Rectangular eigenvalue problems

Behnam Hashemi[1] · Yuji Nakatsukasa[2] · Lloyd N. Trefethen[2]

Abstract

Often the easiest way to discretize an ordinary or partial differential equation is by a *rectangular numerical method,* in which n basis functions are sampled at $m \gg n$ collocation points. We show how eigenvalue problems can be solved in this setting by QR reduction to square matrix generalized eigenvalue problems. The method applies equally in the limit "$m = \infty$" of eigenvalue problems for quasimatrices. Numerical examples are presented as well as pointers to related literature.

Keywords Eigenvalue problems · Quasimatrix · Spectral methods · Method of fundamental solutions · Lightning solver · Vandermonde with Arnoldi · Helmholtz equation · Fourier extension

Mathematics Subject Classification (2010) 47A75 · 65F15 · 65N35

1 Introduction

Problems involving ordinary and partial differential equations (ODEs and PDEs) are traditionally discretized by square matrices. Such methods are effective when a well-conditioned basis is available in which to expand the numerical solution and good

---

✉ Lloyd N. Trefethen
  trefethen@maths.ox.ac.uk

  Behnam Hashemi
  bhashemi@qu.edu.qa

  Yuji Nakatsukasa
  nakatsukasa@maths.ox.ac.uk

[1] Mathematics Program, Department of Mathematics, Statistics and Physics, College of Arts and Sciences, Qatar University, 2713, Doha, Qatar

[2] Mathematical Institute, University of Oxford, Woodstock Rd., Oxford, OX2 6GG, UK

Ⓐ Springer

quadrature or collocation points are known at which to enforce the equations. Sometimes, however, these conditions do not hold, and it becomes advantageous to sample the equation at more data points than there are basis functions and to solve the problem in a least-squares formulation. We call these *rectangular numerical methods.* The aim of this paper is to propose rectangular numerical methods for ODE and PDE eigenvalue problems.

Rectangular numerical methods have appeared in many areas, though they have rarely taken center stage. When Fourier, Chebyshev, or related expansions are involved, one can speak of rectangular spectral methods [11], though Boyd observes that such methods are "relatively uncommon" [8, Section 3.1]. In the finite elements literature there are Least-Squares Finite Element Methods [6, 20, 24]. With expansion functions that satisfy the differential equation but not the boundary conditions, one gets series methods [33] or the Method of Fundamental Solutions (MFS) [3, 13] or lightning or log-lightning methods for PDE problems with corner singularities [16, 26]. Related expansions that do not satisfy the differential equation and hence need fitting in the interior of a domain, not just on the boundary, lead to least-squares methods for radial basis functions (RBFs) or other kernels [10, 14, 21, 29]. RBF methods are an example of the broad category of *meshfree* methods.

Our plan is to set forth some of the simplest methods for solving rectangular eigenproblems and to illustrate them with a sequence of examples. The closest previous contributions we know of on this topic are by Manzhos and coauthors, who have developed what they call "rectangular collocation" methods for eigenvalue problems in quantum chemistry [21, 22], and by the first two authors [17]. The emphasis in [17] is on spectral methods for ODEs, and the linear algebra is carried out by the method of Ito and Murota [19], involving the singular value decomposition (SVD) of a matrix with twice as many columns as there are basis functions. (Important earlier related papers are [7] and [35].) Here we look at a wider range of problems and propose simpler methods of linear algebra based on the QR decomposition of a matrix without the doubled dimension.

We will mainly deal with fully discrete $m \times n$ rectangular matrices, always with $m > n$. As pointed out in [17], however, it makes good sense conceptually to consider the limit in which the columns are functions of one or more continuous variables, so that instead of matrices, we have quasimatrices; see [4, 32] and [12, chap. 6]. Nothing essential changes here, and we shall include quasimatrices in the discussion as the case "$m = \infty$". For spectral ODE problems, the quasimatrices can be realized numerically in Chebfun [12], and the first two of the examples of Section 3 follow this path. After that, our computed examples are fully discrete, though the mathematical derivations apply equally to $m < \infty$ or $m = \infty$.

Rectangular numerical methods for eigenvalue problems are related to ideas going back a century, first associated with Rayleigh, Ritz, and Galerkin, in which square matrix approximations are obtained by quadrature and projection. See [30, Section 4.3], [31, Section 6.4], and the fascinating historical discussion in [15]. In the finite elements literature, Galerkin and Petrov-Galerkin methods can often be interpreted this way. Arnoldi and Jacobi-Davidson iterative methods for computing eigenvalues of large matrices are also of this nature. What is different in the present paper is that no explicit quadrature or projection ideas are employed, just numerical

algorithms applied to eigenvalue problems configured rectangularly. This diminishes the need for case-by-case analysis and permits great flexibility in the choice of basis functions and sample points.

Our eight numerical examples are intended to illustrate the simplicity and potential of rectangular eigenvalue discretizations, not as optimal or guaranteed numerical methods for any particular problem. Their success depends on reasonable choices of parameters such as numbers of expansion terms and numbers of boundary and interior sample points, and they can be made to fail in certain other parameter ranges. Apart from a few comments, our discussion does not explore these parameter dependences systematically, as that would lengthen the presentation greatly.

2 The numerical method, three variants

Let L be a linear operator acting on functions in a univariate or multivariate domain Ω, and suppose we seek eigenvalues λ and nonzero eigenfunctions u such that

$$Lu = \lambda u. \tag{1}$$

We shall consider three variants of this problem, in which (1) is coupled with no boundary conditions, a finite number of boundary conditions, or boundary conditions applied on a continuum. In all three cases we suppose that for some $n \geq 1$, we have a set of functions $g_1, \ldots, g_n$ defined in Ω whose span contains good approximations to the eigenfunctions of interest, and we let G be the "$\infty \times n$ matrix" (also known as a quasimatrix) whose columns are these functions. Setting

$$u = Gx, \tag{2}$$

we seek a coefficient vector $x \in \mathbb{C}^n$ such that

$$LGx = \lambda Gx. \tag{3}$$

This is an $\infty \times n$ generalized eigenvalue problem, which can also be described as the eigenvalue problem for the $\infty \times n$ quasimatrix pencil $LG - \lambda G$. Like most rectangular eigenproblems, it will not have exact solutions in general [7, 19, 35], but if G is well chosen, we expect it to have approximate solutions accurate to many digits of accuracy as measured, for example, in the sense of small residuals. In such contexts G will often be highly ill-conditioned.

