TDB – NLA

Numerical Linear Algebra

Numerical methods for eigenvalue problems

Maya Neytcheva, TDB, Feb-March 2021

Numerical Linear Algebra

Numerical methods for eigenvalue problems

Basics

Background: Vector and matrix norms

Vector norms

Given $\boldsymbol{v} \in V$, where V is $\mathbb{C}$ or $\mathbb{R}$.

$$\|\boldsymbol{v}\|_2 = \left(\sum_{k=1}^{n} |v_k|^2\right)^{1/2} \quad \text{Euclidean norm}$$

$$\|\boldsymbol{v}\|_1 = \sum_{k=1}^{n} |v_k| \quad \text{the absolute sum norm}$$

$$\|\boldsymbol{v}\|_\infty = \max_k |v_k| \quad \text{the maximum norm}$$

$$\|\boldsymbol{v}\|_p = \left(\sum_{k=1}^{n} |v_k|^p\right)^{1/p}, 1 \leq p \quad \text{Hölder (p) norm}$$

Vector norm equivalence

Theorem: For every pair of norms $\|\boldsymbol{v}\|_a$ and $\|\boldsymbol{v}\|_b$ in $\mathbb{C}$, there exist constants $0 < m > M$ such that

$$m\|\boldsymbol{v}\|_a \leq \|\boldsymbol{v}\|_b \leq M\|\boldsymbol{v}\|_a$$

for any $\boldsymbol{v} \in \mathbb{C}$, where m and M do not depend on $\boldsymbol{v}$.

Matrix norms I

- Induced norms
  For any given vector norm $\|\boldsymbol{v}\|$,

$$\|A\| = \sup_{\boldsymbol{v}\neq\boldsymbol{0}} \frac{\|A\boldsymbol{v}\|}{\|\boldsymbol{v}\|}$$

  is said to be the matrix norm, *induced by the vector norm* $\|\boldsymbol{v}\|$ (or the natural norm).
- If $\|A\boldsymbol{v}\| \leq \|A\|\|\boldsymbol{v}\|$, then the matrix norm $\|A\|$ is compatible with the vector norm $\|\boldsymbol{v}\|$.
- The infinity norm (the max row sum of $|A|$)

$$\|A\|_\infty = \max_k \sum_\ell |A_{k\ell}|$$

Matrix norms II

- The '1'-norm (the max column sum of $|A|$)

$$\|A\|_1 = \max_\ell \sum_k |A_{k\ell}|$$

- The Frobenius norm

$$\|A\|_F = \left(\sum_k \sum_\ell |A_{k\ell}|^2 \right)^{1/2}$$

  Note

$$\|A\|_F = trace(A^*A) = \sum_k (\lambda_k(A^*A))$$

  The Weighted Frobenius norm (W is an spd matrix)

$$\|A\|_W = trace((AW)^*(AW)).$$

Norm relations I

- Let $\rho(A)$ be the spectral radius of A. Then there holds

$$\|A\|_2 = (\rho(A^*A))^{1/2} = (\rho(AA^*))^{1/2}$$

  Note: For Hermitian matrices, $\|A\|_2 = \rho(A)$ but the spectral radius is not a norm.
  Example: $A = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}$, $A \neq 0$ but $\rho(a) = 0$.

Norm relations II

- For an arbitrary square matrix we have $\rho(A) \leq \|A\|$ because $A\boldsymbol{v} = \lambda\boldsymbol{v}$ implies

$$|\lambda|\|\boldsymbol{v}\| = \|A\boldsymbol{v}\| \leq \|A\|\,\|\boldsymbol{v}\|.$$

- For arbitrary square matrix, $\|A\|_2^2 \leq \|A\|_1\|A\|_\infty$. The inequality is sharp.
-

$$\|A\|_2 \leq \|A\|_F = \sqrt{tr(A^*A)} \leq \sqrt{n}\|A\|_2,\ A(n,n).$$

Background: Matrix types

- Normal matrix: $A^*A = AA^*$
  Theorem: A matrix is normal if and only if it is unitary similar to a diagonal matrix.

  Theorem: A normal matrix with real eigenvalues is Hermitian.
- Hermitian matrix: $A = A^*$

  Theorem: The eigenvalues of a Hermitian matrix are real.

  Theorem: Any Hermitian matrix is unitary similar to a real diagonal matrix.

- Reducible/Irreducible matrices:
  Theorem: A matrix A is reducible if and only if there exists a permutation matrix P such that

$$PAP^T = \begin{bmatrix} A_{11} & A_{12} \\ 0 & A_{22} \end{bmatrix},$$

  where A_{11}, A_{22} are square blocks.

  Q: Is the diagonal matrix irreducible?

