Krylov subspace methods, minimal polynomials and clustering of eigenvalues: intriguing relationships and challenges

Zdeněk Strakoš
Charles University, Prague
Jindřich Nečas Center for Mathematical Modelling

Workshop on iterative methods, Strathclyde, September 27, 2019

Motto of this talk

A. Einstein,
in Oxford User's Guide to Mathematics, E. Zeidler (ed), OUP (2004), p. 3:

Motto of this talk

A. Einstein,
in Oxford User's Guide to Mathematics, E. Zeidler (ed), OUP (2004), p. 3:

"Everything should be made as simple as possible, but not simpler."

Motto of this talk

A. Einstein,
in Oxford User's Guide to Mathematics, E. Zeidler (ed), OUP (2004), p. 3:

"Everything should be made as simple as possible, but not simpler."

(This talk does not concern uniformly bounded condition and spectral numbers in preconditioning of discretized PDEs.)

Outline

1. (Preconditioned) Krylov subspace methods
2. Krylov subspace method as approximation of the minimal polynomial?
3. Conjugate gradients case study
4. Nonsymmetry and nonnormality
5. Conclusions

1 Hierarchy of linear problems starting at infinite dimension

Problem with bounded invertible operator $\mathcal{G}$ on the infinite dimensional Hilbert space S

$$\mathcal{G}\, u \;=\; f$$

is approximated on a finite dimensional subspace $S_h \subset S$ by a problem with the finite dimensional operator

$$\mathcal{G}_h\, u_h \;=\; f_h\,,$$

represented, using an appropriate basis of S_h, by the (sparse?) matrix problem

$$\mathbf{A}\,\mathbf{x} \;=\; \mathbf{b}\,.$$

1 Approximation by compact operators

Bounded operators in infinite dimensional Hilbert spaces having bounded inversion can not be compact.

1 Approximation by compact operators

Bounded operators in infinite dimensional Hilbert spaces having bounded inversion can not be compact.

They can be approximated by compact (e.g. finite dimensional) operators only in the sense of strong convergence (pointwise limit)

$$\| \mathcal{G}_h \, w - \mathcal{G} \, w \| \;\to\; 0 \quad \text{as} \quad h \to 0 \quad \text{for all } w \in S \, ;$$

The convergence $\mathcal{G}_h \, w \;\to\; \mathcal{G} \, w$ is not uniform w.r.t. w ; which points out the role of the right hand sides in approximate solutions of linear equations.

1 Approximation by compact operators

Bounded operators in infinite dimensional Hilbert spaces having bounded inversion can not be compact.

They can be approximated by compact (e.g. finite dimensional) operators only in the sense of strong convergence (pointwise limit)

$$\| \mathcal{G}_h \, w - \mathcal{G} \, w \| \;\to\; 0 \quad \text{as} \quad h \to 0 \quad \text{for all } w \in S \, ;$$

The convergence $\mathcal{G}_h \, w \;\to\; \mathcal{G} \, w$ is not uniform w.r.t. w ; which points out the role of the right hand sides in approximate solutions of linear equations.

Another non-uniformity wrt right hand sides: Krylov subspace methods.

1 Polynomial (Krylov subspace) methods

(Infinite dimensional) Krylov subspace methods at the step n implicitly construct the finite dimensional approximation $\mathcal{G}_n$ of $\mathcal{G}$ which determines the desired approximate solution $u_n \in u_0 + \mathcal{K}_n(\mathcal{G}, r), \quad r = f - \mathcal{G}u_0$

$$u_n := u_0 + p_{n-1}(\mathcal{G})\, r \approx u = \mathcal{G}^{-1} f\,.$$

Here $p_{n-1}(\lambda)$ is the associated polynomial of degree at most $n-1$ and $\mathcal{G}_n$ is obtained by restricting and projecting $\mathcal{G}$ onto the nth Krylov subspace

$$\mathcal{K}_n(\mathcal{G}, r) := \operatorname{span}\left\{r, \mathcal{G}r, \ldots, \mathcal{G}^{n-1}r\right\}.$$

A.N. Krylov (1931), Gantmakher (1934), Hestenes and Stiefel (1952), Lanczos (1952-53); Karush (1952), Hayes (1954), Stesin (1954), Vorobyev (1958)

From

$$r_n^{\mathrm{M}} \;=\; f \;-\; \mathcal{G}\, u_n^{\mathrm{M}} \;=\; r \;-\; \mathcal{G}\, p_{n-1}^{\mathrm{M}}(\mathcal{G})\, r \;=:\; \varphi_n^{\mathrm{M}}(\mathcal{G})\, r$$

we get the approximation polynomial

$$\varphi_n^{\mathrm{M}}(\lambda) = 1 \;-\; \lambda\, p_{n-1}^{\mathrm{M}}(\lambda)\,,$$

which is nonlinear both in $\mathcal{G}$ (obvious) and f (through the orthogonality/optimality property defining the particular method M). Clearly

$$\varphi_n^{\mathrm{M}}(0) \;=\; 1\,.$$

1 Now consider the discretized problem $\mathbf{Ax} = \mathbf{b}$

It is often stated that the main idea behind Krylov subspace methods is approximating the minimal polynomial of the system matrix (more accurately, of the system matrix wrt the initial residual). This polynomial approximation idea is subsequently combined with an argument about clustering of eigenvalues.

1 Now consider the discretized problem $\mathbf{Ax} = \mathbf{b}$

It is often stated that the main idea behind Krylov subspace methods is approximating the minimal polynomial of the system matrix (more accurately, of the system matrix wrt the initial residual). This polynomial approximation idea is subsequently combined with an argument about clustering of eigenvalues.

Reaching the minimal polynomial of $\mathbf{A}$ wrt $\mathbf{r}_0 = \mathbf{b} - \mathbf{A}\mathbf{x}_0$ means reaching the (exact) solution.

1 Now consider the discretized problem $\mathbf{Ax} = \mathbf{b}$

It is often stated that the main idea behind Krylov subspace methods is approximating the minimal polynomial of the system matrix (more accurately, of the system matrix wrt the initial residual). This polynomial approximation idea is subsequently combined with an argument about clustering of eigenvalues.

Reaching the minimal polynomial of $\mathbf{A}$ wrt $\mathbf{r}_0 = \mathbf{b} - \mathbf{A}\mathbf{x}_0$ means reaching the (exact) solution.

The minimal polynomial approximation idea combined with clustering of eigenvalues can, however, easily leads to incorrect statements.

Outline

1. (Preconditioned) Krylov subspace methods
2. Krylov subspace method as approximation of the minimal polynomial?
3. Conjugate gradients case study
4. Nonsymmetry and nonnormality
5. Conclusions

- If the minimal polynomial of a matrix has degree t,

  then the matrix has at most t distinct eigenvalues.

- If the minimal polynomial of a matrix has degree t,

  then the matrix has at most t distinct eigenvalues.

- If the matrix has t distinct eigenvalues,

  then without additional assumption we can say no more than the minimal polynomial has degree at least t. This degree can be as large as the size of the matrix.

- If the minimal polynomial of a matrix has degree t,

  then the matrix has at most t distinct eigenvalues.

- If the matrix has t distinct eigenvalues,

  then without additional assumption we can say no more than the minimal polynomial has degree at least t. This degree can be as large as the size of the matrix.

- Krylov subspace methods converge (in exact arithmetic and assuming no breakdown) to the exact solution in the number of steps estimated from above by the degree of the minimal polynomial of the system matrix.

- If the minimal polynomial of a matrix has degree t,

  then the matrix has at most t distinct eigenvalues.

- If the matrix has t distinct eigenvalues,

  then without additional assumption we can say no more than the minimal polynomial has degree at least t. This degree can be as large as the size of the matrix.

- Krylov subspace methods converge (in exact arithmetic and assuming no breakdown) to the exact solution in the number of steps estimated from above by the degree of the minimal polynomial of the system matrix.

  No analogous claim can be made, in general, in relation to the number of distinct eigenvalues.

2 Some perturbation results

- For a general matrix, a small perturbation of the matrix entries can induce a large change of the eigenvalues. Sensitivity of the eigenvalues depends of the structure of the invariant subspaces. Sensitivity of a simple eigenvalue is inversely proportional to the cosine of the angle between its left and right eigenvectors. Sensitivity of multiple eigenvalues can be worse.

- For a general matrix, a small perturbation of the matrix entries can induce a large change of the eigenvalues. Sensitivity of the eigenvalues depends of the structure of the invariant subspaces. Sensitivity of a simple eigenvalue is inversely proportional to the cosine of the angle between its left and right eigenvectors. Sensitivity of multiple eigenvalues can be worse.

- For normal matrices, however, a change of the eigenvalues is bounded by the size of the matrix perturbation (due to orthogonality of the eigenvectors that form the basis of the whole space). Still a fundamental difficulty remains!

