P}_{\text{bad}}\mathbf{v}_1\|/\|\mathbf{P}_{\text{good}}\mathbf{v}_1\|$. Elementary geometric considerations yield an alternate expression for C_1 :

$$C_1 = \sqrt{\left(\frac{1}{\|\mathbf{P}_{\text{good}}\|} \frac{\sin \Theta(\mathcal{U}_{\text{good}}, \mathbf{v}_1)}{\cos \Theta(\mathcal{U}_{\text{bad}}^{\perp}, \mathbf{v}_1)}\right)^2 + \left(1 - \frac{1}{\|\mathbf{P}_{\text{good}}\|} \frac{\cos \Theta(\mathcal{U}_{\text{good}}, \mathbf{v}_1)}{\cos \Theta(\mathcal{U}_{\text{bad}}^{\perp}, \mathbf{v}_1)}\right)^2},$$

where $\Theta(\mathcal{U}_{\text{good}}, \mathbf{v}_1)$ and $\Theta(\mathcal{U}_{\text{bad}}^\perp, \mathbf{v}_1)$ are the smallest angles between \mathbf{v}_1 and the one-dimensional subspaces $\mathcal{U}_{\text{good}}$ and $\mathcal{U}_{\text{bad}}^\perp$, respectively. This special case is stated as Proposition 2.1 of [18].³

Our next step is to reduce the conclusion of Theorem 3.4 to an approximation problem in the complex plane. Let \mathcal{U} be an invariant subspace of \mathbf{A} associated with a compact subset $\Omega \subset \mathbb{C}$ (that is, Ω contains only those eigenvalues of \mathbf{A} associated with \mathcal{U} and no others). Define $\kappa(\Omega)$ as the smallest constant for which the inequality

$$(3.6) \quad \|f(\mathbf{A})\|_{\mathcal{U}} \leq \kappa(\Omega) \max_{z \in \Omega} |f(z)|$$

holds uniformly over all $f \in H(\Omega)$, where $H(\Omega)$ denotes the functions analytic on Ω .⁴

Evidently, the value of the constant $\kappa(\Omega)$ depends on the particular choice of Ω (a set containing, in any case, those eigenvalues of \mathbf{A} associated with \mathcal{U}). The following properties of $\kappa(\Omega)$ are shared by the generalized Kreiss constant $\tilde{\mathcal{K}}(\Omega)$ of Toh and Trefethen [44] (defined for $\mathcal{U} = \mathbb{C}^n$). $\kappa(\Omega)$ is monotone decreasing with respect to set inclusion on Ω . Indeed, if $\Omega_1 \subseteq \Omega_2$, then for each function f analytic on Ω_2 ,

$$\frac{\|f(\mathbf{A})\|_{\mathcal{U}}}{\max\{|f(z)| : z \in \Omega_1\}} \geq \frac{\|f(\mathbf{A})\|_{\mathcal{U}}}{\max\{|f(z)| : z \in \Omega_2\}}.$$

Thus, $\Omega_1 \subseteq \Omega_2$ implies $\kappa(\Omega_1) \geq \kappa(\Omega_2)$.

Since the constant functions are always among the available analytic functions on Ω , $\kappa(\Omega) \geq 1$. If \mathbf{A} is normal, $\kappa(\Omega) = 1$. Indeed, if \mathbf{A} is normal and Σ denotes the set of eigenvalues of \mathbf{A} associated with the invariant subspace \mathcal{U} , then

$$1 \leq \kappa(\Omega) = \sup_{f \in H(\Omega)} \frac{\|f(\mathbf{A})\|_{\mathcal{U}}}{\max\{|f(z)| : z \in \Omega\}} = \sup_{f \in H(\Omega)} \frac{\max\{|f(\lambda)| : \lambda \in \Sigma\}}{\max\{|f(z)| : z \in \Omega\}} \leq 1.$$

If any eigenvalue associated with the invariant subspace \mathcal{U} is defective, then some choices of Ω will not yield a finite value for $\kappa(\Omega)$. For example, let

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

and take $\mathcal{U} = \mathbb{C}^2$ as an invariant subspace associated with the defective eigenvalue $\lambda = 0$. If Ω consists of the single point $\{0\}$ and $f(z) = z$ then evidently $\|f(\mathbf{A})\|_{\mathcal{U}} = 1$ but $\max_{z \in \Omega} |f(z)| = 0$. So, no finite value of $\kappa(\Omega)$ is possible (see [33, p. 440]). More generally, if Ω is the spectrum of a defective matrix \mathbf{A} , then the monic polynomial consisting of a single linear factor for each distinct eigenvalue of \mathbf{A} is zero on Ω but cannot annihilate \mathbf{A} , as it has lower degree than the minimum polynomial of \mathbf{A} .

We now use κ to adapt Theorem 3.4 into a more approachable approximation problem. In particular, if Ω_{good} is a compact subset of \mathbb{C} containing all the good eigenvalues of \mathbf{A} but none of the bad, then

$$\begin{aligned} \|\phi(\mathbf{A})\Psi_{\nu p}(\mathbf{A})^{-1}\|_{\mathcal{U}_{\text{good}}} &\leq \kappa(\Omega_{\text{good}}) \max\{|\phi(z)\Psi_{\nu p}(z)^{-1}| : z \in \Omega_{\text{good}}\} \\ &= \frac{\kappa(\Omega_{\text{good}})}{\min\{|\phi(z)\Psi_{\nu p}(z)| : z \in \Omega_{\text{good}}\}}. \end{aligned}$$

³[18] contains an error amounting to the tacit assumption that \mathbf{P}_{good} is an orthogonal projection, which is true only if $\mathcal{U}_{\text{good}} \perp \mathcal{U}_{\text{bad}}$. Thus the results coincide only in this special case (note $C_0 = 1$).

⁴For given $k \geq 1$, the sets Ω that (i) contain all eigenvalues of \mathbf{A} , and (ii) satisfy $\kappa(\Omega) \leq k$ are called *k-spectral sets* and figure prominently in dilation theory of operators [31].

Applying a similar bound to $\|\phi(\mathbf{A})\Psi_{\nu p}(\mathbf{A})\|_{\Omega_{\text{bad}}}$, we obtain the following result, the centerpiece of our development.

THEOREM 3.5. *Suppose \mathbf{A} and \mathbf{v}_1 satisfy the conditions of Theorem 3.3. Let Ω_{good} and Ω_{bad} be disjoint compact subsets of \mathbb{C} that contain, respectively, the good and bad eigenvalues of \mathbf{A} , and suppose that none of the filter shifts $\{\mu_{jk}\}$ lies in Ω_{good} . Then, for $\ell \geq m$,*

$$\delta(\mathcal{U}_{\text{good}}, \mathcal{K}_\ell(\mathbf{A}, \mathbf{v}_1^{(\nu)})) \leq C_0 C_1 C_2 \min_{\phi \in \mathcal{P}_{\ell-m}} \frac{\max\{|\Psi_{\nu p}(z)\phi(z)| : z \in \Omega_{\text{bad}}\}}{\min\{|\Psi_{\nu p}(z)\phi(z)| : z \in \Omega_{\text{good}}\}},$$

where C_0 and C_1 are the constants introduced in Theorems 3.3 and 3.4, respectively, and $C_2 \equiv \kappa(\Omega_{\text{good}})\kappa(\Omega_{\text{bad}})$.

Evidently, Theorem 3.5 can be implemented with a variety of choices for Ω_{good} and Ω_{bad} , which affects both the polynomial approximation problem (discussed in the next section) as well as the constant C_2 (considered in §5.3). A key feature of this development, which becomes more evident in the next section, is the limited role the location of the good eigenvalues play in the approximation problem.

4. The Polynomial Approximation Problem. Theorem 3.5 suggests the gap between a Krylov subspace and an invariant subspace will converge to zero at a rate determined by how small polynomials of increasing degree can become on Ω_{bad} while maintaining a minimal uniform magnitude on Ω_{good} . How can this manifest as a linear convergence rate? Consider the ansatz

$$\min_{\phi \in \mathcal{P}_{\ell^*}} \frac{\max\{|\phi(w)| : w \in \Omega_{\text{bad}}\}}{\min\{|\phi(z)| : z \in \Omega_{\text{good}}\}} = r^{\ell^*},$$

for some $0 < r \leq 1$. Pick a fixed $\phi \in \mathcal{P}_{\ell^*}$, say with exact degree ℓ^* . Then

$$(4.1) \quad \log \left(\frac{\max\{|\phi(w)| : w \in \Omega_{\text{bad}}\}}{\min\{|\phi(z)| : z \in \Omega_{\text{good}}\}} \right) \geq \ell^* \log(r).$$

Introducing $U_\phi(z, \Omega_{\text{bad}}) \equiv \frac{1}{\ell^*} \log \left(\frac{|\phi(z)|}{\max\{|\phi(w)| : w \in \Omega_{\text{bad}}\}} \right)$, (4.1) is equivalent to

$$\min_{z \in \Omega_{\text{good}}} U_\phi(z, \Omega_{\text{bad}}) \leq -\log(r).$$

Evidently, the size of r will be related to how large $U_\phi(z, \Omega_{\text{bad}})$ can be made uniformly throughout Ω_{good} ; larger U_ϕ values allow smaller r (faster rates). $U_\phi(z, \Omega_{\text{bad}})$ has the following properties

- $U_\phi(z, \Omega_{\text{bad}})$ is harmonic at z where $\phi(z) \neq 0$;
- $U_\phi(z, \Omega_{\text{bad}}) = \log|z| + c + o(1)$ for a finite constant c as $|z| \rightarrow \infty$;
- $U_\phi(z, \Omega_{\text{bad}}) \leq 0$ for all $z \in \partial\Omega_{\text{bad}}$.

Potential theory provides a natural setting for studying such approximation problems. It is central to the analysis of Krylov subspace methods for linear systems [10, 29], and has been used by Calvetti, Reichel, and Sorensen to analyze the Hermitian Lanczos algorithm with restarts [6]. We will apply similar techniques to study $U_\phi(z, \Omega_{\text{bad}})$, and thus begin by presenting some fundamental background.

4.1. Potential Theory Background. Consider the Dirichlet problem:

Suppose \mathcal{G} is an open region with compact closure in \mathbb{C} and suppose $f(z)$ is a continuous function on the boundary $\partial\mathcal{G}$. Find $u(z)$ such that

$$(4.2) \quad \begin{aligned} \Delta u &= 0 \quad \text{on } \mathcal{G}, \\ u &= f \quad \text{on } \partial\mathcal{G}. \end{aligned}$$

Solutions to this problem exist for each continuous f on $\partial\mathcal{G}$ provided the set \mathcal{G} is a “Dirichlet region” [8, §X.4]. This holds, for example, when \mathcal{G} has a piecewise smooth boundary with no isolated points. For our purposes here, such sets are sufficient; the effect of isolated points is addressed in §4.3.

Suppose, then, that \mathcal{G} is a bounded Dirichlet region containing the origin, $0 \in \mathcal{G}$. Assign $f(z) = \log|z|$ in (4.2) and let $u(z)$ denote the associated solution to the Dirichlet problem. Define the Green’s function of \mathcal{G} with pole at 0 by

$$\dot{g}(z, \mathcal{G}) = u(z) - \log|z| \quad \text{for } z \in \mathcal{G}.$$

The following relevant properties of $\dot{g}(z, \mathcal{G})$ are straightforward:

- (i) \dot{g} is harmonic in $\mathcal{G} \setminus \{0\}$;
- (ii) $\lim_{z \rightarrow 0} \dot{g}(z, \mathcal{G}) + \log|z| = \text{finite constant}$;
- (iii) $\lim_{z \rightarrow \hat{z}} \dot{g}(z, \mathcal{G}) = 0$ for all $\hat{z} \in \partial\mathcal{G}$;
- (iv) $\dot{g}(z, \mathcal{G}) > 0$ for all $z \in \mathcal{G}$.

Properties (i) and (iii) are elementary. Property (ii) is a consequence of the continuity of $u(z)$ at 0. Property (iv) follows from (i), (ii), the fact that (ii) implies that $\dot{g} > 0$ in any sufficiently small neighborhood of 0, and the maximum principle for harmonic functions. The maximum principle also shows that $\dot{g}(z, \mathcal{G})$ is the only function satisfying (i)–(iv).

The Green’s function of an *unbounded* region with pole at infinity can be defined and understood in terms of \dot{g} . Let $\mathcal{D} \subset \mathbb{C}$ be a compact set whose complement $\mathbb{C} \setminus \mathcal{D}$ is a connected Dirichlet region (say, \mathcal{D} has a piecewise smooth boundary and no isolated points). The conformal change of variable $z \mapsto z^{-1}$ maps $\mathbb{C} \setminus \mathcal{D}$ to a bounded region \mathcal{G} with $0 \in \mathcal{G}$. The Green’s function of $\mathbb{C} \setminus \mathcal{D}$ with pole at ∞ then is then defined as

$$g[z, \mathcal{D}] = \dot{g}(z^{-1}, \mathcal{G}) \quad \text{for } z \in \mathbb{C} \setminus \mathcal{D}.$$

Harmonicity is preserved under this change of variable, and one can see that properties (i)–(iv) hold for $g[z, \mathcal{D}]$, replacing 0 by ∞ , $\log|z|$ by $-\log|z|$, and \mathcal{G} by $\mathbb{C} \setminus \mathcal{D}$.