- Rayleigh quotient:
  Let $(\lambda, \boldsymbol{v})$ be an eigenpair of a matrix A. Then there holds

$$\lambda = \frac{(A\boldsymbol{v}, \boldsymbol{v})}{(\boldsymbol{v}, \boldsymbol{v})}.$$

  For any nonzero vector $\boldsymbol{u} \in \mathbb{C}$ define the so-called Rayleigh quotient

$$\frac{(A\boldsymbol{u}, \boldsymbol{u})}{(\boldsymbol{u}, \boldsymbol{u})}$$

Gershgorin type eigenvalue estimates

Theorem: The spectrum $S(A)$ of a matrix $A = [a_{ij}]$ is enclosed in the union of the disks

$$C_i = \left\{ z \in \mathbb{C} : |z - a_{ii}| \leq \sum_{j \neq i} |a_{ij}|,\ 1 \leq i \leq n \right\}$$

and in the union of the disks

$$C_i' = \left\{ z \in \mathbb{C} : |z - a_{ii}| \leq \sum_{j \neq i} |a_{ji}|,\ 1 \leq i \leq n \right\}$$

That is, $S(A) \in (\cup C_i) \cap (\cup C_i')$. Recall that $S(A) = S(A^T)$.

Min-Max theorem (Courant-Fisher, Poincaré-Weyl)

Theorem: Let A be Hermitian of order n with eigenvalues $\lambda_1 \leq \lambda_2 \leq \cdots \leq \lambda_n$ and $\boldsymbol{v}_1, \boldsymbol{v}_2, \cdots, \boldsymbol{v}_n$ their corresponding eigenvectors. Let $\boldsymbol{p}_1, \cdots, \boldsymbol{p}_n$ be orthonormal vectors, Then, assuming that $\boldsymbol{x}^*\boldsymbol{x} = 1$, we have

$$(a) \quad \min_{\boldsymbol{x} \perp \boldsymbol{p_1} \cdots \boldsymbol{p_{s-1}}} \boldsymbol{x}^* A \boldsymbol{x} \leq \lambda_s$$

$$(b) \quad \max_{\boldsymbol{x} \perp \boldsymbol{p_n}, \boldsymbol{p_{n-1}} \cdots \boldsymbol{p_{s+1}}} \boldsymbol{x}^* A \boldsymbol{x} \geq \lambda_s$$

$$(c) \quad \max_{\substack{\boldsymbol{p}_i,\ 1 \leq i \leq s-1 \\ \boldsymbol{p}_i^* \boldsymbol{p}_j = \delta_{ij}}} \ \min_{\boldsymbol{x} \perp \boldsymbol{p_1} \cdots \boldsymbol{p_{s-1}}} \boldsymbol{x}^* A \boldsymbol{x} = \lambda_s$$

$$(d) \quad \min_{\substack{\boldsymbol{p}_i,\ s+1 \leq i \leq n \\ \boldsymbol{p}_i^* \boldsymbol{p}_j = \delta_{ij}}} \ \max_{\boldsymbol{x} \perp \boldsymbol{p_n}, \boldsymbol{p_{n-1}} \cdots \boldsymbol{p_{s+1}}} \boldsymbol{x}^* A \boldsymbol{x} = \lambda_s$$

and the extreme values are attained for $\boldsymbol{p}_i = \boldsymbol{v}_i, i = 1, 2, \cdots, s-1$ and $\boldsymbol{p}_i = \boldsymbol{v}_i, i = n, n-1, \cdots, s+1$, respectively.

Min-Max theorem (Courant-Fisher, Poincaré-Weyl)

$$\lambda_k \leq \min_{U, dim(U) = n-k+1} \ \max_{\boldsymbol{u} \in U, \boldsymbol{u} \neq \boldsymbol{0}} \frac{(A\boldsymbol{u}, \boldsymbol{u})}{(\boldsymbol{u}, \boldsymbol{u})}$$

or

$$\lambda_k \leq \max_{U, dim(U) = k} \ \min_{\boldsymbol{u} \in U, \boldsymbol{u} \neq \boldsymbol{0}} \frac{(A\boldsymbol{u}, \boldsymbol{u})}{(\boldsymbol{u}, \boldsymbol{u})}$$

Corollaries

Let A, B be symmetric matrices and consider the eigenvalues of $A, A + B, AB$ ordered increasingly. Then

1. $\lambda_i(A) + \lambda_{min}(B) \leq \lambda(A + B) \leq \lambda_i(A) + \lambda_{max}(B)$
2. If A positive definite and $\lambda_{max}(B)$ in nonnegative, then
$$\lambda_i(AB) \leq \lambda_i(A)\lambda_{max}(B)$$
3. If A positive definite and $\lambda_{min}(B)$ in nonnegative, then
$$\lambda_i(AB) \geq \lambda_i(A)\lambda_{min}(B)$$

Proof: Let $\boldsymbol{v}_i$ be the eigenvectors of A, orthonormal; $\boldsymbol{x}$, $\|\boldsymbol{x}\| = 1$, arbitrary. Then

$$\begin{aligned} \lambda_i(A+B) &\geq \min_{\boldsymbol{x} \perp \boldsymbol{v_1}, \cdots, \boldsymbol{v_{i-1}}} \boldsymbol{x}^T (A+B) \boldsymbol{x} \\ &\geq \min_{\boldsymbol{x} \perp \boldsymbol{v_1}, \cdots, \boldsymbol{v_{i-1}}} \boldsymbol{x}^T A \boldsymbol{x} + \min_{\boldsymbol{x}^T \boldsymbol{x} = 1} \boldsymbol{x}^T B \boldsymbol{x} \\ &= \lambda_i(A) + \lambda_{min}(B). \end{aligned}$$

Cauchy interlace theorem

Theorem: Let $B = \begin{bmatrix} A & \boldsymbol{w} \\ \boldsymbol{w}^* & \alpha \end{bmatrix}$, $A(n-1, n-1), B(n, n)$, Hermitian.
The eigenvalues of A and B interlace,

$$\lambda_1(B) \leq \lambda_1(A) \leq \lambda_2(B) \leq \lambda_2(A) \leq \cdots \leq \lambda_{n-1}(B) \leq \lambda_{n-1}(A) \leq \lambda_n(B).$$

Reformulated: the characteristic polynomial of a Hermitian matrix is interlaced by the characteristic polynomial of any principle submatrix.