2 Some perturbation results

- For a general matrix, a small perturbation of the matrix entries can induce a large change of the eigenvalues. Sensitivity of the eigenvalues depends of the structure of the invariant subspaces. Sensitivity of a simple eigenvalue is inversely proportional to the cosine of the angle between its left and right eigenvectors. Sensitivity of multiple eigenvalues can be worse.

- For normal matrices, however, a change of the eigenvalues is bounded by the size of the matrix perturbation (due to orthogonality of the eigenvectors that form the basis of the whole space). Still a fundamental difficulty remains!

- What can be claimed regarding the behaviour of Krylov subspace methods for the perturbed problems?

2 Some perturbation results

- For a general matrix, a small perturbation of the matrix entries can induce a large change of the eigenvalues. Sensitivity of the eigenvalues depends of the structure of the invariant subspaces. Sensitivity of a simple eigenvalue is inversely proportional to the cosine of the angle between its left and right eigenvectors. Sensitivity of multiple eigenvalues can be worse.

- For normal matrices, however, a change of the eigenvalues is bounded by the size of the matrix perturbation (due to orthogonality of the eigenvectors that form the basis of the whole space). Still a fundamental difficulty remains!

- What can be claimed regarding the behaviour of Krylov subspace methods for the perturbed problems?

- E.g., for the method of conjugate gradients and SPD matrices with large outlying mutliple eigenvalues?

2 Backward stability

- A method provides a backward stable approximation to the solution providing that the computed approximation solves exactly a nearby problem (i.e. a problem within a multiple of machine precision from the original one).

2 Backward stability

- A method provides a backward stable approximation to the solution providing that the computed approximation solves exactly a nearby problem (i.e. a problem within a multiple of machine precision from the original one).

- If the problem is sensitive to perturbations, backward stability does not imply small error of the computed approximation to the solution (small forward error). Nonsymmetric eigenvalue problems can provide illustrative examples.

2 Backward stability

- A method provides a backward stable approximation to the solution providing that the computed approximation solves exactly a nearby problem (i.e. a problem within a multiple of machine precision from the original one).

- If the problem is sensitive to perturbations, backward stability does not imply small error of the computed approximation to the solution (small forward error). Nonsymmetric eigenvalue problems can provide illustrative examples.

- How a backward error argument can be applied to Krylov subspace methods?

2 Backward stability

- A method provides a backward stable approximation to the solution providing that the computed approximation solves exactly a nearby problem (i.e. a problem within a multiple of machine precision from the original one).

- If the problem is sensitive to perturbations, backward stability does not imply small error of the computed approximation to the solution (small forward error). Nonsymmetric eigenvalue problems can provide illustrative examples.

- How a backward error argument can be applied to Krylov subspace methods?

- The case is clear only regarding the maximum attainable accuracy.

2 Backward stability

- A method provides a backward stable approximation to the solution providing that the computed approximation solves exactly a nearby problem (i.e. a problem within a multiple of machine precision from the original one).
- If the problem is sensitive to perturbations, backward stability does not imply small error of the computed approximation to the solution (small forward error). Nonsymmetric eigenvalue problems can provide illustrative examples.
- How a backward error argument can be applied to Krylov subspace methods?
- The case is clear only regarding the maximum attainable accuracy.
- The meaning is not clear regarding the convergence behaviour.

2 Correct statement (assuming exact arithmetic)

The low degree minimal polynomial argument
in Murphy, Golub and Wathen, SISC (2000):

A sufficient condition for a good preconditioner is that the preconditioned matrix $\mathbf{T} = \mathbf{P}^{-1}\mathbf{A}$ *has a low degree minimal polynomial.*

2 Correct statement (assuming exact arithmetic)

The low degree minimal polynomial argument
in Murphy, Golub and Wathen, SISC (2000):

A sufficient condition for a good preconditioner is that the preconditioned matrix $\mathbf{T} = \mathbf{P}^{-1}\mathbf{A}$ *has a low degree minimal polynomial. This condition is more usually expressed in terms of* $\mathbf{T}$ *having only a few distinct eigenvalues: in this form we must insist that* $\mathbf{T}$ *is not degenerate (.....) or at least that its Jordan canonical form has Jordan blocks of only small dimension.*

2 Correct statement (assuming exact arithmetic)

The low degree minimal polynomial argument
in Murphy, Golub and Wathen, SISC (2000):

A sufficient condition for a good preconditioner is that the preconditioned matrix $\mathbf{T} = \mathbf{P}^{-1}\mathbf{A}$ *has a low degree minimal polynomial. This condition is more usually expressed in terms of* $\mathbf{T}$ *having only a few distinct eigenvalues: in this form we must insist that* $\mathbf{T}$ *is not degenerate (.....) or at least that its Jordan canonical form has Jordan blocks of only small dimension.*

Here we skipped *(derogatory)* which can make the text difficult to understand.

? *If the preconditioned matrix has k distinct clusters of eigenvalues, then the backward stability of the algorithm in finite precision arithmetic together with the polynomial convergence bound based on eigenvalues ensures that, computationally, there will be a large error reduction after k steps if the algorithm is applied to the preconditioned system.*

2 Incorrect statements

? *If the preconditioned matrix has k distinct clusters of eigenvalues, then the backward stability of the algorithm in finite precision arithmetic together with the polynomial convergence bound based on eigenvalues ensures that, computationally, there will be a large error reduction after k steps if the algorithm is applied to the preconditioned system.*

? Let the eigenvalues of the diagonalizable preconditioned matrix belong into a few clusters, say t of them. If the diameters of the clusters are small enough, then the preconditioned matrix behaves numerically like a matrix with t distinct eigenvalues. As a result, we would expect t iterations of a Krylov subspace method to produce a reasonably accurate approximation.

Outline

1. (Preconditioned) Krylov subspace methods
2. Krylov subspace method as approximation of the minimal polynomial?
3. Conjugate gradients case study
4. Nonsymmetry and nonnormality
5. Conclusions

3 Minimization of the energy functional ($\mathcal{G}$ coercive self-adjoint)

Defining the energy functional

$$J(v) := \frac{1}{2}(\mathcal{G}v, v) - (f, v)\,, \qquad v \in S,$$

the solution is equivalently given by the condition

$$u \in S \text{ minimizes the functional } J \text{ over } S\,.$$

The Galerkin solution (of the discretized problem) then solves

$$u_h \in S_h \text{ minimizes the functional } J \text{ over } S_h\,.$$

3 Minimization of the energy functional ($\mathcal{G}$ coercive self-adjoint)

Defining the energy functional

$$J(v) := \frac{1}{2}(\mathcal{G}v, v) - (f, v)\,, \qquad v \in S,$$

the solution is equivalently given by the condition

$$u \in S \text{ minimizes the functional } J \text{ over } S\,.$$

The Galerkin solution (of the discretized problem) then solves

$$u_h \in S_h \text{ minimizes the functional } J \text{ over } S_h\,.$$

There is a flexibility in considering S_h Minimization of the energy functional over the sequence of Krylov subspaces defines the iterates of the conjugate gradient method in infinite dimensional Hilbert space.

- Finite dimensional self-adjoint operators (finite Hermitian matrices)

$$\begin{aligned}\mathcal{G} &= \frac{1}{2\pi\iota}\int_{\Gamma}\lambda\,(\lambda I_N-\mathcal{G})^{-1}\,d\lambda \;=\; \frac{1}{2\pi\iota}\sum_{\ell=1}^{N}\int_{\Gamma_\ell}\lambda\,(\lambda I_N-\mathcal{G})^{-1}\,d\lambda\\ &= \sum_{\ell=1}^{N} Y\,\mathrm{diag}\left(\frac{1}{2\pi\iota}\int_{\Gamma_\ell}\frac{\lambda}{\lambda-\lambda_\ell}\,d\lambda\right)Y^* \;=\; \sum_{\ell=1}^{N}\lambda_\ell\, y_\ell y_\ell^*\\ &= \int \lambda\, dE(\lambda)\,.\end{aligned}$$

- Compact infinite dimensional self-adjoint operators
- Bounded infinite dimensional self-adjoint operators
- Generalization to bounded normal and non-normal operators

3 Conjugate Gradient (CG) method for $Ax = b$ with A SPD (1952)

$r_0 = b - Ax_0$, $p_0 = r_0$. For $n = 1, \ldots, n_{\max}$:

$$
\begin{aligned}
\alpha_{n-1} &= \frac{r_{n-1}^* r_{n-1}}{p_{n-1}^* A p_{n-1}} \\
x_n &= x_{n-1} + \alpha_{n-1} p_{n-1}\,, \quad \text{stop when the stopping criterion is satisfied} \\
r_n &= r_{n-1} - \alpha_{n-1} A p_{n-1} \\
\beta_n &= \frac{r_n^* r_n}{r_{n-1}^* r_{n-1}} \\
p_n &= r_n + \beta_n p_{n-1}
\end{aligned}
$$

