

FEAST FOR DIFFERENTIAL EIGENVALUE PROBLEMS*

ANDREW HORNING[†] AND ALEX TOWNSEND[‡]

Abstract. An operator analogue of the FEAST matrix eigensolver is developed to compute the discrete part of the spectrum of a differential operator in a region of interest in the complex plane. Unbounded search regions are handled with a novel rational filter for the right half-plane. If the differential operator is normal or self-adjoint, then the operator analogue preserves that structure and robustly computes eigenvalues to near machine precision accuracy. The algorithm is particularly adept at computing high-frequency modes of differential operators that possess self-adjoint structure with respect to weighted Hilbert spaces.

Key words. FEAST, filtered subspace, differential eigenvalue problems, spectral methods

AMS subject classifications. 34L16, 65F15

DOI. 10.1137/19M1238708

1. Introduction. In this paper, we consider differential eigenvalue problems posed on the interval $[-1, 1]$, i.e.,

$$(1.1) \quad \mathcal{L}u = \lambda u, \quad u(\pm 1) = \dots = u^{(N/2)}(\pm 1) = 0.$$

Here, \mathcal{L} is a linear, ordinary differential operator of even order N . A complex number λ and a function u satisfying (1.1) are called an eigenvalue and eigenfunction of \mathcal{L} , respectively. We focus on computing the eigenvalues of \mathcal{L} contained in a simply connected region $\Omega \subset \mathbb{C}$. Throughout, we assume that the boundary $\partial\Omega$ is a rectifiable, simple closed curve and that the spectrum $\lambda(\mathcal{L})$ of \mathcal{L} is discrete, does not intersect $\partial\Omega$, and only finitely many eigenvalues counting multiplicities are in Ω . To simplify discussion about the eigenfunctions of (1.1), we assume that there are eigenfunctions of \mathcal{L} that form a basis for the invariant subspace of \mathcal{L} associated with Ω .

Since the development of the QR algorithm in the 1960s, the standard methods for solving (1.1) have adopted a “discretize-then-solve” paradigm. These algorithms first discretize \mathcal{L} to obtain a finite matrix eigenvalue problem and then solve the matrix eigenvalue problem with algorithms from numerical linear algebra [17, 21, 25, 39]. Motivated by mathematical software for highly adaptive computations with functions [18], we propose an alternative strategy: an algorithm that solves (1.1) by directly manipulating \mathcal{L} at the continuous level and only discretizes functions, not operators. By designing an eigensolver for \mathcal{L} rather than intermediate discretizations, we are able to leverage spectrally accurate approximation schemes for functions while avoiding several pitfalls that plague spectral discretizations of (1.1) (see [60], [22, Chap. 2], and [59, Chap. 30]). For this reason, we view our proposed algorithms as adopting a “solve-then-discretize” paradigm. This paradigm has been applied to Krylov methods [23], iterative eigensolvers [28], and contour integral projection eigensolvers [7] for differential operators. Related techniques for computing with operators on infinite-dimensional spaces have been proposed and studied in [14, 37].

*Received by the editors January 14, 2019; accepted for publication (in revised form) December 4, 2019; published electronically April 20, 2020.

<https://doi.org/10.1137/19M1238708>

Funding: The work of the authors was supported by the National Science Foundation grant 1818757.

[†]Center for Applied Mathematics, Cornell University, Ithaca, NY 14853 (ajh326@cornell.edu).

[‡]Department of Mathematics, Cornell University, Ithaca, NY 14853 (townsend@cornell.edu).

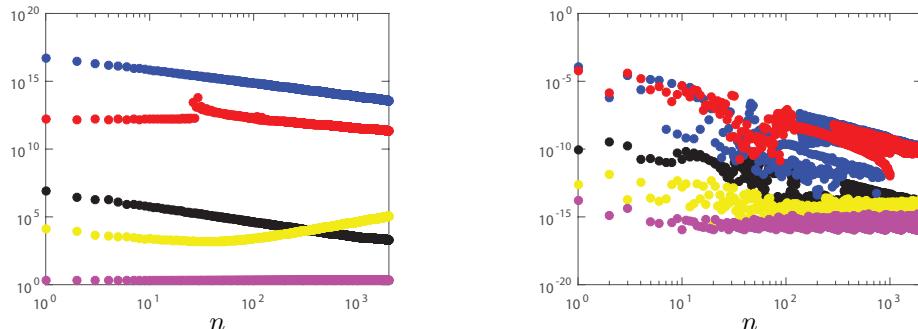


FIG. 1.1. Left: The eigenvalue condition numbers [5] for 4000×4000 discretizations of (1.2) obtained by collocation (blue dots), tau (red dots), Chebyshev–Galerkin (black dots), and ultraspherical (yellow dots) spectral methods are compared to the eigenvalue condition numbers (magenta dots) of (1.2), which are preserved by the operator analogue of FEAST. Right: The relative errors in the first 2000 eigenvalues of each spectral discretization of (1.2), computed with a backward stable eigensolver [24, p. 385]. We observe fluctuations in the relative errors due to the ill-conditioning introduced by using nonsymmetric spectral discretizations of \mathcal{L} . In contrast, the relative errors (magenta dots) in the eigenvalues computed by *contFEAST*, a practical implementation of the operator analogue of FEAST (see section 4), are on the order of machine precision.

As an example of the advantages of our methodology, consider the simplest possible differential eigenvalue problem given by

$$(1.2) \quad -\frac{d^2u}{dx^2} = \lambda u, \quad u(\pm 1) = 0.$$

The eigenvalues of (1.2) are $\lambda_k = (k\pi/2)^2$, for $k \geq 1$, and are well-conditioned due to the fact that the eigenfunctions form a complete orthonormal set in the Hilbert space $L^2([-1, 1])$ [31, p. 382]. However, spectral discretizations of (1.2) lead to highly non-normal matrices with eigenvalues that are far more ill-conditioned than expected. Due to this ill-conditioning, the accuracy in the computed eigenvalues can be extremely variable and difficult to predict, ranging from a few digits to nearly full precision (see Figure 1.1).

It is possible to use structure-preserving spectral discretizations to solve (1.2) accurately [13, 49]. However, there is a lack of literature on designing spectral discretizations of (1.1) when \mathcal{L} is self-adjoint or normal with respect to an inner product other than $L^2([-1, 1])$. On the other hand, our solve-then-discretize methodology automatically preserves the normality or self-adjointness of \mathcal{L} with respect to a relevant Hilbert space \mathcal{H} , provided that the inner product $(\cdot, \cdot)_{\mathcal{H}}$ can be evaluated.

At the heart of our approach is an operator analogue of the FEAST matrix eigensolver, which we briefly outline:

- (1) We construct a basis for the eigenspace \mathcal{V} corresponding to Ω by sampling the range of the associated spectral projector $\mathcal{P}_{\mathcal{V}}$.
- (2) We extract an \mathcal{H} -orthonormal basis for \mathcal{V} with a continuous analogue of the QR factorization [57].
- (3) We perform a Rayleigh–Ritz projection [46, p. 98] of \mathcal{L} onto \mathcal{V} with the orthonormal basis in (2). We solve the resulting matrix eigenvalue problem to obtain approximations to the eigenvalues of \mathcal{L} in Ω .

As with the FEAST matrix eigensolver, the spectral projector $\mathcal{P}_{\mathcal{V}}$ is applied approximately via a quadrature rule approximation. For matrices, this involves solving

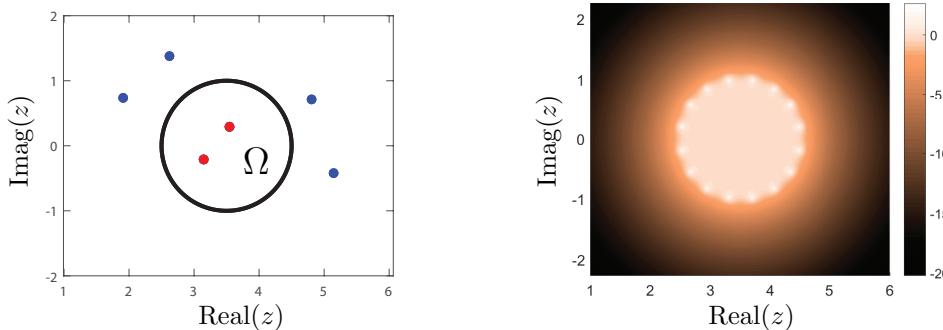


FIG. 2.1. Left: FEAST uses an approximation to the spectral projector to compute the eigenvalues that lie inside Ω (red dots) and project away the eigenvalues outside of Ω (blue dots). Right: The rational map in (2.2) that approximates the characteristic function on Ω .

shifted linear systems, while for differential operators one needs to solve shifted linear differential equations. We solve these differential equations with the ultraspherical spectral method, which is a well-conditioned spectral method that is capable of resolving solutions that exhibit layers, rapid oscillations, and weak corner singularities [36].

Critically, we discretize basis functions for \mathcal{V} as opposed to discretizing the differential operator \mathcal{L} when solving (1.1). While discretizations of a normal operator \mathcal{L} can lead to nonnormal matrices, the Rayleigh–Ritz projection described in (3) always leads to a normal matrix eigenvalue problem when \mathcal{L} is normal (see Theorems 3.1 and 3.2). In fact, we prove that using a sufficiently good approximate basis for \mathcal{V} does not significantly increase the sensitivity of the eigenvalues when \mathcal{L} is normal (see subsection 5.2 for a precise statement). The result is a highly accurate eigensolver for normal differential operators \mathcal{L} , requiring $\mathcal{O}(mMN \log(N) + m^2N + m^3)$ floating point operations, where $m = \dim(\mathcal{V})$ and M and N are the polynomial degrees used to resolve the variable coefficients in \mathcal{L} and the eigenfunctions in \mathcal{V} , respectively.

The eigensolver we develop is competitive in the high-frequency regime because it efficiently resolves oscillatory basis functions in \mathcal{V} . Furthermore, it handles operators that are self-adjoint or normal with respect to nonstandard Hilbert spaces. Finally, our algorithm is parallelizable like the FEAST matrix eigensolver [43]. This work is a step towards closing the gap between the frequency regimes that are accessible to computational techniques and asymptotic methods for differential eigenvalue problems posed on higher-dimensional domains [9, 10].

This paper is organized as follows. We begin in section 2 by reviewing FEAST for matrix eigenvalue problems. In section 3 we introduce an analogue of FEAST for differential operators and show that the operator analogue preserves eigenvalue sensitivity. In section 4 we discuss a practical implementation of the operator analogue and provide two examples from Sturm–Liouville theory to illustrate its capabilities in the high-frequency regime. We analyze the convergence and stability of this implementation in section 5. Sections 6 and 7 develop further applications of the solve-then-discretize paradigm, including an operator analogue of the Rayleigh Quotient Iteration and an extension of FEAST to unbounded search regions.

2. The FEAST matrix eigensolver. The FEAST matrix eigensolver uses approximate spectral projection to compute the eigenvalues of a matrix $A \in \mathbb{C}^{n \times n}$ in a region of interest $\Omega \subset \mathbb{C}$ [32] (see Figure 2.1). It is usually more computationally efficient than standard eigensolvers when the number of eigenvalues in Ω is much

Algorithm 2.1 The FEAST algorithm for matrix eigenvalue problems [43]. This is often viewed as a single iteration that is repeated to improve the accuracy of the computed eigenvalues and eigenvectors [53].

Input: $A \in \mathbb{C}^{n \times n}$, $\Omega \subset \mathbb{C}$ containing m eigenvalues of A , $Y : \mathbb{C}^{n \times m}$.

- 1: Compute $V = \mathcal{P}_{\mathcal{V}}Y$.
- 2: Compute the QR factorization $V = QR$.
- 3: Compute $A_Q = Q^*AQ$ and solve the eigenvalue problem $A_QX = \Lambda X$ for $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_m)$ and $X \in \mathbb{C}^{m \times m}$.

Output: Eigenvalues $\lambda_1, \dots, \lambda_m$ in Ω and eigenfunctions $U = QX$.

smaller than n . The dominating computational cost of FEAST is solving several independent shifted linear systems, but these can be performed in parallel [32].

There are three essential ingredients to FEAST:

- (i) **Spectral projector.** Let $\lambda_1, \dots, \lambda_m$ be the eigenvalues of A in Ω and let \mathcal{V} be the associated invariant subspace of A , i.e., $A\mathcal{V} = \mathcal{V}$. The *spectral projector* onto \mathcal{V} is defined as

$$(2.1) \quad \mathcal{P}_{\mathcal{V}} = \frac{1}{2\pi i} \int_{\partial\Omega} (zI - A)^{-1} dz.$$

The important fact here is that $\text{range}(\mathcal{P}_{\mathcal{V}}) = \mathcal{V}$ and so $\mathcal{P}_{\mathcal{V}}$ is a projection onto the invariant subspace of A [31].

- (ii) **Basis for \mathcal{V} .** FEAST uses the spectral projector to construct a basis for \mathcal{V} . It begins with a matrix $Y \in \mathbb{C}^{n \times m}$ with linearly independent columns that are not in $\ker(\mathcal{P}_{\mathcal{V}})$, then it computes $Z = \mathcal{P}_{\mathcal{V}}Y$. The columns of Z span \mathcal{V} , and a QR factorization of Z provides an orthonormal basis, Q , for \mathcal{V} .
- (iii) **Rayleigh–Ritz projection.** Having obtained an orthonormal basis for \mathcal{V} , FEAST solves $A_Qx = \lambda x$ using a dense eigensolver [43], where $A_Q = Q^*AQ$. Since $\text{range}(Q) = \mathcal{V}$, the eigenvalues of A_Q are the eigenvalues of A that lie inside Ω . When A_Q is diagonalizable, the eigenvectors of A are given by $u_i = Qx_i$, for $i = 1, \dots, m$, where x_1, \dots, x_m are the corresponding eigenvectors of A_Q .