Background: Canonical forms of matrices

Definitions

1. **Def.:** An eigenvalue λ of A is said to have algebraic multiplicity m if it is a root of multiplicity m of the characteristic polynomial of A.

Definitions

1. **Def.:** An eigenvalue λ of A is said to have algebraic multiplicity m if it is a root of multiplicity m of the characteristic polynomial of A.
2. **Def.:** An eigenvalue λ of A of algebraic multiplicity 1 is said to be simple.

Definitions

1. **Def.:** An eigenvalue λ of A is said to have algebraic multiplicity m if it is a root of multiplicity m of the characteristic polynomial of A.
2. **Def.:** An eigenvalue λ of A of algebraic multiplicity 1 is said to be simple.
3. **Def.:** An eigenvalue λ of A is said to have geometric multiplicity μ if the maximum number of independent eigenvectors, associated with it is μ. Thus, the geometric multiplicity i s the dimension of the eigenspace $Null(A - \lambda I)$.

Definitions

1. **Def.:** An eigenvalue λ of A is said to have algebraic multiplicity m if it is a root of multiplicity m of the characteristic polynomial of A.
2. **Def.:** An eigenvalue λ of A of algebraic multiplicity 1 is said to be simple.
3. **Def.:** An eigenvalue λ of A is said to have geometric multiplicity μ if the maximum number of independent eigenvectors, associated with it is μ. Thus, the geometric multiplicity i s the dimension of the eigenspace $Null(A - \lambda I)$.
4. **Def.:** Semi-simple eigenvalue, if $m = \mu$.

Definitions

Def.: The matrices A and B are said to be similar if there is a nonsingular matrix X, such that

$$A = XBX^{-1}$$

(similarity transformation).

A and B have the same spectrum including the algebraic multiplicity.

Jordan canonical form

Theorem: For any square matrix A there exists a nonsingular matrix X that reduces A to a block-diagonal form $XAX^{-1} = diag(J_1, J_2, \cdots, J_p)$, where J_k of order n_k is either $J_k = \lambda_k$ or

$$J_k = \begin{bmatrix} \lambda_k & 1 & & & \\ & \lambda_k & 1 & & \\ & & \ddots & & \\ & & & \lambda_k & 1 \\ 0 & & & & \lambda_k \end{bmatrix}, \text{ if } n_k \geq 2.$$

A is diagonalizable iff $n_k = 1, k = 1, 2, \cdots, p$.

Schur canonical form

Theorem: Any square matrix A is unitary similar to an upper-triangular matrix, where the diagonal elements are the eigenvalues of A:

$$U^{-1}AU = \begin{bmatrix} \lambda_1 & b_{12} & \cdots & b_{1n} \\ 0 & \lambda_2 & \cdots & b_{2n} \\ & & \ddots & \\ 0 & \cdots & 0 & \lambda_n \end{bmatrix}$$

Eigenvalue-revealing !

Schur canonical form I

$$A = URU^* \quad \text{equivalently} \quad AU = UR$$

Note 1: The first column of U $\boldsymbol{u}_1 = U\boldsymbol{e}_1$ is an eigenvector of A with eigenvalue $\lambda_1 = \boldsymbol{e}_1^T R \boldsymbol{e}_1$.
Note 2: The last column of U $\boldsymbol{u}_n = U\boldsymbol{e}_n$ is an eigenvector of A with eigenvalue $\lambda_n = \boldsymbol{e}_n^T R \boldsymbol{e}_n$.
Note 3: The second column $\boldsymbol{u}_2$ of U is an eigenvalue of $A' = (I - \boldsymbol{u}_1\boldsymbol{u}_1^*)A(I - \boldsymbol{u}_1\boldsymbol{u}_1^*)$ with eigenvalue $\lambda_2 = \boldsymbol{e}_2^T R \boldsymbol{e}_2$.
Proof: R is upper-triangular. $R\boldsymbol{e}_2 = R_{12}\boldsymbol{e}_1 + \lambda_2\boldsymbol{e}_2$. Proof is completed by expanding $A'\boldsymbol{u}_e$,
In A' the eigenvalue λ_1 is deflated.

Corollaries

1. Every selfadjoint matrix is unitary equivalent to a real diagonal matrix.

Corollaries

1. Every selfadjoint matrix is unitary equivalent to a real diagonal matrix.
2. Every real symmetric matrix is orthogonally equivalent to a real diagonal matrix.

Corollaries

1. Every selfadjoint matrix is unitary equivalent to a real diagonal matrix.
2. Every real symmetric matrix is orthogonally equivalent to a real diagonal matrix.
3. A matrix of order n is similar to a diagonal matrix if and only if its eigenvalues form a basis for $\mathbb{C}^n$, i.e., the eigenvector space of A is complete.