Here α_{n-1} ensures the minimization of the energy norm $\|x - x_n\|_A$ along the line

$$z(\alpha) = x_{n-1} + \alpha p_{n-1}\,.$$

Provided that

$$p_i \perp_A p_j, \quad i \neq j,$$

the one-dimensional line minimizations at the individual steps 1 to n result in the n-dimensional minimization over the whole shifted Krylov subspace

$$x_0 + \mathcal{K}_n(A, r_0) \;=\; x_0 + \mathrm{span}\{p_0, p_1, \ldots, p_{n-1}\}\,.$$

Indeed,

$$x - x_0 \;=\; \sum_{\ell=0}^{N-1} \alpha_\ell p_\ell \;=\; \sum_{\ell=0}^{n-1} \alpha_\ell p_\ell \;+\; x - x_n\,,$$

where

$$x - x_n \;\perp_A\; K_n(A, r_0)\,, \quad \text{or, equivalently,} \quad r_n \;\perp\; K_n(A, r_0)\,.$$

$$w_1^* A w_1 \;=\; w_1^* \left(\sum_{\ell=1}^{N} \lambda_\ell \, y_\ell y_\ell^* \right) w_1 \;\equiv\; w_1^* \left(\int \lambda \, dE(\lambda) \right) w_1$$

$$= \sum_{\ell=1}^{N} \lambda_\ell \, |(y_\ell, w_1)|^2 \;=\; \sum_{\ell=1}^{N} \lambda_\ell \, \omega_\ell = \int \lambda \, d\omega(\lambda) \,,$$

where the spectral function $E(\lambda)$ of A is understood to be a nondecreasing family of projections with increasing λ, symbolically $dE(\lambda_\ell) \equiv y_\ell y_\ell^*$ and

$$I \;=\; \sum_{\ell=1}^{N} y_\ell y_\ell^* \;\equiv\; \int dE(\lambda) \,.$$

Hilbert (1906, 1912, 1928), Von Neumann (1927, 1932), Wintner (1929) .

3 Distribution function

Distribution function $\omega(\lambda)$ associated with $Ax = b,\ r_0 = b - Ax_0,\ A$ SPD,

λ_i, y_i are the eigenpairs of A, $\omega_i = |(y_i, w_1)|^2$, $(w_1 = r_0/\|r_0\|)$

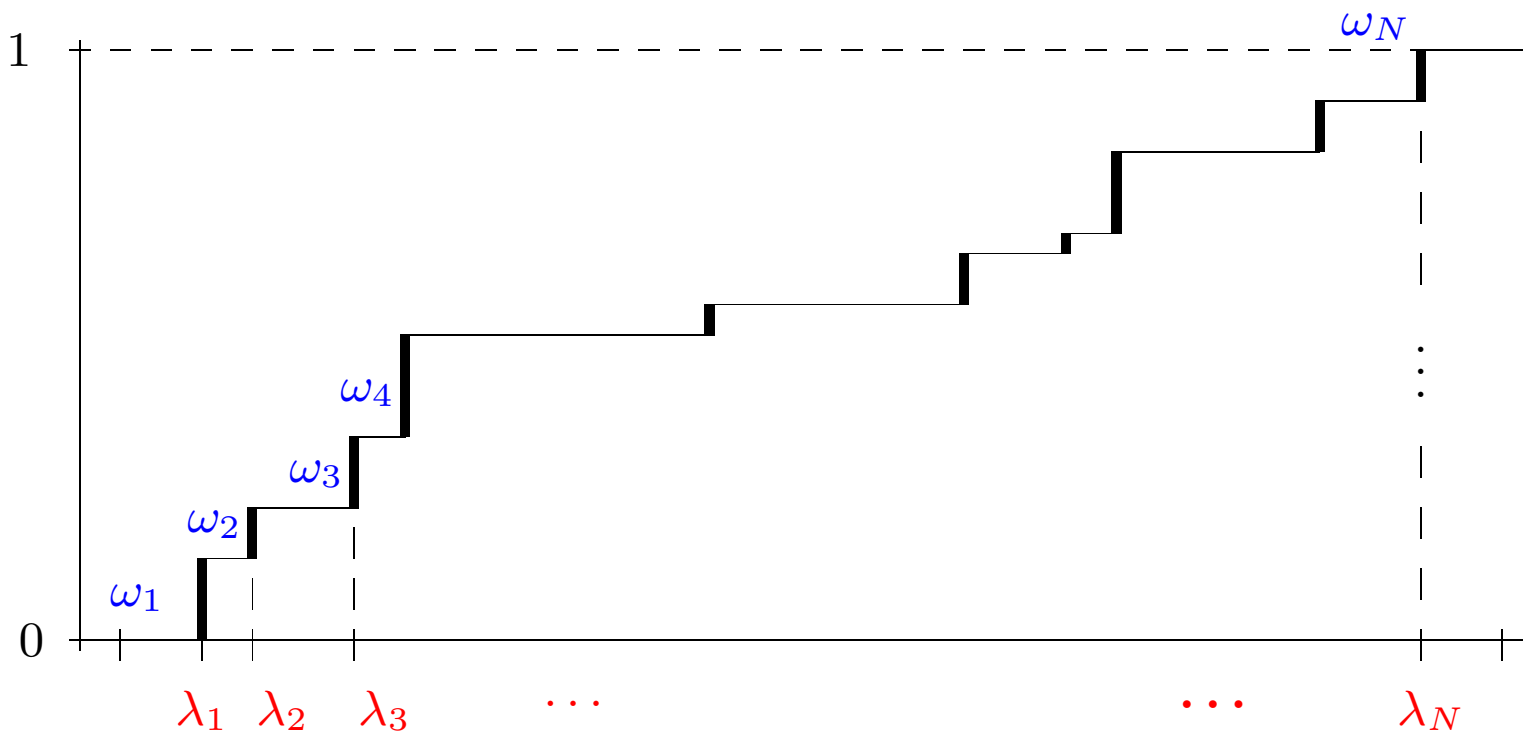

- With $\omega(\lambda)$ determined by the SPD A and r_0, solve the finite Stieltjes moment problem, i.e., determine the distribution function $\omega^{(n)}(\lambda)$ with the n points of increase such that the first $2n$ moments are matched, i.e.,

$$m_\ell = \int_0^\infty \lambda^\ell \, d\omega(\lambda) = \int_0^\infty \lambda^\ell \, \omega^{(n)}(\lambda) \,, \qquad \ell = 0, 1, 2, \ldots, 2n-1 \,, \quad n < N \,.$$

- With $\omega(\lambda)$ determined by the SPD A and r_0, solve the finite Stieltjes moment problem, i.e., determine the distribution function $\omega^{(n)}(\lambda)$ with the n points of increase such that the first $2n$ moments are matched, i.e.,

$$m_\ell = \int_0^\infty \lambda^\ell \, d\omega(\lambda) = \int_0^\infty \lambda^\ell \, \omega^{(n)}(\lambda) \,, \qquad \ell = 0, 1, 2, \ldots, 2n-1 \,, \quad n < N \,.$$

- Equivalently, compute the n-point Gauss quadrature of the Riemann-Stieltjes integral with the distribution function $\omega(\lambda)$.

- With $\omega(\lambda)$ determined by the SPD A and r_0, solve the finite Stieltjes moment problem, i.e., determine the distribution function $\omega^{(n)}(\lambda)$ with the n points of increase such that the first $2n$ moments are matched, i.e.,

$$m_\ell = \int_0^\infty \lambda^\ell \, d\omega(\lambda) = \int_0^\infty \lambda^\ell \, \omega^{(n)}(\lambda) \,, \qquad \ell = 0, 1, 2, \ldots, 2n-1 \,, \quad n < N \,.$$

- Equivalently, compute the n-point Gauss quadrature of the Riemann-Stieltjes integral with the distribution function $\omega(\lambda)$.

- Apply the Vorobyev method of moments to the vectors $z_0, z_1, \ldots, z_n$ given by $r_0, Ar_0, \ldots, A^n r_0$, which corresponds to matching the $2n$ moments ($w_1 = r_0/\|r_0\|$)

$$w_1^* A^\ell w_1 = w_1^* A_n^\ell w_1 \,, = e_1^* T_n^\ell e_1 \,, \qquad \ell = 0, 1, \ldots, 2n-1 \,.$$

- With $\omega(\lambda)$ determined by the SPD A and r_0, solve the finite Stieltjes moment problem, i.e., determine the distribution function $\omega^{(n)}(\lambda)$ with the n points of increase such that the first $2n$ moments are matched, i.e.,

$$m_\ell = \int_0^\infty \lambda^\ell \, d\omega(\lambda) = \int_0^\infty \lambda^\ell \, \omega^{(n)}(\lambda) \,, \qquad \ell = 0, 1, 2, \ldots, 2n-1 \,, \quad n < N \,.$$

- Equivalently, compute the n-point Gauss quadrature of the Riemann-Stieltjes integral with the distribution function $\omega(\lambda)$.