Usually we will discretize the problem by sampling in a subset $\Omega_m \subseteq \Omega$ of m points of Ω, so that the quasimatrices become $m \times n$ matrices. Let S, of dimensions "$m \times \infty$," denote the sampling operator. Then SG is the $m \times n$ matrix whose (j, k) entry is g_k sampled in Ω_m, and SLG is the $m \times n$ matrix whose (j, k) entry is Lg_k sampled in Ω_m. We will now attempt to solve the discretized variant of (3),

$$SLGx = \lambda SGx. \tag{4}$$

In cases where we wish to solve a problem in the original quasimatrix form (3), without discretization, we can speak of "$m = \infty$," and (4) and the following formulas remain valid with S equal to the identity operator.

Variant 1 No explicit boundary conditions. Suppose first that (1) is the whole problem: as in [7] and [19], there are no explicit boundary conditions. An example would be the harmonic oscillator $-u'' + x^2 u = \lambda u$ defined on the real axis, with eigenvalues $1, 3, 5, \ldots$. In this case our proposed numerical method begins by computing the QR factorization of SG,

$$SG = QR, \tag{5}$$

where Q is $m \times n$ and R is $n \times n$ and upper-triangular. (For details of QR factorization in the quasimatrix case $m = \infty$, see [32].) Premultiplying (4) by Q^* then gives

$$Q^*(SLG)x = \lambda Q^* SGx = \lambda Rx. \tag{6}$$

This equation enforces the condition that the residual $(SLG)x - \lambda SGx$ is orthogonal to the range of SG. (Alternatively one could require

$$(SG)^*(SLG)x = \lambda(SG)^*(SG)x, \tag{7}$$

though with less numerical stability, in analogy to the normal equations for least-squares fitting problems; compare [21, eq. (6)] and the earlier [25, eq. (5)].) Equation (6) is a square matrix generalized eigenvalue problem with dimensions $n \times n$, which we solve by the standard QZ algorithm. Note that (4) implies (6). Conversely, (6) implies (4) if the columns of SLG lie in the column space of SG. This may or may not hold exactly, but in many applications it will hold to high accuracy, making (4) and (6) effectively equivalent.

In the computation above, as in Variants 2 and 3 below, an alternative (mathematically equivalent) possibility is to use the SVD instead of the QR decomposition to construct an orthonormal basis of the columns of SG. In our experience this may improve the accuracy slightly, typically by less than one digit, at the cost of a slight increase in computing time. We have not investigated the matter carefully.

In the continuous case where S is the identity operator, the formulation (6) is particularly close to the Rayleigh-Ritz method, as mentioned in the introduction. Here one has $G = QR$, hence $Gx = QRx$, so (6) reduces to the Rayleigh-Ritz form

$$Q^* LQv = \lambda v, \quad v = Rx. \tag{8}$$

Before finishing with Variant 1 we should comment on what it means for a differential equations problem to have "no boundary conditions." On a domain with no boundary, such as a periodic interval or a torus, this may be strictly accurate. On the real axis as in the example of our harmonic oscillator, however, it would be more correct to say that there are implicit boundary conditions involving decay as $|x| \to \infty$, which practitioners may neglect to specify explicitly and which numerical methods may enforce implicitly. In Example 1 of the next section, the implicit enforcement results from L^2-boundedness of computed eigenfunctions.

Variant 2 Finite set of boundary conditions. Suppose next that (1) is coupled with a finite set of $\mu > 0$ homogeneous linear boundary conditions, as is considered (along with other possibilities) in [17]. An example would be $-(4/\pi^2)u'' = \lambda u$ on $[-1, 1]$ with boundary conditions $u(\pm 1) = 0$, with eigenvalues $1, 4, 9, \ldots$. We can write the boundary conditions in the form

$$Bu = \mathbf{0}, \tag{9}$$

where B ("boundary") is a $\mu \times \infty$ quasimatrix and $\mathbf{0}$ is the $\mu \times 1$ zero vector. Each row of B is a linear functional, which might, for example, evaluate u or u' at a boundary point. Applying (2), this becomes the $\mu \times n$ set of equations

$$(BG)x = \mathbf{0}. \tag{10}$$

We can now combine (10) with (4) to get the $(m + \mu) \times n$ rectangular generalized eigenvalue problem

$$\begin{bmatrix} SLG \\ \hline BG \end{bmatrix} x = \lambda \begin{bmatrix} SG \\ \hline 0 \end{bmatrix} x, \tag{11}$$

where 0 denotes the zero matrix of dimensions $\mu \times n$. Various methods can be employed to make this equation square, as discussed in [17] in the context of the Ito-Murota formulation. The simplest, analogous to what is called the "tau method" of imposing boundary conditions in spectral methods [8], is to let Q_- denote the $m \times (n - \mu)$ matrix or quasimatrix consisting of Q with its final μ columns removed and then consider

$$\begin{bmatrix} Q_-^* SLG \\ \hline BG \end{bmatrix} x = \lambda \begin{bmatrix} Q_-^* SG \\ \hline 0 \end{bmatrix} x. \tag{12}$$

This equation enforces the boundary conditions exactly while requiring the residual $(SLG)x - \lambda SGx$ to be orthogonal to the range of the first $n - \mu$ columns of SG. It is a square matrix generalized eigenvalue problem of dimensions $n \times n$, which again we solve by standard methods. For essentially the same structure but not based on a QR factorization, see [11, Section 5] and [1, Section 5].