To summarize:

A is diagonalizable if there exists a diagonal matrix D and a nonsingular matrix X, such that $A = XDX^{-1}$.
In other words, if and only if there exists a basis of the whole space consisting of eigenvectors of A only. Matrices that cannot be diagonalized are called defective.

1. A is real symmetric $\Longrightarrow$ A is Hermitian $\Longrightarrow$ A is normal $\Longrightarrow$ A is diagonalizable,
2. A is defective $\Longrightarrow$ A is non-normal $\Longrightarrow$ A is non-Hermitian.

None of the above implications are an equivalence.

Singular Value decomposition

Let $A \in \mathbb{C}^{n\times k}$. Then A can be factorized as

$$A = U\Sigma V^T \quad \text{or} \quad AV = U\Sigma,$$

where U, V are unitary and Σ is diagonal with nonnegative entries. Rank-revealing factorization!

(Considered separately.)

Proposition: Let A be normal. Then for each $\varepsilon > 0$ there exists a non-normal matrix A_ε such that

$$\|A - A_\varepsilon\| \leq \varepsilon.$$

The set of non-normal matrices is dense in the set of all matrices.
Proof: (Note, that since all norms on finite dimensional spaces are equivalent, we need not specify the topology implied by the word 'dense'.) Let A be normal and $AU = UD$ with $U^*U = I$ and D is diagonal. Define matrices A_ε as follows

$$A_\varepsilon = U_\varepsilon^* \Lambda U_\varepsilon, \text{ where } U_\varepsilon = U \begin{bmatrix} cos(\varepsilon) & 0 & 0 \\ sin(\varepsilon) & 1 & 0 \\ 0 & 0 & I \end{bmatrix}.$$

Then $U_0 = U$ and $A_0 = A$, and for all $\varepsilon \in (0, \pi)$, we have that U_ε is non-unitary. But the above is a diagonalization of A_ε. Therefore, $\{A_\varepsilon\}$ with $\varepsilon \to 0$ is a sequence of non-normal matrices converging to A.

Proposition: Let A be defective. Then for each $\varepsilon > 0$ there exists a diagonalizable matrix A_ε such that

$$\|A - A_\varepsilon\| \leq \varepsilon.$$

The set of diagonalizable matrices is dense in the set of all matrices.

Gram-Schmidt orthogonalization

1: Compute $r_{11} = \|\boldsymbol{v}_1\|_2$. If $r_{11} = 0$, stop, else $\boldsymbol{q}_1 = \boldsymbol{v}_1 / r_{11}$
2: **for** $j = 2, \cdots n$ **do**
3: **for** $i = 1, 2, \cdots, j-1$ **do**
4: $r_{ij} = (\boldsymbol{v}_j, \boldsymbol{q}_i)$
5: **end for**
6: $\widehat{\boldsymbol{q}} = \boldsymbol{v}_j - \sum_{i=1}^{j-1} r_{ij} \boldsymbol{q}_i$
7: $r_{jj} = \|\widehat{\boldsymbol{q}}\|_2$
8: If $r_{jj} = 0$ then stop, else $\boldsymbol{q}_j = \widehat{\boldsymbol{q}} / r_{jj}$
9: **end for**

Modified Gram-Schmidt orthogonalization

1: Compute $r_{11} = \|\boldsymbol{v}_1\|_2$. If $r_{11} = 0$, stop, else $\boldsymbol{q}_1 = \boldsymbol{v}_1 / r_{11}$
2: **for** $j = 2, \cdots n$ **do**
3: Set $\widehat{\boldsymbol{q}} = \boldsymbol{v}_j$
4: **for** $i = 1, 2, \cdots, j-1$ **do**
5: $r_{ij} = (\widehat{\boldsymbol{q}}, \boldsymbol{q}_i)$
6: $\widehat{\boldsymbol{q}} = \widehat{\boldsymbol{q}} - r_{ij} \boldsymbol{q}_i$
7: **end for**
8: $r_{jj} = \|\widehat{\boldsymbol{q}}\|_2$
9: If $r_{jj} = 0$ then stop, else $\boldsymbol{q}_j = \widehat{\boldsymbol{q}} / r_{jj}$
10: **end for**

Numerical Linear Algebra

Perturbation theory and error analysis

Formulations of the eigenvalue problem: I

Given a square matrix A of size n.

- Find all complex numbers λ for which there exists a nonzero solution $\boldsymbol{v} \in \mathbb{C}$ such that $A\boldsymbol{v} = \boldsymbol{v}\lambda$.
  (λ is seen as a matrix of size 1.)
- Find λ, such that $A - \lambda I$ vanishes on a nontrivial subset $V_\lambda \subset \mathbb{C}^n$. Note that $dim(V_\lambda) \geq 1$.