EXAMPLE 4.1. If $\mathbb{C} \setminus \mathcal{D}$ is simply connected, one is assured (from the Riemann mapping theorem; see, e.g., [8, §VII.4]) of the existence of a function $F(z)$ that maps $\mathbb{C} \setminus \mathcal{D}$ conformally onto the exterior of the closed unit disk $\mathbb{C} \setminus \mathcal{B}_1 = \{z : |z| > 1\}$ such that $F(\infty) = \infty$. Such an F must behave asymptotically as $\alpha z + \mathcal{O}(1)$ as $|z| \rightarrow \infty$ for some constant α , since it must remain one-to-one in any neighborhood of ∞ . Since $\log|z|$ is harmonic for any $z \neq 0$, one may check that $u(z) = \log|F(z)|$ is a solution to (4.2) with $f = 0$ and $u(\infty) = \infty$. Notice that $\log|z|$ itself is the Green’s function with pole at infinity for $\mathbb{C} \setminus \mathcal{B}_1$. Thus, $\log|F(z)|$ is the Green’s function with pole at infinity for $\mathbb{C} \setminus \mathcal{D}$. Evidently, $\lim_{|z| \rightarrow \infty} u(z) - \log|z| \rightarrow \log|\alpha|$. \square

Even for more complicated compact sets \mathcal{D} , the condition that $g[z, \mathcal{D}]$ is harmonic everywhere outside \mathcal{D} with a pole at ∞ restricts the rate of growth of $g[z, \mathcal{D}]$ near ∞ . Loosely speaking, as $|z|$ becomes very large, the compact set \mathcal{D} becomes less and less

distinguishable from a disk centered at 0 (say, with radius γ), and so $g[z, \mathcal{D}]$ becomes less and less distinguishable from $g[z, \mathcal{B}_\gamma] = \log |z/\gamma| = \log |z| - \log \gamma$, which is the Green's function with pole at infinity for $\mathbb{C} \setminus \mathcal{B}_\gamma = \{z : |z| > \gamma\}$. From property (ii) (with ∞ now replacing 0), we conclude that $g[z, \mathcal{D}]$ has growth at infinity satisfying

$$(4.3) \quad \lim_{|z| \rightarrow \infty} g[z, \mathcal{D}] - \log |z| = -\log \gamma$$

for some constant $\gamma > 0$ known as the *logarithmic capacity* of the set \mathcal{D} . This γ can be thought of as the effective radius of \mathcal{D} in the sense we've just described.

EXAMPLE 4.2. Suppose $\Phi_\ell(z)$ is a monic polynomial of degree ℓ and let

$$\mathcal{D}_\varepsilon(\Phi_\ell) = \{z \in \mathbb{C} : |\Phi_\ell(z)| \leq \varepsilon\}$$

be a family of regions whose boundaries are the ε -lemniscates of $\Phi_\ell(z)$. $\mathcal{D}_\varepsilon(\Phi_\ell)$ is compact for each $\varepsilon > 0$, though it need not be a connected region. With an easy calculation one may verify that $\mathcal{D}_\varepsilon(\Phi_\ell)$ has the Green's function (cf. [38, p. 164])

$$g[z, \mathcal{D}_\varepsilon(\Phi_\ell)] = \frac{1}{\ell} \log \left(\frac{|\Phi_\ell(z)|}{\varepsilon} \right). \quad \square$$

Equipped with the Green's function $g[z, \mathcal{D}]$, we can return our attention to the function $U_\phi(z, \mathcal{D})$ describing the error in our approximation problem. The following result is a simplified version of the Bernstein–Walsh lemma (see [38, §III.2]).

PROPOSITION 4.3. *Let \mathcal{D} be a compact set with piecewise smooth boundary $\partial\mathcal{D}$. Suppose u is harmonic outside \mathcal{D} and that $u(z) \leq 0$ for $z \in \partial\mathcal{D}$. If $u(z) = \log |z| + c + o(1)$ for some constant c as $|z| \rightarrow \infty$, then $u(z) \leq g[z, \mathcal{D}]$. In particular, if $\phi(z)$ is any polynomial of degree ℓ , then*

$$(4.4) \quad U_\phi(z, \mathcal{D}) = \frac{1}{\ell} \log \left(\frac{|\phi(z)|}{\max\{|\phi(w)| : w \in \mathcal{D}\}} \right) \leq g[z, \mathcal{D}]$$

for each $z \in \mathbb{C} \setminus \mathcal{D}$.

Proof. Define $v(z) \equiv u(z) - g[z, \mathcal{D}]$ for $z \in \mathbb{C} \setminus \mathcal{D}$. Note that v is harmonic in $\mathbb{C} \setminus \mathcal{D}$, $\lim_{z \rightarrow \partial\mathcal{D}} v(z) \leq 0$, and v is bounded at infinity. We will show this implies $v \leq 0$ throughout $\mathbb{C} \setminus \mathcal{D}$. Consider any $w \in \mathbb{C} \setminus \mathcal{D}$. Without loss of generality, suppose \mathcal{D} contains the unit disk centered at the origin. For any $\varepsilon > 0$, $v_\varepsilon(z) \equiv v(z) - \varepsilon \log |z|$ is harmonic in $\mathbb{C} \setminus \mathcal{D}$ with $\lim_{z \rightarrow \partial\mathcal{D}} v_\varepsilon(z) \leq 0$. Since v is bounded in $\mathbb{C} \setminus \mathcal{D}$, there exists $R > |w|$ such that $v_\varepsilon(z) \leq 0$ on the circle $|z| = R$. By the maximum principle, $v_\varepsilon(w) \leq 0$. Letting $\varepsilon \rightarrow 0$, we see that $v(w) \leq 0$, and so $u(w) \leq g[w, \mathcal{D}]$.

When ϕ has all its roots in \mathcal{D} and is of exact degree ℓ , then $U_\phi(z, \mathcal{D})$ satisfies the hypotheses on $u(z)$, giving (4.4). If ϕ has exact degree ℓ but roots outside \mathcal{D} , then $U_\phi(z, \mathcal{D})$ has singularities in $\mathbb{C} \setminus \mathcal{D}$. Enclose these singularities in a set \mathcal{B}_γ consisting of the union of closed disks with radius γ , with $\gamma > 0$ sufficiently small that $U_\phi < 0$ throughout the interior of \mathcal{B}_γ . Now $\lim_{z \rightarrow \partial\mathcal{D} \cup \partial\mathcal{B}_\gamma} U_\phi(z, \mathcal{D}) - g[z, \mathcal{D}] \leq 0$ and $U_\phi(z, \mathcal{D}) - g[z, \mathcal{D}]$ is harmonic on $\mathbb{C} \setminus \{\mathcal{D} \cup \mathcal{B}_\gamma\}$. By the above arguments, (4.4) holds throughout $\mathbb{C} \setminus \{\mathcal{D} \cup \mathcal{B}_\gamma\}$ and remains true as $\gamma > 0$ is made arbitrarily small.

If $\deg(\phi) < \ell$, replace ℓ in the definition of U_ϕ by $\deg(\phi)$. Apply the argument above to prove this new function is bounded by $g[z, \mathcal{D}]$; since this function is larger than U_ϕ , (4.4) holds. \square

For certain special choices of $\mathcal{D} = \Omega_{\text{bad}}$, the polynomial approximation problem of Theorem 3.5 can be solved exactly.

THEOREM 4.4. *Suppose $\Phi_{\ell^*}(z)$ is a monic polynomial of degree ℓ^* . Let $\Omega_{\text{bad}} = \mathcal{D}_\varepsilon(\Phi_{\ell^*})$ be an associated ε -lemniscatic set as defined in Example 4.2 and suppose Ω_{good} is a compact subset of \mathbb{C} such that $\Omega_{\text{good}} \cap \mathcal{D}_\varepsilon(\Phi_{\ell^*}) = \emptyset$. Then*

$$\min_{\phi \in \mathcal{P}_{\ell^*}} \frac{\max\{|\phi(w)| : w \in \Omega_{\text{bad}}\}}{\min\{|\phi(z)| : z \in \Omega_{\text{good}}\}} = \frac{\varepsilon}{\min\{|\Phi_{\ell^*}(z)| : z \in \Omega_{\text{good}}\}}.$$

Proof. Using the Green's function for $\mathcal{D}_\varepsilon(\Phi_{\ell^*})$ described in Example 4.2, we can rearrange (4.4) to show that for any $\phi \in \mathcal{P}_{\ell^*}$,

$$\frac{|\phi(z)|}{\max\{|\phi(w)| : w \in \mathcal{D}_\varepsilon(\Phi_{\ell^*})\}} \leq \frac{|\Phi_{\ell^*}(z)|}{\varepsilon}$$

holds for all $z \in \Omega_{\text{good}}$. Equality is attained for every $z \in \mathbb{C}$ whenever $\phi = \Phi_{\ell^*}$. Minimizing over $z \in \Omega_{\text{good}}$ and then maximizing over $\phi \in \mathcal{P}_{\ell^*}$ yields

$$(4.5) \quad \max_{\phi \in \mathcal{P}_{\ell^*}} \frac{\min\{|\phi(z)| : z \in \Omega_{\text{good}}\}}{\max\{|\phi(w)| : w \in \mathcal{D}_\varepsilon(\Phi_{\ell^*})\}} \leq \frac{\min\{|\Phi_{\ell^*}(z)| : z \in \Omega_{\text{good}}\}}{\varepsilon}.$$

In fact, equality must hold in (4.5) since $\phi = \Phi_{\ell^*}$ is included in the class of functions over which the maximization occurs. The conclusion then follows by taking the reciprocal of both sides. \square

More general choices of $\mathcal{D} = \Omega_{\text{bad}}$ will not typically yield an exactly solvable polynomial approximation problem, at least for fixed (finite) polynomial degree. However, the following asymptotic result can be obtained as the polynomial degree is allowed to increase.

THEOREM 4.5. *Let Ω_{bad} and Ω_{good} be two disjoint compact sets in the complex plane such that $\mathbb{C} \setminus \Omega_{\text{bad}}$ is a Dirichlet region. Then*

$$(4.6) \quad \lim_{\ell^* \rightarrow \infty} \min_{\phi \in \mathcal{P}_{\ell^*}} \left(\frac{\max\{|\phi(w)| : w \in \Omega_{\text{bad}}\}}{\min\{|\phi(z)| : z \in \Omega_{\text{good}}\}} \right)^{1/\ell^*} = e^{-\min\{g[z, \Omega_{\text{bad}}] : z \in \Omega_{\text{good}}\}}$$

where $g[z, \Omega_{\text{bad}}]$ is the Green's function of $\mathbb{C} \setminus \Omega_{\text{bad}}$ with pole at infinity.

Proof. The theorem is proved in [27, p. 236], where the left hand side of (4.6) is referred to as the $(\ell^*, 0)$ Zolotarjov number. We give here a brief indication of the proof sufficient to support later discussion. The inequality (4.4) can be manipulated to yield

$$\left(\frac{|\Phi_{\ell^*}(z)|}{\max\{|\Phi_{\ell^*}(w)| : w \in \Omega_{\text{bad}}\}} \right)^{1/\ell^*} \leq e^{g[z, \Omega_{\text{bad}}]},$$

which in turn implies

$$\left(\frac{\max\{|\Phi_{\ell^*}(w)| : w \in \Omega_{\text{bad}}\}}{\min\{|\Phi_{\ell^*}(z)| : z \in \Omega_{\text{good}}\}} \right)^{1/\ell^*} \geq e^{-\min\{g[z, \Omega_{\text{bad}}] : z \in \Omega_{\text{good}}\}}.$$

Furthermore, one may construct polynomials L_k that have as their zeros points distributed on the boundary $\partial\Omega_{\text{bad}}$, the *Leja points* $\{\mu_1, \mu_2, \dots, \mu_k\}$, defined recursively so that

$$\mu_{k+1} = \arg \max \left\{ \prod_{j=1}^k |z - \mu_j| : z \in \Omega_{\text{bad}} \right\};$$

see [38, §V.1]. This sequence of *Leja polynomials* satisfies asymptotic optimality,

$$(4.7) \quad \lim_{k \rightarrow \infty} \left(\frac{|L_k(z)|}{\max\{|L_k(w)| : w \in \Omega_{\text{bad}}\}} \right)^{1/k} = e^{g[z, \Omega_{\text{bad}}]}$$

for each $z \in \mathbb{C} \setminus \Omega_{\text{bad}}$. Convergence is uniform on compact subsets of $\mathbb{C} \setminus \Omega_{\text{bad}}$. Thus we can reverse the order of the limit with respect to polynomial degree and minimization with respect to $z \in \Omega_{\text{good}}$, then take reciprocals to find

$$(4.8) \quad \lim_{k \rightarrow \infty} \left(\frac{\max\{|L_k(w)| : w \in \Omega_{\text{bad}}\}}{\min\{|L_k(z)| : z \in \Omega_{\text{good}}\}} \right)^{1/k} = e^{-\min\{g[z, \Omega_{\text{bad}}] : z \in \Omega_{\text{good}}\}}.$$