Variant 3 Continuum of boundary conditions. Finally, suppose (1) is coupled with a continuum of homogeneous linear boundary conditions. Specifically, suppose we have a PDE in a domain Ω of dimension $d \geq 2$ and a boundary condition applied on the boundary $\partial\Omega$ of dimension $d - 1$. An example would be $-\Delta u = \lambda u$ on the unit disk with boundary condition $u = 0$ on the unit circle, whose first eigenvalue is $5.7831859629\ldots$, the square of the smallest root of the Bessel function $J_0(x)$.

In this continuous case equations (9)–(11) continue to apply, but now the lower part of (11) needs discretization too. Let S_b be the $\mu \times \infty$ operator that samples Ω in a subset of μ points on the boundary. (As before we could take "$\mu = \infty$" for a quasimatrix formulation.) The discretized variant of (11) becomes

$$\begin{bmatrix} SLG \\ \hline S_b BG \end{bmatrix} x = \lambda \begin{bmatrix} SG \\ \hline 0 \end{bmatrix} x. \tag{13}$$

Now that μ is large or infinite, it is no longer normally appropriate to attempt to enforce the boundary conditions exactly. Instead, the natural thing to do is to treat

all of (13), both the upper and lower parts, in a least-squares fashion. We do this as follows. Let **G** denote the $(m + \mu) \times n$ structure

$$\mathbf{G} = \left[\begin{array}{c} SG \\ \hline S_b G \end{array} \right]. \tag{14}$$

Thus each column of **G** is an object whose upper part is the m-point discretization of a function of d dimensions and whose lower part is the μ-point discretization of the same function along the boundary $\partial \Omega$. We now compute a QR factorization of **G**,

$$\mathbf{G} = \mathbf{Q}R, \tag{15}$$

in which **Q** has the same structure as **G**,

$$\mathbf{Q} = \left[\begin{array}{c} Q \\ \hline \partial Q \end{array} \right], \tag{16}$$

and R is an $n \times n$ upper-triangular matrix. The columns of **Q** are orthonormal with respect to an inner product that combines sums (or integrals, when m or μ is infinite) associated with both Ω and $\partial \Omega$. For discussions of QR factorization and singular value decomposition of such mixed objects, see [17]. One could analyze what relative weighting is most appropriate in balancing the two halves, but the expectation is that in applications it will not make much difference since we are aiming for residuals close to zero. In our fully discrete computed examples we give equal weights to all sample points, both the m points in the interior and the μ points on the boundary. At another extreme, if one gave infinitely more weight to boundary points than interior ones, Variant 3 would reduce to Variant 2, except with the "finite set of boundary conditions" corresponding to the discretization rather than to the underlying continuous mathematical problem.

To square up the eigenvalue problem, we left-multiply (13) by $\mathbf{Q}^*$ to obtain

$$\left[Q^*(SLG) + (\partial Q)^*(S_b BG) \right] x = \lambda Q^* SGx. \tag{17}$$

This equation enforces the condition that a combined boundary-interior residual is orthogonal in the mixed inner product to the basis vectors (columns of SG) and their boundary traces (columns of $S_b G$). Like (6) and (12), (17) is an $n \times n$ generalized eigenvalue problem, and again we solve it by standard numerical methods.

We now turn to computed examples. The ODE problems of Section 3 illustrate variants 1 and 2, and the PDE problems of Section 4 illustrate variant 3.

3 One-dimensional examples (ODEs)

Example 1 Harmonic oscillator with no boundary conditions. We begin with the harmonic oscillator mentioned on p. 4, $-u'' + x^2 u = \lambda u$ on the real axis. Using

Chebfun for the quasimatrices, and approximating the real axis by $[-8, 8]$, we can compute eigenvalues based on an $\infty \times 40$ rectangular Chebyshev spectral discretization with the code below, just six lines long. The first three eigenvalues come out as 1.0000000008, 3.0000000113, and 5.0000005634, and this accuracy can be improved by increasing n up to about $n = 50$ in standard floating-point arithmetic. (After about $n = 75$, spurious eigenvalues appear for reasons related to the Chebfun implementation of quasimatrices.)

```
n = 40;
L = chebop(@(x,u) -diff(u,2) + x^2*u,[-8,8]);
G = chebpoly(0:n-1,[-8,8]);
[Q,R] = qr(G);
A = Q'*(L*G); C = R;
lam = sort(eig(A,C))
```

By adjusting a few of the commands we get a code for the corresponding fully discrete computation with 100×40 matrices, 100 being the default number of points in the `linspace` command. Chebfun is still used in this code segment, but only because it offers a convenient way to construct a matrix of sampled Chebyshev polynomials scaled to $[-8, 8]$ and their second derivatives. The first three eigenvalues come out with approximately the same accuracy as before as 1.0000000004, 3.0000000050, and 5.0000002819. (This time there is no problem with much larger values of n, so long as $m \gg n$ points are used in the discretization.)

```
n = 40;
L = chebop(@(x,u) -diff(u,2) + x^2*u,[-8,8]);
G = chebpoly(0:n-1,[-8,8]); LG = L*G;
X = linspace(-8,8)'; SG = G(X); SLG = LG(X);
[Q,R] = qr(SG,0);
A = Q'*SLG; C = R;
lam = sort(eig(A,C))
```

For both of the computations just presented, the accuracy of the computed eigenvalues is undiminished if the formulation (7) without the QR factorization is used instead of (6). This makes sense since G is a matrix of Chebyshev polynomials on $[-1, 1]$, hence well-conditioned.