Formulations of the eigenvalue problem: II

- The eigenvalue problem can be seen as an eigenvector problem:

  $$A\boldsymbol{v} = \boldsymbol{v}\lambda \Rightarrow \boldsymbol{v}^* A\boldsymbol{v} = \boldsymbol{v}^*\boldsymbol{v}\lambda \Rightarrow \boldsymbol{v}\boldsymbol{v}^* A\boldsymbol{v} = \boldsymbol{v}\boldsymbol{v}^*\boldsymbol{v}\lambda$$
  $$\boldsymbol{v}\boldsymbol{v}^* A\boldsymbol{v} = \boldsymbol{v}\boldsymbol{v}^*\boldsymbol{v}\lambda \Rightarrow \boldsymbol{v}\boldsymbol{v}^* A\boldsymbol{v} = \boldsymbol{v}^*\boldsymbol{v} A\boldsymbol{v} \Rightarrow A\boldsymbol{v}\boldsymbol{v}^*\boldsymbol{v} = \boldsymbol{v}\boldsymbol{v}^* A\boldsymbol{v}.$$

  Thus, the eigenvectors are the nonzero roots of the nonlinear functional $f : \mathbb{C}^n \to \mathbb{C}^n : A\boldsymbol{v}\boldsymbol{v}^*\boldsymbol{v} - \boldsymbol{v}\boldsymbol{v}^* A\boldsymbol{v} \equiv (I - \boldsymbol{v}\boldsymbol{v}^*)A\boldsymbol{v}$ (if $\|\boldsymbol{v}\| = 1$).

Despite of the fact that the eigenproblems are classified as Linear Algebra problems, they are highly nonlinear - f is a quadratic functional and the problem could be solved using the Newton method, for instance.

Perturbation theory

Motivation to analyse the effect of perturbations I

$$A\boldsymbol{v} = \boldsymbol{v}\lambda$$

As we compute the eigensolutions inexactly, up to a machine precision, we are interested to see what happens with the eigenvalues and the eigenvectors when we slightly perturb the matrix.

Def.: Let A be of size n, μ - scalar and $\boldsymbol{v}$ - a vector in $\mathbb{C}^n$. The eigenvalue residual for the pair $(\mu, \boldsymbol{v})$ is defined as

$$\boldsymbol{r} = A\boldsymbol{v} - \boldsymbol{v}\mu.$$

Theorem: Under the above setting, $(\mu, \boldsymbol{v})$ is an eigenpair of a matrix $\widehat{A} = A + E$, where $E = -\boldsymbol{v}\boldsymbol{r}^*$.

Proof:

$\widehat{A}\boldsymbol{v} = (A + E)\boldsymbol{v} = (A - \boldsymbol{v}\boldsymbol{r}^*)\boldsymbol{v} = A\boldsymbol{v} - \boldsymbol{r} = A\boldsymbol{v} - A\boldsymbol{v} + \boldsymbol{v}\mu = \boldsymbol{v}\mu.$

Motivation to analyse the effect of perturbations II

Note: $\|E\| = \|\boldsymbol{r}\|$.
Theorem: There exists no matrix F with $\|F\| < \|\boldsymbol{r}\|$ such that $(A+F)\boldsymbol{v} = \boldsymbol{v}\mu$. **Proof:** From $(A+F)\boldsymbol{v} = \boldsymbol{v}\mu$ we have

$$\|F\boldsymbol{v}\| = \|A\boldsymbol{v} - \mu\boldsymbol{v}\| = \|\boldsymbol{r}\|, \text{ hence, } \|F\| = \sup_{\boldsymbol{w}\neq\boldsymbol{0}} \frac{\|F\boldsymbol{w}\|}{\|\boldsymbol{w}\|} \geq \frac{\|F\boldsymbol{v}\|}{\|\boldsymbol{v}\|} = \|\boldsymbol{r}\|.$$

In other words, $E = -\boldsymbol{r}\boldsymbol{v}^*$ is the smallest perturbation of A having $(\mu, \boldsymbol{v})$ as eigenpair.

Regarding computational algorithms: they should produce small residuals.
If $\|\boldsymbol{r}\|$ is small, then it depends on the properties of A whether or not this will result in accurate eigenvalues.

Classical perturbation bounds I

Theorem [Bauer-Fike]: Suppose that $AV = V\Lambda$ with V nonsingular and Λ diagonal. Let μ be an eigenvalue of $A+E$. Then, there exists an eigenvalue λ of A such that

$$|\lambda - \mu| \leq \varkappa[V]\|E\|, \text{ where } \varkappa(V) = \|V\|\,\|V^{-1}\|.$$

Corollary: Let A be a normal matrix, μ be an eigenvalue of $A+E$ with $\|E\| < \varepsilon$. Then there exists an eigenvalue of A, λ such that

$$|\lambda - \mu| \leq \varepsilon.$$

(In this case $\|V\| = 1$.)

Classical perturbation bounds II

Theorem [Henrici]: Let $AQ = QR$ be a Schur decomposition of A. Let $R = D + N$, where D is diagonal and N is strictly upper-triangular (nilpotent). Let p be the smallest integer such that $N^p = 0$. Then for each eigenvalue μ of $A+E$ there exists an eigenvalue λ of A such that

$$|\lambda - \mu| \leq \max(\theta, \theta^{1/p}), \text{where } \theta = \|E\| \sum_{k=0}^{p-1} \|N\|^k.$$

The quantity $\nu(A) = \|N\|$ is referred to as departure from normality of A.

Classical perturbation bounds III

Moral: When A is not normal or close to defective, the eigenvalues of $A+E$ may differ significantly from those of A even if $\|E\|$ is small.