Since

$$\begin{aligned} \left(\frac{\max\{|L_{\ell^*}(w)| : w \in \Omega_{\text{bad}}\}}{\min\{|L_{\ell^*}(z)| : z \in \Omega_{\text{good}}\}} \right)^{1/\ell^*} &\geq \min_{\phi \in \mathcal{P}_{\ell^*}} \left(\frac{\max\{|\phi(w)| : w \in \Omega_{\text{bad}}\}}{\min\{|\phi(z)| : z \in \Omega_{\text{good}}\}} \right)^{1/\ell^*} \\ &\geq e^{-\min\{g[z, \Omega_{\text{bad}}] : z \in \Omega_{\text{good}}\}}, \end{aligned}$$

equality must hold throughout and thus (4.6) holds. \square

In the context of Example 4.1, where $F(z)$ was a conformal map taking the exterior of Ω_{bad} to the exterior of the closed unit disk with $F(\infty) = \infty$, Theorem 4.5 reduces to (cf. [10, Thm. 2])

$$\lim_{\ell^* \rightarrow \infty} \min_{\phi \in \mathcal{P}_{\ell^*}} \left(\frac{\max\{|\phi(w)| : w \in \Omega_{\text{bad}}\}}{\min\{|\phi(z)| : z \in \Omega_{\text{good}}\}} \right)^{1/\ell^*} = \max_{z \in \Omega_{\text{good}}} \frac{1}{|F(z)|}.$$

4.2. Effective restart strategies. The usual goal in constructing a restart strategy is to limit the size of the Krylov subspace (restricting the maximum degree of the polynomial ϕ) without degrading the asymptotic convergence rate. Demonstrating equality in (4.6) pivoted on the construction of an optimal family of polynomials—in this case, Leja polynomials. There are other possibilities, however. Fekete polynomials are the usual choice for the construction in Theorem 4.5; see [38, §III.1]. Chebyshev polynomials and Faber polynomials offer familiar alternatives. (For Hermitian matrices, a practical Leja shift strategy has been developed by Calvetti, et al. [3, 6]. Heuveline and Sadkane advocate numerical conformal mapping to determine Faber polynomials for restarting non-Hermitian iterations [18].) Once some optimal family of polynomials is known that solves (4.6), effective restart strategies become evident.

THEOREM 4.6. *Let Ω_{good} and Ω_{bad} be two disjoint compact sets in the complex plane containing, respectively, the good and bad eigenvalues of \mathbf{A} , and such that $\mathbb{C} \setminus \Omega_{\text{bad}}$ is a Dirichlet region. Suppose that $\Psi_{\nu, p}(z)$ is the aggregate restart polynomial*

representing ν restarts each of order p .

(a) If polynomial restarts are performed using roots of optimal polynomials for Ω_{bad} (i.e., $\Psi_{\nu p}(z)$ are optimal polynomials of degree νp), then

$$(4.9) \quad \lim_{\nu \rightarrow \infty} \min_{\phi \in \mathcal{P}_{\ell^*}} \left(\frac{\max\{|\Psi_{\nu p}(w)\phi(w)| : w \in \Omega_{\text{bad}}\}}{\min\{|\Psi_{\nu p}(z)\phi(z)| : z \in \Omega_{\text{good}}\}} \right)^{\frac{1}{\nu p + \ell^*}} = e^{-\min\{g[z, \Omega_{\text{bad}}] : z \in \Omega_{\text{good}}\}},$$

where $g[z, \Omega_{\text{bad}}]$ is the Green's function of Ω_{bad} with pole at infinity.

(b) If the boundary of Ω_{bad} is a lemniscate of $\Psi_{\nu p}\Phi_{\ell^*}$,

$$\Omega_{\text{bad}} = \mathcal{D}_\varepsilon(\Psi_{\nu p}\Phi_{\ell^*}) = \{z \in \mathbb{C} : |\Psi_{\nu p}(z)\Phi_{\ell^*}(z)| \leq \varepsilon\},$$

for some degree- ℓ^* monic polynomial Φ_{ℓ^*} and some $\varepsilon > 0$, then

$$\min_{\phi \in \mathcal{P}_{\ell^*}} \frac{\max\{|\Psi_{\nu p}(w)\phi(w)| : w \in \Omega_{\text{bad}}\}}{\min\{|\Psi_{\nu p}(z)\phi(z)| : z \in \Omega_{\text{good}}\}} = \frac{\varepsilon}{\min\{|\Psi_{\nu p}(z)\Phi_{\ell^*}(z)| : z \in \Omega_{\text{good}}\}}.$$

Proof. Part (b) follows immediately from Theorem 4.4. Part (a) can be seen by observing that since $\Psi_{\nu p}(z)$ is an asymptotically optimal family for Ω_{bad} ,

$$\begin{aligned} \frac{\max\{|\Psi_{\nu p}(w)| : w \in \Omega_{\text{bad}}\}}{\min\{|\Psi_{\nu p}(z)| : z \in \Omega_{\text{good}}\}} &\geq \min_{\phi \in \mathcal{P}_{\ell^*}} \left(\frac{\max\{|\Psi_{\nu p}(w)\phi(w)| : w \in \Omega_{\text{bad}}\}}{\min\{|\Psi_{\nu p}(z)\phi(z)| : z \in \Omega_{\text{good}}\}} \right) \\ &\geq \left(e^{-\min\{g[z, \Omega_{\text{bad}}] : z \in \Omega_{\text{good}}\}} \right)^{\nu p + \ell^*}. \end{aligned}$$

Now fixing p and ℓ^* , the conclusion follows from (4.8) by following the subsequence generated by $\nu = 1, 2, \dots$ \square

Recall that the desired effect of the restart polynomial is to retain the rapid convergence rate of the full (unrestarted) Krylov subspace without requiring the dimension ℓ^* to grow without bound. We have seen here that restarting with optimal polynomials for Ω_{bad} recovers the expected linear convergence rate for Ω_{bad} (presuming one can identify this set, not a trivial matter in practice). Still, the unrestarted process may take advantage of the discrete nature of the spectrum, accelerating convergence beyond the expected linear rate. Designing a restart strategy that yields similar behavior is more elaborate.

4.3. Superlinear effects from assimilation of bad eigenvalues. In a variety of situations, the gap appears to converge superlinearly. True superlinear convergence is an asymptotic phenomenon that has a nontrivial meaning only for nonterminating iterations. Thus one must be cautious about describing superlinear effects relating to (unrestarted) Krylov subspaces, since $\mathcal{U}_{\text{good}}$ is eventually completely captured by the Krylov subspace as discussed in §2. Here our point of view follows that of [49, 51], showing the estimated gap may be bounded by a family of linearly converging processes exhibiting increasingly rapid linear rates. The next result mimics the Ritz value bounds for Hermitian matrices developed by van der Sluis and van der Vorst [50, §6.6]. We assume here that Ω_{bad} consists of the union of s discrete points, potentially with some additional Dirichlet region. That is, some bad eigenvalues (typically those closest to the good eigenvalues, or distant outliers) are treated as discrete points, while any leftovers are collected in the Dirichlet region.

THEOREM 4.7. *Let Ω_{good} and Ω_{bad} be disjoint compact subsets of \mathbb{C} and suppose Ω_{bad} contains s isolated points, z_1, z_2, \dots, z_s . Define a sequence of $s + 1$ nested subsets as $\Omega_k = \Omega_{k+1} \cup \{z_k\}$ for $k = 1, \dots, s$ with $\Omega_1 \equiv \Omega_{\text{bad}}$, so that each set $\Omega_k \supset \Omega_{k+1} \neq \emptyset$ differs from adjacent sets in the sequence by single points. Define also the associated diameters*

$$e_k \equiv \max\{|w - z_k| : w \in \Omega_k\} \quad \text{and} \quad d_k \equiv \min\{|z - z_k| : z \in \Omega_{\text{good}}\}.$$

Then for $r = 1, \dots, s$ and each $\ell^* > r$,

$$\min_{\phi \in \mathcal{P}_{\ell^*}} \frac{\max\{|\phi(w)| : w \in \Omega_{\text{bad}}\}}{\min\{|\phi(z)| : z \in \Omega_{\text{good}}\}} \leq \left(\prod_{j=1}^r \frac{e_j}{d_j} \right) \min_{\phi \in \mathcal{P}_{\ell^* - r}} \frac{\max\{|\phi(w)| : w \in \Omega_{r+1}\}}{\min\{|\phi(z)| : z \in \Omega_{\text{good}}\}}.$$

Proof. Fix an integer $k \geq 1$ and observe that

$$\begin{aligned} \min_{\phi \in \mathcal{P}_{\ell^*}} \frac{\max_{w \in \Omega_k} |\phi(w)|}{\min_{z \in \Omega_{\text{good}}} |\phi(z)|} &\leq \min_{\phi \in \mathcal{P}_{\ell^* - 1}} \frac{\max_{w \in \Omega_k} |(w - z_k)\phi(w)|}{\min_{z \in \Omega_{\text{good}}} |(z - z_k)\phi(z)|} \\ &= \min_{\phi \in \mathcal{P}_{\ell^* - 1}} \frac{\max_{w \in \Omega_{k+1}} |(w - z_k)\phi(w)|}{\min_{z \in \Omega_{\text{good}}} |(z - z_k)\phi(z)|} \\ &\leq \frac{e_k}{d_k} \min_{\phi \in \mathcal{P}_{\ell^* - 1}} \frac{\max_{w \in \Omega_{k+1}} |\phi(w)|}{\min_{z \in \Omega_{\text{good}}} |\phi(z)|}. \end{aligned}$$

The conclusion follows by applying the argument repeatedly for $k = 1, 2, \dots, r$. \square

Asymptotically, the discrete points in Ω_{bad} have no effect on the convergence rate.

COROLLARY 4.8. *In the notation of Theorem 4.7, suppose Ω_{s+1} is a Dirichlet region. Then*

$$\lim_{\ell^* \rightarrow \infty} \min_{\phi \in \mathcal{P}_{\ell^*}} \left(\frac{\max\{|\phi(w)| : w \in \Omega_{\text{bad}}\}}{\min\{|\phi(z)| : z \in \Omega_{\text{good}}\}} \right)^{1/\ell^*} \leq e^{-\min\{g[z, \Omega_{s+1}] : z \in \Omega_{\text{good}}\}},$$

where $g[z, \Omega_{s+1}]$ is the Green's function with pole at infinity associated with $\mathbb{C} \setminus \Omega_{s+1}$.

Proof. e Since $\Omega_{r+1} \subseteq \Omega_{r+1}^{\text{conv}}$, The result follows by applying the asymptotic approach of Theorem 4.5 to the result of Theorem 4.7 for $r = s$. \square

To demonstrate such superlinear effects, we consider a parameterized diagonal matrix \mathbf{A}_α having 100 bad eigenvalues spaced uniformly in the unit interval $[-1 - \alpha, -\alpha]$ and 4 good eigenvalues uniformly spaced in $[0, 1]$. Figure 4.1 illustrates convergence of the gap $\delta(\mathcal{U}_{\text{good}}, \mathcal{K}_\ell(\mathbf{A}_\alpha, \mathbf{v}_1))$ for $\alpha = 0.1, 0.01, 0.05$, and 0.001 , always with the starting vector \mathbf{v}_1 having $1/\sqrt{n}$ in each component ($n = 104$). Above each convergence curve are bounds from Theorem 3.5 and Theorem 4.7. (The calculation of C_1 is addressed in §5.1.) For the superlinear bounds, take Ω_{bad} to be the set of bad eigenvalues and set Ω_r to be Ω_{bad} less the $r - 1$ rightmost bad eigenvalues. We approximate the optimal polynomial in Theorem 4.7 by Chebyshev polynomials for $\Omega_{r+1}^{\text{conv}}$ (see [37, §IV.4.1] for details). Notice the envelope produced by the aggregated linear rates creates a superlinear convergence effect to an extent determined by the ‘‘granularity’’ of bad eigenvalues as viewed from the nearest good eigenvalue. Greater granularity (smaller α) causes poor initial rates due to nearby bad eigenvalues, that rapidly dissipate as these eigenvalues are assimilated, yielding to improved rates determined by more remote bad eigenvalues. The same phenomenon is observed in §6.4




FIG. 4.1. Aggregate linear rates produce a superlinear effect. Observed gap convergence (broken line) and aggregate bounds (solid lines) predicted by Theorem 3.5 and Theorem 4.7 for Ω_r with $r = 1, \dots, 50$.

for a Markov chain eigenvalue problem. But assimilation of nearby bad eigenvalues is not the only mechanism for superlinear convergence. In §5.3, we describe how non-normality can also give rise to such behavior, illustrated experimentally in §6.2.

5. Analysis of Constants. This section contains a more detailed discussion of the constants C_1 and C_2 that arise in the convergence bounds given in Theorems 3.4 and 3.5. The magnitude of these constants controls the predicted start of the linear phase of convergence: larger constants suggest delayed linear convergence. Thus we seek an appreciation of those matrix and starting vector properties that lead to more or less favorable convergence bounds.