For practical computation, FEAST approximates the contour integral in (2.1) with a quadrature rule. Given a quadrature rule with nodes z_1, \dots, z_{ℓ} and weights w_1, \dots, w_{ℓ} , one can approximate $\mathcal{P}_{\mathcal{V}}Y$ by

$$(2.2) \quad \mathcal{P}_{\mathcal{V}}Y \approx \frac{1}{2\pi i} \sum_{k=1}^{\ell} w_k (z_k I - A)^{-1} Y.$$

In this case, the eigenpairs of A_Q provide approximations to the eigenpairs of A , known as Ritz values and vectors [53]. To refine the accuracy of the Ritz values and vectors, a more accurate quadrature rule can be used to compute $\mathcal{P}_{\mathcal{V}}Y$ [53]. FEAST also refines the approximate eigenvalues and eigenvectors by applying $\mathcal{P}_{\mathcal{V}}$ to the $n \times m$ block of approximate eigenvectors using a quadrature rule and iterating (ii) and (iii) until convergence. To fully understand this refinement process, one must examine FEAST through the lens of rational subspace iteration [53].

When the dimension m of the invariant subspace \mathcal{V} is unknown, there are several techniques for estimating m and selecting an appropriate value [32, 35, 53]. Most of these techniques can be incorporated into the operator analogue of FEAST. Conse-

quently, we assume that m is known throughout the paper and focus on the algorithmic and theoretical aspects of FEAST that are relevant in the operator setting.

Curiously, the originally proposed FEAST algorithm does not compute an orthonormal basis for \mathcal{V} before performing the Rayleigh–Ritz projection [32, 43].¹ However, when Q has orthonormal columns and $\text{range}(Q)$ is an invariant subspace of A , then the eigenvalues of the small matrix Q^*AQ are no more sensitive to perturbations than the original eigenvalues of A . This highly desirable property follows from an examination of the structure of the left and right invariant subspaces of Q^*AQ or, alternatively, from the ϵ -pseudospectra of Q^*AQ [59, p. 382].

3. An operator analogue of FEAST. The FEAST matrix algorithm provides a natural starting point for an operator analogue because it provides a recipe to construct a small matrix Q^*AQ whose eigenvalues coincide with those of A inside Ω and have related invariant subspaces. Moreover, the eigenstructure of Q^*AQ reflects the eigenstructure of A when the columns of Q are orthonormal. As the sensitivity of the eigenvalues of A depends intimately on the structure of the associated eigenvectors, Q^*AQ may be used to compute the desired eigenvalues of A efficiently without sacrificing accuracy. Here, we generalize FEAST so that it constructs a matrix whose eigenvalues coincide with those of a differential operator inside Ω .

3.1. FEAST for differential operators. In place of a matrix A acting on vectors from \mathbb{C}^n , we now consider a differential operator \mathcal{L} acting on functions from a Hilbert space \mathcal{H} . As described in section 2, the FEAST recipe prescribes a spectral projection to compute a basis for \mathcal{V} , which is then used for the Rayleigh–Ritz projection to construct a matrix representation on \mathcal{V} . Throughout, we require that \mathcal{L} be a closed operator² and that its domain $\mathcal{D}(\mathcal{L})$ is dense in \mathcal{H} .

- (i) **Spectral projector.** Although \mathcal{L} is unbounded, the resolvent $(z\mathcal{I} - \mathcal{L})^{-1}$ is bounded when $z \notin \lambda(\mathcal{L})$ and the spectral projector onto \mathcal{V} may be defined via contour integral [31, p. 178]. It is given by

$$(3.1) \quad \mathcal{P}_{\mathcal{V}} = \frac{1}{2\pi i} \int_{\partial\Omega} (z\mathcal{I} - \mathcal{L})^{-1} dz.$$

- (ii) **Basis for \mathcal{V} .** With the spectral projector at our disposal, we apply $\mathcal{P}_{\mathcal{V}}$ to functions f_1, \dots, f_m in $\mathcal{H} \setminus \ker(\mathcal{P}_{\mathcal{V}})$ to obtain a basis of functions v_1, \dots, v_m for \mathcal{V} . Orthonormalizing v_1, \dots, v_m with respect to the inner product $(\cdot, \cdot)_{\mathcal{H}}$ on \mathcal{H} gives us an \mathcal{H} -orthonormal basis q_1, \dots, q_m for \mathcal{V} .
- (iii) **Rayleigh–Ritz projection.** To compute a matrix representation L of \mathcal{L} on \mathcal{V} , the Rayleigh–Ritz projection is performed using the inner product on \mathcal{H} . The elements of L are given by $L_{ij} = (q_i, \mathcal{L}q_j)_{\mathcal{H}}$ for $1 \leq i, j \leq m$. The eigenvalues of L are precisely the eigenvalues $\lambda_1, \dots, \lambda_m$ of \mathcal{L} that lie inside Ω . The eigenfunctions of \mathcal{L} are recovered from the eigenvectors x_1, \dots, x_m of L by computing $u_i = \sum_{k=1}^m x_i^{(k)} q_k$, for $i = 1, \dots, m$, where $x_i^{(k)}$ is the k th component of x_i .

To avoid a clutter of indices, we employ the notation of quasimatrices.³ If Q is the quasimatrix with columns q_1, \dots, q_m , then the matrix L whose elements are

¹The FEAST algorithm for non-Hermitian matrices utilizes dual bases for the left and right eigenspaces to improve stability [32].

²An operator $\mathcal{A} : \mathcal{D}(\mathcal{A}) \rightarrow \mathcal{H}$ is *closed* if its graph is a closed linear subspace of $\mathcal{H} \times \mathcal{H}$ [31, p. 165].

³A *quasimatrix* is a matrix whose columns (or rows) are functions defined on an interval $[a, b]$, in contrast to matrices whose columns (or rows) are vectors [18, Chap. 6].

Algorithm 3.1 An operator analogue of FEAST for differential operators.

Input: $\mathcal{L} : \mathcal{D}(\mathcal{L}) \rightarrow \mathcal{H}$, $\Omega \subset \mathbb{C}$ containing m eigenvalues of \mathcal{L} , $F : \mathbb{C}^m \rightarrow \mathcal{H}$.

- 1: Compute $V = \mathcal{P}_{\mathcal{V}} F$.
- 2: Compute $V = QR$, where $Q : \mathbb{C}^m \rightarrow \mathcal{D}(\mathcal{L}) \subset \mathcal{H}$ has \mathcal{H} -orthonormal columns and $R \in \mathbb{C}^{m \times m}$ is upper triangular.
- 3: Compute $L = Q^* \mathcal{L} Q$ and solve $LX = \Lambda X$ for $\Lambda = \text{diag}[\lambda_1, \dots, \lambda_m]$ and $X \in \mathbb{C}^{m \times m}$.

Output: Eigenvalues $\lambda_1, \dots, \lambda_m$ in Ω and eigenfunctions $U = QX$.

$L_{ij} = (q_i, \mathcal{L}q_j)_{\mathcal{H}}$ in (iii) is expressed compactly in quasimatrix notation as $L = Q^* \mathcal{L} Q$. Here, Q^* is the conjugate transpose of the quasimatrix Q so its rows are complex conjugates of the functions q_1, \dots, q_m .

The analogue of FEAST for differential operators is summarized in Algorithm 3.1 using quasimatrix notation so that it resembles its matrix counterpart.

Keep in mind that Algorithm 3.1 is a formal algorithm. In general, we cannot apply the spectral projector exactly, nor represent the basis \mathcal{V} exactly with finite memory. A practical implementation is discussed in section 4.

3.2. Condition number of the Ritz values. As illustrated in Figure 1.1, the eigenvalues of matrix discretizations of \mathcal{L} can be more sensitive to perturbations than the eigenvalues of \mathcal{L} . The advantage of our FEAST approach in subsection 3.1 is that the Ritz values, i.e., the eigenvalues of $Q^* \mathcal{L} Q$, are no more sensitive to perturbations than the original eigenvalues of \mathcal{L} when $\text{range}(Q)$ is an invariant subspace of \mathcal{L} .

To see this, let λ be a simple eigenvalue of a differential operator \mathcal{L} . Let $u, w \in \mathcal{H}$ satisfy $\mathcal{L}u = \lambda u$ and $\mathcal{L}^*w = \bar{\lambda}w$, where $\bar{\lambda}$ denotes the complex conjugate of λ . The *condition number*⁴ of λ is given by [5, Theorem 2.3]

$$(3.2) \quad \kappa_{\mathcal{H}}(\lambda) = \frac{\|u\|_{\mathcal{H}}\|w\|_{\mathcal{H}}}{(w, u)_{\mathcal{H}}}.$$

The condition number $\kappa_{\mathcal{H}}(\lambda)$ quantifies the worst-case first-order sensitivity of λ to perturbations of \mathcal{L} . For instance, if we compute λ using a backward stable algorithm in floating point arithmetic, we expect to achieve an accuracy of at least $\kappa_{\mathcal{H}}(\lambda)\epsilon_{\text{mach}}$, where ϵ_{mach} is machine precision [58, Theorem 15.1].

THEOREM 3.1. *Let $\mathcal{L} : \mathcal{D}(\mathcal{L}) \rightarrow \mathcal{H}$ be a closed and densely defined operator on a Hilbert space \mathcal{H} , $Q : \mathbb{C}^m \rightarrow \mathcal{H}$ be an invariant subspace of \mathcal{L} satisfying $Q^* Q = I$, and $L = Q^* \mathcal{L} Q$. Suppose that $u \in \text{range}(Q)$ satisfies $\mathcal{L}u = \lambda u$ and w satisfies $\mathcal{L}^*w = \bar{\lambda}w$, where \mathcal{L}^* denotes the adjoint of \mathcal{L} and λ is a simple eigenvalue of \mathcal{L} with condition number $\kappa_{\mathcal{H}}(\lambda)$. Then,*

- (1) $LQ^*u = \lambda Q^*u$ and $L^*Q^*w = \bar{\lambda}Q^*w$,
- (2) $(Q^*w, Q^*u)_{\mathbb{C}^m} = (w, u)_{\mathcal{H}}$, and
- (3) $\kappa_{\mathbb{C}^m}(\lambda) \leq \kappa_{\mathcal{H}}(\lambda)$.

Proof. Denote $x = Q^*u$ and $y = Q^*w$. We prove the statements of the theorem in order. (1) Since $u \in \text{range}(Q)$, we can write $u = Qx$. Then, $\mathcal{L}(Qx) = \lambda(Qx)$ implies that $Q^* \mathcal{L} Qx = \lambda x$ using the fact that $Q^* Q = I$. For the left eigenvector, we write $w = Qy + v$ for some $v \in \text{range}(Q)^\perp$. Rewriting the adjoint equation for w , we find

⁴Although this formula is usually associated with the condition number for a simple eigenvalue of a matrix, its proof extends to our general setting [54, Theorem 5].

that $\mathcal{L}^*(Qy+v) = \bar{\lambda}(Qy+v)$ and multiplying by Q^* on both sides yields $Q^*\mathcal{L}^*Qy = \bar{\lambda}y$. Here, we have used the fact that $Q^*\mathcal{L}^*v = 0$, which holds because $v^*\mathcal{L}Q = 0$. (2) By calculating $(w, u)_{\mathcal{H}} = (Qy + v, Qx)_{\mathcal{H}}$, we find that $(w, u)_{\mathcal{H}} = (Qy, Qx)_{\mathcal{H}}$ because $v \in \text{range}(Q)^\perp$. Moreover, since $Q^*Q = I$ we conclude that $(w, u)_{\mathcal{H}} = (y, Q^*Qx)_{\mathbb{C}^m} = (y, x)_{\mathbb{C}^m}$. (3) We know that $\|u\|_{\mathcal{H}} = (Qx, Qx)_{\mathcal{H}} = (x, x)_{\mathbb{C}^m} = \|x\|_{\mathbb{C}^m}$ and $\|w\|_{\mathcal{H}} = (Qy + v, Qy + v)_{\mathcal{H}} = \|y\|_{\mathbb{C}^m} + \|v\|_{\mathcal{H}}$. Therefore,

$$\|u\|_{\mathcal{H}}\|w\|_{\mathcal{H}} = \|x\|_{\mathbb{C}^m}(\|y\|_{\mathbb{C}^m} + \|v\|_{\mathcal{H}}) \geq \|x\|_{\mathbb{C}^m}\|y\|_{\mathbb{C}^m}.$$

Referring to (2) for equality of the inner products in the denominator, we have

$$\kappa_{\mathbb{C}^m}(\lambda) = \frac{\|x\|_{\mathbb{C}^m}\|y\|_{\mathbb{C}^m}}{(y, x)_{\mathbb{C}^m}} \leq \frac{\|u\|_{\mathcal{H}}\|w\|_{\mathcal{H}}}{(w, u)_{\mathcal{H}}} = \kappa_{\mathcal{H}}(\lambda),$$

which concludes the proof. \square

Theorem 3.1 shows that if \mathcal{L} is a normal operator, then $u = w$ and we have $\kappa_{\mathbb{C}^m}(\lambda) = \kappa_{\mathcal{H}}(\lambda) = 1$. For nonnormal operators, item (3) of Theorem 3.1 may seem to erroneously indicate that ill-conditioning in the eigenvalues of \mathcal{L} can be overcome by a Rayleigh–Ritz projection. However, when \mathcal{L} is nonnormal the spectral projector $\mathcal{P}_{\mathcal{V}}$ is an oblique projection and computing the basis Q may be itself an ill-conditioned problem. Theorem 3.1 also illustrates why the operator analogue of FEAST leads to a well-conditioned matrix eigenvalue problem when the differential eigenvalue problem is well-conditioned. By computing an \mathcal{H} -orthonormal basis for the Rayleigh–Ritz projection, the relevant structure in the eigenspaces of \mathcal{L} and \mathcal{L}^* is preserved. However, the first-order analysis above is limited to simple eigenvalues.