Example 2 Wave oscillator with two boundary conditions. Our second example, mentioned on p. 5, is $-(4/\pi^2)u'' = \lambda u$ on $[-1, 1]$ with $u(\pm 1) = 0$. The following code implements an $\infty \times 30$ quasimatrix discretization, computing the first ten eigenvalues $1, 4, 9, \ldots, 100$ to 11–14 digits of relative accuracy.

```
n = 30;
L = chebop(@(x,u) -(4/pi^2)*diff(u,2));
G = chebpoly(0:n-1);
[Q,R] = qr(G);
A = [Q(:,1:n-2)'*(L*G); G(-1); G(1)];
```

```
C = [R(1:n-2,:); zeros(2,n)];
lam = sort(eig(A,C))
```

Here is the adjustment needed for a fully discrete 200×30 discretization using equi-spaced points in $[-1, 1]$. (It makes little difference if Chebyshev points are used instead, since we are in the regime $m \gg n$ with plenty of sample points. With $m \approx n$, it would be important to be careful about the distribution of sample points, but rectangular numerical methods make it unnecessary for m to be small.) The relative accuracy of the first ten eigenvalues is now 9–14 digits, which returns to 11–14 digits if n is increased to 34.

```
n = 30;
L = chebop(@(x,u) -(4/pi^2)*diff(u,2));
G = chebpoly(0:n-1); LG = L*G;
X = linspace(-1,1,200)'; SG = G(X); SLG = LG(X);
[Q,R] = qr(SG,0);
A = [Q(:,1:n-2)'*SLG; SG(1,:); SG(end,:)];
C = [R(1:n-2,:); zeros(2,n)];
lam = sort(eig(A,C))
```

As with the last pair of computations, there is again little difference in accuracy here if one bypasses the QR factorization and uses (7) instead of (6).

These examples are of a Chebyshev spectral flavor, and in such cases, at least in simple domains, square discretizations are often readily available. Now we turn to problems related to the Method of Fundamental Solutions or RBF or other meshfree discretizations, where the need for rectangular formulations is more pressing. The reason is that the representation of the solution involves n points that do not lie in the domain, hence have no naturally associated grid for interpolation or quadrature.

Example 3 Wave oscillator, method of fundamental solutions. Our third example is the problem $-(4/\pi^2)u'' = \lambda u$ on $[-1, 1]$ with $u(\pm 1) = 0$ again, but now solved by a kind of method of fundamental solutions, with the solution represented as a linear combination of point charges. The following code implements a 150×35 matrix discretization involving a constant term plus 34 point charge potentials $\log |x - p_j|$ with p_j equally spaced from $-1.5 + 0.5i$ to $1.5 + 0.5i$. The first ten eigenvalues $1, 4, 9, \ldots, 100$ are computed to 12–15 digits of relative accuracy. A similar 150×35 discretization based on 17 dipoles equally spaced from $-1.5 + 0.5i$ to $1.5 + 0.5i$, that is, real and imaginary parts of complex poles $1/(x - p_j)$ (not shown), gives 10–11 digits.

```
n = 35;
pts = linspace(-1.5+.5i,1.5+.5i,n-1);
X = linspace(-1,1,150)';
SG = [X.^0 log(abs(X-pts))];
SGpp = [0*X -real(1./(X-pts).^2)];
SLG = -(4/pi^2)*SGpp;
[Q,R] = qr(SG,0); A = [Q(:,1:n-2)'*SLG;
```

```
SG(1,:); SG(end,:)];
C = [R(1:n-2,:); zeros(2,n)];
lam = sort(eig(A,C));
```

For this problem, the QR factorization makes a big difference. If we use (7) instead of (6), some spurious eigenvalues appear and the first ten nonspurious computed eigenvalues fall to 2–10 digits of accuracy.

Example 4 Quantum oscillator with singularity: lightning discretization. We now look at a problem with a singularity,

$$-0.01u'' + |x|^{1/2}u = \lambda u, \quad u(\pm 1) = 0, \tag{18}$$

posed on the interval $[-1, 1]$. This is a Schrödinger equation with the singular potential $V(x) = |x|^{1/2}$. Smooth discretizations will have difficulty achieving more than 2 or 3 digits of accuracy, but we can do much better with a "lightning discretization" involving poles on the imaginary x-axis exponentially clustered near the singular point $x = 0$. Specifically, following eq. (3.2) of [16], a formula that is justified in [34], we fix a number $n_{\text{poles}} \geq 0$ and define

$$d_j = \exp(4(\sqrt{j} - \sqrt{n_{\text{poles}}})), \quad 1 \leq j \leq n_{\text{poles}}. \tag{19}$$

The columns of G will include both the real and imaginary parts of the simple pole functions $d_j^3/(x - id_j)$, making $2n_{\text{poles}}$ columns all together. (The constant d_j^3 is included for scaling, since the second derivative of this function is $2d_j^3/(x - id_j)^3$.) In addition we fix a number $n_{\text{poly}} \geq 0$ and include the Chebyshev polynomials $T_k(x)$ with $0 \leq k \leq n_{\text{poly}}$ as further columns of the matrix.

Figure 1 shows results for this scheme for $0 \leq n_{\text{poly}} \leq 20$ and $0 \leq n_{\text{poles}} \leq 25$. The interval $[-1, 1]$ is discretized by 1000 points exponentially spaced from 10^{-10} to 1 and their negatives, so the matrices have 2000 rows and between 1 and 71 columns. (This space discretization could undoubtedly be improved.) Neither poles nor smooth polynomials alone give good accuracy, but in combination they achieve up to 12 digits on this difficult problem.

If Fig. 1 is recomputed based on the formulation (7) without QR factorization, the accuracy falls to 2 digits and spurious eigenvalues appear. For our further examples we will report results only from the stable formulation (6).

4 Two-dimensional examples (PDEs)

Now we move to two-dimensional (2D) domains and PDE eigenvalue problems. Though problems without boundaries can certainly be considered (such as the 2D harmonic oscillator in the x-y plane), we shall look at examples where Ω has a boundary $\partial\Omega$ with explicit boundary conditions, leading to rectangular discretizations in the Variant 3 form (17).