- Apply the Vorobyev method of moments to the vectors $z_0, z_1, \ldots, z_n$ given by $r_0, Ar_0, \ldots, A^n r_0$, which corresponds to matching the $2n$ moments ($w_1 = r_0/\|r_0\|$)

$$w_1^* A^\ell w_1 = w_1^* A_n^\ell w_1 \,, = e_1^* T_n^\ell e_1 \,, \qquad \ell = 0, 1, \ldots, 2n-1 \,.$$

- Equivalently, consider the system $A\,x = b$ with the initial approximation x_0 and compute n iterations of the conjugate gradient (Lanczos) method.

Let $W_n = [w_1, \ldots, w_n]$, $A W_n = W_n T_n + \delta_{n+1} w_{n+1} e_n^T$, form the Lanczos orthonormal basis of the Krylov subspace $K_n(A, r_0)$. Here the Jacobi matrix of the orthonormalization coefficients

$$T_n = \begin{pmatrix} \gamma_1 & \delta_2 & & & \\ \delta_2 & \ddots & \ddots & & \\ & \ddots & \ddots & \ddots & \\ & & \ddots & \ddots & \delta_n \\ & & & \delta_n & \gamma_n \end{pmatrix}$$

represents, at the same time, the matrix of the restricted and orthogonally projected operator $A_n = W_n W_n^* A$ on $K_n(A, r_0)$ in the basis W_n.
The CG approximation is determined by

$$T_n t_n = \|r_0\| e_1 , \quad x_n = x_0 + W_n t_n .$$

3 Summary of the relationship

$$
\begin{array}{ccc}
A,\ v_1 = r_0/\|r_0\| & \longleftrightarrow & \omega(\lambda),\ \int f(\lambda)\, d\omega(\lambda) \\
\downarrow & & \downarrow \\
\text{Hermitian Lanczos/CG} & & \text{Gauss-Christoffel quadrature} \\
\downarrow & & \downarrow \\
T_n,\ e_1 & \longleftrightarrow & \omega^{(n)}(\lambda),\ \sum_{j=1}^{n} \omega_j^{(n)} f(\theta_j^{(n)})
\end{array}
$$

The quadrature nodes $\lambda_j^{(n)}$ are the eigenvalues $\theta_j^{(n)}$ of T_n and the weights $\omega_j^{(n)}$ are the squared first components of the associated normalized eigenvectors.

At any iteration step n, CG represents the matrix formulation of the n-point Gauss quadrature of the Riemann-Stieljes integral determined by A and r_0,

$$\int_0^\infty f(\lambda)\, d\omega(\lambda) \;=\; \sum_{i=1}^{n} \omega_i^{(n)} f(\theta_i^{(n)}) \;+\; R_n(f)\,.$$

For $f(\lambda) \equiv \lambda^{-1}$,

$$\frac{\|x - x_0\|_A^2}{\|r_0\|^2} \;=\; n\text{-th Gauss quadrature} \;+\; \frac{\|x - x_n\|_A^2}{\|r_0\|^2}\,.$$

3 In contrast, linear "description" of the nonlinear CG method

- The CG optimality property

$$\|x - x_n\|_A = \min_{z \in x_0 + \mathcal{K}_n(A, r_0)} \|x - z\|_A = \min_{p \in \mathcal{P}_n(0)} \|p(A)(x - x_0)\|_A$$

yields in two derivation steps the (worst case) linear convergence bound valid and relevant for the Chebyshev method

$$\frac{\|x - x_n\|_A}{\|x - x_0\|_A} \leq \min_{p \in \mathcal{P}_n(0)} \max_{1 \leq j \leq N} |p(\lambda_j)| \leq \min_{p \in \mathcal{P}_n(0)} \max_{\lambda \in [\lambda_1, \lambda_N]} |p(\lambda)|$$

$$\leq 2 \left(\frac{\sqrt{\kappa(A)} - 1}{\sqrt{\kappa(A)} + 1} \right)^n, \quad \kappa(A) = \frac{\lambda_N}{\lambda_1}.$$

- The worst-case nonlinear bound is completely determined by the distribution of the eigenvalues of A.

3 In contrast, linear "description" of the nonlinear CG method

- The CG optimality property

$$\|x - x_n\|_A = \min_{z \in x_0 + \mathcal{K}_n(A, r_0)} \|x - z\|_A = \min_{p \in \mathcal{P}_n(0)} \|p(A)(x - x_0)\|_A$$

yields in two derivation steps the (worst case) linear convergence bound valid and relevant for the Chebyshev method

$$\frac{\|x - x_n\|_A}{\|x - x_0\|_A} \leq \min_{p \in \mathcal{P}_n(0)} \max_{1 \leq j \leq N} |p(\lambda_j)| \leq \min_{p \in \mathcal{P}_n(0)} \max_{\lambda \in [\lambda_1, \lambda_N]} |p(\lambda)|$$

$$\leq 2 \left(\frac{\sqrt{\kappa(A)} - 1}{\sqrt{\kappa(A)} + 1} \right)^n, \quad \kappa(A) = \frac{\lambda_N}{\lambda_1}.$$

- The worst-case nonlinear bound is completely determined by the distribution of the eigenvalues of A.

3 Remarkable history, rarely quoted

- Markov (1890)
- Flanders and Shortley (1950)
- Lanczos (1953), Kincaid (1947), Young (1954, ...)
- Stiefel (1958), Rutishauser (1959)
- Meinardus (1963), Kaniel (1966)
- Daniel (1967a, 1967b)
- Luenberger (1969)

Derivations are repeated in recent textbooks and monographs and the resulting bound is identified with the convergence of CG without noticing severe limitations.

3 Example: Composite bounds considering large outliers

The condition-number-based bound should be used with a great care in connection with the behaviour of CG unless $\kappa(A) = \lambda_N/\lambda_1$ is really small or unless the (very special) distribution of eigenvalues makes the bound tight.

The condition-number-based bound should be used with a great care in connection with the behaviour of CG unless $\kappa(A) = \lambda_N/\lambda_1$ is really small or unless the (very special) distribution of eigenvalues makes the bound tight.

In particular, one should be very careful while using it as a part of a composite bound in the presence of the large outlying eigenvalues

$$\min_{\substack{p(0)=1 \\ \deg(p)\le n-s}} \max_{1\le j\le N} |\, q_s(\lambda_j)\, p(\lambda_j)\,| \quad \le \quad \max_{1\le j\le N} |q_s(\lambda_j)| \left| \frac{T_{n-s}(\lambda_j)|}{T_{n-s}(0)} \right|$$

$$< \quad \max_{1\le j\le N-s} \left| \frac{T_{n-s}(\lambda_j)}{T_{n-s}(0)} \right| .$$

This Chebyshev method bound on the interval $[\lambda_1, \lambda_{N-s}]$ is then valid after s initial steps.

3 Polynomial $q_s(\lambda)$ has the desired root, but look at $T_{4-5}(\lambda)$

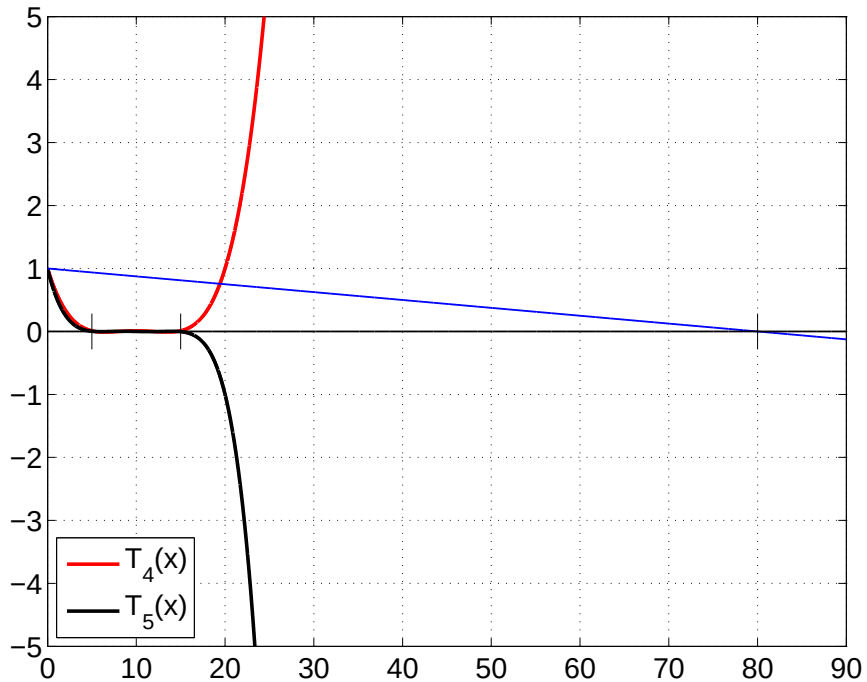

The Chebyshev polynomials $T_4(\lambda), T_5(\lambda)$, and the polynomial $q_1(\lambda)$, $q_1(0) = 1$ having the single root at the large outlying eigenvalue.