3.3. Pseudospectra of $\mathbf{Q}^*\mathcal{L}\mathbf{Q}$. To go beyond first-order sensitivity analysis, we compare the ϵ -pseudospectra of \mathcal{L} and $\mathbf{Q}^*\mathcal{L}\mathbf{Q}$. Fix any $\epsilon > 0$ and let $\mathcal{L} : \mathcal{D}(\mathcal{L}) \rightarrow \mathcal{H}$ be a closed operator with a domain $\mathcal{D}(\mathcal{L})$ that is dense in \mathcal{H} . The ϵ -pseudospectrum of \mathcal{L} is defined as the set [59, p. 31]

$$(3.3) \quad \lambda_{\epsilon}(\mathcal{L}) = \{z \in \mathbb{C} : \|(z\mathcal{I} - \mathcal{L})^{-1}\|_{\mathcal{H}} > 1/\epsilon\}.$$

Here, we adopt the usual convention that $\|(z\mathcal{I} - \mathcal{L})^{-1}\|_{\mathcal{H}} = \infty$ when $z \in \lambda(\mathcal{L})$ so that $\lambda(\mathcal{L}) \subset \lambda_{\epsilon}(\mathcal{L})$. The ϵ -pseudospectrum set of \mathcal{L} bounds the region in which the eigenvalues of the perturbed operator $\mathcal{L} + \mathcal{E}$ with $\|\mathcal{E}\|_{\mathcal{H}} < \epsilon$ can be found [59, p. 31]. This means that $\lambda(\mathcal{L} + \mathcal{E}) \subset \lambda_{\epsilon}(\mathcal{L})$. In fact, there is an equivalence so that [59, p. 31]

$$(3.4) \quad \bigcup_{\|\mathcal{E}\|_{\mathcal{H}} < \epsilon} \lambda(\mathcal{L} + \mathcal{E}) = \lambda_{\epsilon}(\mathcal{L}).$$

This allows us to relate the sensitivity of the eigenvalues of \mathcal{L} and $\mathbf{Q}^*\mathcal{L}\mathbf{Q}$ by comparing the resolvent norms $\|(z\mathcal{I} - \mathcal{L})^{-1}\|_{\mathcal{H}}$ and $\|(z\mathcal{I} - \mathbf{Q}^*\mathcal{L}\mathbf{Q})^{-1}\|_{\mathbb{C}^m}$, respectively.

A useful generalization of Theorem 3.1 is that the ϵ -pseudospectrum of $\mathbf{Q}^*\mathcal{L}\mathbf{Q}$ is contained in the ϵ -pseudospectrum of \mathcal{L} . Since this holds for any $\epsilon > 0$, it demonstrates that the eigenvalues (even those with multiplicity) of $\mathbf{Q}^*\mathcal{L}\mathbf{Q}$ are no more sensitive to perturbations than those of \mathcal{L} . This inclusion result is well known in the matrix case where projection methods are a popular method for approximating the ϵ -pseudospectra of large data-sparse matrices [59, p. 381].

THEOREM 3.2. *Let $\mathcal{L} : \mathcal{D}(\mathcal{L}) \rightarrow \mathcal{H}$ be a closed and densely defined operator on a Hilbert space \mathcal{H} . For a fixed $\epsilon > 0$, suppose that $Q : \mathbb{C}^m \rightarrow \mathcal{H}$ satisfies $Q^*Q = I$ and that $\text{range}(Q)$ is an invariant subspace of \mathcal{L} . Then, $\lambda_{\epsilon}(\mathbf{Q}^*\mathcal{L}\mathbf{Q}) \subset \lambda_{\epsilon}(\mathcal{L})$.*

Proof. We follow the proof of Proposition 40.1 in [59, p. 382] for matrices, but with a closed operator. Since $Qx \in \mathcal{H}$ for any $x \in \mathbb{C}^m$, we have that

$$\|(zI - \mathcal{L})^{-1}\|_{\mathcal{H}} = \sup_{f \in \mathcal{H}, \|f\|_{\mathcal{H}}=1} \|(z\mathcal{I} - \mathcal{L})^{-1}f\|_{\mathcal{H}} \geq \max_{x \in \mathbb{C}^m, \|x\|_{\mathbb{C}^m}=1} \|(z\mathcal{I} - \mathcal{L})^{-1}Qx\|_{\mathcal{H}}.$$

Now, when $\text{range}(Q)$ is an invariant subspace of \mathcal{L} and $Q^*Q = I$, we have that $\|(z\mathcal{I} - \mathcal{L})^{-1}Qx\|_{\mathcal{H}} = \|Q^*(z\mathcal{I} - \mathcal{L})^{-1}Qx\|_{\mathbb{C}^m}$. Because $QQ^*f = f$ for all $f \in \mathcal{V}$, we can check that $Q^*(z\mathcal{I} - \mathcal{L})^{-1}Q = (Q^*(z\mathcal{I} - \mathcal{L})Q)^{-1}$. Since $Q^*Q = I$, it follows that

$$\|(zI - \mathcal{L})^{-1}\|_{\mathcal{H}} \geq \|(Q^*(z\mathcal{I} - \mathcal{L})Q)^{-1}\|_{\mathbb{C}^m} = \|(zI - Q^*\mathcal{L}Q)^{-1}\|_{\mathbb{C}^m}.$$

Therefore, $z \in \lambda_{\epsilon}(\mathcal{L})$ whenever $z \in \lambda_{\epsilon}(Q^*\mathcal{L}Q)$. \square

The inclusion in Theorem 3.2 may be strict, indicating that the eigenvalues of $Q^*\mathcal{L}Q$ are less sensitive than those of \mathcal{L} . For example, this may occur when the projection onto $\text{range}(Q)$ targets a subset of well-conditioned eigenvalues of \mathcal{L} . However, we emphasize that ill-conditioning in the eigenvalues of \mathcal{L} cannot be overcome by a Rayleigh–Ritz projection: in general, the situation is complicated [59, Chap. 40].

Theorem 3.2 is useful for studying the stability of Algorithm 3.1. If an approximate eigenvalue $\hat{\lambda}$ of $Q^*\mathcal{L}Q$ is computed with an error tolerance of $\epsilon > 0$, then

$$\hat{\lambda} \in \lambda_{\epsilon}(Q^*\mathcal{L}Q) \subset \lambda_{\epsilon}(\mathcal{L}).$$

From this, we know by (3.4) that $\hat{\lambda}$ is an eigenvalue of a perturbed operator $\mathcal{L} + \mathcal{E}$ with $\|\mathcal{E}\|_{\mathcal{H}} < \epsilon$. In other words, the operator analogue of FEAST, Algorithm 3.1, is backward stable. As we see in section 5, Theorem 3.2 is also the starting point for a stability analysis when the spectral projection is no longer exact and the Rayleigh–Ritz projection is performed with a matrix \hat{Q} that only approximates a basis for an invariant subspace of \mathcal{L} .

4. A practical differential eigensolver based on an operator analogue of FEAST. The operator analogue of FEAST requires the manipulation of objects such as differential operators, functions, and contour integrals (see Algorithm 3.1). For a practical implementation, these objects must be discretized; however, we avoid discretizing \mathcal{L} directly. Instead, we construct polynomial approximants to the basis for \mathcal{V} by approximately solving shifted linear ODEs. These polynomial approximants are used in the Rayleigh–Ritz projection to compute the eigenvalues of \mathcal{L} in Ω .

Let z_1, \dots, z_{ℓ} and w_1, \dots, w_{ℓ} be a set of quadrature nodes and weights to approximate the integral in (3.1). As FEAST does in the matrix case, we approximate $\mathcal{P}_{\mathcal{V}}$ in (3.1) with a quadrature rule as follows:

$$(4.1) \quad \hat{\mathcal{P}}_{\mathcal{V}} = \frac{1}{2\pi i} \sum_{k=1}^{\ell} w_k (z_k \mathcal{I} - \mathcal{L})^{-1}.$$

If F is a quasimatrix with columns $f_1, \dots, f_m \in \mathcal{H}$, then $\mathcal{P}_{\mathcal{V}}F$ is replaced by the approximation $\hat{\mathcal{P}}_{\mathcal{V}}F = \frac{1}{2\pi i} \sum_{k=1}^{\ell} w_k (z_k \mathcal{I} - \mathcal{L})^{-1}F$. Therefore, to compute $\hat{\mathcal{P}}_{\mathcal{V}}F$ we need to solve ℓ shifted linear ODEs, each with m right-hand sides, i.e.,

$$(4.2) \quad (z_k \mathcal{I} - \mathcal{L})g_{i,k} = f_i, \quad g_{i,k}(\pm 1) = \dots = g_{i,k}^{(N/2)}(\pm 1) = 0, \quad 1 \leq i \leq m.$$

If the quasimatrix with columns $g_{1,k}, \dots, g_{m,k}$ is denoted by G_k for $k = 1, \dots, \ell$, then we have $\hat{\mathcal{P}}_{\mathcal{V}}F = \sum_{k=1}^{\ell} w_k G_k$.

Algorithm 4.1 A practical algorithm for computing the eigenvalues of a differential operator \mathcal{L} , which we refer to as `contFEAST`.

Input: $\mathcal{L} : \mathcal{D}(\mathcal{L}) \rightarrow \mathcal{H}$, $z_1, \dots, z_\ell \in \partial\Omega$, $w_1, \dots, w_\ell \in \mathbb{C}$, $F : \mathbb{C}^m \rightarrow \mathcal{H}$, $\epsilon > 0$.

1: **repeat**

2: Solve $(z_k \mathcal{I} - \mathcal{L})G_k = F$, $G_k(\pm 1) = 0, \dots, G_k^{(N/2)}(\pm 1) = 0$, for $k = 1, \dots, \ell$.

3: Set $\hat{V} = \sum_{k=1}^{\ell} w_k G_k$.

4: Compute $\hat{V} = \hat{Q}\hat{R}$, where $\hat{Q} : \mathbb{C}^m \rightarrow \mathcal{D}(\mathcal{L}) \subset \mathcal{H}$ has \mathcal{H} -orthonormal columns and $\hat{R} \in \mathbb{C}^{m \times m}$ is upper triangular.

5: Compute $\hat{L} = \hat{Q}^* \mathcal{L} \hat{Q}$ and solve $\hat{L}\hat{X} = \hat{X}\hat{\Lambda}$ for $\hat{\Lambda} = \text{diag}[\hat{\lambda}_1, \dots, \hat{\lambda}_m]$ and $\hat{X} \in \mathbb{C}^{m \times m}$. Set $F = \hat{Q}\hat{X}$.

6: **until** $\|\mathcal{L}F - F\hat{\Lambda}\|_{\mathcal{H}} \leq \epsilon \|\hat{\Lambda}\|_{\mathbb{C}^m}$.

Output: $\hat{\Lambda}$, $\hat{U} = \hat{Q}\hat{X}$.

To construct a basis for \mathcal{V} , it is important to choose F so that the columns of $\hat{V} = \hat{\mathcal{P}}_{\mathcal{V}}F$ are linearly independent and, if possible, well-conditioned. Analogous to the implementation of matrix FEAST [32, 43], we obtain the columns of F by selecting m band-limited random functions⁵ on $[-1, 1]$ [20]. When \mathcal{L} is a normal operator, this typically yields a well-conditioned basis \hat{V} .

We now outline the key implementation details of our differential eigensolver:

- (i) **Approximate spectral projection.** To compute $\hat{V} = \hat{\mathcal{P}}_{\mathcal{V}}F$, we solve the shifted linear ODEs in (4.2) using the ultraspherical spectral method [36]. The ultraspherical spectral method leads to well-conditioned linear systems and is capable of accurately resolving the functions $g_{i,k}$ even when they are highly oscillatory or have boundary layers. Moreover, an adaptive QR factorization automatically determines the degree of the polynomial interpolants needed to approximate the functions $g_{i,k}$ to near machine precision [37, 38]. After accurately resolving the functions $g_{i,k}$, we can accurately compute a basis for \mathcal{V} provided that both the spectral projector is well-conditioned (i.e., \mathcal{L} is not highly nonnormal) and the quadrature rule is sufficiently accurate.
- (ii) **Orthonormal basis.** To compute an orthonormal basis \hat{Q} for the columns of \hat{V} , we compute a QR factorization of the quasimatrix \hat{V} by Householder triangularization [57]. The Householder reflectors are constructed with respect to the inner product $(\cdot, \cdot)_{\mathcal{H}}$ so that the columns of \hat{Q} are \mathcal{H} -orthonormal.
- (iii) **Computing $\hat{Q}^* \mathcal{L} \hat{Q}$.** To construct the matrix $\hat{L} = \hat{Q}^* \mathcal{L} \hat{Q}$, we apply \mathcal{L} to the columns of \hat{Q} and then evaluate the action of \hat{Q}^* on $\mathcal{L}\hat{Q}$. Multiplying \hat{Q}^* with $\mathcal{L}\hat{Q}$ involves taking the inner products

$$(4.3) \quad \hat{L}_{ij} = (\hat{q}_i, \mathcal{L}\hat{q}_j)_{\mathcal{H}}, \quad 1 \leq i, j \leq m,$$

where \hat{q}_i denotes the i th column of \hat{Q} . The eigenvalues $\hat{\lambda}_1, \dots, \hat{\lambda}_m$ and eigenvectors $\hat{x}_1, \dots, \hat{x}_m$ of the matrix \hat{L} are computed using the QR algorithm [24, p. 385].