Pseudoeigenvalues I

Given A ,$z \in \mathbb{C}$, the eigenvalues of A are those z for which $A - zI$ is singular.
Consider subsets of the complex plane that are close to these eigenvalues.
Def.: The **ε-pseudospectrum** $\Lambda_\varepsilon(A)$ of A is the set of points z such that z is an eigenvalue of $A + E$ with $\|E\| \leq \varepsilon$.
In this way the pseudospectra define the exact bounds for how far and eigenvalue of A can move under a perturbation E.
Def.: **Resolvent** For any $z \in \mathbb{C}$

$R(A, z) = (A - zI)^{-1}$ $(Y.Saad)$ $R(A, z) = (zI - A)^{-1}$ $(N.Trefethen)$

Pseudoeigenvalues II

Equivalent definitions for the **ε-pseudospectrum** of A, $\Lambda_\varepsilon(A)$:
Def.: $\Lambda_\varepsilon(A) = \{z \in \mathbb{C} : z \in \Lambda(A + E) \text{ for some } E : \|E\| \leq \varepsilon\}$.
Def.:

$$\Lambda_\varepsilon(A) = \left\{z \in \mathbb{C} : \|(zI - A)^{-1}\| \geq \varepsilon^{-1}\right\}.$$

$\Lambda_\varepsilon(A)$ is the subset of the complex plane, bounded by the ε^{-1} level curve of the resolvent norm.
Def.: $\Lambda_\varepsilon(A) = \{z \in \mathbb{C} : z \in \Lambda(A + E) \text{ for some } E : \|E\| \leq \varepsilon\}$
Def.: $\Lambda_\varepsilon(A) = \{z \in \mathbb{C} : \sigma_{min}(zI - A) \leq \varepsilon\}$, where $\sigma_{min}(A)$ denotes the minimum singular value of a generic matrix A.
$\Lambda_\varepsilon(A)$ are the sets in the complex plane bounded by the level curves of $\sigma_{min}(zI - A)$.

The resolvent I

The spectrum of A is defined as the set in $\mathbb{C}$ where the inverse of $A - zI$ does not exist.
$R(A, z)$ admits singularities at the eigenvalues of A. Away, the resolvent is analytic w.r.t z:

$$\begin{aligned} R(A, z) \equiv (A - zI)^{-1} &= ((A - z_0 I) - (z - z_0)I)^{-1} \\ &= R(z_0)(I - (z - z_0)R(z_0))^{-1} \end{aligned}$$

... many properties of $R(A, z)$ ...

The resolvent I

Proposition: $\Lambda_\varepsilon(A)$ is the set of points $z \in \mathbb{C}$, for which

- $\sigma_{min} \leq \varepsilon$, with σ_{min} the minimal singular value of $A - zI$,
- $\|(A - zI)^{-1}\| \geq \varepsilon^{-1}$

Theorem: Let $(\mu, \boldsymbol{v})$ be an approximate eigenpair of A and assume $\|\boldsymbol{v}\| = 1$. Then $\mu \in \Lambda_\varepsilon(A)$ with $\varepsilon = \|\boldsymbol{r}\|$, $\boldsymbol{r} = A\boldsymbol{v} - \boldsymbol{v}\mu$.

The resolvent II

Theorem: Let $(\mu, \boldsymbol{v})$ be an approximate eigenpair of A and assume $\|\boldsymbol{v}\| = 1$. Let θ be the positive acute angle between $\boldsymbol{r}$ and $\boldsymbol{v}$. Let $\widehat{\mu}$ be such that $A\boldsymbol{v} - \boldsymbol{v}\widehat{\mu} \perp \boldsymbol{v}$. Then

$$\widehat{\mu} = \boldsymbol{v}^* A \boldsymbol{v} = \Lambda_{\sin(\theta)\|\boldsymbol{r}\|}(A).$$

The above result shows a possibility to improve the eigenvalue once an approximate eigenvector has been computed.

Example (Trefethen)

Consider the time-reduced 1D Schrödinger operator with a potential function that is complex:

$$Au(x) = u'' + (cx^2 - dx^4), \; c = 3 + 3i, d = 1/16.$$

The operator is highly non-normal, invariant with respect to x or $-x$, thus if $u(x)$ is an eigenfinction, then $u(-x)$ is also an eigenfunction.
Shown: all eigenvalues of A are single and each eigenfunction is either even or odd (B. Davies, 1998).
Discretize: Chebyshev collocation spectral method on a finite interval $[-L, L]$.

Example (Trefethen)

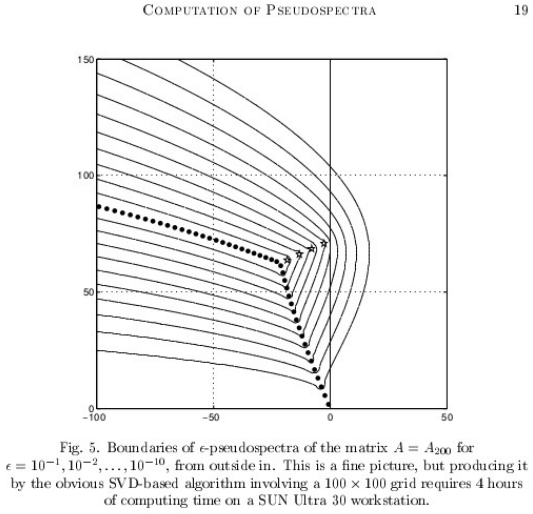

Fig. 5. Boundaries of ϵ-pseudospectra of the matrix $A = A_{200}$ for $\epsilon = 10^{-1}, 10^{-2}, \ldots, 10^{-10}$, from outside in. This is a fine picture, but producing it by the obvious SVD-based algorithm involving a 100×100 grid requires 4 hours of computing time on a SUN Ultra 30 workstation.