5.1. Bounding C_1 . Notice that

$$C_1 = \max_{\psi \in \mathcal{P}_{m-1}} \frac{\|\psi(\mathbf{A})\mathbf{P}_{\text{bad}}\mathbf{v}_1\|}{\|\psi(\mathbf{A})\mathbf{P}_{\text{good}}\mathbf{v}_1\|} = \max_{\mathbf{v} \in \mathcal{K}_m(\mathbf{A}, \mathbf{v}_1)} \frac{\|\mathbf{P}_{\text{bad}}\mathbf{v}\|}{\|\mathbf{P}_{\text{good}}\mathbf{v}\|} = \max_{\mathbf{x} \in \mathbb{C}^m} \frac{\|\mathbf{P}_{\text{bad}}\mathbf{V}_m\mathbf{x}\|}{\|\mathbf{P}_{\text{good}}\mathbf{V}_m\mathbf{x}\|},$$

where the columns of \mathbf{V}_m form a basis for $\mathcal{K}_m(\mathbf{A}, \mathbf{v}_1)$. This last expression for C_1 is simply the largest generalized singular value of the pair of matrices $\mathbf{P}_{\text{bad}}\mathbf{V}_m$ and $\mathbf{P}_{\text{good}}\mathbf{V}_m$ (see, e.g., [14, §8.7.3]). This is how we determine C_1 for our computational examples.

The dependence of C_1 on the starting vector \mathbf{v}_1 is critical. If \mathbf{v}_1 is biased against $\mathcal{U}_{\text{good}}$, then C_1 will be large and our bounds predict a delay in convergence. Likewise, a good starting vector accelerates convergence as expected.⁵ We will investigate this

⁵Though our bounds explicitly incorporate restart effects into the polynomial approximation problem, an alternative approach could instead handle restarts via the constant C_1 , which we expect to shrink as restarts enrich the starting vector in $\mathcal{U}_{\text{good}}$.

behavior with an illustrative example, but first give bounds for C_1 that relate its magnitude to the orientation of $\mathcal{K}_m(\mathbf{A}, \mathbf{v}_1)$ relative to the invariant subspaces $\mathcal{U}_{\text{good}}$ and \mathcal{U}_{bad} .

PROPOSITION 5.1. *Under the conditions of Theorem 3.4,*

$$\frac{1}{\|\mathbf{P}_{\text{good}}\|} \frac{\delta(\mathcal{K}_m(\mathbf{A}, \mathbf{v}_1), \mathcal{U}_{\text{good}})}{\delta(\mathcal{K}_m(\mathbf{A}, \mathbf{v}_1), \mathcal{U}_{\text{bad}})} \leq C_1 \leq \frac{\|\mathbf{P}_{\text{good}}\| \delta(\mathcal{K}_m(\mathbf{A}, \mathbf{v}_1), \mathcal{U}_{\text{good}})}{1 - \|\mathbf{P}_{\text{good}}\| \delta(\mathcal{K}_m(\mathbf{A}, \mathbf{v}_1), \mathcal{U}_{\text{good}})},$$

where the second inequality holds provided $\|\mathbf{P}_{\text{good}}\| \delta(\mathcal{K}_m(\mathbf{A}, \mathbf{v}_1), \mathcal{U}_{\text{good}}) < 1$.

Proof. If $\mathbf{\Pi}_{\text{good}}$ denotes the orthogonal projection onto $\mathcal{U}_{\text{good}}$ then $\mathbf{I} - \mathbf{\Pi}_{\text{good}} = (\mathbf{I} - \mathbf{\Pi}_{\text{good}})(\mathbf{I} - \mathbf{P}_{\text{good}})$, and so

$$\|(\mathbf{I} - \mathbf{\Pi}_{\text{good}})\psi(\mathbf{A})\mathbf{v}_1\| \leq \|(\mathbf{I} - \mathbf{P}_{\text{good}})\psi(\mathbf{A})\mathbf{v}_1\| = \|\psi(\mathbf{A})\mathbf{P}_{\text{bad}}\mathbf{v}_1\|.$$

Thus,

$$\begin{aligned} \delta(\mathcal{K}_m(\mathbf{A}, \mathbf{v}_1), \mathcal{U}_{\text{good}}) &= \max_{\psi \in \mathcal{P}_{m-1}} \min_{\mathbf{u} \in \mathcal{U}_{\text{good}}} \frac{\|\mathbf{u} - \psi(\mathbf{A})\mathbf{v}_1\|}{\|\psi(\mathbf{A})\mathbf{v}_1\|} \\ &= \max_{\psi \in \mathcal{P}_{m-1}} \frac{\|(\mathbf{I} - \mathbf{\Pi}_{\text{good}})\psi(\mathbf{A})\mathbf{v}_1\|}{\|\psi(\mathbf{A})\mathbf{v}_1\|} \\ &= \max_{\psi \in \mathcal{P}_{m-1}} \frac{\|\psi(\mathbf{A})\mathbf{P}_{\text{good}}\mathbf{v}_1\|}{\|\psi(\mathbf{A})\mathbf{v}_1\|} \frac{\|(\mathbf{I} - \mathbf{\Pi}_{\text{good}})(\mathbf{I} - \mathbf{P}_{\text{good}})\psi(\mathbf{A})\mathbf{v}_1\|}{\|\psi(\mathbf{A})\mathbf{P}_{\text{good}}\mathbf{v}_1\|} \\ &\leq \max_{\psi \in \mathcal{P}_{m-1}} \frac{\|(\mathbf{I} - \mathbf{P}_{\text{bad}})(\mathbf{I} - \mathbf{\Pi}_{\text{bad}})\psi(\mathbf{A})\mathbf{v}_1\|}{\|\psi(\mathbf{A})\mathbf{v}_1\|} \frac{\|\mathbf{P}_{\text{bad}}\psi(\mathbf{A})\mathbf{v}_1\|}{\|\psi(\mathbf{A})\mathbf{P}_{\text{good}}\mathbf{v}_1\|} \\ &\leq \|\mathbf{I} - \mathbf{P}_{\text{bad}}\| \max_{\psi \in \mathcal{P}_{m-1}} \frac{\|(\mathbf{I} - \mathbf{\Pi}_{\text{bad}})\psi(\mathbf{A})\mathbf{v}_1\|}{\|\psi(\mathbf{A})\mathbf{v}_1\|} \frac{\|\psi(\mathbf{A})\mathbf{P}_{\text{bad}}\mathbf{v}_1\|}{\|\psi(\mathbf{A})\mathbf{P}_{\text{good}}\mathbf{v}_1\|} \\ &\leq \|\mathbf{P}_{\text{good}}\| \delta(\mathcal{K}_m(\mathbf{A}, \mathbf{v}_1), \mathcal{U}_{\text{bad}}) C_1. \end{aligned}$$

This gives the first inequality. For the second, note that for any $\psi \in \mathcal{P}_{m-1}$,

$$\begin{aligned} \frac{\|\psi(\mathbf{A})\mathbf{P}_{\text{bad}}\mathbf{v}_1\|}{\|\psi(\mathbf{A})\mathbf{P}_{\text{good}}\mathbf{v}_1\|} &= \frac{\|(\mathbf{I} - \mathbf{P}_{\text{good}})\psi(\mathbf{A})\mathbf{v}_1\|}{\|\psi(\mathbf{A})\mathbf{v}_1\|} \frac{\|\psi(\mathbf{A})\mathbf{v}_1\|}{\|\psi(\mathbf{A})\mathbf{P}_{\text{good}}\mathbf{v}_1\|} \\ &= \frac{\|(\mathbf{I} - \mathbf{P}_{\text{good}})(\mathbf{I} - \mathbf{\Pi}_{\text{good}})\psi(\mathbf{A})\mathbf{v}_1\|}{\|\psi(\mathbf{A})\mathbf{v}_1\|} \frac{\|\psi(\mathbf{A})(\mathbf{P}_{\text{good}} + \mathbf{P}_{\text{bad}})\mathbf{v}_1\|}{\|\psi(\mathbf{A})\mathbf{P}_{\text{good}}\mathbf{v}_1\|} \\ &\leq \|\mathbf{I} - \mathbf{P}_{\text{good}}\| \frac{\|(\mathbf{I} - \mathbf{\Pi}_{\text{good}})\psi(\mathbf{A})\mathbf{v}_1\|}{\|\psi(\mathbf{A})\mathbf{v}_1\|} \left(1 + \frac{\|\psi(\mathbf{A})\mathbf{P}_{\text{bad}}\mathbf{v}_1\|}{\|\psi(\mathbf{A})\mathbf{P}_{\text{good}}\mathbf{v}_1\|}\right). \end{aligned}$$

(A more frugal inequality leads to a sharper but rather intricate upper bound for C_1 .) Maximizing over $\psi \in \mathcal{P}_{m-1}$ and noting that $\|\mathbf{I} - \mathbf{P}_{\text{good}}\| = \|\mathbf{P}_{\text{good}}\|$ [22] yields

$$C_1 \leq \|\mathbf{P}_{\text{good}}\| \delta(\mathcal{K}_m(\mathbf{A}, \mathbf{v}_1), \mathcal{U}_{\text{good}})(1 + C_1).$$

When $\|\mathbf{P}_{\text{good}}\| \delta(\mathcal{K}_m(\mathbf{A}, \mathbf{v}_1), \mathcal{U}_{\text{good}}) < 1$, this expression can be rearranged to give the desired upper bound. \square

The bounds given in Proposition 5.1 can be disparate when $\|\mathbf{P}_{\text{good}}\|$ is large or $\delta(\mathcal{K}_m(\mathbf{A}, \mathbf{v}_1), \mathcal{U}_{\text{good}})$ is close to one. To obtain alternative lower bounds, approximate

the maximizing polynomial ψ in (3.5). Some intuitively appealing choices for the roots of $\psi \in \mathcal{P}_{m-1}$ include the Ritz values or harmonic Ritz values generated from $\mathcal{K}_{m-1}(\mathbf{A}, \mathbf{P}_{\text{good}}\mathbf{v}_1)$. (This is motivated by the fact that taking ψ to be a degree- m polynomial with the m Ritz values from $\mathcal{K}_m(\mathbf{A}, \mathbf{P}_{\text{good}}\mathbf{v}_1)$ as roots would zero the denominator of the expression (3.5) for C_1 .)

5.2. An Illustration of Starting Vector Influence. Consider a Hermitian matrix $\mathbf{A} \in \mathbb{C}^{128 \times 128}$ with eight good eigenvalues uniformly distributed in the interval $[1, 2]$. The remaining eigenvalues uniformly fill the interval $[-1, 0]$. Since \mathbf{A} is normal, the constants C_0 and C_2 are trivial, $C_0 = C_2 = 1$. Theorem 3.5 thus bounds gap convergence as the product of the constant C_1 , which depends on the starting vector, and a polynomial approximation problem, which is independent of it. Taking $\Omega_{\text{bad}} = [-1, 0]$ and $\Omega_{\text{good}} = [1, 2]$, Theorem 4.5 yields an asymptotic convergence factor of $3 - \sqrt{2} \approx 0.1716$, an expedient rate due to the good separation of Ω_{good} from Ω_{bad} . To study the role of C_1 , we construct seven different starting vectors \mathbf{v}_1 that form angles of $\theta = 10^{-15}, 10^{-12}, 10^{-9}, 10^{-6}, 10^{-3}, 1$ radians with $\mathcal{U}_{\text{good}}$. (Each starting vector has equal components in each unwanted eigenvector direction.) Figure 5.1 shows the result of this experiment. The gap convergence curves are solid lines; the dotted lines show bounds from Theorem 3.5. For the finite-degree polynomial approximation problem in Theorem 3.5, we use Chebyshev polynomials for $\Omega_{\text{bad}} = [-1, 0]$. (Since $\delta(\mathcal{U}_{\text{good}}, \mathcal{K}_\ell(\mathbf{A}, \mathbf{v}_1)) = 1$ when $\ell < m = \dim \mathcal{U}_{\text{good}} = 8$, we show the complementary measure $\delta(\mathcal{K}_\ell(\mathbf{A}, \mathbf{v}_1), \mathcal{U}_{\text{good}})$ for the first seven iterations.) As predicted by our bounds, the asymptotic convergence rate appears largely independent of the orientation of \mathbf{v}_1 . Interestingly, even a considerable starting vector bias toward $\mathcal{U}_{\text{good}}$ yields only a modest improvement in convergence, which may appear even less significant for problems with slower convergence rates.




FIG. 5.1. The effect of a biased starting vector on gap convergence. The solid lines denote the computed gap convergence curves for starting vectors \mathbf{v}_1 that form angles of θ radians with $\mathcal{U}_{\text{good}}$. The dotted lines show the bound derived from Theorem 3.5 for each value of θ . The black dots denote the values of C_1 . In the vertical axis label, \mathcal{K}_ℓ is a shorthand for $\mathcal{K}_\ell(\mathbf{A}, \mathbf{v}_1)$.

5.3. Bounding C_2 . In contrast to C_1 , which was strongly linked to the orientation of the starting vector \mathbf{v}_1 with respect to the good invariant subspace, the constant C_2 has a somewhat more diffuse interpretation. C_2 captures the effect of the non-normality of \mathbf{A} , yet ambiguity in the selection of Ω_{good} and Ω_{bad} injects wide variability to the values C_2 can achieve. Generally speaking, choosing the sets Ω_{good} and Ω_{bad} to be overly large yields a small constant C_2 at the expense of a slow convergence rate for the polynomial approximation problem. Shrinking these sets increases the constant but improves the predicted convergence rate. The smallest possible sets that can be chosen for Ω_{good} and Ω_{bad} are the sets of good and bad eigenvalues, respectively. If \mathbf{A} is diagonalizable, it is possible to pose the approximation problem over these discrete point sets, at the expense of a potentially large C_2 term arising from eigenvector conditioning.