Critically, the inner product $(\cdot, \cdot)_{\mathcal{H}}$ used in the QR factorization of \hat{V} and the construction of $\hat{Q}^* \mathcal{L} \hat{Q}$ depends on the choice of the Hilbert space \mathcal{H} . As long as we are able to evaluate $(\cdot, \cdot)_{\mathcal{H}}$, we can exploit the fact that \mathcal{L} is self-adjoint or a normal

⁵A periodic band-limited random function on $[-L, L]$ is a periodic function defined by a truncated Fourier series with random (e.g., standard Gaussian distributed) coefficients. In the nonperiodic setting, the Fourier series is defined on a larger interval $[-L', L']$ and the domain is then truncated [20].

operator with respect to $(\cdot, \cdot)_{\mathcal{H}}$ so that we can accurately compute the eigenvalues of \mathcal{L} in Ω (see Theorem 5.2). For this reason, our algorithm is able to accurately compute the eigenvalues and eigenfunctions of differential operators that are self-adjoint with respect to nonstandard Hilbert spaces (see subsection 4.1).

Evaluating the inner product $(\cdot, \cdot)_{\mathcal{H}}$ usually means computing an integral, which we approximate with a quadrature rule. For example, if $\mathcal{H} = L^2([-1, 1])$,

$$(f, g)_{L^2([-1, 1])} = \int_{-1}^1 \overline{f(x)} g(x) dx.$$

Given the Gauss–Legendre quadrature nodes x_1, \dots, x_p and weights w_1, \dots, w_p on $[-1, 1]$, then one uses the approximation [11]

$$(f, g)_{L^2([-1, 1])} \approx \sum_{k=1}^p w_k \overline{f(x_k)} g(x_k).$$

A practical implementation of the operator analogue of FEAST is presented in Algorithm 4.1. As with matrix FEAST, there are two approaches for improving the accuracy of the Ritz values $\hat{\lambda}_1, \dots, \hat{\lambda}_m$ and vectors $\hat{Q}\hat{x}_1, \dots, \hat{Q}\hat{x}_m$. The first is to improve the accuracy of the quadrature rule in (4.1). The second is to iterate the algorithm by replacing F by the quasimatrix \hat{U} with columns $\hat{u}_i = \hat{Q}\hat{x}_i$ for $1 \leq i \leq m$, repeating the process if necessary.⁶ For normal operators, this iteration generates a sequence of quasimatrices \hat{Q}_k with \mathcal{H} -orthonormal columns that converge to an \mathcal{H} -orthonormal basis for the invariant subspace \mathcal{V} as $k \rightarrow \infty$. This can be viewed as a rational subspace iteration, and geometric convergence of the Ritz pairs is typical (see section 5).

With either refinement strategy, the accuracy of the Ritz pairs may be monitored using the residual norm (see step 6 of Algorithm 4.1) as a proxy, just as in the matrix case. For normal operators, the error in the eigenvalues and eigenvectors computed by Algorithm 4.1 is typically $\mathcal{O}(\epsilon)$, where ϵ is the threshold for the residual norm in step 6. We defer a discussion of the convergence and stability of Algorithm 4.1 to section 5. Additional resources on residual norm bounds for eigenvalues and eigenvectors of matrices and extensions to closed linear operators are found in [8, 12, 51].

In practice, when \mathcal{L} is nonnormal it may be beneficial to use a dual Rayleigh–Ritz projection $\hat{Q}_L^* \mathcal{L} \hat{Q}_R$, where the columns of \hat{Q}_R approximate an orthonormal basis for the target eigenspace of \mathcal{L} and the columns of \hat{Q}_L approximate an orthonormal basis for the associated eigenspace of the adjoint \mathcal{L}^* . In the case of matrix FEAST, the use of the dual projection leads to a nonnormal matrix eigensolver with improved robustness [32]. Although it is not difficult to adapt Algorithm 4.1 to an operator analogue of FEAST that uses dual projection, we focus on the implementation and analysis of the one-sided iteration.

Typically, solving the ODEs in (4.2) dominates the computational cost of Algorithm 4.1. With the ultraspherical spectral method, the computational complexity of solving the linear ODEs with m distinct right-hand sides is $\mathcal{O}(mMN \log(N))$ floating point operations (flops) [36]. Here, N and M are, respectively, the degrees of

⁶When \mathcal{L} is nonnormal the Ritz vectors $\hat{Q}\hat{x}_1, \dots, \hat{Q}\hat{x}_m$ may become numerically linearly dependent, which can lead to an ill-conditioned basis \hat{V} in subsequent iterations. The robustness of Algorithm 4.1 may be improved by computing the Schur vectors v_1, \dots, v_m of \hat{L} and using the orthonormal basis $\hat{Q}\hat{v}_1, \dots, \hat{Q}\hat{v}_m$ to seed the next iteration [52].

the truncated Chebyshev series needed to resolve the columns of G_k and the variable coefficients in \mathcal{L} to within the tolerance ϵ specified in Algorithm 4.1. In addition to the ODE solve, the QR factorization in (ii) requires $\mathcal{O}(m^2N)$ flops [57], while the dense eigenvalue computation with a small $m \times m$ matrix in (iii) takes $\mathcal{O}(m^3)$ flops [24, p. 391]. The complexity of one iteration of Algorithm 4.1 is therefore $\mathcal{O}(mMN \log(N) + m^2N + m^3)$ flops. In practice, convergence to machine precision usually occurs within two or three iterations.

4.1. Computing high-frequency eigenmodes. Algorithm 4.1 adaptively and accurately resolves basis functions for highly oscillatory eigenmodes and preserves the sensitivity of the eigenvalues of the differential operator \mathcal{L} , so it is well-suited to computing high-frequency eigenmodes when \mathcal{L} is self-adjoint or normal with respect to $(\cdot, \cdot)_{\mathcal{H}}$. We provide two examples from Sturm–Liouville theory to illustrate the effectiveness of the solve-then-discretize methodology in the high-frequency regime.

4.1.1. A regular Sturm–Liouville eigenvalue problem. First, consider a regular Sturm–Liouville eigenvalue problem (SLEP) given by

$$(4.4) \quad -\frac{d^2u}{dx^2} + x^2u = \lambda \cosh(x)u, \quad u(\pm 1) = 0.$$

This defines a self-adjoint differential operator with respect to the inner product

$$(4.5) \quad (v, u)_w = \int_{-1}^1 \bar{v}u \cosh(x) dx.$$

Consequently, (4.4) possesses a complete $(\cdot, \cdot)_w$ -orthonormal basis of eigenfunctions u_1, u_2, u_3, \dots for the weighted Hilbert space $\mathcal{H}_w = \{u : \|u\|_w = \sqrt{(u, u)_w} < \infty\}$ and an unbounded set of real eigenvalues $\lambda_1 \leq \lambda_2 \leq \lambda_3 \leq \dots$.

Asymptotics for the large eigenvalues of (4.4) are given by [2]

$$(4.6) \quad \sqrt{\lambda_n} \sim \frac{n\pi}{\int_{-1}^1 \sqrt{\cosh(x)} dx}, \quad n \rightarrow \infty.$$

To accurately compute the large eigenvalues of (4.4) with Algorithm 4.1, we prescribe circular search regions with unit radius centered at the values given by the asymptotic formula in (4.6) (see Figure 4.1). Each search region contains one eigenvalue.

4.1.2. An indefinite Sturm–Liouville eigenvalue problem. Next, we consider the following indefinite SLEP:

$$(4.7) \quad -\frac{d^2u}{dx^2} = \lambda x^3u, \quad u(\pm 1) = 0,$$

which is closely related to models of light propagation in a nonhomogeneous material [2, 61]. Since the weight function x^3 changes sign at $x = 0$, (4.7) has a bi-infinite sequence of eigenvalues [6]. We index them in order as $\dots \leq \lambda_{-2} \leq \lambda_{-1} < 0 < \lambda_1 \leq \lambda_2 \leq \dots$. The asymptotics for the positive eigenvalues are given by [3]

$$(4.8) \quad \sqrt{\lambda_n} \sim \frac{(n - 1/4)\pi}{\int_0^1 x^{3/2} dx}, \quad \lambda_n > 0, \quad n \rightarrow \infty.$$

A similar expansion holds for the negative eigenvalues [3].

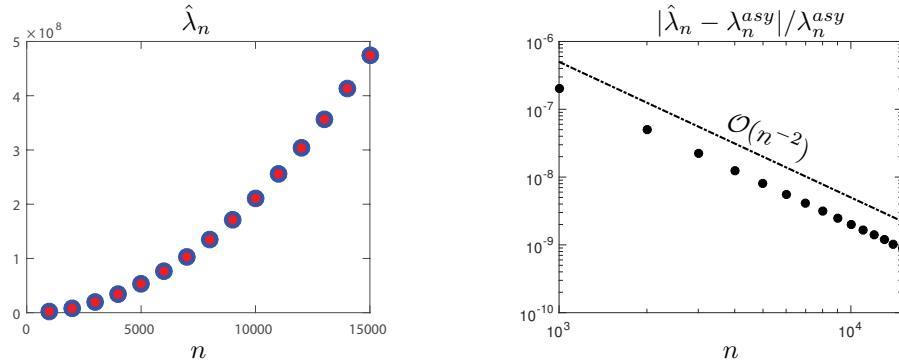


FIG. 4.1. Left: The large eigenvalues of (4.4) are computed by *contFEAST* (see Algorithm 4.1) using search regions given by asymptotic estimates for the eigenvalues (4.6). Right: The relative difference $|\hat{\lambda}_n - \lambda_n^{asy}|/\lambda_n^{asy}$ between the eigenvalues $\hat{\lambda}_n$ computed by *contFEAST* and the asymptotic values λ_n^{asy} from (4.6). The difference is compared to an $\mathcal{O}(n^{-2})$ relative error estimate [2].

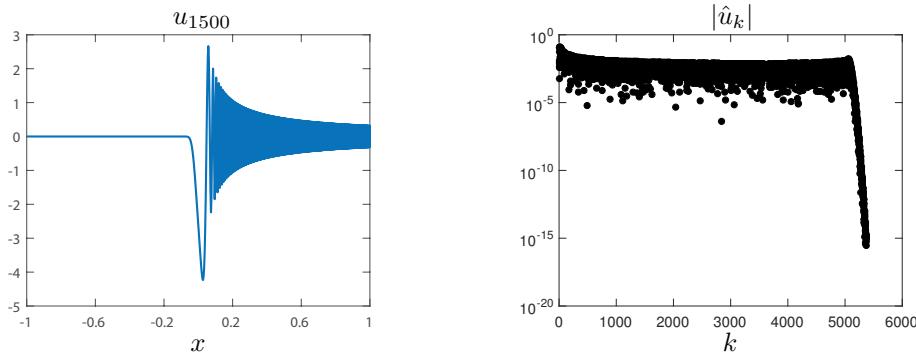


FIG. 4.2. Left: The high-frequency eigenfunction associated to λ_{1500} of the indefinite SLEP (4.7) computed by *contFEAST* (see Algorithm 4.1). Right: The Chebyshev coefficients $\{\hat{u}_k\}$ in a series expansion used to represent the eigenfunction. About 5371 Chebyshev coefficients are needed to accurately resolve the eigenfunction. The rapid decay in the coefficients to essentially machine precision is a good indication that the solution is fully resolved.

In contrast to the previous example, the indefinite weight function x^3 means that (4.7) is not immediately associated with a self-adjoint operator on a Hilbert space. Instead, (4.7) is usually studied through the lens of a Krein space and the eigenfunctions form a Riesz basis for the Hilbert space with the inner product [15]

$$(4.9) \quad (v, u)_{|w|} = \int_{-1}^1 \bar{v}u|x|^3 dx.$$

We use the leading order asymptotics in (4.8) to identify search regions that are likely to contain an eigenvalue of (4.7). Because the ultraspherical spectral method used to solve the ODEs in step 2 of Algorithm 4.1 is efficient when applied to ODEs with smooth variable coefficients, it is convenient to treat (4.7) as a generalized eigenvalue problem, i.e., as $\mathcal{L}_1 u = \lambda \mathcal{L}_2 u$, where $\mathcal{L}_1 u = -\frac{d^2 u}{dx^2}$ and $\mathcal{L}_2 u = x^3 u$. The eigenvalues of the pencil $z\mathcal{L}_2 - \mathcal{L}_1$ are then computed with a straightforward generalization of Algorithm 4.1 that is based on the spectral projector for the generalized eigenvalue problem, i.e.,

$$(4.10) \quad \mathcal{P}_V = \frac{1}{2\pi i} \int_{\partial\Omega} (z\mathcal{L}_2 - \mathcal{L}_1)^{-1} \mathcal{L}_2 dz.$$

The eigenvalues and eigenfunctions are automatically resolved to essentially machine precision because of the use of the adaptive QR solver (see Figure 4.2).

5. Convergence and stability. The primary consequence of the approximations introduced in Algorithm 4.1 is that the spectral projector is no longer applied exactly. Therefore, the basis \hat{Q} computed for the Rayleigh–Ritz projection is not an exact basis for the invariant subspace \mathcal{V} of \mathcal{L} and may require further refinement. Here, we view the iterative refinement procedure used in Algorithm 4.1 as a rational subspace iteration applied to a normal differential operator \mathcal{L} in order to provide a preliminary analysis of the stability of the iteration and the sensitivity of the Ritz values. The main results may be summarized as follows.