Consider the desired accuracy ϵ, $\kappa_s(A) \equiv \lambda_{N-s}/\lambda_1$. Then, assuming exact arithmetic, n CG steps, where

$$n = \mathbf{s} + \left\lceil \frac{\ln(2/\epsilon)}{2} \sqrt{\kappa_s(A)} \right\rceil ,$$

will produce the approximate solution x_n satisfying

$$\|x - x_n\|_A \leq \epsilon \, \|x - x_0\|_A .$$

This statement has been applied to explain superlinear convergence of CG at the presence of large outliers in the spectrum. Due to rounding errors, this concept can not be applied to practical computations.

3 Why is consideration of rounding errors so fundamental?

1. CG should be used when it has a chance to accelerate its convergence due to adaptation to the information (hidden) in data, i.e., when the eigenvalues are not almost uniformly distributed in the spectral interval. Presence of large outlying eigenvalues seem the most favourable case.

3 Why is consideration of rounding errors so fundamental?

1. CG should be used when it has a chance to accelerate its convergence due to adaptation to the information (hidden) in data, i.e., when the eigenvalues are not almost uniformly distributed in the spectral interval. Presence of large outlying eigenvalues seem the most favourable case.

2. Hovewer, apart from the trivial situation mentioned next, in such cases CG convergence behaviour is then substantially affected by rounding errors due to the loss of orthogonality among the direction vectors (and residuals).

3 Why is consideration of rounding errors so fundamental?

1. CG should be used when it has a chance to accelerate its convergence due to adaptation to the information (hidden) in data, i.e., when the eigenvalues are not almost uniformly distributed in the spectral interval. Presence of large outlying eigenvalues seem the most favourable case.

2. Hovewer, apart from the trivial situation mentioned next, in such cases CG convergence behaviour is then substantially affected by rounding errors due to the loss of orthogonality among the direction vectors (and residuals).

3. This is the generic situation. If CG behaviour is not affected by rounding errors, then either we are lucky because convergence is so fast that rounding errors have not enough iterations to amplify (trivial cases), or CG convergence is hopelessly linear with no chance to accelerate. In the latter case linear methods would with a high probability be much more efficient in terms of computing time (as well as energy consumption).

3 Why is consideration of rounding errors so fundamental?

1. CG should be used when it has a chance to accelerate its convergence due to adaptation to the information (hidden) in data, i.e., when the eigenvalues are not almost uniformly distributed in the spectral interval. Presence of large outlying eigenvalues seem the most favourable case.

2. Hovewer, apart from the trivial situation mentioned next, in such cases CG convergence behaviour is then substantially affected by rounding errors due to the loss of orthogonality among the direction vectors (and residuals).

3. This is the generic situation. If CG behaviour is not affected by rounding errors, then either we are lucky because convergence is so fast that rounding errors have not enough iterations to amplify (trivial cases), or CG convergence is hopelessly linear with no chance to accelerate. In the latter case linear methods would with a high probability be much more efficient in terms of computing time (as well as energy consumption).

4. The presented facts do not allow any escape.

3 Explanation

3 In which sense $\varphi_n^{\mathrm{CG}}(\lambda)$ approximates the minimal polynomial?

$$\begin{aligned} \|\mathbf{x}-\mathbf{x}_n\|_{\mathbf{A}}^2 &= \min_{\varphi\in\Pi_n} \|\varphi(\mathbf{A})(\mathbf{x}-\mathbf{x}_0)\|_{\mathbf{A}}^2 \\ &= \sum_{j=1}^{N} \lambda_j \,|\varphi_n^{\mathrm{CG}}(\lambda_j)\,\zeta_j|^2, \quad j=1,2,\dots \end{aligned}$$

Here

$$\varphi_n^{\mathrm{CG}}(\lambda) = \frac{(\lambda-\theta_1^{(n)})\cdots(\lambda-\theta_n^{(n)})}{(-1)^n\,\theta_1^{(n)}\cdots\theta_n^{(n)}}$$

is determined by the eigenvalues of the orthogonally restricted operator (Ritz values) $\theta_1^{(n)},\dots,\theta_n^{(n)}$.