LEMMA 5.2. *Suppose Σ is a subset of the spectrum of \mathbf{A} consisting only of non-defective eigenvalues, and let \mathcal{U} denote the maximal invariant subspace associated with eigenvalues in Σ . If the columns of \mathbf{X} are eigenvectors of \mathbf{A} forming a basis for \mathcal{U} , then*

$$\kappa(\Sigma) \leq \text{cond}_2(\mathbf{X}).$$

(The condition number $\text{cond}_2(\cdot)$ is the ratio of the maximum to the minimum nonzero singular value.)

Proof. Observe that $\mathbf{\Pi} \equiv \mathbf{X}(\mathbf{X}^*\mathbf{X})^{-1}\mathbf{X}^*$ defines an orthogonal projection onto \mathcal{U} , and suppose $\mathbf{\Lambda}$ is a diagonal matrix with entries in Σ such that $\mathbf{A}\mathbf{X} = \mathbf{X}\mathbf{\Lambda}$. Then for any function f that is analytic on Σ , $f(\mathbf{A})\mathbf{X} = \mathbf{X}f(\mathbf{\Lambda})$, and

$$\begin{aligned} \|f(\mathbf{A})\|_{\mathcal{U}} &= \|f(\mathbf{A})\mathbf{X}(\mathbf{X}^*\mathbf{X})^{-1}\mathbf{X}^*\| \\ &= \|\mathbf{X}f(\mathbf{\Lambda})(\mathbf{X}^*\mathbf{X})^{-1}\mathbf{X}^*\| \\ &\leq \|\mathbf{X}\| \|(\mathbf{X}^*\mathbf{X})^{-1}\mathbf{X}^*\| \|f(\mathbf{\Lambda})\| \\ &= \text{cond}_2(\mathbf{X}) \max_{\lambda \in \Sigma} |f(\lambda)|. \quad \square \end{aligned}$$

Now if Ω_{good} and Ω_{bad} in Theorem 3.5 are precisely the sets of good and bad eigenvalues of \mathbf{A} , respectively, Lemma 5.2 leads to a bound on C_2 .

FIRST COROLLARY TO THEOREM 3.5. *To the conditions of Theorem 3.5, add the assumption that \mathbf{A} is diagonalizable,*

$$\mathbf{A}[\mathbf{X}_{\text{good}}, \mathbf{X}_{\text{bad}}] = [\mathbf{X}_{\text{good}}, \mathbf{X}_{\text{bad}}] \text{diag}(\mathbf{\Lambda}_{\text{good}}, \mathbf{\Lambda}_{\text{bad}}).$$

Then

$$(5.1) \quad \delta(\mathcal{U}_{\text{good}}, \mathcal{K}_\ell(\mathbf{A}, \mathbf{v}_1^{(\nu)})) \leq C_0 C_1 \widehat{C}_2 \min_{\phi \in \mathcal{P}_\ell^*} \frac{\max_{j=L+1, \dots, N} |\phi(\lambda_j) \Psi_{\nu p}(\lambda_j)|}{\min_{k=1, \dots, L} |\phi(\lambda_k) \Psi_{\nu p}(\lambda_k)|},$$

where C_0 and C_1 are as defined in Theorems 3.3 and 3.4 and

$$\widehat{C}_2 \equiv \text{cond}_2(\mathbf{X}_{\text{good}}) \text{cond}_2(\mathbf{X}_{\text{bad}}).$$

When \mathbf{A} is far from normal, the constant \widehat{C}_2 will typically be large; it grows infinite as \mathbf{A} tends towards a defective matrix. However, such extreme situations are

not necessarily associated with severe degradation in convergence behavior, and so the bound (5.1) will be most appropriate when \mathbf{A} is either normal or nearly so.

Non-normality can complicate invariant subspace computation in a variety of ways. The good eigenvalues can be individually ill-conditioned, with $\text{cond}_2(\mathbf{X}_{\text{good}}) \gg 1$, while the associated invariant subspace is perfectly conditioned. In other cases, one may find the desired eigenvalues are well-conditioned, while the bad eigenvalues are highly non-normal (as when $\text{cond}_2(\mathbf{X}_{\text{bad}}) \gg \text{cond}_2(\mathbf{X}_{\text{good}}) \approx 1$).⁶ In either case, the good invariant subspace may still have physical significance, and we would like to understand how this ill-conditioning affects the rate at which we can compute it.

Since non-normal matrices are of special interest, consideration of pseudospectra yields a natural approach that often can provide sharper, more descriptive convergence bounds. Recall that the ε -pseudospectrum [45, 46] is the set

$$\Lambda_\varepsilon(\mathbf{A}) \equiv \{z \in \mathbb{C} : \|(z - \mathbf{A})^{-1}\| \geq \varepsilon^{-1}\},$$

or, equivalently, $\Lambda_\varepsilon(\mathbf{A}) = \{z \in \Lambda(\mathbf{A} + \mathbf{E}) : \|\mathbf{E}\| \leq \varepsilon\}$, where $\Lambda(\mathbf{M})$ denotes the set of eigenvalues of a matrix \mathbf{M} .

For a fixed ε , $\Lambda_\varepsilon(\mathbf{A})$ is a closed set in the complex plane consisting of the union of no more than N connected sets, each of which must contain at least one eigenvalue. As $\varepsilon \rightarrow 0$, $\Lambda_\varepsilon(\mathbf{A})$ tends to N disjoint disks (whose radii depend on eigenvalue conditioning and defectiveness) centered at and shrinking around the N distinct eigenvalues.

LEMMA 5.3. *Let \mathcal{U} be an invariant subspace of \mathbf{A} and suppose Σ is the set of eigenvalues associated with \mathcal{U} .*

(a) *Let Ω be a set containing Σ but no eigenvalues of \mathbf{A} outside Σ , and suppose the boundary $\partial\Omega$ is the finite union of positively-oriented Jordan curves. Then*

$$(5.2) \quad \kappa(\Omega) \leq \frac{1}{2\pi} \int_{\partial\Omega} \|(z - \mathbf{A})^{-1}\|_{\mathcal{U}} |dz|.$$

(b) *Let Σ_ε contain the union of those connected components of $\Lambda_\varepsilon(\mathbf{A})$ that include $\lambda \in \Sigma$, and suppose further that Σ_ε contains no eigenvalues outside of Σ and its boundary $\partial\Sigma_\varepsilon$ is the finite union of positively-oriented Jordan curves. Then*

$$(5.3) \quad \kappa(\Sigma_\varepsilon) \leq \frac{\mathcal{L}(\partial\Sigma_\varepsilon)}{2\pi\varepsilon},$$

where $\mathcal{L}(\partial\Sigma_\varepsilon)$ is the length of the boundary of Σ_ε .

Proof. For part (a), let $\mathbf{\Pi}$ be the orthogonal projector onto the given invariant subspace \mathcal{U} and let \mathbf{P} be the spectral projector for \mathbf{A} associated with \mathcal{U} . For any function f analytic on Ω , $\|f(\mathbf{A})\|_{\mathcal{U}} = \|f(\mathbf{A})\mathbf{\Pi}\| = \|f(\mathbf{A})\mathbf{P}\mathbf{\Pi}\| \leq \|f(\mathbf{A})\mathbf{P}\|$. Now,

$$f(\mathbf{A})\mathbf{P} = \frac{1}{2\pi i} \int_{\partial\Omega} f(z)(z - \mathbf{A})^{-1} dz.$$

Thus for any vector $\mathbf{x} \in \mathcal{U}$,

$$\begin{aligned} \|f(\mathbf{A})\mathbf{x}\| &\leq \frac{1}{2\pi} \int_{\partial\Omega} |f(z)| \|(z - \mathbf{A})^{-1}\mathbf{x}\| |dz| \\ &\leq \left(\frac{1}{2\pi} \int_{\partial\Omega} \|(z - \mathbf{A})^{-1}\|_{\mathcal{U}} |dz| \right) \max_{z \in \partial\Omega} |f(z)| \|\mathbf{x}\|. \end{aligned}$$

⁶This is the case for the Markov chain example described in §6.4. Trefethen describes another example, the Gauss–Seidel iteration matrix for the centered difference discretization of the second derivative [46, Example 10].

But since f is analytic on Ω , $\max_{z \in \partial\Omega} |f(z)| = \max_{z \in \Omega} |f(z)|$. Part (b) follows from (a) by assigning $\Omega = \Sigma_\varepsilon$. \square

Pseudospectral bounds were developed by Trefethen to bound the GMRES residual norm [45], and Simoncini has used a similar approach to analyze block-Arnoldi convergence [39]. In the single eigenvector case, her Theorem 3.1 closely resembles our (5.6) below. (Lemma 5.3 could easily be sharpened to instead involve $\Lambda_\varepsilon(\mathbf{U}^* \mathbf{A} \mathbf{U})$, where the columns of \mathbf{U} form an orthonormal basis for $\mathcal{U}_{\text{good}}$; note that $\Lambda_\varepsilon(\mathbf{U}^* \mathbf{A} \mathbf{U}) \subseteq \Lambda_\varepsilon(\mathbf{A})$ [43].)

The pseudospectral approach leads to a robust alternative to the eigenvector-based bound (5.1).⁷ Suppose ε is sufficiently small that the components of the ε -pseudospectrum enclosing the good eigenvalues are disjoint from those components enclosing the bad eigenvalues. $\Lambda_\varepsilon(\mathbf{A})$ can then be contained in the two disjoint sets $\Sigma_\varepsilon^{\text{good}}$ and $\Sigma_\varepsilon^{\text{bad}}$, leading to an alternative bound.

SECOND COROLLARY TO THEOREM 3.5. *Assume the conditions of Theorem 3.5 and suppose that $\varepsilon > 0$ is sufficiently small that $\Sigma_\varepsilon^{\text{good}} \cap \Sigma_\varepsilon^{\text{bad}} = \emptyset$. Then, provided $\Psi_{\nu p}(z)$ has no roots in $\Sigma_\varepsilon^{\text{good}}$, and the boundaries of $\Sigma_\varepsilon^{\text{good}}$ and $\Sigma_\varepsilon^{\text{bad}}$ are finite unions of positively-oriented Jordan curves,*

$$(5.4) \quad \delta(\mathcal{U}_{\text{good}}, \mathcal{K}_\ell(\mathbf{A}, \mathbf{v}_1^{(\nu)})) \leq C_0 C_1 \tilde{C}_2(\varepsilon) \min_{\phi \in \mathcal{P}_{\ell^*}} \frac{\max\{|\phi(z) \Psi_{\nu p}(z)| : z \in \Sigma_\varepsilon^{\text{bad}}\}}{\min\{|\phi(z) \Psi_{\nu p}(z)| : z \in \Sigma_\varepsilon^{\text{good}}\}},$$

where C_0 and C_1 are as defined in Theorems 3.3 and 3.4, and

$$(5.5) \quad \tilde{C}_2(\varepsilon) \equiv \frac{\mathcal{L}(\partial \Sigma_\varepsilon^{\text{good}}) \mathcal{L}(\partial \Sigma_\varepsilon^{\text{bad}})}{4\pi^2 \varepsilon^2}.$$

$\mathcal{L}(\partial \Sigma_\varepsilon^{\text{good}})$ and $\mathcal{L}(\partial \Sigma_\varepsilon^{\text{bad}})$ are the boundary lengths of $\Sigma_\varepsilon^{\text{good}}$ and $\Sigma_\varepsilon^{\text{bad}}$, respectively.

This pseudospectral bound holds for a range of ε -values, providing a natural mechanism for adjusting the sets Ω_{good} and Ω_{bad} . As ε gets smaller, $\tilde{C}_2(\varepsilon)$ generally increases, but the convergence rate induced by the polynomial approximation problem improves, since the sets on which the approximation problem is posed recede from one another. For the most descriptive convergence bound, take the envelope of individual bounds corresponding to a variety of ε -values; see Figures 6.1 and 6.3. Of course, the bound (5.4) is only meaningful when ε is sufficiently small that $\Sigma_\varepsilon^{\text{good}} \cap \Sigma_\varepsilon^{\text{bad}} = \emptyset$. The need to take ε particularly small to satisfy this condition may signal an ill-conditioned problem; consider enlarging the set of good eigenvalues.

In some situations, one may wish to use different values of ε for the good and bad pseudospectra, in which case (5.4) changes in the obvious way. Furthermore, when the good eigenvalues are normal (i.e., one can take $\text{cond}_2(\mathbf{X}_{\text{good}}) = 1$), it is best to combine the pseudospectra and eigenvector approaches to obtain

$$(5.6) \quad \delta(\mathcal{U}_{\text{good}}, \mathcal{K}_\ell(\mathbf{A}, \mathbf{v}_1^{(\nu)})) \leq \frac{C_0 C_1 \mathcal{L}(\Sigma_\varepsilon^{\text{bad}})}{2\pi\varepsilon} \min_{\phi \in \mathcal{P}_{\ell^*}} \frac{\max\{|\phi(z) \Psi_{\nu p}(z)| : z \in \Sigma_\varepsilon^{\text{bad}}\}}{\min_{k=1, \dots, L} |\phi(\lambda_k) \Psi_{\nu p}(\lambda_k)|}.$$

We close this section by pointing out one non-normal situation where the eigenvector-based bound (5.1) can be dramatically superior to the pseudospectral bound (5.4).