- (i) Algorithm 4.1 yields a sequence of quasimatrizes $\hat{Q}_1, \dots, \hat{Q}_k$ that (generically) converge geometrically to an orthonormal basis for the eigenspace \mathcal{V} (see Theorem 5.1).
- (ii) If \hat{Q}_k is a sufficiently good approximation to an orthonormal basis for \mathcal{V} , then the ϵ -psuedospectrum of $\hat{Q}_k^* \mathcal{L} \hat{Q}_k$ is contained in the 2ϵ -psuedospectrum of \mathcal{L} itself (see Theorem 5.2).
- (iii) Under mild conditions on the initial quasimatrix F in Algorithm 4.1, the sequence $\|\mathcal{L}(\hat{Q}_k - Q)\|_{\mathbb{C}^m \rightarrow \mathcal{H}}$ is uniformly bounded as $k \rightarrow \infty$ (see Lemma 5.3).

Taken together, these results demonstrate that each iteration of Algorithm 4.1 yields uniformly consistent Ritz pairs that converge linearly to the desired eigenpair and that the unboundedness of \mathcal{L} does not lead to instability. Note that this analysis does not take into account the impact of finite-precision arithmetic or the fact that the shifted differential equations in (4.2) are not solved exactly at each iteration (see the discussion at the end of subsection 5.1). However, (ii) ensures that the eigenvalues of the small matrix $\hat{Q}_k^* \mathcal{L} \hat{Q}_k$ are not much more sensitive than the eigenvalues of \mathcal{L} . Therefore, provided that the eigenvalue problem for \mathcal{L} is well-conditioned and we compute a sufficiently accurate approximation to a basis for \mathcal{V} , then we expect that the eigenvalues computed with Algorithm 4.1 provide an accurate approximation to the desired eigenvalues of \mathcal{L} .

5.1. Rational subspace iteration for differential operators. Analogous to the matrix case [53], Algorithm 4.1 may be interpreted as a filtered subspace iteration. Filtered subspace iteration is a variant of standard subspace iteration for computing a target subset of eigenvalues of a matrix A [46, Chap. 5]. The main idea is to choose a filter function $s(\cdot)$ that is large on the targeted eigenvalues of A and small on the unwanted eigenvalues of A . Applying the spectral transformation $s(A)$,⁷ one uses standard subspace iteration to compute a basis for the eigenspace of $s(A)$ corresponding to its largest eigenvalues, i.e., the targeted eigenvalues of A . With an approximate basis for the eigenspace available, the eigenvalues and eigenvectors can be extracted with a Rayleigh–Ritz step.

From this perspective, Algorithm 4.1 computes the eigenvalues of \mathcal{L} in Ω with the aid of a rational filter function induced by the quadrature rule in (4.1), i.e., $s(\cdot)$ is given by

$$(5.1) \quad s(z) = \sum_{k=1}^{\ell} \frac{w_k}{z_k - z}, \quad z \in \mathbb{C} \setminus \{z_1, \dots, z_l\}.$$

⁷A spectral transformation $s(\cdot)$ may be applied to A via the eigendecomposition of A , or more generally the Jordan decomposition. For example, if A has eigendecomposition $A = X\Lambda X^{-1}$ with $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_n)$, then $s(A) = X s(\Lambda) X^{-1}$, where $s(\Lambda) = \text{diag}(s(\lambda_1), \dots, s(\lambda_n))$.

The functional calculus for unbounded normal operators ensures that if λ_i is an eigenvalue of \mathcal{L} with eigenfunction u_i , then $s(\lambda_i)$ is an eigenvalue of $s(\mathcal{L})$ with eigenfunction u_i [44, VIII.5].⁸ As the degree of the quadrature rule is increased, the rational function becomes an increasingly good approximation to the Cauchy integral

$$(5.2) \quad \chi(z) = \frac{1}{2\pi i} \int_{\partial\Omega} \frac{dw}{w-z}, \quad z \in \mathbb{C} \setminus \partial\Omega.$$

Therefore, the eigenvalues of \mathcal{L} in Ω are usually $\mathcal{O}(1)$ in size under the spectral transformation $s(\cdot)$ while the eigenvalues outside of Ω are much smaller.

We now turn to the convergence of the iteration described in Algorithm 4.1, which we interpret as a subspace iteration applied to the bounded linear operator $\hat{\mathcal{P}}_{\mathcal{V}} = s(\mathcal{L})$. It is helpful to introduce the notions of the *spectral radius* and a *dominant eigenspace* of a bounded linear operator \mathcal{B} . The spectral radius of a bounded linear operator \mathcal{B} on a Hilbert space \mathcal{H} is defined as [16, p. 99]

$$(5.3) \quad \rho(\mathcal{B}) = \max\{|z| : z \in \lambda(\mathcal{B})\}.$$

The spectral radius is useful because it characterizes the asymptotic behavior of $\|\mathcal{B}^k\|_{\mathcal{H}}$, in the sense that $\rho(\mathcal{B}) = \lim_{k \rightarrow \infty} \|\mathcal{B}^k\|_{\mathcal{H}}^{1/k}$ [16, Theorem 4.1.3]. Let \mathcal{V} be an invariant subspace of \mathcal{B} associated with eigenvalues $\lambda_1 \geq \dots \geq \lambda_m$ and a spectral projector $\mathcal{P}_{\mathcal{V}}$. We say that \mathcal{B} has dominant eigenspace \mathcal{V} if

$$(5.4) \quad \rho((\mathcal{I} - \mathcal{P}_{\mathcal{V}})\mathcal{B}) < |\lambda_m|.$$

The following theorem is an extension of a convergence analysis [46, p. 119] for matrix subspace iteration to the setting of bounded linear operators with a dominant eigenspace. We omit the details of the proof, as they are identical to those found in the proof of Lemmas 3.1 and 3.2 of [26].

THEOREM 5.1. *Let \mathcal{B} be a bounded linear operator on a Hilbert space \mathcal{H} with dominant eigenspace \mathcal{V} , defined in (5.4), having $\dim(\mathcal{V}) = m$. Select a quasimatrix $F : \mathbb{C}^m \rightarrow \mathcal{H}$ such that the columns of $\mathcal{P}_{\mathcal{V}}F$ are linearly independent, and suppose the columns of the quasimatrix $\hat{Q}_k : \mathbb{C}^m \rightarrow \mathcal{H}$ form an orthonormal basis for $\text{range}(\mathcal{B}^k F)$, for $k = 1, 2, 3, \dots$. If $u \in \mathcal{V}$ is an eigenvector of \mathcal{B} with eigenvalue λ , then there is a function $\hat{u}_k \in \text{range}(\hat{Q}_k)$ such that*

$$\|\hat{u}_k - u\|_{\mathcal{H}} \leq (|\rho/\lambda| + \epsilon_k)^k \|(I - \mathcal{P}_{\mathcal{V}})Fx\|_{\mathcal{H}}, \quad k = 1, 2, 3, \dots,$$

where $\rho = \rho((\mathcal{I} - \mathcal{P}_{\mathcal{V}})\mathcal{B})$, $\epsilon_k \rightarrow 0$ as $k \rightarrow \infty$, and $u = \mathcal{P}_{\mathcal{V}}Fx$.

Although we have neglected the effects of approximately solving the ODEs in (4.2) and the impact of round-off errors in our brief analysis of rational subspace iteration for normal differential operators, we mention two recent results for rational subspace iteration with matrices [47] and self-adjoint differential operators [26, 27].

- For matrices, small errors made during application of the spectral projector generally do not alter the convergence behavior of subspace iteration [47]. In this case, the sequence \hat{Q}_k no longer converges to an exact basis for \mathcal{V} . However, the matrices \hat{Q}_k approximate a basis for \mathcal{V} and the approximation error converges geometrically to a constant determined by the sizes of the errors introduced at each iteration [47].

⁸The result [44, VIII.5] is stated for closed self-adjoint operators on \mathcal{H} ; however, it extends immediately to closed normal operators on \mathcal{H} if the spectral decomposition of a closed normal operator [45, Theorem 13.33] is used in place of the spectral decomposition of a self-adjoint operator. For information on the spectral decomposition of unbounded normal operators and the associated functional calculus, see [45, Chap. 13] and [19].

- For self-adjoint differential operators (closed and densely defined on \mathcal{H}), rational subspace iteration converges to a subspace even when the resolvent operator is discretized to solve the ODEs in (4.2) [26, 27]. The distance between the computed subspace and the target eigenspace (in a distance metric between subspaces) is proportional to the approximation error in the discretized resolvent [26, 27].

We expect that similar statements hold for normal operators on \mathcal{H} , but a rigorous and detailed convergence analysis is more subtle and beyond the scope of this paper.

5.2. A pseudospectral inclusion theorem. As $\text{range}(\hat{Q})$ is not an invariant subspace of \mathcal{L} , the ϵ -pseudospectrum of $\hat{Q}^* \mathcal{L} \hat{Q}$ is not, in general, contained in the ϵ -pseudospectrum of \mathcal{L} . However, if $\|\hat{Q} - Q\|_{\mathbb{C}^m \rightarrow \mathcal{H}}$ is sufficiently small, then the ϵ -pseudospectrum of $\hat{Q}^* \mathcal{L} \hat{Q}$ is contained in the 2ϵ -pseudospectrum of \mathcal{L} .

THEOREM 5.2. *Consider a closed operator \mathcal{L} with domain $\mathcal{D}(\mathcal{L})$ that is densely defined on a Hilbert space \mathcal{H} and fix $\epsilon > 0$. Let $Q : \mathbb{C}^m \rightarrow \mathcal{D}(\mathcal{L}) \cap \mathcal{D}(\mathcal{L}^*)$ satisfy $Q^* Q = I$, and let $\text{range}(Q)$ be an m -dimensional invariant subspace of \mathcal{L} . If $\hat{Q} : \mathbb{C}^m \rightarrow \mathcal{D}(\mathcal{L})$ satisfies*

$$\|\hat{Q} - Q\|_{\mathbb{C}^m \rightarrow \mathcal{H}} \left(\|\mathcal{L}^* Q\|_{\mathbb{C}^m \rightarrow \mathcal{H}} + \|\mathcal{L} Q\|_{\mathbb{C}^m \rightarrow \mathcal{H}} + \|\mathcal{L}(\hat{Q} - Q)\|_{\mathbb{C}^m \rightarrow \mathcal{H}} \right) < \frac{\epsilon}{2},$$

then $\lambda_\epsilon(\hat{Q}^* \mathcal{L} \hat{Q}) \subset \lambda_{2\epsilon}(\mathcal{L})$.

Proof. Consider $z \in \lambda_\epsilon(\hat{Q}^* \mathcal{L} \hat{Q})$. If $z \in \lambda_\epsilon(\hat{Q}^* \mathcal{L} \hat{Q}) \cap \lambda_\epsilon(\mathcal{L})$, there is nothing to prove, so assume without loss of generality that $z \notin \lambda_\epsilon(\mathcal{L})$. If we denote $R_Q(z) = (zI - Q^* \mathcal{L} Q)^{-1}$, $R_{\hat{Q}}(z) = (zI - \hat{Q}^* \mathcal{L} \hat{Q})^{-1}$, and $E = \hat{Q} - Q$, then we have that $R_{\hat{Q}}(z) = [R_Q(z)^{-1} - B]^{-1}$, where $B = Q^* \mathcal{L} E + E^* \mathcal{L} Q + E^* \mathcal{L} E$. Employing a formula for the inverse of the sum of two matrices, we obtain $R_{\hat{Q}}(z) = R_Q(z) + R_Q(z)[I - BR_Q(z)]^{-1}BR_Q(z)$ [30].

Now, $\|B\|_{\mathbb{C}^m} \leq \|Q^* \mathcal{L} E\|_{\mathbb{C}^m} + \|E^* \mathcal{L} Q\|_{\mathbb{C}^m} + \|E^* \mathcal{L} E\|_{\mathbb{C}^m}$. Since $\|E^*\|_{\mathcal{H} \rightarrow \mathbb{C}^m} = \|E\|_{\mathbb{C}^m \rightarrow \mathcal{H}}$ and $\|Q^* \mathcal{L}\|_{\mathcal{H} \rightarrow \mathbb{C}^m} = \|\mathcal{L}^* Q\|_{\mathbb{C}^m \rightarrow \mathcal{H}}$ [31, p. 256], our hypothesis indicates that the sum of the three terms comprising B is bounded in norm by

$$\|B\|_{\mathbb{C}^m} \leq \|E\|_{\mathbb{C}^m \rightarrow \mathcal{H}} (\|\mathcal{L}^* Q\|_{\mathbb{C}^m \rightarrow \mathcal{H}} + \|\mathcal{L} Q\|_{\mathbb{C}^m \rightarrow \mathcal{H}} + \|\mathcal{L} E\|_{\mathbb{C}^m \rightarrow \mathcal{H}}) < \frac{\epsilon}{2}.$$

Moreover, since $z \notin \lambda_\epsilon(\mathcal{L})$, we have that $\|R_Q(z)\|_{\mathbb{C}^m} \leq 1/\epsilon$ by Theorem 3.2. Therefore, $\|BR_Q(z)\|_{\mathbb{C}^m} \leq 1/2$.