3 Illustration of $\varphi_n^{\mathrm{CG}}(\lambda)$, $\mathbf{Ax} = \mathbf{b}$, a single large outlier λ_N

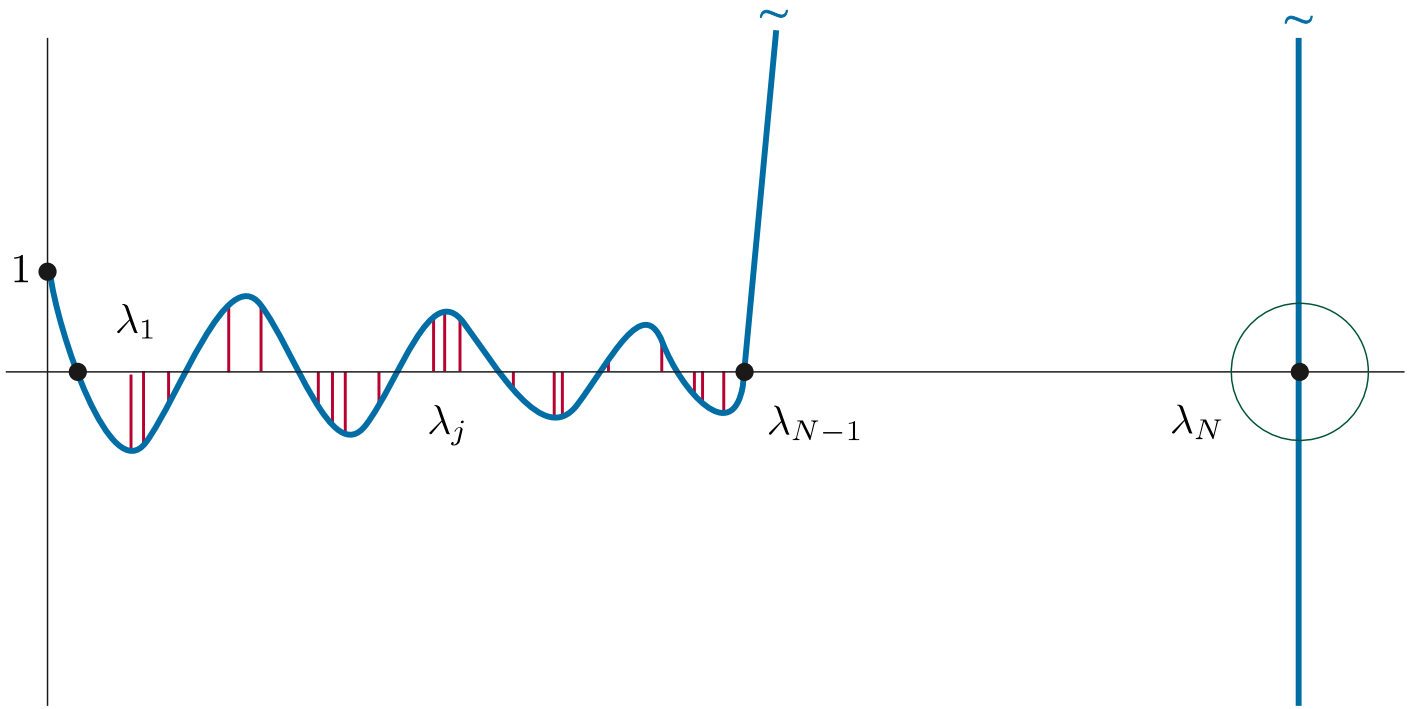

3 Consider replacing λ_N by a tight cluster

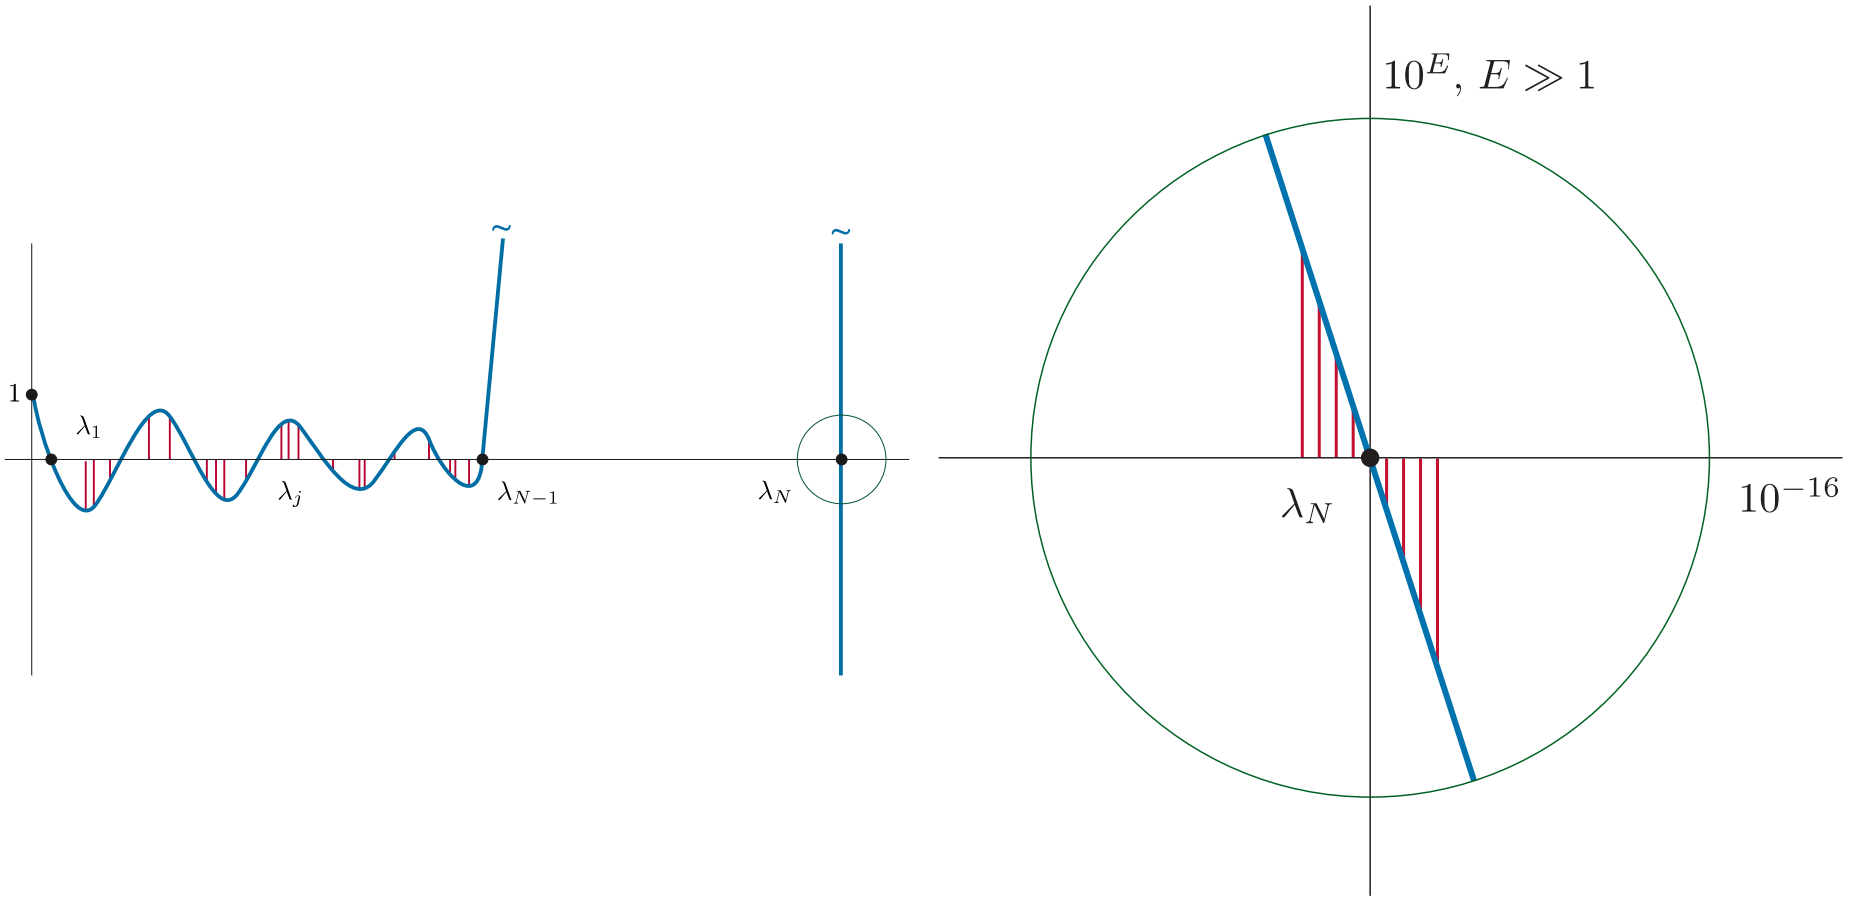

Since E in the illustration of the slope is enormous, many roots close to λ_N are needed.