Example 5 Circular drum, RBF discretization. Consider the planar Laplace problem mentioned on p. 6,

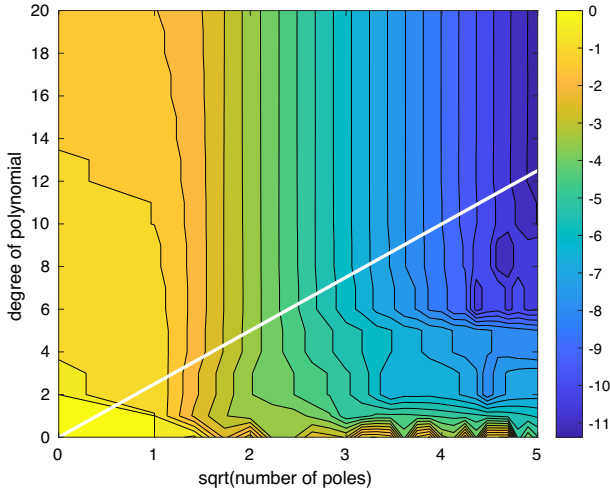$$-\Delta u = \lambda u, \quad |z| < 1, \tag{20}$$

Fig. 1 Error (log base 10) in the first computed eigenvalue $\lambda_1 \approx 0.42185900459401$ of the singular Schrödinger problem (18) as a function of the polynomial degree n_{poly} and the number of exponentially clustered "lightning" poles n_{poles}. The rectangular discretization combining the polynomial and the poles converges rapidly to 12 digits of accuracy with a basis matrix G with 2000 rows and about 50 columns. Note that the horizontal axis is $\sqrt{n_{\text{poles}}}$, to highlight the root-exponential convergence investigated in [34]; the white line marks $n_{\text{poly}} = 2.5\sqrt{n_{\text{poles}}}$. Examination of the data along the two axes shows that neither the polynomial nor the clustered poles alone get better than 2 digits

with $u(z) = 0$ for $|z| = 1$. The eigenvalues are the squares of the zeros of the Bessel functions $J_k(r)$, $k \geq 0$. For $k = 0$, the eigenfunctions are axisymmetric and the eigenvalues are simple, whereas for $k \geq 1$, the eigenfunctions are not axisymmetric and each eigenvalue is of multiplicity 2.

Our first rectangular discretization will be based on RBFs. We follow [28] and take as a radial basis function the multiquadric

$$\phi(r) = \sqrt{c^2 + r^2} \tag{21}$$

for a fixed parameter c, so that each eigenfunction is approximated by a sum

$$u(z) = a_0 + \sum_{k=1}^{n-1} a_k \phi(|z - \zeta_k|), \tag{22}$$

where $\{\zeta_k\}$ is a set of $n - 1$ centers. In [28], as in most RBF literature, the emphasis is on obtaining square discretizations based on interpolation. This requires care in selecting the centers, which must be clustered near the boundary to avoid a Runge phenomenon [27]. In rectangular mode, however, with $m \gg n$ sample points, one can be more relaxed. To illustrate the method, Fig. 2 shows a square grid of sample points in the unit disk $|z| \leq 1$ with spacing 0.04 as well as a sparser square grid of RBF centers in the disk $|\zeta| \leq 1.25$ with spacing 0.08. We take $c = 0.4$ for the constant of (21). Figure 3 shows that the resulting 2241×770 rectangular eigenvalue problem gives 5–8 digit accuracy in the first eight eigenvalues.

Fig. 2 RBF discretization of Example 5. The dots show 1941 interior sample points in the unit disk, and the circles show 769 RBF centers in the disk $|z| \leq 1.25$. The unit circle boundary is discretized by 400 equispaced points. The resulting rectangular eigenproblem is of dimensions 2241×770
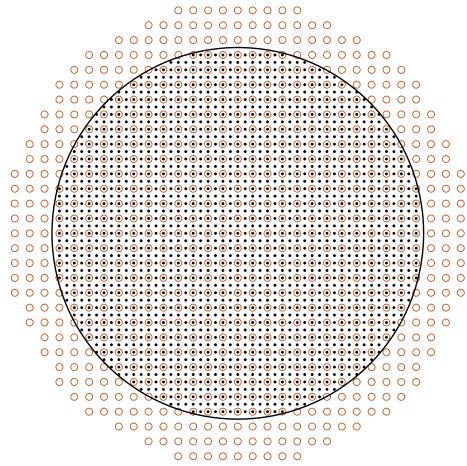
Example 6 Circular drum, Fourier extension discretization. Consider (20) for a circular drum again, but now discretized by a Fourier extension method. For a rectangular discretization of (20), we start from a 41×41 square grid in $[-1, 1]^2$ (i.e., grid spacing 0.05) and discard the points outside the unit disk; the 1245 points that remain are our interior sample points. On the boundary we take 300 equispaced sample points. The basis functions are the products $\cos(kx)\cos(my)$, $\cos(kx)\sin(my)$, $\sin(kx)\cos(my)$, and $\sin(kx)\sin(my)$ with $0 \leq k + m \leq 10$ (discarding those that are exactly zero). This leads to an eigenvalue problem of dimensions 1545×221,
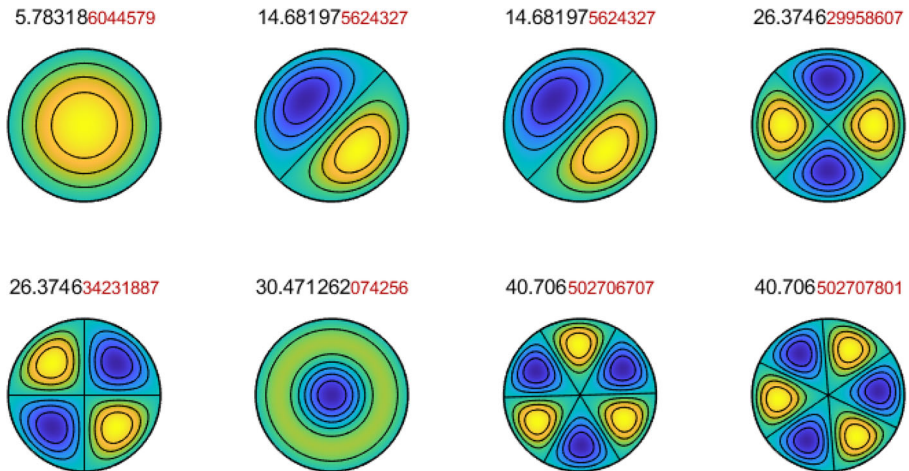
Fig. 3 First 8 computed eigenvalues and eigenfunctions of a disk, based on an RBF discretization defined by (21)–(22) with the RBF and sampling grids of Fig. 2. The rectangular eigenvalue problem is of dimensions 2241×770. Correct digits are printed in black and incorrect ones in a smaller font in red. The eigenvalues of multiplicity 2 are identified correctly, but the associated eigenfunction pairs do not come out orthogonal, reflecting the fact that the numerical method is not self-adjoint