Numerical Linear Algebra

Numerical methods for small/medium eigenvalue problems

Task

Given A, B, C, square of order n. Find pairs $\lambda, \boldsymbol{v}$) such that

$A\boldsymbol{v} = \lambda\boldsymbol{v}$	standard eigenvalue problem
$A\boldsymbol{v} = \lambda B\boldsymbol{v}$	generalized eigenvalue problem
$A\boldsymbol{v} + \lambda B\boldsymbol{v} + \lambda^2 C\boldsymbol{v} = \boldsymbol{0}$	quadratic eigenvalue problem

Usually n is large ($10^5 - 10^8$).
For smaller n ($n \leq 5000$) all eigenpairs are computed using the QR iteration (`Matlab eig`).

The Power iteration

The Power method (Single vector iteration technique) - constructs a sequence of vectors that converges to the dominant eigenvector, the eigenvector, corresponding to the eigenvalue with largest modulus.

1: Choose $\boldsymbol{v}_0 \neq \boldsymbol{0}$
2: **for** $k = 1, \cdots$ until convergence **do**
3: Find $\alpha_k = \max_{1 \leq k \leq n}(|\boldsymbol{v}_k(i)|)$
4: $\boldsymbol{v}_k = \frac{1}{\alpha_k} A\boldsymbol{v}_{k-1}$
5: **end for**

Theorem: Let λ_i be the eigenvalues of A ordered as $|\lambda_1| \geq \lambda_2| \geq \cdots \geq |\lambda_n|$ The convergence factor is proportional to the relative gap between $|\lambda_1|$ and $|\lambda_2|$, $\frac{|\lambda_2|}{|\lambda_1|}$.

The Power iteration, another form

1: Choose $\boldsymbol{v}_0 \neq \boldsymbol{0}$
2: **for** $k = 1, \cdots$ until convergence **do**
3: $\widetilde{\boldsymbol{v}}_k = \frac{1}{\alpha_k} A\boldsymbol{v}_{k-1}$
4: $\alpha_k = \|\widetilde{\boldsymbol{v}}_k\|$
5: $\boldsymbol{v}_k = \frac{1}{\alpha_k}\widetilde{\boldsymbol{v}}_k$
6: **end for**

The Power iteration, yet another form (Rayleigh quotient)

Previous form:

1: Choose $\boldsymbol{v}_0 \neq \boldsymbol{0}$
2: **for** $k = 1, \cdots$ until convergence **do**
3: $\widetilde{\boldsymbol{v}}_k = \frac{1}{\alpha_k} A\boldsymbol{v}_{k-1}$
4: $\alpha_k = \|\widetilde{\boldsymbol{v}}_k\|$
5: $\boldsymbol{v}_k = \frac{1}{\alpha_k}\widetilde{\boldsymbol{v}}_k$
6: **end for**

Alternative form:

1: Choose $\boldsymbol{v}_0 \neq \boldsymbol{0}$
2: **for** $k = 1, \cdots$ until convergence **do**
3: $\widetilde{\boldsymbol{v}}_k = \frac{1}{\alpha_k} A\boldsymbol{v}_{k-1}$
4: $\boldsymbol{v}_k = \frac{\widetilde{\boldsymbol{v}}_k}{\|\widetilde{\boldsymbol{v}}_k\|_2}$
5: $\alpha_k = \boldsymbol{v}^* A\boldsymbol{v}_k$
6: **end for**

The Power iteration, stopping test

- fixed number of iterations
- $\|\boldsymbol{v}_{k+1} - \boldsymbol{v}\| \leq \tau$
- Check on the change in the eigenvalue

The Inverse Power iteration

1: Choose $\boldsymbol{v}_0 \neq \boldsymbol{0}$
2: **for** $k = 1, \cdots$ until convergence **do**
3: $\quad A\widetilde{\boldsymbol{v}}_k = \boldsymbol{v}_{k-1}$
4: $\quad \alpha_k = \frac{1}{\|\widetilde{\boldsymbol{v}}_k\|}$
5: $\quad \boldsymbol{v}_k = \alpha_k \widetilde{\boldsymbol{v}}_k$
6: **end for**

Converges to the first (smallest) eigenvalue and its eigenvector.

The Shifted Inverse Power iteration

1: Choose $\boldsymbol{v}_0 \neq \boldsymbol{0}$
2: **for** $k = 1, \cdots$ until convergence **do**
3: $\quad (A - \sigma I)\widetilde{\boldsymbol{v}}_k = \boldsymbol{v}_{k-1}$
4: $\quad \alpha_k = \frac{1}{\|\widetilde{\boldsymbol{v}}_k\|}$
5: $\quad \boldsymbol{v}_k = \nu_k \widetilde{\boldsymbol{v}}_k$
6: $\quad \alpha_k = \sigma + \nu_k$
7: **end for**

To increase the convergence one can use variable shifts:
$(A - \sigma_{k-1} I)\widetilde{\boldsymbol{v}}_k = \boldsymbol{v}_{k-1}$.