3 Consider replacing λ_N by a tight cluster

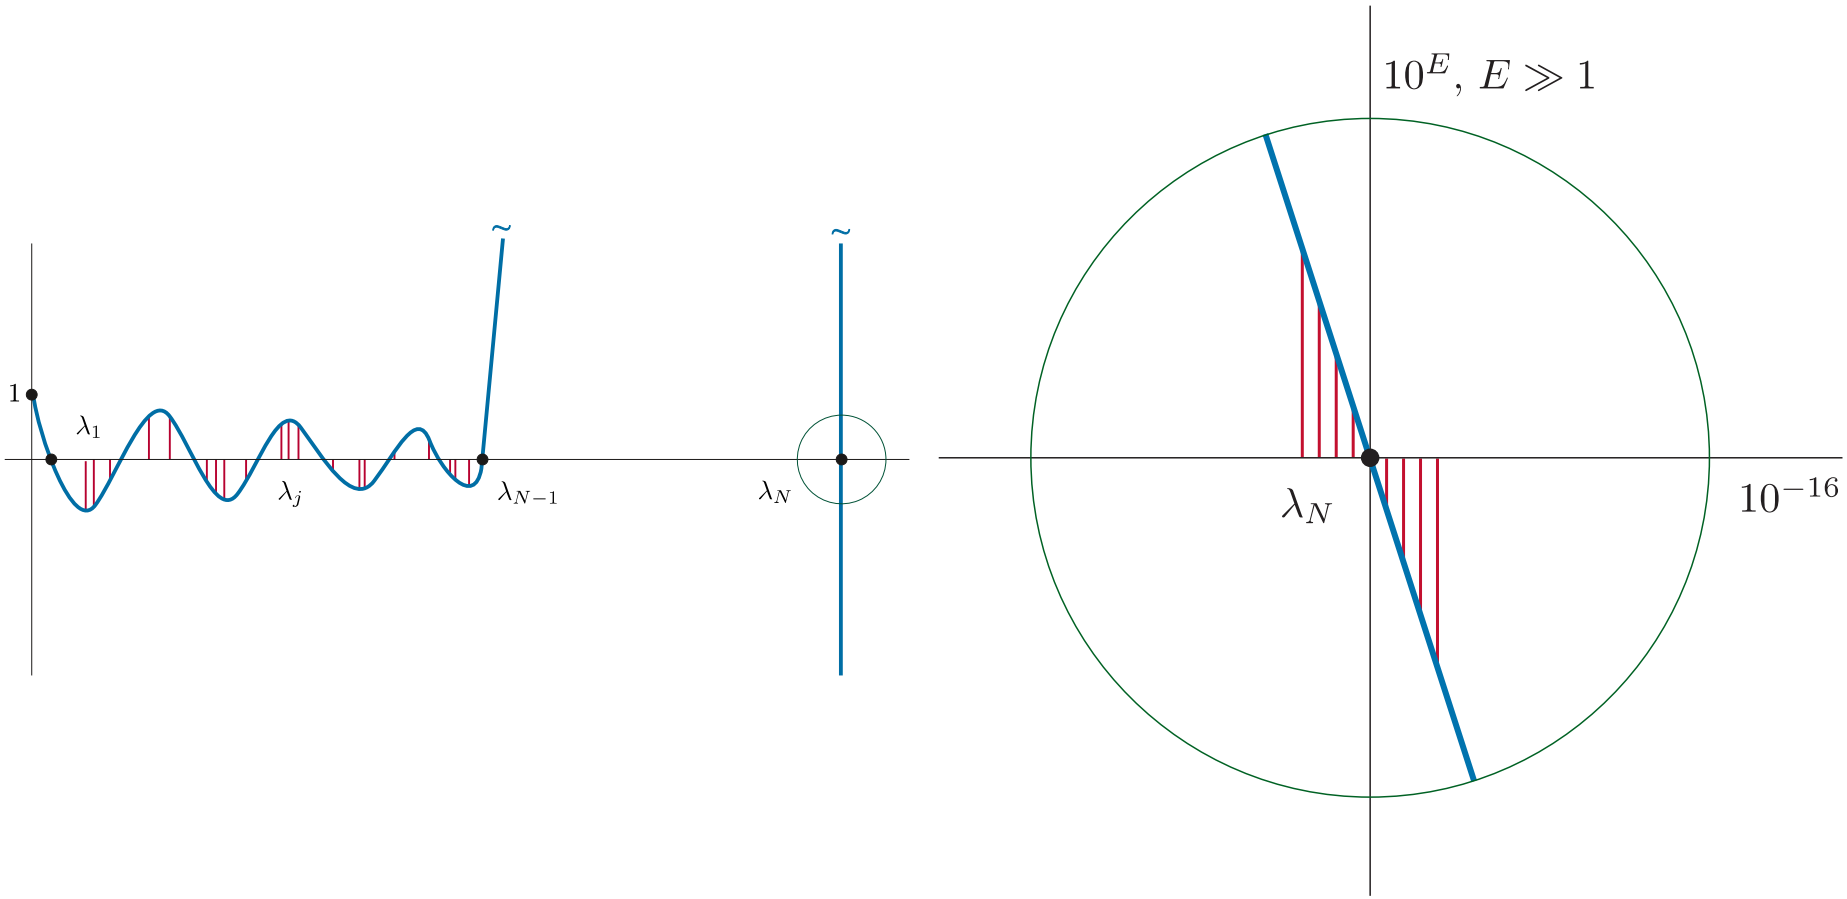

Since E in the illustration of the slope is enormous, many roots close to λ_N are needed. Relationship with propagation of rounding errors was revealed by Paige (1971 - 80); Greenbaum (1989); Greenbaum, S (1992); Golub, S (1994); ...

3 Beautiful idea of Greenbaum (1989) - not a perturbation argument!

Consider the Jacobi matrix $\mathbf{T}_n$ computed in n steps of CG in FP arithmetic. This matrix is constructively extended to a larger Jacobi matrix $\mathbf{T}_{n+m(n)}$ having all its (clusters of) eigenvalues close to the eigenvalues of the matrix $\mathbf{A}$.

3 Beautiful idea of Greenbaum (1989) - not a perturbation argument!

Consider the Jacobi matrix $\mathbf{T}_n$ computed in n steps of CG in FP arithmetic. This matrix is constructively extended to a larger Jacobi matrix $\mathbf{T}_{n+m(n)}$ having all its (clusters of) eigenvalues close to the eigenvalues of the matrix $\mathbf{A}$.

The EXACT Lanczos (CG) applied to $\mathbf{T}_{n+m(n)}$ and the initial residual e_1 thus gives in the first n steps $\mathbf{T}_n$, which was originally COMPUTED in FP arithmetic. Using this, finite precision computation is analyzed as exact computation for the problem having clusters of eigenvalues.

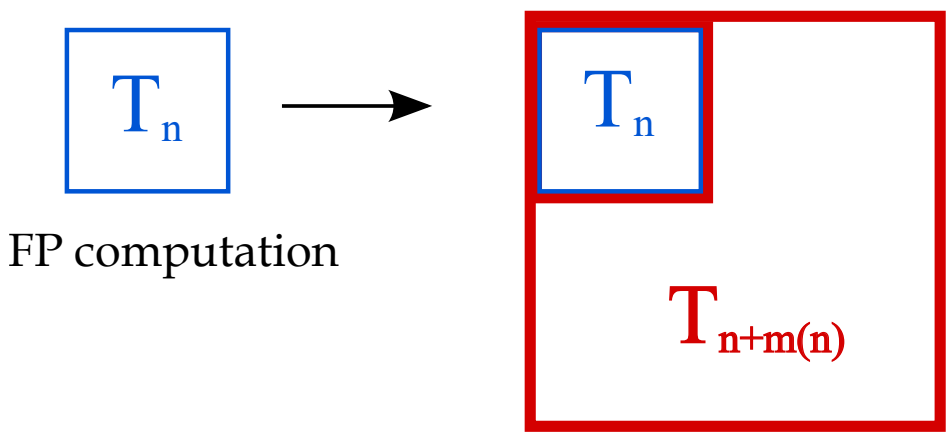

3 The appearance of clusters is not uniform (Carson, S (2019))

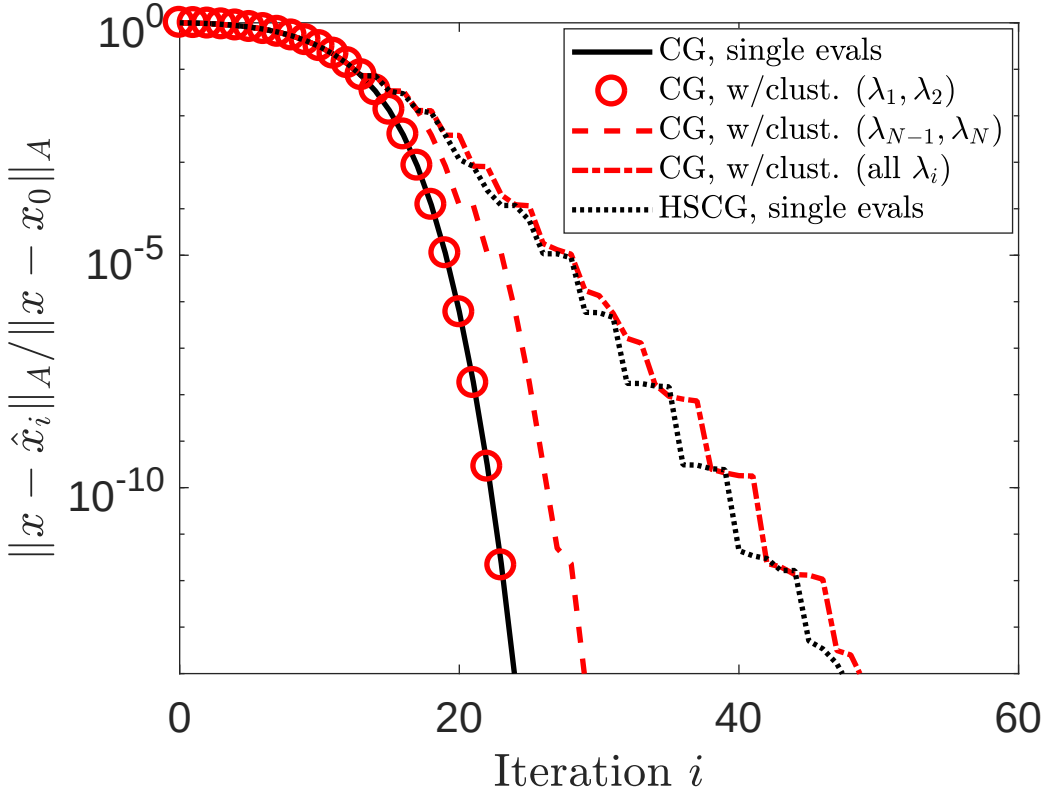

Given λ_1, λ_N and a parameter $0 < \gamma \leq 1$,

$$\lambda_j = \lambda_1 + \frac{j-1}{N-1}(\lambda_N - \lambda_1)\gamma^{N-j}, \quad j = 2, \ldots, N-1.$$

3 Consequence

Argument about replacing clusters by its single number representatives (and vice versa) is applicable to clusters (eigenvalues) positioned between zero and the bulk of the spectrum, i.e., to small outliers.

Van der Vorst, Van der Sluis (1987); Liesen, S (2013); Gergelits, S (2015).

3 Summary of the point

3 Summary of the point

1. The first n steps of finite precision Lanczos (CG) applied to $\mathbf{A}$ behaves, as proved in Greenbaum (1989), as exact Lanczos (CG) applied to a certain $\tilde{\mathbf{A}}_{n+m(n)}$ having clusters of eigenvalues arround the original eigenvalues of $\mathbf{A}$.

3 Summary of the point

1. The first n steps of finite precision Lanczos (CG) applied to $\mathbf{A}$ behaves, as proved in Greenbaum (1989), as exact Lanczos (CG) applied to a certain $\tilde{\mathbf{A}}_{n+m(n)}$ having clusters of eigenvalues arround the original eigenvalues of $\mathbf{A}$.

2. Finite precision Lanczos (CG) applied to $\mathbf{A}$ in general behaves differently from the exact Lanczos (CG) applied to the same matrix $\mathbf{A}$.

3 Summary of the point

1. The first n steps of finite precision Lanczos (CG) applied to $\mathbf{A}$ behaves, as proved in Greenbaum (1989), as exact Lanczos (CG) applied to a certain $\tilde{\mathbf{A}}_{n+m(n)}$ having clusters of eigenvalues arround the original eigenvalues of $\mathbf{A}$.

2. Finite precision Lanczos (CG) applied to $\mathbf{A}$ in general behaves differently from the exact Lanczos (CG) applied to the same matrix $\mathbf{A}$.

3. Therefore, in general, exact Lanczos (CG) applied to the matrix $\mathbf{A}$ can not behave closely to the exact Lanczos (CG) applied to the matrix with the eigenvalues replaced by tight clusters, and vice-versa.