and Fig. 4 shows the first eight computed eigenfunctions, with eigenvalues accurate to 9–11 digits.

Example 7 Elliptical drum, Fourier extension discretization. Figure 5 shows results for another 2D Fourier extension computation, this time involving an elliptical drum of length 1 and width 1/2. Although it doesn't make much difference for this problem, we have switched here to a representation in which the basis functions are orthogonalized by a Vandermonde with Arnoldi process [9, Example 3]. Two-dimensional Vandermonde with Arnoldi has been utilized previously for bivariate and trivariate polynomials in [2, 18], and [36], and here we do it for Fourier extension. (In separate experiments not reported here, we have successfully computed eigenmodes of the ellipse in this manner by bivariate polynomials.) Setting $X = e^{ix}$ and $Y = e^{iy}$, we note that the real part of $X^k Y^{\pm m}$ is $\cos(kx)\cos(my) \mp \sin(kx)\sin(my)$ and the imaginary part is $\pm\cos(kx)\sin(my) + \sin(kx)\cos(my)$, so these real and imaginary parts span the necessary space of bivariate trigonometric polynomials. To be precise, we fix $K \geq 1$ and work with integers k and m with $0 \leq k \leq K$ and $0 \leq m \leq K$ for $k = 0$, $k - K \leq m \leq K - k$ for $k \geq 1$. Arnoldi orthogonalization is carried out in the order $1, Y, X, Y^2, XY, X^2, Y^3, XY^2, X^2Y, X^3, \ldots$ (compare the paragraph after eq. (7) of [2]). The rectangular matrix whose eigenfunctions are shown in the figure is of dimensions 1399×313.

Fourier approximations of analytic functions on analytic domains should converge exponentially, and for Examples 6 and 7, the data confirm this nicely, as shown in Fig. 6. (Here and in the next example, the correct eigenvalues are not known analytically but are determined numerically by higher-resolution calculations.) For more complicated domains, however, especially if they are nonconvex, the exponential rate becomes very slow because the solution can only be analytically continued a short distance outside the boundary. Such effects have been studied by Barnett and Betcke [3], and for more on the theory of analytic continuation of Helmholtz fields,
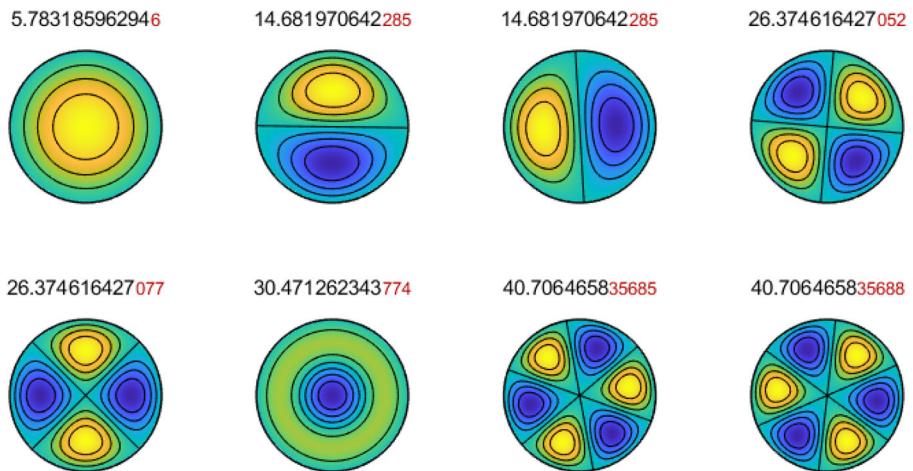

5.783185962946   14.681970642285   14.681970642285   26.374616427052

26.374616427077   30.471262343774   40.706465835685   40.706465835688

Fig. 4 Like Fig. 3 but for a Fourier extension discretization. The rectangular eigenvalue problem is of dimensions 1545×221. The degenerate pairs 2–3, 4–5, and 7–8 again show orientations at arbitrary angles

14.2669064117   25.1017224807   40.1136064818   46.9466617364

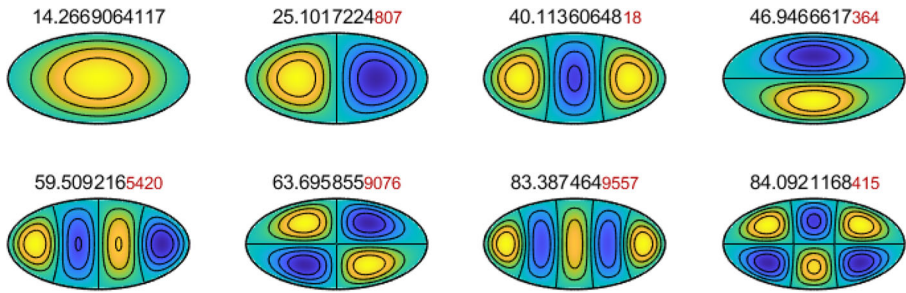59.5092165420   63.6958559076   83.3874649557   84.0921168415

Fig. 5 Like Fig. 4 but for an elliptical drum of axis lengths 1 and $\frac{1}{2}$

see [23]. This difficulty pertains to the choice of expansion functions, not to the method of dealing with them by rectangular eigenvalue problems.