Because $\|BR_Q(z)\|_{\mathbb{C}^m} \leq 1/2$, we may use the Neumann series to compute $(I - BR_Q(z))^{-1} = \sum_{k=0}^{\infty} (BR_Q(z))^k$. We see that $R_{\hat{Q}}(z) = R_Q(z)(I + \sum_{k=1}^{\infty} (BR_Q(z))^k)$ and therefore,

$$\|R_{\hat{Q}}(z)\|_{\mathbb{C}^m} \leq \left(1 + \sum_{k=1}^{\infty} \frac{1}{2^k} \right) \|R_Q(z)\|_{\mathbb{C}^m} = 2 \|R_Q(z)\|_{\mathbb{C}^m}.$$

Now, if $z \in \lambda_\epsilon(\hat{Q}^* \mathcal{L} \hat{Q})$, then $\|R_Q(z)\|_{\mathbb{C}^m} \geq \|R_{\hat{Q}}(z)\|_{\mathbb{C}^m}/2 > 1/(2\epsilon)$. By Theorem 3.2, we have that $\|(zI - \mathcal{L})^{-1}\|_{\mathcal{H}} \geq \|R_Q(z)\|_{\mathbb{C}^m}$. Collecting inequalities yields the result $\|(zI - \mathcal{L})^{-1}\|_{\mathcal{H}} > 1/(2\epsilon)$, i.e., $z \in \lambda_{2\epsilon}(\mathcal{L})$. \square

A consequence of Theorem 5.2 is that Algorithm 4.1 possesses a type of stability provided that \mathcal{L} is uniformly bounded on the sequence E_1, E_2, E_3, \dots , where $E_k = \hat{Q}_k - Q$ for $k \geq 1$. If \mathcal{L} is uniformly bounded on $\{E_k\}_{k=1}^{\infty}$, then there is a $\Lambda \geq 0$ such that $\sup_{k \geq 1} \|\mathcal{L} E_k\|_{\mathbb{C}^m \rightarrow \mathcal{H}} \leq \Lambda$. Applying Theorem 5.2, we see that Algorithm 4.1

computes elements in the 2ϵ -pseudospectrum of \mathcal{L} provided that a basis for \mathcal{V} is resolved to within $\epsilon/(2(\|\mathcal{L}^*Q\|_{\mathbb{C}^m \rightarrow \mathcal{H}} + \|\mathcal{L}Q\|_{\mathbb{C}^m \rightarrow \mathcal{H}} + \Lambda))$.

We now verify, with two mild constraints placed on the choice of the initial quasimatrix F , that \mathcal{L} is uniformly bounded on the sequence $\{\hat{Q}_k\}_{k=1}^\infty$ generated by Algorithm 4.1. Note that this implies that \mathcal{L} is uniformly bounded on $\{E_k\}_{k=1}^\infty$ because $E_k = \hat{Q}_k - Q$ and $\text{range}(Q) \subset \mathcal{D}(\mathcal{L})$. The constraints on F are generically satisfied when F is selected as in section 4. In the statement of the bound on $\|\mathcal{L}\hat{Q}_k\|_{\mathbb{C}^m \rightarrow \mathcal{H}}$, we use the notation $\sigma_{\min}(\mathcal{P}_V F)$ and $\sigma_{\min}((\mathcal{I} - \mathcal{P}_V)F)$ to denote the smallest singular values of the quasimatrices $\mathcal{P}_V F$ and $(\mathcal{I} - \mathcal{P}_V)F$, respectively.⁹

LEMMA 5.3. *Consider a closed, normal operator \mathcal{L} with domain $\mathcal{D}(\mathcal{L})$ that is densely defined on a Hilbert space \mathcal{H} . Let $\hat{\mathcal{P}}_V$ be the bounded operator on \mathcal{H} defined in (4.1) and suppose that $\hat{\mathcal{P}}_V$ has a dominant eigenspace of \mathcal{V} (see (5.4)) with $\dim(\mathcal{V}) = m$. Let F , \mathcal{P}_V , and $\{\hat{Q}_k\}_{k=1}^\infty$ be as in Theorem 5.1 with $\mathcal{B} = \hat{\mathcal{P}}_V$. Suppose that $\hat{\mathcal{P}}_V^k F$ (for each $k \geq 1$) and $(\mathcal{I} - \mathcal{P}_V)F$ each have linearly independent columns and that $\text{range}(F) \subset \mathcal{D}(\mathcal{L})$. Then, we have that*

$$\|\mathcal{L}\hat{Q}_k\|_{\mathbb{C}^m \rightarrow \mathcal{H}} \leq 2M\|\mathcal{L}F\|_{\mathbb{C}^m \rightarrow \mathcal{H}}, \quad k = 1, 2, 3, \dots,$$

where $M = \max\{1/\sigma_{\min}(\mathcal{P}_V F), 1/\sigma_{\min}((\mathcal{I} - \mathcal{P}_V)F)\}$.

Proof. Since \hat{Q}_k is an orthonormal basis for $\hat{\mathcal{P}}_V^k F$, there is a matrix $R_k \in \mathbb{C}^{m \times m}$ such that $\hat{\mathcal{P}}_V^k F = \hat{Q}_k R_k$. By the assumption that $\hat{\mathcal{P}}_V^k F$ has linearly independent columns, we know that R_k is invertible. We obtain that

$$(5.5) \quad \hat{Q}_k = \hat{\mathcal{P}}_V^k F R_k^{-1}.$$

We use the spectral projector \mathcal{P}_V to rewrite (5.5) as

$$(5.6) \quad \hat{Q}_k = \hat{\mathcal{P}}_V^k (\mathcal{P}_V F + (\mathcal{I} - \mathcal{P}_V)F) R_k^{-1}.$$

Now, $\text{range}(\mathcal{P}_V F)$ and $\text{range}((\mathcal{I} - \mathcal{P}_V)F)$ are invariant under $\hat{\mathcal{P}}_V$ [31, p. 178]. Consequently, there are matrices $D_1, D_2 \in \mathbb{C}^{m \times m}$ such that

$$(5.7) \quad \hat{\mathcal{P}}_V^k \mathcal{P}_V F = \mathcal{P}_V F D_1^k, \quad \hat{\mathcal{P}}_V^k (\mathcal{I} - \mathcal{P}_V)F = (\mathcal{I} - \mathcal{P}_V)F D_2^k.$$

Substituting (5.7) into (5.6) yields the following useful equation for \hat{Q}_k :

$$(5.8) \quad \hat{Q}_k = (\mathcal{P}_V F D_1^k + (\mathcal{I} - \mathcal{P}_V)F D_2^k) R_k^{-1}.$$

Applying \mathcal{L} to both sides of (5.8) and commuting with the spectral projectors \mathcal{P}_V and $\mathcal{I} - \mathcal{P}_V$ [31, p. 179], we obtain

$$(5.9) \quad \mathcal{L}\hat{Q}_k = (\mathcal{P}_V \mathcal{L} F D_1^k + (\mathcal{I} - \mathcal{P}_V) \mathcal{L} F D_2^k) R_k^{-1}.$$

Since $\text{range}(F) \subset \mathcal{D}(\mathcal{L})$, we have that $\|\mathcal{L}F\|_{\mathbb{C}^m \rightarrow \mathcal{H}} < \infty$. Additionally, since \mathcal{L} is normal, the spectral projectors have norms equal to 1 [31, p. 277]. Therefore, it remains to find a uniform bound for $\|D_1^k R_k^{-1}\|_{\mathbb{C}^m}$ and $\|D_2^k R_k^{-1}\|_{\mathbb{C}^m}$ as $k \rightarrow \infty$.

⁹The singular value decomposition of a quasimatrix $A : \mathbb{C}^m \rightarrow \mathcal{H}$ is the decomposition $A = U\Sigma V^*$, where $U : \mathbb{C}^m \rightarrow \mathcal{H}$ is a quasimatrix with \mathcal{H} -orthonormal columns, $\Sigma \in \mathbb{C}^{m \times m}$ is a diagonal matrix with nonnegative entries $\sigma_1 \geq \dots \geq \sigma_m$, and $V \in \mathbb{C}^{m \times m}$ is a unitary matrix [56].

For brevity, we prove uniform boundedness of $\|D_1^k R_k^{-1}\|_{\mathbb{C}^m}$ and note that the proof for $\|D_2^k R_k^{-1}\|_{\mathbb{C}^m}$ is essentially identical. We begin by commuting $\hat{\mathcal{P}}_{\mathcal{V}}$ with the spectral projectors in (5.7) and substituting the QR factorization of $\hat{\mathcal{P}}_{\mathcal{V}}^k F$ to see that

$$(5.10) \quad \mathcal{P}_{\mathcal{V}} \hat{Q}_k = (\mathcal{P}_{\mathcal{V}} F) D_1^k R_k^{-1}.$$

Using the pseudoinverse $(\mathcal{P}_{\mathcal{V}} F)^+$ of the quasimatrix¹⁰ $\mathcal{P}_{\mathcal{V}} F$ and noting that $\mathcal{P}_{\mathcal{V}} F$ has linearly independent columns, (5.10) implies that

$$(5.11) \quad D_1^k R_k^{-1} = (\mathcal{P}_{\mathcal{V}} F)^+ \mathcal{P}_{\mathcal{V}} \hat{Q}_k.$$

Now, we know that $\|\mathcal{P}_{\mathcal{V}} \hat{Q}_k\|_{\mathcal{H}} \leq 1$, because $\|\mathcal{P}_{\mathcal{V}}\|_{\mathcal{H}} = 1$ and \hat{Q}_k has orthonormal columns. We conclude that

$$(5.12) \quad \|D_1^k R_k^{-1}\|_{\mathbb{C}^m} \leq \frac{1}{\sigma_{\min}(\mathcal{P}_{\mathcal{V}} F)}.$$

A similar argument shows that

$$(5.13) \quad \|D_2^k R_k^{-1}\|_{\mathbb{C}^m} \leq \frac{1}{\sigma_{\min}((\mathcal{I} - \mathcal{P}_{\mathcal{V}}) F)}.$$

Taking norms in (5.9) and substituting the bounds from (5.12) and (5.13), we find

$$(5.14) \quad \|\mathcal{L} \hat{Q}_k\|_{\mathbb{C}^m \rightarrow \mathcal{H}} \leq \|\mathcal{L} F\|_{\mathbb{C}^m \rightarrow \mathcal{H}} \left(\frac{1}{\sigma_{\min}(\mathcal{P}_{\mathcal{V}} F)} + \frac{1}{\sigma_{\min}((\mathcal{I} - \mathcal{P}_{\mathcal{V}}) F)} \right).$$

The lemma follows immediately from (5.14). \square

Theorem 5.1, Theorem 5.2, and Lemma 5.3 provide a preliminary analysis to explain why Algorithm 4.1 accurately computes the eigenvalues of normal operators with a dominant eigenspace \mathcal{V} . Theorem 5.1 allows us to accurately resolve an orthonormal basis Q for \mathcal{V} by refining the quasimatrix \hat{Q}_k with subspace iteration. Lemma 5.3 confirms that $\mathcal{L} \hat{Q}_k$ does not grow without bound as \hat{Q}_k is refined. Finally, Theorem 5.2 demonstrates that the eigenvalues are computed to the expected accuracy, provided that the basis for \mathcal{V} has been resolved.

6. An operator analogue of the Rayleigh Quotient Iteration. It is useful to have operator analogues for other eigensolvers too; particularly, when the eigenvalues of interest are difficult to target with a pre-selected search region $\Omega \subset \mathbb{C}$. The Rayleigh Quotient Iteration (RQI) is a generalization of the inverse iteration that incorporates dynamic shifting to obtain cubic (for Hermitian problems) or quadratic (non-Hermitian problems) convergence [42]. Given a matrix $A \in \mathbb{C}^{n \times n}$ and an initial vector $\tilde{y}_0 \in \mathbb{C}^n$, RQI computes the iterates

$$(6.1) \quad \tilde{y}_{k+1} = (A - \beta_k I)^{-1} y_k, \quad \beta_k = y_k^* A y_k, \quad y_k = \frac{\tilde{y}_k}{\|\tilde{y}_k\|_2}, \quad k = 0, 1, 2, \dots$$

The vectors y_k typically converge to a nearby eigenvector of A , while the sequence β_k converges to the associated eigenvalue of A [40]. In the matrix setting, (6.1)

¹⁰The pseudoinverse of a quasimatrix $A : \mathbb{C}^m \rightarrow \mathcal{H}$ may be defined via the SVD as $A^+ = V \Sigma^+ U^*$, where Σ^+ is the diagonal matrix with entries $\Sigma_{ii}^+ = 1/\sigma_i$ if $\sigma_i \neq 0$ and 0 otherwise. It is easy to verify familiar properties from the matrix case [24, p. 290], i.e., if A has linearly independent columns, then $A^+ A = I$ and $\|A^+\|_{\mathcal{H} \rightarrow \mathbb{C}^m} = 1/\sigma_{\min}(A)$.

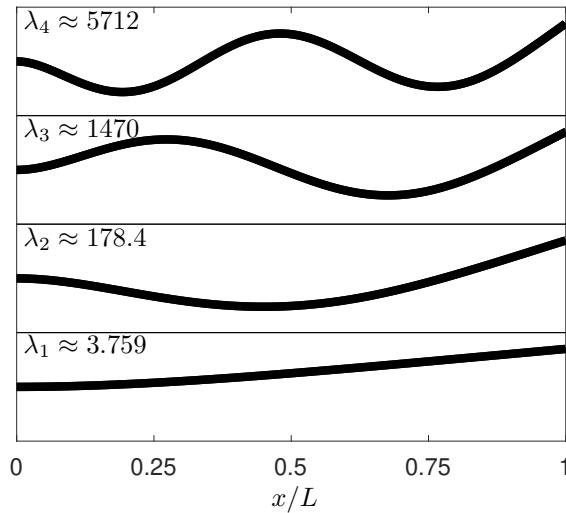


FIG. 6.1. Selected free-vibration modes of an airplane wing modeled by (6.3).

is often used to compute interior eigenvalues or refine an estimate of an invariant subspace [41, 42].