⁷Note that Greenbaum has demonstrated how more clever use of eigenvector information can sometimes be superior to estimating integrals of the resolvent norm [15].

Suppose for simplicity that $\dim \mathcal{U}_{\text{good}} = \dim \mathcal{U}_{\text{bad}}$ with $\mathcal{U}_{\text{good}} \approx \mathcal{U}_{\text{bad}}$ for some diagonalizable \mathbf{A} . It is possible for the basis vectors in \mathbf{X}_{good} and \mathbf{X}_{bad} to be perfectly conditioned on their own, but terribly conditioned if taken together, e.g.,

$$\mathbf{X}_{\text{good}} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \\ 0 & 0 \\ 0 & 0 \end{bmatrix}, \quad \mathbf{X}_{\text{bad}} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \\ \gamma & 0 \\ 0 & \gamma \end{bmatrix},$$

with $0 < |\gamma| \ll 1$. This results in $\widehat{C}_2 = 1$ but $\widetilde{C}_2(\varepsilon) \gg 1$ for usefully small values of ε . (This can be remedied by considering the pseudospectra of \mathbf{A} orthogonally projected onto $\mathcal{U}_{\text{good}}$ and \mathcal{U}_{bad} .) What is happening here? The more alike $\mathcal{U}_{\text{good}}$ and \mathcal{U}_{bad} are, the more prominent their general orientation is in the Krylov subspace, possibly resulting in an initial period of rapid sublinear convergence. Discriminating the fine difference between $\mathcal{U}_{\text{good}}$ and \mathcal{U}_{bad} may still be challenging.

6. Some Examples. How well does the machinery constructed in the previous sections work? Here we demonstrate our bounds for a variety of examples. These test problems are contrived to illustrate the effects we have described as cleanly as possible. Eigenvalue problems from applications inevitably involve more complicated spectral structure.

6.1. Influence of Non-normality on Predicted Rates. We begin with two examples involving non-diagonalizable matrices where pseudospectral convergence bounds can be used to good effect. (While the examples in this subsection and the next are defective, we emphasize that the pseudospectral bound can also be useful for diagonalizable matrices with large values of \widehat{C}_2 .) Define

$$(6.1) \quad \mathbf{A} = \begin{bmatrix} \mathbf{D}_{\text{good}} & \mathbf{0} \\ \mathbf{0} & \mathbf{J}_{58}(-1) \end{bmatrix},$$

where \mathbf{D}_{good} is a 6×6 diagonal matrix containing good eigenvalues uniformly distributed in $[1, 2]$, and $\mathbf{J}_{58}(-1)$ is a Jordan block of dimension 58 with the bad eigenvalue $\lambda = -1$ on the main diagonal and ones on the first superdiagonal. Note that $\mathcal{U}_{\text{good}} \perp \mathcal{U}_{\text{bad}}$, so $C_0 = 1$. Since the good eigenvalues are normal, we apply the hybrid pseudospectral bound (5.6). The pseudospectra of a direct sum of matrices is the union of the pseudospectra of each component matrix [48], so we need only focus on the pseudospectra of the Jordan block, which are circular disks for all $\varepsilon > 0$ [32]; see Figure 6.1. It follows that $\widetilde{C}_2(\varepsilon) = r_\varepsilon/\varepsilon$, where r_ε is the radius of $\Sigma_\varepsilon^{\text{bad}} = \Lambda_\varepsilon(\mathbf{J}_{58}(-1))$, determined numerically. For $\phi \in \mathcal{P}_{\ell^*}$ we take the Chebyshev polynomial for $\Sigma_\varepsilon^{\text{bad}}$, $\phi(z) = (z + 1)^{\ell^*}$. For all ε such that $r_\varepsilon < 2$, (5.6) gives

$$(6.2) \quad \delta(\mathcal{U}_{\text{good}}, \mathcal{K}_\ell(\mathbf{A}, \mathbf{v}_1)) \leq \frac{C_1 r_\varepsilon}{\varepsilon} \left(\frac{r_\varepsilon}{2}\right)^{\ell^*},$$

where we have used the fact that $|\phi(\lambda)| \geq 2$ for all good eigenvalues λ . The convergence curve and corresponding bounds are shown in Figure 6.1 for the starting vector \mathbf{v}_1 with $1/\sqrt{n}$ in each component; no restarting is performed. Interestingly, for small values of ε the bound (5.6) accurately captures the finite termination that must occur when $\ell = n = 64$, a trait exhibited by pseudospectral bounds in other contexts.




FIG. 6.1. On the left, good eigenvalues (\times) and pseudospectral boundaries $\partial\Sigma_\varepsilon^{\text{bad}}$ for $\varepsilon = 10^{-2}$, 10^{-5} , 10^{-15} , and 10^{-100} , where \mathbf{A} is given by (6.1). (The bad eigenvalue (\cdot) is obscured by the $\varepsilon = 10^{-100}$ boundary.) On the right, gap convergence (solid line) together with the bound (6.2) (dotted lines) for each of the pseudospectral curves shown on the left. For small values of ε , (6.2) captures the finite termination that must occur at the 64th iteration.

Our second example is the same, except the good eigenvalues are now replaced with a Jordan block,

$$(6.3) \quad \mathbf{A} = \begin{bmatrix} \mathbf{J}_6(\frac{3}{2}) & \mathbf{0} \\ \mathbf{0} & \mathbf{J}_{58}(-1) \end{bmatrix},$$

where $\mathbf{J}_6(\frac{3}{2})$ is a 6×6 Jordan block with $\frac{3}{2}$ on the main diagonal and ones on the first superdiagonal; $\mathbf{J}_{58}(-1)$ is as before. Again note that $\mathcal{U}_{\text{good}} \perp \mathcal{U}_{\text{bad}}$, implying $C_0 = 1$. Since both the good and bad eigenvalues are defective, apply the pseudospectral bound (5.4). Recalling the pseudospectra of Jordan blocks are circular disks, let $r_\varepsilon^{\text{bad}}$ and $r_\varepsilon^{\text{good}}$ denote the radii of $\Sigma_\varepsilon^{\text{bad}} = \Lambda_\varepsilon(\mathbf{J}_{58}(-1))$ and $\Sigma_\varepsilon^{\text{good}} = \Lambda_\varepsilon(\mathbf{J}_6(\frac{3}{2}))$, respectively; see the left plot of Figure 6.2. The Second Corollary to Theorem 3.5 holds whenever $r_\varepsilon^{\text{bad}} + r_\varepsilon^{\text{good}} < \frac{5}{2}$. For such ε , $\tilde{C}_2(\varepsilon) = r_\varepsilon^{\text{bad}} r_\varepsilon^{\text{good}} / \varepsilon^2$ and

$$(6.4) \quad \delta(\mathcal{U}_{\text{good}}, \mathcal{K}_\ell(\mathbf{A}, \mathbf{v}_1)) \leq C_1 \frac{r_\varepsilon^{\text{bad}} r_\varepsilon^{\text{good}}}{\varepsilon^2} \left(\frac{r_\varepsilon^{\text{bad}}}{\frac{5}{2} - r_\varepsilon^{\text{good}}} \right)^{\ell^*},$$

where again we have taken for $\phi \in \mathcal{P}_{\ell^*}$ the Chebyshev polynomial for $\Sigma_\varepsilon^{\text{bad}}$, $\phi(z) = (z+1)^{\ell^*}$. The convergence curve and corresponding bounds are shown in Figure 6.2 for the starting vector \mathbf{v}_1 with $1/\sqrt{n}$ in each component; no restarting is performed.

6.2. Superlinear Effects Due to Non-normality. Our final example of pseudospectral bounds addresses the matrix

$$(6.5) \quad \mathbf{A} = \begin{bmatrix} 0 & \mathbf{0} \\ \mathbf{0} & \mathbf{F} \end{bmatrix},$$

where there is a single good eigenvalue $\lambda = 0$ (with multiplicity one) and a bad eigenvalue $\lambda = -\frac{1}{3}$ associated with the 63×63 bidiagonal matrix \mathbf{F} , which has $-\frac{1}{3}$ in

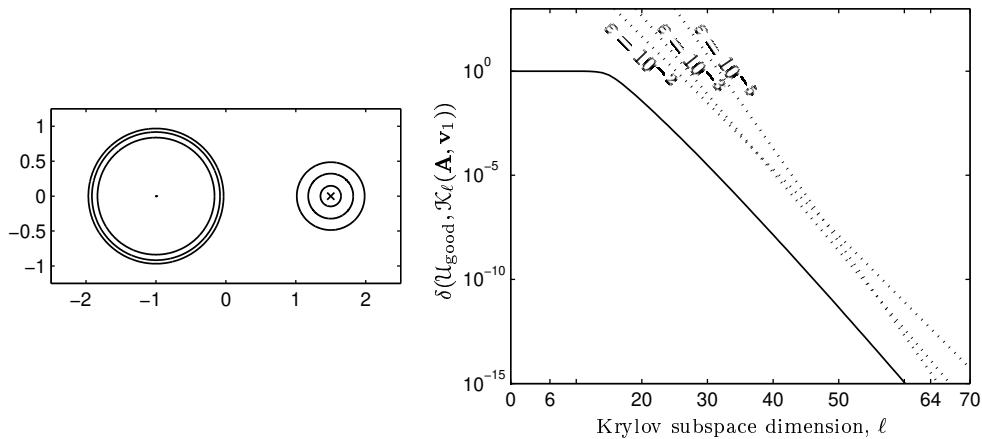


FIG. 6.2. On the left, bad eigenvalue (\cdot), good eigenvalue (\times), and pseudospectral boundaries $\partial\Sigma_\epsilon^{\text{bad}}$ and $\partial\Sigma_\epsilon^{\text{good}}$ for \mathbf{A} given by (6.3) and $\epsilon = 10^{-2}$, 10^{-3} , and 10^{-5} . On the right, gap convergence (solid line) with the bound (6.4) (dotted lines) for the three ϵ values used in the left plot.

the main diagonal entries and $1/j$ in the $(j, j + 1)$ entry of the superdiagonal. Like the Jordan blocks described before, the pseudospectra of \mathbf{F} are circular disks [32], but the radii of these disks shrink much more rapidly as ϵ decreases than observed for the Jordan block. As a result, the convergence rate steadily improves as ϵ gets smaller; this is compensated by growing $\tilde{C}_2(\epsilon)$ values. Taking $\phi(z) = (z + \frac{1}{3})^{\ell^*}$, we obtain

$$(6.6) \quad \delta(\mathcal{U}_{\text{good}}, \mathcal{K}_\ell(\mathbf{A}, \mathbf{v}_1)) \leq \frac{C_1 r_\epsilon}{\epsilon} (3r_\epsilon)^{\ell^*},$$

provided $r_\epsilon < \frac{1}{3}$, where r_ϵ is the radius of $\Sigma_\epsilon^{\text{bad}}$. Figure 6.3 shows the spectrum of \mathbf{A} and pseudospectra of \mathbf{F} . As ϵ gets smaller, the bound (6.6) traces out an envelope that predicts early stagnation followed by improving linear convergence rates. This is “superlinear” convergence, but of a different nature from that described in §4.3. Figure 6.3 shows these bounds along with the gap convergence curve for a vector \mathbf{v}_1 with real entries drawn from the standard normal distribution. Pseudospectral bounds for GMRES exhibit similar superlinear behavior for matrices like \mathbf{F} [10, 12]. Although all the examples here have used defective matrices, these bounds are also appropriate for diagonalizable matrices with a large eigenvector condition number.

6.3. Shift Selection for Restarted Algorithms. The results of §4 indicate that effective restart strategies can be constructed using optimal polynomials associated with sets containing the bad eigenvalues. In this section, we give some examples of how choices for $\Psi_{\nu,p}$ based on partial information (or misinformation) about bad eigenvalue location affects the observed convergence rates and illustrate how well our bounds can predict this.

Consider the 200×200 upper triangular matrix

$$\mathbf{A} = \begin{bmatrix} \mathbf{D}_{\text{good}} & \mathbf{C} \\ \mathbf{0} & \mathbf{D}_{\text{bad}} \end{bmatrix},$$

where \mathbf{D}_{good} is a 16×16 diagonal matrix of good eigenvalues, distributed uniformly around the circle in the complex plane centered at 3 with radius 1; \mathbf{D}_{bad} is a diagonal




FIG. 6.3. On the left, bad eigenvalue (\cdot), good eigenvalue (\times), and pseudospectral boundaries $\partial\Sigma_\varepsilon^{\text{bad}}$ for \mathbf{A} given by (6.5) and $\varepsilon = 10^{-2}, \dots, 10^{-12}$. On the right, gap convergence (solid line) with the bound (6.4) (dotted lines) for the eleven ε values shown in the left plot.

matrix containing the bad eigenvalues distributed uniformly along the line segment (designated \mathcal{J}_{bad}) parallel to the imaginary axis connecting the points $-1 \pm 5i$; \mathbf{C} is a full (row) rank matrix scaled so that $\|\mathbf{P}_{\text{good}}\| \approx 1000$. The starting vector, \mathbf{v}_1 , has normally distributed random complex entries. (The same \mathbf{v}_1 was used for all experiments shown in this subsection.)