3 Summary of the point

1. The first n steps of finite precision Lanczos (CG) applied to $\mathbf{A}$ behaves, as proved in Greenbaum (1989), as exact Lanczos (CG) applied to a certain $\tilde{\mathbf{A}}_{n+m(n)}$ having clusters of eigenvalues arround the original eigenvalues of $\mathbf{A}$.

2. Finite precision Lanczos (CG) applied to $\mathbf{A}$ in general behaves differently from the exact Lanczos (CG) applied to the same matrix $\mathbf{A}$.

3. Therefore, in general, exact Lanczos (CG) applied to the matrix $\mathbf{A}$ can not behave closely to the exact Lanczos (CG) applied to the matrix with the eigenvalues replaced by tight clusters, and vice-versa.

   Therefore the argument persistently used in literature about representing clusters by single eigenvalues and vice-versa is mathematically incorrect.

3 Summary of the point

1. The first n steps of finite precision Lanczos (CG) applied to $\mathbf{A}$ behaves, as proved in Greenbaum (1989), as exact Lanczos (CG) applied to a certain $\tilde{\mathbf{A}}_{n+m(n)}$ having clusters of eigenvalues arround the original eigenvalues of $\mathbf{A}$.

2. Finite precision Lanczos (CG) applied to $\mathbf{A}$ in general behaves differently from the exact Lanczos (CG) applied to the same matrix $\mathbf{A}$.

3. Therefore, in general, exact Lanczos (CG) applied to the matrix $\mathbf{A}$ can not behave closely to the exact Lanczos (CG) applied to the matrix with the eigenvalues replaced by tight clusters, and vice-versa.

   Therefore the argument persistently used in literature about representing clusters by single eigenvalues and vice-versa is mathematically incorrect.

   Here we compare two mathematical approximation problems. There is no effect of rounding errors involved!

Outline

$$x - x_n \;=\; p_n(A)\,(x - x_0) \;=\; W p_n(J)\,W^{-1}\,(x - x_0)$$

does not offer an insight unless W is close to unitary and the spectrum of A has some particular structure. Separating the operator from the initial error and using Cauchy integral representation

$$p_n(A) \;=\; \frac{1}{2\pi i}\int_\Gamma p(\lambda)\,(\lambda - A)^{-1}\,d\lambda\,.$$

we get with denoting $\mathcal{L}(\Gamma)$ the length of the integrating curve

$$\|p_n(A)\| \;=\; \frac{\mathcal{L}(\Gamma)}{2\pi}\,\max_{\lambda\in\Gamma}\|(\lambda - A)^{-1}\|\,\max_{\lambda\in\Gamma}|p(\lambda)|\,.$$

Considering the curves Γ_ϵ on which the resolvent norm is constant $\|(\lambda - A)^{-1}\| = 1/\epsilon$, i.e., the boundaries of the ϵ-pseudospectra of A,

$$\|p_n(A)\| \;=\; \frac{\mathcal{L}(\Gamma_\epsilon)}{2\pi\epsilon}\,\max_{\lambda\in\Gamma_\epsilon}|p(\lambda)|\,.$$

4 Spectral information alone is not sufficient

Given any spectrum and any sequence of the nonincreasing residual norms, a complete parametrization is known of the set of all matrices and right-hand sides for which GMRES generates the prescribed convergence curve while the matrix has the prescribed spectrum.

The set of problems for which the distribution of eigenvalues alone does not correspond to convergence behavior is not of measure zero and it is not pathological.

4 Spectral information alone is not sufficient

Given any spectrum and any sequence of the nonincreasing residual norms, a complete parametrization is known of the set of all matrices and right-hand sides for which GMRES generates the prescribed convergence curve while the matrix has the prescribed spectrum.

The set of problems for which the distribution of eigenvalues alone does not correspond to convergence behavior is not of measure zero and it is not pathological.

- Widespread eigenvalues alone can not be identified with poor convergence.
- Clustered eigenvalues alone can not be identified with fast convergence.

Here the link between the matrix and the right-hand side is crucial.

Theorem

1° The spectrum of A is given by $\{\lambda_1, \ldots, \lambda_N\}$ and GMRES(A, r_0) yields residuals with the nonincreasing sequence

$$\|r_0\| \geq \|r_1\| \geq \cdots \geq \|r_{N-1}\| > \|r_N\| = 0\,.$$

2° Let C be the spectral companion matrix, $h = (h_1, \ldots, h_N)^T$, $h_i^2 = \|r_{i-1}\|^2 - \|r_i\|^2$, $i = 1, \ldots, N$. Let R be a nonsingular upper triangular matrix such that $Rs = h$ with s being the first column of C^{-1}, and let W be a unitary matrix. Then

$$A = W R C R^{-1} W^* \quad \text{and} \quad b = W h\,.$$

Greenbaum, Pták, Arioli and S (1994 - 98); Liesen (1999); Eiermann and Ernst (2001); Meurant (2012); Meurant and Tebbens (2012, 2014); .....

Theorem

Let the spectrum of A be given by $\{\lambda_1, \ldots, \lambda_N\}$ and GMRES(A, r_0) yields residuals with the non-increasing sequence

$$\|r_0\| \geq \|r_1\| \geq \cdots \geq \|r_{N-1}\| > \|r_N\| = 0\,.$$

1° There is always a GMRES-equivalent unitary matrix B.

2° Let zero is out of the field of values of the matrix A.
Then there is always a GMRES-equivalent
Hermitian positive definite matrix B.

Unfortunately, there is no known relationship between the spectra of A and B.
Greenbaum and S (1994)

4 Positive results and unexplored research challenges

- When the operators (matrices arising from discretization) are far from normal and the spectral information is descriptive for convergence behavior of Krylov subspace methods, this points out to some fundamental mathematical, physical, ... properties of the problem.

- When the operators (matrices arising from discretization) are far from normal and the spectral information is descriptive for convergence behavior of Krylov subspace methods, this points out to some fundamental mathematical, physical, ... properties of the problem.

- In such cases there must be some special inner structure of invariant subspaces and/or special right hand side (in BVP that means boundary conditions and outer forces).

4 Positive results and unexplored research challenges

- When the operators (matrices arising from discretization) are far from normal and the spectral information is descriptive for convergence behavior of Krylov subspace methods, this points out to some fundamental mathematical, physical, ... properties of the problem.

- In such cases there must be some special inner structure of invariant subspaces and/or special right hand side (in BVP that means boundary conditions and outer forces).

- Very little has unfortunately been done in that much needed direction of research. The challenge is rarely mentioned.

Outline

1. (Preconditioned) Krylov subspace methods
2. Krylov subspace method as approximation of the minimal polynomial?
3. Conjugate gradients case study
4. Nonsymmetry and nonnormality
5. Conclusions

5 Myths in Mathematics?

Myth:
A belief given uncritical acceptance by the members of a group especially in support of existing or traditional practices and institutions.

Webster's Third New International Dictionary, Enc. Britannica Inc., Chicago (1986)

5 Myths concerning Krylov subspace methods include

- Minimal polynomials and finite termination property
- Chebyshev bounds and CG
- Spectral information and clustering of eigenvalues
- Operator-based bounds and functional analysis arguments on convergence rate
- Finite precision computations seen as a minor modification of the exact considerations
- Linearization of nonlinear phenomenon
- Considering CG in matrix computations as a simplification of CG in general nonlinear optimization
- Well conditioned basis and short recurrences (look-ahead)
- Sparsity as an ultimate positive feature
- Discretization and algebraic errors in numerical PDEs

5 Subtle points that can not be avoided

5 Subtle points that can not be avoided

- For normal matrices any argumentation using clusters must include position of the clusters.

5 Subtle points that can not be avoided

- For normal matrices any argumentation using clusters must include position of the clusters.
- For non-normal matrices the step from low degree minimal polynomials to small number of distinct eigenvalues requires substantial restrictive assumptions.

5 Subtle points that can not be avoided

- For normal matrices any argumentation using clusters must include position of the clusters.
- For non-normal matrices the step from low degree minimal polynomials to small number of distinct eigenvalues requires substantial restrictive assumptions.
- Perturbation results, backward stability and analysis of the effect of rounding errors should be well thought when using in investigation of iterative computations.

Carson, S (2019)

5 Lanczos, Why Mathematics (1966)

"We will go on pondering and meditating, the great mysteries still ahead of us, we will err and stumble on the way, and if we win a little victory, we will be jubilant and thankful, without claiming, however, that we have done something that can eliminate the contribution of all the millenia before us."

5 Wallis, Arithmetica Infinitorium (1656) (see Khruschev 2008)

"There remains this: we beech the skilled in these things, that we thought worth showing, they will think openly receiving, an whatever it hides, worth imparting more properly by themselves to the wider mathematical community."

Thank you very much for your kind patience!