Replacing a matrix A by a differential operator $\mathcal{L} : \mathcal{D}(\mathcal{L}) \rightarrow \mathcal{H}$, as in (1.1), and the vectors \tilde{y}_k by functions $f_k \in \mathcal{D}(\mathcal{L})$, we obtain an operator analogue of RQI. One needs to select an initial function $f_0 \in \mathcal{D}(\mathcal{L})$ and solve a sequence of ODEs, i.e.,

$$(6.2) \quad (\mathcal{L} - \beta_k I) f_{k+1} = f_k, \quad f_{k+1}(\pm 1) = \dots = f_{k+1}^{(N/2)}(\pm 1) = 0.$$

At each iteration, the shift β_k is computed from the Rayleigh quotient $(f_k, \mathcal{L}f_k)_{\mathcal{H}}$ (in strong form) and the solution f_{k+1} is normalized after each iteration. Analogous to the matrix setting, we observe that the operator analogue of RQI converges cubically for self-adjoint operators and quadratically otherwise [28].

We note that block generalizations of RQI (e.g., RSQR and GRQI [1]) are also easily extended to the differential operator setting. In this case, a sequence of quasi-matrices \hat{Q}_k with \mathcal{H} -orthonormal columns are generated to approximate an invariant subspace of \mathcal{L} and a Rayleigh–Ritz projection is performed to compute approximate eigenvalues and eigenvectors. As with the operator analogue of FEAST, Theorem 5.2 implies that the iteration (6.2) accurately computes eigenvalues of normal differential operators when the basis for the target eigenspace is sufficiently resolved.

6.1. Free vibrations of an airplane wing. The improved convergence rate of RQI can offer much faster computation time than subspace iteration, often requiring only three or four ODE solves to reach an accuracy of essentially machine precision [28]. We now employ (6.2) for the rapid computation of vibrational modes of an airplane wing.

An airplane wing may be crudely modeled as a thin, cantilevered beam of length L with a linear taper. The governing equation for free vibrations is [29]

$$(6.3) \quad \frac{d^2}{dx^2} \left((1+x) \frac{d^2u}{dx^2} \right) = \lambda u, \quad u(0) = u'(0) = 0, \quad u''(L) = u'''(L) = 0.$$

The variable coefficient $1+x$ accounts for the linear taper of the wing, while the

boundary conditions on u'' and u''' at $x = 1$ express the natural requirement that the bending moment and shear force vanish at the endpoint.

To compute a few of the smoothest modes of (6.3) we use the eigenfunctions w_n of the cantilevered beam equation with constant coefficients, given in closed form by [29]

$$(6.4) \quad w_n(x) = \cosh \beta_n x - \cos \beta_n x + \frac{\cos \beta_n L + \cosh \beta_n L}{\sin \beta_n L + \sinh \beta_n L} (\sin \beta_n x + \sinh \beta_n x).$$

Here β_n is the n th root of $g(\beta) = \cosh(\beta L) \cos(\beta L) + 1$ [29]. We target a mode of (6.3) by setting $f_0(x) = w_n(x)$. Figure 6.1 shows the modes that are computed using initial guesses w_1, \dots, w_4 , corresponding to the smallest four positive roots of $g(\beta)$.

7. Computing eigenvalues in unbounded regions. The stability analysis of solutions to time-dependent partial differential equations (PDEs) provides an abundant source of differential eigenvalue problems. Consider the initial boundary value problem (IBVP) with periodic boundary conditions

$$(7.1) \quad u_t = \mathcal{L}u + \mathcal{N}(u), \quad u_t(x, 0) = g(x), \quad u(-1, t) = u(1, t).$$

Here, \mathcal{L} and \mathcal{N} are linear and nonlinear ordinary differential operators (with respect to the variable x), respectively. In many instances, (7.1) supports steady-states, traveling wave states, or other phenomena whose stability is of critical importance in the physical problem under study [4, 34, 48]. When \mathcal{L} is self-adjoint or normal, the stability analysis often reduces to determining whether or not the eigenvalues of \mathcal{L} are contained in one half-plane [4, 33, 48, 59]. We now show how to modify the spectral projector in (3.1) to derive a practical rational filter to compute (finitely many) eigenvalues of \mathcal{L} in the right half-plane.

7.1. A rational filter for the half-plane. Let \mathcal{L} be a closed linear operator that is densely defined on a Hilbert space \mathcal{H} . Suppose that \mathcal{L} is a normal operator with a spectrum in the left half-plane $\text{Re}(z) < 0$ except for finitely many eigenvalues $\lambda_1, \dots, \lambda_m$ (including multiplicities) such that $\text{Re}(\lambda_i) > 0$ for $1 \leq i \leq m$. Denote the eigenspace associated with $\lambda_1, \dots, \lambda_m$ by \mathcal{V} and consider search regions that are semicircles of radius R , i.e.,

$$(7.2) \quad \Omega_R = \{z \in \mathbb{C} : |z| < R, \text{Re}(z) > 0\}, \quad R > \max_{1 \leq i \leq m} |\lambda_i|.$$

To construct a computable spectral projector onto the right half-plane we consider taking $R \rightarrow \infty$. We adopt the following strategy:

- (i) Introduce a $1/R$ decay into the integrand of the spectral projector (3.1) as $R \rightarrow \infty$, while preserving the projection onto \mathcal{V} .
- (ii) Split the projector into an integral over the vertical part of $\partial\Omega_R$ and an integral over the circular arc of $\partial\Omega_R$. By taking $R \rightarrow \infty$, we observe that the contribution from the circular arc goes to 0 due to the additional $1/R$ decay in the integrand.
- (iii) Map the imaginary axis to the interval $[-1, 1]$ and approximate the spectral projector by a quadrature rule.

Select $a \in \mathbb{R}^+$ and consider the family of functions that are analytic in the right half-plane defined by

$$(7.3) \quad \mathcal{P}_R(\lambda) = \frac{1}{2\pi i} \int_{\partial\Omega_R} (z + a)^{-1} (z - \lambda)^{-1} dz.$$

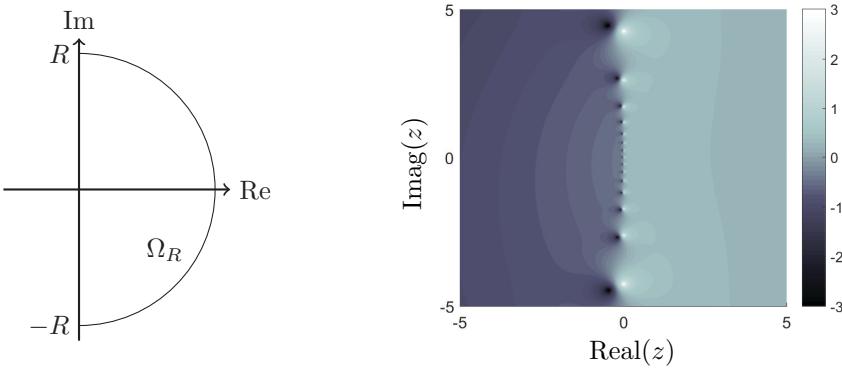


FIG. 7.1. Left: The region Ω_R from (7.2) used in the derivation of the rational filter over the right half-plane. Right: The constructed rational filter (7.6) for the right half-plane with $\ell = 20$.

By Cauchy's Integral Formula, we know that $\mathcal{P}_R(\lambda) = (\lambda + a)^{-1}$ if $\lambda \in \Omega_R$ and is zero otherwise [50]. Taking the limit $R \rightarrow \infty$, we obtain

$$(7.4) \quad \mathcal{P}(\lambda) = \lim_{R \rightarrow \infty} \mathcal{P}_R(\lambda) = \frac{1}{2\pi} \int_{-\infty}^{\infty} (iy + a)^{-1} (iy - \lambda)^{-1} dy.$$

Using functional calculus for unbounded normal operators we can extend $\mathcal{P}(\lambda)$ to an operator-valued function $\mathcal{P}(\mathcal{L})$ [45, Theorem 13.24].¹¹ Moreover, we have that $\mathcal{P}(\mathcal{L})u = \mathcal{P}(\lambda)u$ when $\mathcal{L}u = \lambda u$ [44, VIII.5]. Consequently, $\text{range}(\mathcal{P}(\mathcal{L})) = \mathcal{V}$.

Now, take the change-of-variables $x = \frac{2}{\pi} \tan^{-1} y$ in (7.4) to obtain

$$(7.5) \quad \mathcal{P}(\mathcal{L}) = \frac{1}{4} \int_{-1}^1 \left(i \tan\left(\frac{\pi x}{2}\right) + a \right)^{-1} \left(i \tan\left(\frac{\pi x}{2}\right) \mathcal{I} - \mathcal{L} \right)^{-1} \sec^2\left(\frac{\pi x}{2}\right) dx.$$

Using Gauss-Legendre quadrature nodes x_1, \dots, x_ℓ and weights w_1, \dots, w_ℓ on $[-1, 1]$, we can approximate $\mathcal{P}(\mathcal{L})$ by

$$(7.6) \quad \hat{\mathcal{P}}(\mathcal{L}) = \frac{1}{4} \sum_{k=1}^{\ell} w_k \frac{1 - z_k^2}{z_k + a} (z_k \mathcal{I} - \mathcal{L})^{-1}, \quad z_k = i \tan\left(\frac{\pi x_k}{2}\right).$$

Figure 7.1 (right) shows the derived rational filter $\hat{\mathcal{P}}(\lambda)$ in the complex plane.

7.2. Stability of thin fluid films. To demonstrate the utility of the filter in (7.6), we assess the stability of the steady-state solutions to a PDE governing the motion of a thin film of fluid supported below by a flat substrate. The PDE is

$$(7.7) \quad u_t = \partial_x^4 u + \partial_x(u \partial_x u),$$

where $u(x, t)$ is a positive, periodic function representing the thickness of the fluid [34]. The nonlinear term models gravitational effects and substrate-fluid interactions [34].

A droplet steady-state $u_{ss}(x)$ of (7.7), rescaled so that it is supported on $[0, l]$ with contact angle $\pi/4$, is stable if all the eigenvalues of a fourth-order differential operator are in the left half-plane. The associated differential eigenproblem is [33]

$$(7.8) \quad \frac{d^4 u}{dx^4} + \frac{d}{dx} \left(u_{ss} \frac{du}{dx} \right) = \lambda u, \quad u(0) = u(l) = 0, \quad u''(0) = u''(l) = 0.$$

¹¹In [45, Theorem 13.24], $E_{x,y}$ is the spectral measure of \mathcal{L} [45, Theorem 13.33].

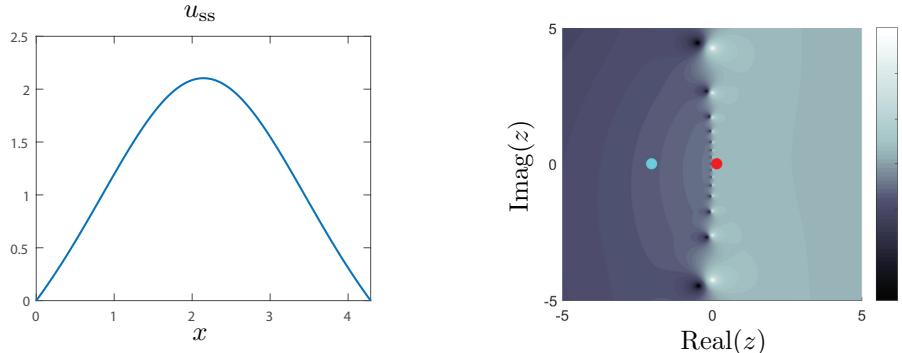


FIG. 7.2. Left: A droplet, u_{ss} , which is a steady-state solution to (7.7), computed from the nonlinear ODE in (7.9). Right: Two rightmost eigenvalues (blue and red dots) of (7.8) together with a log-scale colormap of the rational filter in (7.6) with $\ell = 20$, which is used in place of (4.1). The eigenvalue with a positive real part (red dot) indicates that this steady-state droplet is unstable.

We compute the steady-state $u_{ss}(x)$ by solving the second-order nonlinear ODE [34]

$$(7.9) \quad \frac{du_{ss}}{dx} + \frac{1}{2}u_{ss}^2 - \delta = 0, \quad u_{ss}(0) = 0, \quad u'_{ss}(0) = 1.$$

Here, δ is a dimensionless quantity relating the rescaled problem to the original contact angle [34]. The length l of the droplet's base and δ may be calculated analytically [34].

In Figure 7.2, we show an approximation to the rescaled steady-state u_{ss} along with the rightmost eigenvalues of (7.8). Using the rational filter in (7.6) with $\ell = 20$ (the degree of the quadrature rule defining the filter) to perform the approximate spectral projection in Algorithm 4.1, we are able to identify an eigenvalue of (7.8) in the right half-plane, which indicates that the droplet (see Figure 7.2 (left)) is unstable.

Techniques for selecting the dimension m of the subspace \mathcal{V} [32, 53] are important in stability analysis as one is trying to determine the number of eigenvalues in the right half-plane. To select m , we monitor the singular values of the matrix $\hat{V}^* \hat{V}$ after each iteration and adjust the number of basis functions by removing columns of \hat{V} associated with singular values that are close to machine precision (relative to the largest singular value) [53]. This procedure usually allows us to capture the dominant eigenspace of the filtered operator $\hat{\mathcal{P}}(\mathcal{L})$ that includes the target eigenspace as well as any eigenvalues clustered near the imaginary axis. We then determine whether there are any eigenvalues in the right half-plane by sorting through the computed eigenvalues. However, this strategy may break down, for instance, if there is an eigenvalue close to a quadrature node. Additionally, the sharp decay of the filter (7.6) across the imaginary axis is softened as $|\text{Im}(z)| \rightarrow \infty$, which can lead to difficulties when there are clusters of eigenvalues near the imaginary axis with large imaginary part. In this case, one may need to take a large number of basis functions to accurately resolve the dominant eigenvalues of $\hat{\mathcal{P}}(\mathcal{L})$.