Figure 6.4 compares the predicted and observed convergence curves for the un-restarted iteration, where the Krylov subspace grows without bound. The left plot displays the equipotentials of $g(z, \mathcal{J}_{\text{bad}})$ —the physical analog is the potential field generated by a continuous (line) charge distribution spread over \mathcal{J}_{bad} . The color bar is calibrated to show $\exp(-g(z, \mathcal{J}_{\text{bad}}))$, giving the predicted convergence rates at locations in the complex plane if good eigenvalues were present there. In particular, the lowest equipotential contour passing through a good eigenvalue is shown; it leads via (4.6) to a predicted convergence rate of ≈ 0.566 . The right plot shows the iteration history of $\delta(\mathcal{U}_{\text{good}}, \mathcal{K}_\ell(\mathbf{A}, \mathbf{v}_1))$ versus the iteration index ℓ . After an early sublinear surge that flattens out near $1/\|\mathbf{P}_{\text{good}}\|$, an observed linear rate of ≈ 0.539 emerges. In separate experiments (not shown), we have varied the magnitude of $\|\mathbf{C}\|$ (in effect changing $\|\mathbf{P}_{\text{good}}\|$) and have observed variations in the sublinear stagnation level roughly proportional to $1/\|\mathbf{P}_{\text{good}}\|$, consistent with the discussion surrounding Figure 2.1. The convergence bound is derived from the First Corollary to Theorem 3.5, using for ϕ Chebyshev polynomials for \mathcal{J}_{bad} . (For all experiments in this subsection, $C_0 = \sqrt{2}$, $C_1 \approx 4.4325 \times 10^{11}$, $\widehat{C}_2 \approx 1.2439 \times 10^3$.)

Figure 6.5 shows results for polynomial restarts using fast Leja points [3] associated with \mathcal{J}_{bad} . These appear as a dense line of white dots atop the black band of bad eigenvalues. The base dimension is 20 and restarts are each of order 5. (The Krylov subspace dimension never exceeds 25.) The left plot displays the effective potential, $g(z, \Omega_{\text{bad}})$, generated by 180 fast Leja points— Ω_{bad} is the smallest polynomial lemniscate generated by the aggregate filter polynomial that contains all bad eigenvalues. The lowest equipotential contour passing through a good eigenvalue is shown; it leads

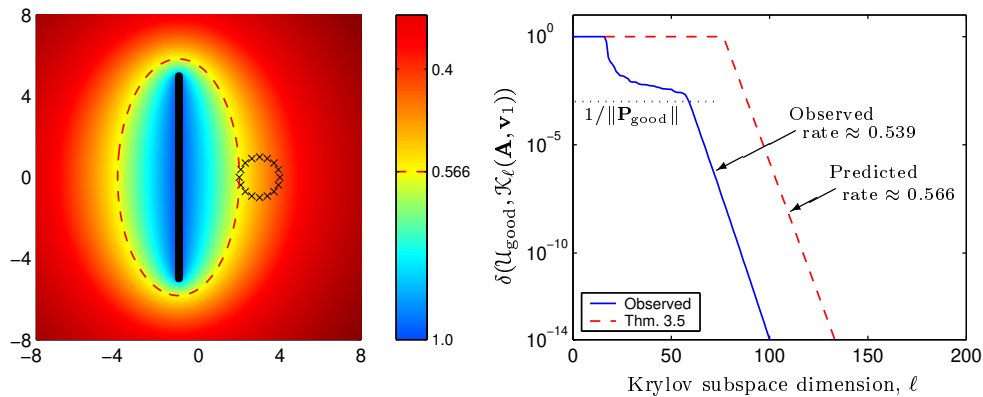


FIG. 6.4. *Unrestarted Subspace.* On the left, good and bad eigenvalues are shown in the “potential field” generated by the bad eigenvalues. The colorbar is calibrated to show effective convergence rates for different components of $\mathcal{U}_{\text{good}}$. The right plot shows the observed gap history (solid line) together with a bound (dashed) derived from the First Corollary to Theorem 3.5.

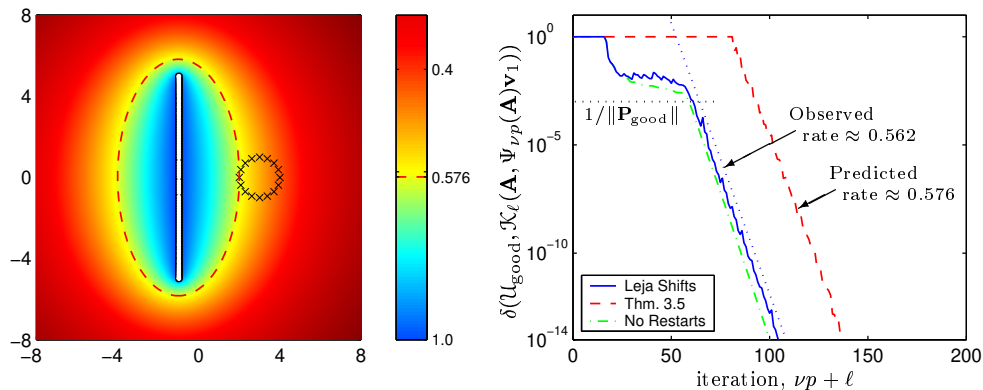


FIG. 6.5. *Polynomial restarts at fast Leja points of \mathcal{J}_{bad} (white dots).* The base dimension is 20 and restarts are each of degree $p = 5$ (so the subspace dimension never exceeds 25).

via (4.6) and Example 4.2 to a predicted convergence rate of ≈ 0.576 . The bound on the right was obtained from the First Corollary to Theorem 3.5, using Chebyshev polynomials for \mathcal{J}_{bad} up to the base dimension, then including the shift polynomials.

The next two figures show the effect of poorer choices for the filter shifts. Suppose we mistakenly believe the bad eigenvalues to be concentrated toward the ends of the interval \mathcal{J}_{bad} and choose filter shifts accordingly grouped in two subintervals that omit the central portion of \mathcal{J}_{bad} (which we believe to be devoid of bad eigenvalues). We use fast Leja points again but this time for pairs of disjoint intervals that in fact cover only 60% and 20%, respectively, of the bad eigenvalues. These are asymptotically optimal filter shifts for misguided guesses of the bad eigenvalue distribution. Ω_{bad} is again the smallest polynomial lemniscate generated by 180 fast Leja points that contains all bad eigenvalues. Here it takes on a more pronounced dumb-bell appearance, reflecting the absence of zeros from the middle of \mathcal{J}_{bad} . As before, the base dimension is 20 and restarts are each of order 5. The convergence rate is seen to deteriorate to

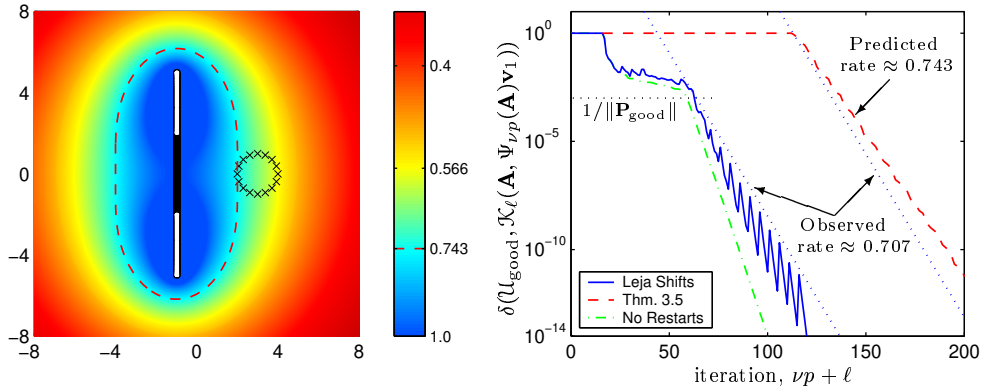


FIG. 6.6. Polynomial restarts with fast Leja points (white dots) for two subintervals covering only 60% of the bad eigenvalues. The subspace dimensions are as in Figure 6.5.

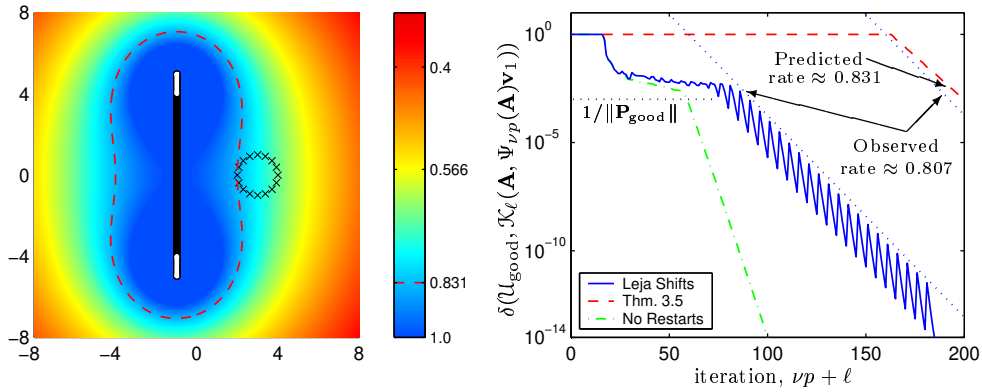


FIG. 6.7. Polynomial restarts with fast Leja points (white dots) for two subintervals covering only 20% of the bad eigenvalues. The subspace dimensions are as in Figure 6.5.

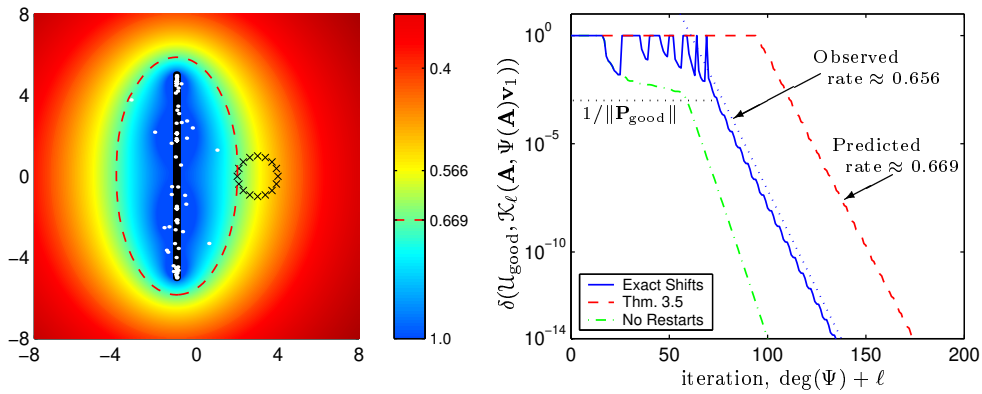


FIG. 6.8. Polynomial restarts using exact shifts (white dots) determined by choosing Ritz values with real part smaller than 1. The subspace dimension never exceeds 20.

≈ 0.707 and ≈ 0.807 , respectively, and is predicted to within an accuracy of roughly 3%–5.2%. By comparing the equipotential contours of Figures 6.4 and 6.5 with those of Figures 6.6 and 6.7, notice the filter shifts in the latter cases create a potential significantly different from what either the bad eigenvalues or optimal filter shifts would generate. Figure 6.8 shows the result of using Sorensen’s exact shifts. The subspace dimension is limited to be no larger than 20, and a Ritz value is used as a shift if it has real part smaller than 1. (The early convergence plateaus occur when the subspace is compressed to have dimension smaller than the number of good eigenvalues.) The potential plot on the left is based on 180 exact shifts. Although these shifts fall outside the convex hull of the bad eigenvalues, they effectively recover the potential generated by those eigenvalues. The convergence rate is predicted to within 2% of the observed rate. The use of exact shifts yields a convergence rate within 25% of the rate for the unrestarted iteration (Figure 6.4) at a lower computational cost and without requiring *a priori* localization of bad eigenvalues to determine optimal shifts (as in Figure 6.5 for good localization and Figures 6.6 and 6.7 for poor localization).

6.4. Markov Chain Example. We close by examining a more realistic eigenvalue problem, taking \mathbf{A} to be the transition matrix for a Markov chain that describes a random walk on a triangular lattice. See Saad [37, §II.5.1] for details of this example, a common test problem for iterative eigenvalue algorithms. Since all the rows of a transition matrix sum to one, \mathbf{A} must have an eigenvalue $\lambda = 1$, and the Perron–Frobenius theorem assures this eigenvalue is simple (see, e.g., [5, Thm. 1.4]). The *left* eigenvector corresponding to $\lambda = 1$ determines a stationary distribution of the Markov chain, so we are interested in the convergence of $\delta(\mathcal{U}_{\text{good}}, \mathcal{K}_\ell(\mathbf{A}^*, \mathbf{v}_1))$, where $\mathcal{U}_{\text{good}}$ is the invariant subspace of \mathbf{A}^* for $\lambda = 1$. Here we consider a lattice with a base and height of 50 nodes, yielding a transition matrix of dimension $n = 1275$. This matrix exhibits a significant degree of non-normality, mostly associated with ill-conditioned eigenvalues far from $\lambda = 1$, as one can infer from the pseudospectra illustrated in Figure 6.9. Unlike the previous examples in this section, the good eigenvalue is quite close to bad eigenvalues, as highlighted by the close-up on the right of Figure 6.9.