Example 8 L-shaped region with singular terms. Solutions to PDEs in regions with corners usually have corner singularities, which make it challenging to get high accuracy. In the context of the global representations explored in this paper, a natural idea for such problems would be to combine a general purpose set of basis functions to capture the "smooth part" of the solution with additional singular terms near the corners. For Laplace Dirichlet or Neumann problems, representations of this kind led to the lightning and log-lightning solvers introduced in [16] and [26]. Here we illustrate that such an approach may be effective for eigenvalue problems too. This is a PDE analogue of Example 4 for ODEs.

Our example, shown in Fig. 7, is the planar drum (20) in the form of the L-shaped region well known from the MATLAB logo, the square $[-1, 1]^2$ with one corner removed. (For numerical eigenvalues of this and other drums calculated by a more specialized method to an accuracy of 8 digits, see [5].) The boundary has been discretized by 420 points exponentially clustered near the reentrant corner, and the interior by a square grid of spacing 1/20. This gives a 1617 × 313 matrix,
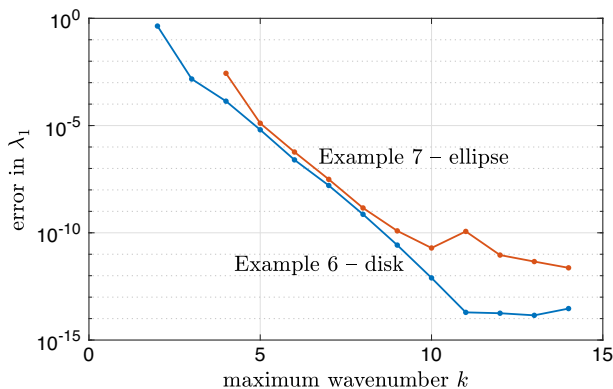

Fig. 6 Convergence curves for the two examples involving Fourier extension discretizations. Exponential convergence is observed to around 14 digits for the disk and 11 digits for the ellipse
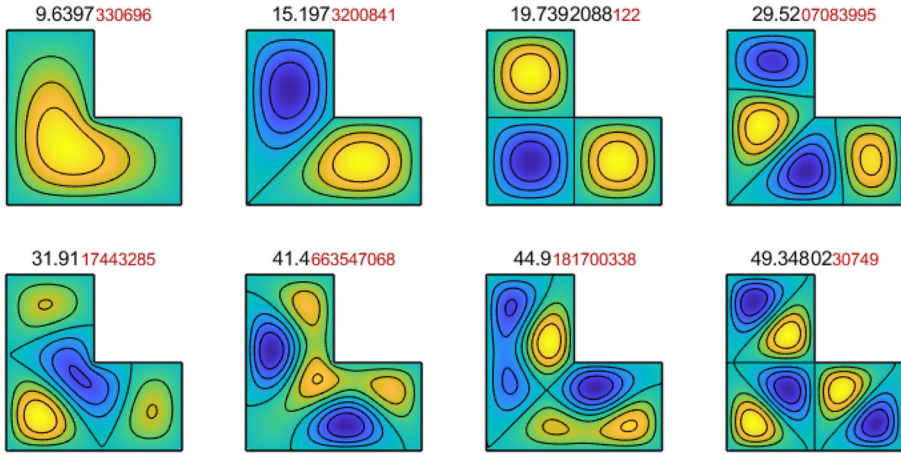
Fig. 7 Eigenfunctions of an L-shaped region computed by a 1617×343 eigenvalue problem, with 30 of the matrix columns devoted to resolving the singularity at the reentrant corner

to which a further 30 columns are added corresponding to the corner singularity. These 30 terms are chosen to capture the dominant behavior of the functions $J_{(2/3)j}(\sqrt{\lambda} r) \sin((2/3)ij\theta)$ that arise in series expansions of eigenfunctions near reentrant right-angle corners, where r is the distance from the corner and θ is the angle measured from one of the adjacent sides. Specifically, we used discretizations of the 30 functions $r^a \sin(b\theta)$ with a and b given by

$$a = 2/3, 4/3, \ \ 8/3, 10/3, \ \ 14/3, 16/3, \ \ 20/3, 22/3, \ \ 26/3, 28/3$$

and

$$b = a, \, a - 2, \, a - 4, \cdots > 0.$$

The figure shows that this rectangular discretization computes the first 8 eigenvalues to about 4 digits of accuracy. So far as we know, discretizations of this kind have not been considered before for eigenvalue problems, and we hope to present them more fully, and improve them, in a future publication.

5 Discussion

The most robust discretizations of differential equation eigenvalue problems, and the ones with the strongest theoretical support, often involve square matrices, especially in the self-adjoint case. The theory of finite element methods has brought such discretizations to an advanced state.

For some problems, however, whether because of irregular geometry, nonself-adjointness, or the presence of singularities, a good square matrix discretization may not be readily available. The aim of this paper has been to show that in such cases rectangular matrices may offer an eminently practical alternative, often making possible high accuracy solutions with a global representation (hence perfectly smooth in the interior, and very fast to evaluate). We make no claim of guaranteed success, and

indeed, in most of our experiments, which are based on new kinds of discretizations with little previous literature, it has been necessary to try several parameter choices to get good accuracy and avoid spurious modes. With further work, more may be learned about these matters and rectangular eigenvalue methods may be developed with guarantees of robustness and accuracy. These methods are easy and flexible and deserve ongoing attention.

Declarations

Conflict of interest The authors declare no competing interests.