Conclusions. An operator analogue of the FEAST matrix eigensolver is derived to solve differential eigenvalue problems without discretizing the operator. This approach leads to an algorithm that can exploit spectrally accurate techniques for computing with functions while preserving the structure of \mathcal{L} . The result is an efficient, automated, and accurate eigensolver for normal and self-adjoint differential operators. This eigensolver is adept in the high-frequency regime and may provide a new direction towards robust high-frequency eigenvalue computations.

The implementation described in section 4 extends to higher dimensions in a straightforward way for simple geometries where spectral methods apply [55]. In the case of more complicated geometries, one may still benefit from the advantages of the “solve-then-discretize” framework outlined in Algorithm 4.1 provided that one is able to accurately compute inner products and solutions to the shifted linear differential equations. Although we have focused on the strong form of the eigenvalue problem in (1.1), it may be necessary to work with the weak form. In this case, one can still follow Algorithm 4.1 provided that the shifted linear systems are solved in weak form and the Rayleigh–Ritz projection is performed with the associated sesquilinear form.

Acknowledgments. We thank Anthony Austin for providing detailed feedback on the results in section 3 and on our rational filter in section 7. We thank David Bindel for helpful discussions about spectral discretizations of differential eigenvalue problems. Finally, we thank Marcus Webb, Sheehan Olver, Yuji Nakatsukasa, Heather Wilber, and Matthew Colbrook for providing comments on a draft.

REFERENCES

- [1] P.-A. ABSIL, R. SEPULCHRE, P. VAN DOOREN, AND R. MAHONY, *Cubically convergent iterations for invariant subspace computation*, SIAM J. Matrix Anal. Appl., 26 (2004), pp. 70–96, <https://doi.org/10.1137/S0895479803422002>.
- [2] A. J. AKBARFAM AND A. MINGARELLI, *Higher order asymptotics of the eigenvalues of Sturm–Liouville problems with a turning point of arbitrary order*, Can. Appl. Math. Q., 12 (2004), pp. 275–301.
- [3] A. J. AKBARFAM AND A. B. MINGARELLI, *Higher order asymptotic distribution of the eigenvalues of nondefinite Sturm–Liouville problems with one turning point*, J. Comput. Appl. Math., 149 (2002), pp. 423–437.
- [4] N. ALIKAKOS, P. BATES, AND X. CHEN, *Periodic traveling waves and locating oscillating patterns in multidimensional domains*, Trans. Amer. Math. Soc., 351 (1999), pp. 2777–2805.
- [5] L. M. ANGUAS, M. I. BUENO, AND F. M. DOPICO, *A comparison of eigenvalue condition numbers for matrix polynomials*, Linear Algebra Appl., 564 (2019), pp. 170–200, <https://doi.org/https://doi.org/10.1016/j.laa.2018.11.031>.
- [6] F. V. ATKINSON AND A. B. MINGARELLI, *Asymptotics of the number of zeros and of the eigenvalues of general weighted Sturm–Liouville problems*, J. Reine Angew. Math., 375 (1987), pp. 380–393.
- [7] A. AUSTIN, *Eigenvalues of Differential Operators by Contour Integral Projection*, <http://www.chebfun.org/examples/ode-eig/ContourProjEig.html>, May 2013.
- [8] Z. BAI, J. DEMMEL, J. DONGARRA, A. RUHE, AND H. VAN DER VORST, *Templates for the Solution of Algebraic Eigenvalue Problems: A Practical Guide*, SIAM, Philadelphia, 2000, <https://doi.org/10.1137/1.9780898719581>.
- [9] A. BARNETT AND A. HASSELL, *Fast computation of high-frequency Dirichlet eigenmodes via spectral flow of the interior Neumann-to-Dirichlet map*, Commun. Pure Appl. Math., 67 (2014), pp. 351–407.
- [10] D. BINDEL AND M. ZWORSKI, *Theory and Computation of Resonances in 1D Scattering*, <http://www.cs.cornell.edu/~bindel/cims/resonant1d/theo2.html>, October 2006.
- [11] I. BOGAERT, *Iteration-free computation of Gauss–Legendre quadrature nodes and weights*, SIAM J. Sci. Comput., 36 (2014), pp. A1008–A1026, <https://doi.org/10.1137/140954969>.
- [12] F. CHATELIN, *Spectral Approximation of Linear Operators*, SIAM, Philadelphia, 2011, <https://doi.org/10.1137/1.9781611970678>.
- [13] L. CHEN AND H.-P. MA, *Approximate solution of the Sturm–Liouville problems with Legendre–Galerkin–Chebyshev collocation method*, Appl. Math. Comput., 206 (2008), pp. 748–754.
- [14] M. J. COLBROOK AND A. C. HANSEN, *On the infinite-dimensional QR algorithm*, Numer. Math., 143 (2019), pp. 17–83.
- [15] B. ĆURGUS, A. FLEIGE, AND A. KOSTENKO, *The Riesz basis property of an indefinite Sturm–Liouville problem with non-separated boundary conditions*, Integral Equations Operator Theory, 77 (2013), pp. 533–557.
- [16] E. B. DAVIES, *Linear Operators and Their Spectra*, Cambridge Stud. Adv. Math. 106, Cambridge University Press, Cambridge, 2007.

- [17] J. J. DONGARRA, B. STRAUGHAN, AND D. W. WALKER, *Chebyshev tau-QZ algorithm methods for calculating spectra of hydrodynamic stability problems*, Appl. Numer. Math., 22 (1996), pp. 399–434.
- [18] T. A. DRISCOLL, N. HALE, AND L. N. TREFETHEN, *Chebfun Guide*, <https://www.chebfun.org/docs/guide/>, 2014.
- [19] N. DUNFORD, *A survey of the theory of spectral operators*, Bull. Amer. Math. Soc., 64 (1958), pp. 217–274.
- [20] S. FILIP, A. JAVEED, AND L. N. TREFETHEN, *Smooth random functions, random ODEs, and Gaussian processes*, SIAM Rev., 61 (2019), pp. 185–205, <https://doi.org/10.1137/17M1161853>.
- [21] J. GARY, *Computing eigenvalues of ordinary differential equations by finite differences*, Math. Comp., 19 (1965), pp. 365–379.
- [22] C.-I. GHEORGHIU, *Spectral Methods for Non-standard Eigenvalue Problems: Fluid and Structural Mechanics and Beyond*, Springer, Cham, 2014.
- [23] M. A. GILLES AND A. TOWNSEND, *Continuous analogues of Krylov subspace methods for differential operators*, SIAM J. Numer. Anal., 57 (2019), pp. 899–924, <https://doi.org/10.1137/18M1177810>.
- [24] G. H. GOLUB AND C. F. VAN LOAN, *Matrix Computations*, Johns Hopkins Ser. Math. Sci. 3, Johns Hopkins University Press, Baltimore, 2012.
- [25] T. R. GOODMAN, *The numerical solution of eigenvalue problems*, Math. Comp., 19 (1965), pp. 462–466.
- [26] J. GOPALAKRISHNAN, L. GRUBIŠIĆ, AND J. OVALL, *Filtered Subspace Iteration for Selfadjoint Operators*, preprint, <https://arxiv.org/abs/1709.06694v1>, 2017.
- [27] J. GOPALAKRISHNAN, L. GRUBIŠIĆ, AND J. OVALL, *Spectral discretization errors in filtered subspace iteration*, Math. Comp., 89 (2020), pp. 203–228.
- [28] N. HALE AND Y. NAKATSUKASA, *Rayleigh Quotient Iteration for an Operator*, <http://www.chebfun.org/examples/ode-eig/RayleighQuotient.html>, March 2017.
- [29] S. M. HAN, H. BENAROYA, AND T. WEI, *Dynamics of transversely vibrating beams using four engineering theories*, J. Sound Vibration, 225 (1999), pp. 935–988.
- [30] H. V. HENDERSON AND S. R. SEARLE, *On deriving the inverse of a sum of matrices*, SIAM Rev., 23 (1981), pp. 53–60, <https://doi.org/10.1137/1023004>.
- [31] T. KATO, *Perturbation Theory for Linear Operators*, Grundlehren der Mathematischen Wissenschaften, Band 132, Springer-Verlag, Berlin, New York, 1976.
- [32] J. KESTYN, E. POLIZZI, AND P. T. P. TANG, *FEAST eigensolver for non-Hermitian problems*, SIAM J. Sci. Comput., 38 (2016), pp. S772–S799, <https://doi.org/10.1137/15M1026572>.
- [33] R. S. LAUGESEN AND M. C. PUGH, *Linear stability of steady states for thin film and Cahn–Hilliard type equations*, Arch. Ration. Mech. Anal., 154 (2000), pp. 3–51.
- [34] R. S. LAUGESEN AND M. C. PUGH, *Properties of steady states for thin film equations*, European J. Appl. Math., 11 (2000), pp. 293–351.
- [35] E. D. NAPOLI, E. POLIZZI, AND Y. SAAD, *Efficient estimation of eigenvalue counts in an interval*, Numer. Linear Algebra Appl., 23 (2016), pp. 674–692.
- [36] S. OLVER AND A. TOWNSEND, *A fast and well-conditioned spectral method*, SIAM Rev., 55 (2013), pp. 462–489, <https://doi.org/10.1137/120865458>.
- [37] S. OLVER AND A. TOWNSEND, *A practical framework for infinite-dimensional linear algebra*, in Proceedings of the 1st Workshop for High Performance Technical Computing in Dynamic Languages, IEEE Press, 2014, pp. 57–62.
- [38] S. OLVER ET AL., *ApproxFun v0.10.8 Julia Package*, <https://github.com/JuliaApproximation/ApproxFun.jl>, 2018.
- [39] S. A. ORSZAG, *Accurate solution of the Orr–Sommerfeld stability equation*, J. Fluid Mech., 50 (1971), pp. 689–703, <https://doi.org/10.1017/S0022112071002842>.
- [40] A. M. OSTROWSKI, *On the convergence of the Rayleigh quotient iteration for the computation of the characteristic roots and vectors. I*, Arch. Rational Mech. Anal., 1 (1957), pp. 233–241.
- [41] R. D. PANTAZIS AND D. B. SZYLD, *Regions of convergence of the Rayleigh quotient iteration method*, Numer. Linear Algebra Appl., 2 (1995), pp. 251–269.
- [42] B. N. PARLETT, *The Rayleigh quotient iteration and some generalizations for nonnormal matrices*, Math. Comp., 28 (1974), pp. 679–693.
- [43] E. POLIZZI, *Density-matrix-based algorithm for solving eigenvalue problems*, Phys. Rev. B, 79 (2009), 115112, <https://doi.org/10.1103/PhysRevB.79.115112>.
- [44] M. REED AND B. SIMON, *Methods of Modern Mathematical Physics. I. Functional Analysis*, Academic Press, San Diego, 1980.
- [45] W. RUDIN, *Functional Analysis*, Internat. Ser. Pure Appl. Math., 2nd ed., McGraw-Hill, New York, 1991.

- [46] Y. SAAD, *Numerical Methods for Large Eigenvalue Problems*, Manchester University Press, Manchester, 1992.
- [47] Y. SAAD, *Analysis of subspace iteration for eigenvalue problems with evolving matrices*, SIAM J. Matrix Anal. Appl., 37 (2016), pp. 103–122, <https://doi.org/10.1137/141002037>.
- [48] N. SANFORD, K. KODAMA, J. D. CARTER, AND H. KALISCH, *Stability of traveling wave solutions to the Whitham equation*, Phys. Lett. A, 378 (2014), pp. 2100–2107.
- [49] J. SHEN AND L.-L. WANG, *Fourierization of the Legendre–Galerkin method and a new space–time spectral method*, Appl. Numer. Math., 57 (2007), pp. 710–720.
- [50] E. M. STEIN AND R. SHAKARCHI, *Complex Analysis*, Princeton Lect. Anal. II, Princeton University Press, Princeton, NJ, 2003.
- [51] G. W. STEWART, *Error bounds for approximate invariant subspaces of closed linear operators*, SIAM J. Numer. Anal., 8 (1971), pp. 796–808, <https://doi.org/10.1137/0708073>.
- [52] G. W. STEWART, *Simultaneous iteration for computing invariant subspaces of non-Hermitian matrices*, Numer. Math., 25 (1976), pp. 123–136.
- [53] P. T. P. TANG AND E. POLIZZI, *FEAST as a subspace iteration eigensolver accelerated by approximate spectral projection*, SIAM J. Matrix Anal. Appl., 35 (2014), pp. 354–390, <https://doi.org/10.1137/13090866X>.
- [54] F. TISSEUR, *Backward error and condition of polynomial eigenvalue problems*, Linear Algebra Appl., 309 (2000), pp. 339–361.
- [55] A. TOWNSEND AND S. OLVER, *The automatic solution of partial differential equations using a global spectral method*, J. Comput. Phys., 299 (2015), pp. 106–123.
- [56] A. TOWNSEND AND L. N. TREFETHEN, *Continuous analogues of matrix factorizations*, Proc. A, 471 (2015), 20140585.
- [57] L. N. TREFETHEN, *Householder triangularization of a quasimatrix*, IMA J. Numer. Anal., 30 (2009), pp. 887–897.
- [58] L. N. TREFETHEN AND D. BAU III, *Numerical Linear Algebra*, SIAM, Philadelphia, 1997.
- [59] L. N. TREFETHEN AND M. EMBREE, *Spectra and Pseudospectra: The Behavior of Nonnormal Matrices and Operators*, Princeton University Press, Princeton, NJ, 2005.
- [60] L. N. TREFETHEN AND M. R. TRUMMER, *An instability phenomenon in spectral methods*, SIAM J. Numer. Anal., 24 (1987), pp. 1008–1023, <https://doi.org/10.1137/0724066>.
- [61] W. WASOW, *Linear Turning Point Theory*, Appl. Math. Sci. 54, Springer-Verlag, New York, 1985.