The eigenvalues of \mathbf{A} appear to be real with $\lambda = 0$ having algebraic and geometric




FIG. 6.9. Eigenvalues and pseudospectra for the random walk transition matrix for a triangular lattice with 1275 nodes. The left plot shows the spectrum and boundaries of ϵ -pseudospectra for $\epsilon = 10^{-1}, \dots, 10^{-8}$. The right plot zooms around $\lambda = 1$, indicating ϵ -pseudospectra for $\epsilon = 10^{-2}, 10^{-3}$.




FIG. 6.10. Gap convergence for the random walk example, $n = 1275$ (solid line). The dashed lines represent the bound (5.1). The best result is obtained when the bad eigenvalues are treated as a discrete point set for the approximation problem, while a slower rate is predicted when the bad eigenvalues are treated as an interval. The dotted lines utilize the superlinear bounds of Theorem 4.7 for $r = 1, \dots, 10$.

multiplicity 25. (Though we formally stipulate that \mathbf{A} be non-derogatory in §3, our proofs only require the good eigenvalues be non-derogatory.) The bound (5.1) based on the conditioning of the matrices of good and bad eigenvectors is simplest to evaluate. We have $C_0 = \sqrt{2}$, and compute $\widehat{C}_2 \approx 3.546 \times 10^9$; for a particular starting vector with normally distributed real random entries, $C_1 \approx 9.933$. Labeling the eigenvalues from right to left, the polynomial approximation problem in (5.1) reduces in this single eigenvector case to a minimax approximation on $\Lambda_{\text{bad}} = \{\lambda_2, \dots, \lambda_n\}$ subject to normalization at $\lambda_1 = 1$. Bounding this approximation problem using Chebyshev polynomials on $[\lambda_n, \lambda_2]$ gives a pessimistic result, as can be seen in the convergence plot in Figure 6.10. The superlinear bounds of Theorem 4.7 yield a marked improvement. In the language of Theorem 4.7, we take $\Omega_k = \{\lambda_j\}_{j=k+1}^n$ and reduce to an approximation problem over Ω_{r+1} for $r = 1, \dots, 10$, for which we use Chebyshev polynomials on $[\lambda_n, \lambda_r]$. An even better bound is obtained by treating Λ_{bad} completely as a discrete point set. One approachable way of doing this is to take $\Lambda_{\text{good}} = \{\lambda_1\}$ and note that

$$(6.7) \quad \min_{\phi \in \mathcal{P}_{\ell^*}} \frac{\max\{|\phi(\lambda)| : \lambda \in \Lambda_{\text{bad}}\}}{\min\{|\phi(\lambda)| : \lambda \in \Lambda_{\text{good}}\}} = \min_{\substack{\phi \in \mathcal{P}_{\ell^*} \\ \phi(\lambda_1)=1}} \max_{\lambda \in \Lambda_{\text{bad}}} |\phi(\lambda)| \leq \min_{\substack{\phi \in \mathcal{P}_{\ell^*} \\ \phi(0)=1}} \|\phi(\mathbf{S})\mathbf{r}\|,$$

where $\mathbf{S} = \text{diag}(\lambda_2 - \lambda_1, \dots, \lambda_n - \lambda_1)$ and $\mathbf{r} = [1, 1, \dots, 1]^T$. The right hand side of (6.7) can be computed as the residual norm of the GMRES algorithm applied to the matrix \mathbf{S} with initial residual \mathbf{r} ; this is no more than a factor of \sqrt{n} worse than the left hand side of (6.7). The resultant bound is shown in Figure 6.10. Alternatively, the minimax problem on the left hand side of (6.7) could be solved directly via a linear program.

Acknowledgements. We thank Dan Sorensen for many constructive comments during the early stages of this work, and Joe Ball, Nick Trefethen and Henk van der Vorst for helpful discussions. The pseudospectral computations in Section 6 were based on software developed by Trefethen and Wright [47, 52].

REFERENCES

- [1] M. ARIOLI, V. PTÁK, AND Z. STRAKOŠ, *Krylov sequences of maximal length and convergence of GMRES*, BIT, 38 (1998), pp. 636–643.
- [2] W. E. ARNOLDI, *The principle of minimized iterations in the solution of the matrix eigenvalue problem*, Quart. Appl. Math., 9 (1951), pp. 17–29.
- [3] J. BAGLAMA, D. CALVETTI, AND L. REICHEL, *Fast Leja points*, Elec. Trans. Numer. Anal., 7 (1998), pp. 124–140.
- [4] Z. BAI, J. DEMMEL, J. DONGARRA, A. RUHE, AND H. VAN DER VORST, eds., *Templates for the Solution of Algebraic Eigenvalue Problems: A Practical Introduction*, SIAM, Philadelphia, 2000.
- [5] A. BERMAN AND R. J. PLEMMONS, *Nonnegative Matrices in the Mathematical Sciences*, SIAM, Philadelphia, 1994.
- [6] D. CALVETTI, L. REICHEL, AND D. C. SORENSEN, *An implicitly restarted Lanczos method for large symmetric eigenvalue problems*, Elec. Trans. Numer. Anal., 2 (1994), pp. 1–21.
- [7] F. CHATELIN, *Eigenvalues of Matrices*, Wiley, Chichester, 1993.
- [8] J. B. CONWAY, *Functions of One Complex Variable*, Springer-Verlag, New York, second ed., 1978.
- [9] G. DE SAMBLANX AND A. BULTHEEL, *Nested Lanczos: implicitly restarting an unsymmetric Lanczos algorithm*, Numer. Algorithms, 18 (1998), pp. 31–50.
- [10] 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.
- [11] N. DUNFORD AND J. SCHWARTZ, *Linear Operators, Part I: General Theory*, Wiley, New York, 1971.
- [12] M. EMBREE, *How descriptive are GMRES convergence bounds?*, Tech. Rep. 99/08, Numerical Analysis Group, Oxford University Computing Laboratory, 1999.
- [13] F. R. GANTMACHER, *Theory of Matrices*, vol. 1, Chelsea, New York, 2nd ed., 1959.
- [14] G. H. GOLUB AND C. F. VAN LOAN, *Matrix Computations*, Johns Hopkins, Baltimore, third ed., 1996.
- [15] A. GREENBAUM, *Using the Cauchy integral formula and the partial fractions decomposition of the resolvent to estimate $\|f(A)\|$* , submitted, 2000.
- [16] E. J. GRIMME, D. C. SORENSEN, AND P. VAN DOOREN, *Model reduction of state space systems via an implicitly restarted Lanczos method*, Numer. Algorithms, 12 (1995), pp. 1–31.
- [17] M. HAVIV AND Y. RITOV, *Bounds on the error of an approximate invariant subspace for non-self-adjoint matrices*, Numer. Math., 67 (1994), pp. 491–500.
- [18] V. HEUVELINE AND M. SADKANE, *Arnoldi-Faber method for large nonhermitian eigenvalue problems*, Elec. Trans. Numer. Anal., 5 (1997), pp. 62–76.
- [19] R. A. HORN AND C. R. JOHNSON, *Topics in Matrix Analysis*, Cambridge University Press, Cambridge, 1991.
- [20] Z. JIA, *The convergence of generalized Lanczos methods for large unsymmetric eigenproblems*, SIAM J. Matrix Anal. Appl., 16 (1995), pp. 843–862.
- [21] Z. JIA AND G. W. STEWART, *An analysis of the Rayleigh–Ritz method for approximating eigenspaces*, Math. Comp, (2000).
- [22] T. KATO, *Estimation of iterated matrices, with application to the von Neumann condition*, Numer. Math., 2 (1960), pp. 22–29.
- [23] ———, *Perturbation Theory for Linear Operators*, Springer-Verlag, Berlin, corrected second ed., 1980.
- [24] L. KNIZHNERMAN, *Error bounds for the Arnoldi method: a set of extreme eigenpairs*, Linear Algebra Appl., 296 (1999), pp. 191–211.
- [25] C. LANCZOS, *An iteration method for the solution of the eigenvalue problem of linear differential and integral operators*, J. Res. Nat. Bur. Stand., 45 (1950), pp. 255–282.
- [26] R. B. LEHOUCQ, *Implicitly restarted Arnoldi methods and subspace iteration*. To appear in *SIAM J. Matrix Anal.*

- [27] A. L. LEVIN AND E. B. SAFF, *Optimal ray sequences of rational functions connected with the Zolotarev problem*, *Constr. Approx.*, 10 (1994), pp. 235–273.
- [28] R. B. MORGAN, *On restarting the Arnoldi method for large nonsymmetric eigenvalue problems*, *Math. Comp.*, 65 (1996), pp. 1213–1230.
- [29] O. NEVANLINNA, *Convergence of Iterations for Linear Equations*, Birkhäuser, Basel, 1993.
- [30] B. N. PARLETT, D. R. TAYLOR, AND Z. A. LIU, *A look-ahead Lanczos algorithm for unsymmetric matrices*, *Math. Comp.*, 44 (1985), pp. 105–124.
- [31] V. I. PAULSEN, *Completely Bounded Maps and Dilations*, Longman Scientific and Technical, Harlow, Essex, 1986.
- [32] L. REICHEL AND L. N. TREFETHEN, *Eigenvalues and pseudo-eigenvalues of Toeplitz matrices*, *Linear Algebra Appl.*, 162–164 (1992), pp. 153–185.
- [33] F. RIESZ AND B. SZ.-NAGY, *Functional Analysis*, Frederick Ungar, New York, 1955.
- [34] Y. SAAD, *Variations on Arnoldi's method for computing eigenvalues of large unsymmetric matrices*, *Linear Algebra Appl.*, 34 (1980), pp. 269–295.
- [35] ———, *Projection methods for solving large sparse eigenvalue problems*, in *Matrix Pencils: Proceedings, Pite Havsbad, 1982*, B. Kågström and A. Ruhe, eds., vol. 973 of *Lecture Notes in Mathematics*, Berlin, 1983, Springer-Verlag, pp. 121–144.
- [36] ———, *Chebyshev acceleration techniques for solving nonsymmetric eigenvalue problems*, *Math. Comp.*, 42 (1984), pp. 567–588.
- [37] ———, *Numerical Methods for Large Eigenvalue Problems*, Manchester University Press, Manchester, 1992.
- [38] E. B. SAFF AND V. TOTIK, *Logarithmic Potentials with External Fields*, Springer-Verlag, Berlin, 1997.
- [39] V. SIMONCINI, *Ritz and pseudo-Ritz values using matrix polynomials*, *Linear Algebra Appl.*, 241–243 (1996), pp. 787–801.
- [40] D. C. SORENSEN, *Implicit application of polynomial filters in a k-step Arnoldi method*, *SIAM J. Matrix. Anal. Appl.*, 13 (1992), pp. 357–385.
- [41] G. W. STEWART, *Error and perturbation bounds for subspaces associated with certain eigenvalue problems*, *SIAM Review*, 15 (1973), pp. 727–764.
- [42] G. W. STEWART AND J.-G. SUN, *Matrix Perturbation Theory*, Academic Press, Boston, 1990.
- [43] K.-C. TOH AND L. N. TREFETHEN, *Calculation of pseudospectra by the Arnoldi iteration*, *SIAM J. Sci. Comput.*, 17 (1996), pp. 1–15.
- [44] ———, *The Kreiss matrix theorem on a general complex domain*, *SIAM J. Matrix Anal. Appl.*, 21 (1999), pp. 145–165.
- [45] L. N. TREFETHEN, *Approximation theory and numerical linear algebra*, in *Algorithms for Approximation II*, J. C. Mason and M. G. Cox, eds., London, 1990, Chapman and Hall.
- [46] ———, *Pseudospectra of matrices*, in *Numerical Analysis 1991*, D. F. Griffiths and G. A. Watson, eds., Harlow, Essex, 1992, Longman Scientific and Technical, pp. 234–266.
- [47] ———, *Computation of pseudospectra*, in *Acta Numerica 8*, Cambridge University Press, Cambridge, 1999, pp. 247–295.
- [48] ———, *Spectra and pseudospectra: The behavior of non-normal matrices and operators*, in *The Graduate Student's Guide to Numerical Analysis '98*, M. Ainsworth, J. Levesley, and M. Marletta, eds., Springer-Verlag, Berlin, 1999, pp. 217–250.
- [49] A. VAN DER SLUIS AND H. A. VAN DER VORST, *The rate of convergence of conjugate gradients*, *Numer. Math.*, 48 (1986), pp. 543–560.
- [50] ———, *The convergence behavior of Ritz values in the presence of close eigenvalues*, *Linear Algebra Appl.*, 88/89 (1987), pp. 651–694.
- [51] H. A. VAN DER VORST AND C. VUIK, *The superlinear convergence behaviour of GMRES*, *J. Comp. Appl. Math.*, 48 (1993), pp. 327–341.
- [52] T. G. WRIGHT, *MATLAB pseudospectra GUI*, 2000–2001. Available online from <http://www.comlab.ox.ac.uk/pseudospectra/psagui>.