References

1. Aurentz, J.L., Trefethen, L.N.: Block operators and spectral discretizations. SIAM Rev. **59**, 423–446 (2017)
2. Austin, A.P., Krishnamoorthy, M., Leyffer, S., Mrenna, S., Müller, J., Schulz, H.: Practical algorithms for multivariate rational approximation. Comput. Phys. Commun. **261**, 107663 (2021)
3. Barnett, A.H., Betcke, T.: Stability and convergence of the method of fundamental solutions for Helmholtz problems on analytic domains. J. Comp. Phys. **14**, 7003–7026 (2008)
4. Battles, Z., Trefethen, L.N.: An extension of MATLAB to continuous functions and operators. SIAM J. Sci. Comput. **25**, 1743–1770 (2004)
5. Betcke, T., Trefethen, L.N.: Computed eigenmodes of planar regions. Contemp. Math. **412**, 297–314 (2006)
6. Bochev, P.B., Gunzburger, M.D.: Least-Squares Finite Element Methods. Springer (2009)
7. Boutry, G., Elad, M., Golub, G.H., Milanfar, P.: The generalized eigenvalue problem for nonsquare pencils using a minimal perturbation approach. SIAM J. Matrix Anal. Appl. **27**, 582–601 (2005)
8. Boyd, J.P.: Chebyshev and Fourier Spectral Methods, 2nd edn. Dover (2001)
9. Brubeck, P.D., Nakatsukasa, Y., Trefethen, L.N.: Vandermonde with Arnoldi. SIAM Rev. **63**, 405–415 (2021)
10. Buhmann, M.D.: Radial Basis Functions: Theory and Implementations. Cambridge U Press (2003)
11. Driscoll, T.A., Hale, N.: Rectangular spectral collocation. IMA J. Numer. Anal. **36**, 108–132 (2016)
12. Driscoll, T.A., Hale, N., Trefethen, L.N.: Chebfun Guide. Pafnuty Press, Oxford (2014). see also www.chebfun.org
13. Fairweather, G., Karageorghis, A.: The method of fundamental solutions for elliptic boundary value problems. Adv. Comp. Math. **9**, 69–95 (1998)
14. Fasshauer, G.: Meshfree Approximation Methods with MATLAB. World Scientific (2007)
15. Gander, M.J., Wanner, G.: From Euler, Ritz, and Galerkin to modern computing. SIAM Rev. **54**, 627–666 (2012)

16. Gopal, A., Trefethen, L.N.: Solving Laplace problems with corner singularities via rational functions. SIAM J. Numer. Anal. **57**, 2074–2094 (2019)
17. Hashemi, B., Nakatsukasa, Y.: Least-squares spectral methods for ODE eigenvalue problems. SIAM J. Sci. Comput. **44**, A3244–A3264 (2022)
18. Hokanson, J.M.: Multivariate rational approxmation using a stabilized Sanathanan–Koerner iteration arXiv:2009.10803v1 (2020)
19. Ito, S., Murota, K.: An algorithm for the generalized eigenvalue problem for nonsquare matrix pencils by minimal perturbation approach. SIAM J. Matrix Anal. Appl. **37**, 409–419 (2016)
20. Jiang, B.: The Least-Squares Finite Element Method: Theory and Applications in Computational Fluid Dynamics and Electromagnetism. Springer (1998)
21. Kamath, A., Manzhos, S.: Inverse multiquadratic functions as the basis for the rectangular collocation method to solve the vibrational Schrödinger equation. Mathematics **6**, 253 (2018)
22. Manzhos, S., Yamashita, K., Carrington Jr., T.: On the advantages of a rectangular matrix collocation equation for computing vibrational spectra from small basis sets. Chem. Phys. Lett. **511**, 434–439 (2011)
23. Millar, R.F.: Singularities and the Rayleigh hypothesis for soutions to the Helmholtz equation. IMA J. Appl. Math. **37**, 155–171 (1986)
24. Monk, P., Wang, D.-Q.: A least-squares method for the Helmholtz equation. Comput. Methods Appl. Mech. Engrg. **175**, 121–136 (1999)
25. Nakatsuji, H., Nakashima, H., Kurokawa, Y., Ishikawa, A.: Solving the Schrödinger equation of atoms and molecules without analytical integration based on the free iterative-complement-interaction wave function. Phys. Rev. Lett. **99**, 240402 (2007)
26. Nakatsukasa, Y., Trefethen, L.N.: Reciprocal-log approximation and planar PDE solvers. SIAM J. Numer. Anal. **59**, 2801–2822 (2021)
27. Platte, R.B.: How fast do radial basis function interpolants of analytic functions converge IMA J. Numer. Anal. **31**, 1578–1597 (2011)
28. Platte, R.B., Driscoll, T.A.: Computing eigenmodes of elliptic operators using radial basis functions. Computers and Math. with Applics. **48**, 561–576 (2004)
29. Platte, R.B., Driscoll, T.A.: Eigenvalue stability of radial basis function discretizations for time-dependent problems. Comput. Math. Applics. **51**, 1251–1268 (2006)
30. Saad, Y.: Numerical Methods for Large Eigenvalue Problems, revised edn. SIAM (2011)
31. Shen, J., Tang, T., Wang, L.-L.: Spectral Methods: Algorithms, Analysis and Applications. Springer (2011)
32. Trefethen, L.N.: Householder triangularization of a quasimatrix. IMA J. Numer. Anal. **30**, 887–897 (2010)
33. Trefethen, L.N.: Series solution of Laplace problems. ANZIAM J. **60**, 1–26 (2018)
34. Trefethen, L.N., Nakatsukasa, Y., Weideman, J.A.C.: Exponential node clustering at singularities for rational approximation, quadrature, and PDEs. Numer. Math. **147**, 227–254 (2021)
35. Wright, T.G., Trefethen, L.N.: Pseudospectra of rectangular matrices. IMA J. Numer. Anal. **22**, 501–519 (2002)
36. Zhu, K.W.: Multivariate Least-Squares Approximations in Irregular Domains via Vandermonde with Arnoldi. MSc diss. University of Oxford (2021)

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.