The Shifted Power iteration I

1: Choose $\boldsymbol{v}_0 \neq \boldsymbol{0}$
2: **for** $k = 1, \cdots$ until convergence **do**
3: $\quad \widetilde{\boldsymbol{v}}_k = (A - \sigma I)\boldsymbol{v}_{k-1}$
4: $\quad \boldsymbol{v}_k = \nu_k \widetilde{\boldsymbol{v}}_k$
5: $\quad \alpha_k = \frac{\boldsymbol{v}_k^* A \boldsymbol{v}_k}{\boldsymbol{v}_k^* \boldsymbol{v}_k}$
6: **end for**

Note: The same eigenvectors, better eigenvalue distribution.
Based of the fact that $f(A)\boldsymbol{v}_k = f(\lambda_k)\boldsymbol{v}_k$ we can do shifts in various ways:
- $f(A) = I - \sigma A; \quad f(A) = A - \sigma(A)$
-
$f(A) = I + \gamma_1 A + \gamma_2 A^2 + \cdots + \gamma_\ell A^\ell = (I - \beta_1 A)(I - \beta_2 A) \cdots (I - \beta_\ell A)$

The Shifted Power iteration II

- $f(A) = (A - \sigma I)^{-1}$
- $f(A) = (A - I)^{-1}(I - A)$ (Cayley transform) or $f(A) = (A - \sigma I)^{-1}(I - \sigma A)$ combined with a shift.

We aim to achieve:
- faster convergence to an eigenvalue close to the shift,
- diminish unwanted components,
- improve eigenvalue distribution,
- amplify wanted components

Demo: `/home/maya/matlab/Eigenvalues/Power_method/Main_Eig_small.m`

The QR iteration

QR factorization , $A(n, n)$

$$A = QR,$$

Q - orthogonal, R - upper-triangular.

Method	Complexity
Gram-Schmidt	n^3
Modified GS	n^3
Householder reflections	n^3
Givens rotations	n^3

QR decomposition of an upper-Hessenberg matrix

$$H = \begin{bmatrix} * & * & * & \cdots & * \\ * & * & * & \cdots & * \\ & * & * & \cdots & * \\ & & \ddots & \ddots & * \\ & & & * & * \end{bmatrix}$$

The QR factorization of H can be done efficiently (in $O(n^2)$ flops).

$$H \equiv H_1 = \begin{bmatrix} a_1 & * & * & * & * & * \\ b_1 & * & * & * & * & * \\ 0 & * & * & * & * & * \\ 0 & 0 & * & * & * & * \\ 0 & 0 & 0 & * & * & * \\ 0 & 0 & 0 & 0 & * & * \end{bmatrix}$$

Construct the Givens rotations ($r_1 = \sqrt{a_1^2 + b_1^2}$)

$$Q_1 = \begin{bmatrix} a_1/r_1 & b_1/r_1 & 0 & 0 & 0 & \\ -b_1/r_1 & a_1/r_1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}$$

$$Q_1 H_1 = \begin{bmatrix} * & * & * & * & * & * \\ 0 & * & * & * & * & * \\ 0 & * & * & * & * & * \\ 0 & 0 & * & * & * & * \\ 0 & 0 & 0 & * & * & * \\ 0 & 0 & 0 & 0 & * & * \end{bmatrix}$$

Repeat I

$$H_2 \equiv Q_1 H_1 = \begin{bmatrix} * & * & * & * & * & * \\ 0 & a_2 & * & * & * & * \\ 0 & b_2 & * & * & * & * \\ 0 & 0 & * & * & * & * \\ 0 & 0 & 0 & * & * & * \\ 0 & 0 & 0 & 0 & * & * \end{bmatrix}$$

Construct the Givens rotations ($r_2 = \sqrt{a_2^2 + b_2^2}$)

Repeat II

$$Q_2 = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & a_2/r_2 & b_2/r_2 & 0 & 0 & \\ 0 & -b_2/r_2 & a_2/r_2 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}$$

$$H_3 \equiv Q_2 H_2 = Q_2 Q_1 H_1 = \begin{bmatrix} * & * & * & * & * & * \\ 0 & * & * & * & * & * \\ 0 & 0 & * & * & * & * \\ 0 & 0 & * & * & * & * \\ 0 & 0 & 0 & * & * & * \\ 0 & 0 & 0 & 0 & * & * \end{bmatrix}$$

QR factorization of H

1: **for** $j = 1, \cdots n-1$ **do**
2: Let QR be the QR factorization of the $2 \times n$ matrix $H(j : j+1, :)$
3: Store the 2×2 matrix $Q_j = Q$ and overwrite $H(j : j+1, :) = R(j : j+1, :)$
4: **end for**

Each Q_j is a plane rotation with $\sin(\theta_j) = a_j/r_j$ and $\cos(\theta_j) = b_j/r_j$.

QR factorization of H

If H is Hessenberg and $H = QR$, then
- Q is Hessenberg and
- RQ is also Hessenberg.

The basic QR iteration

Recall that Schur decomposition is eigenvalue-revealing.
Idea: Compute approximations of eigenvalues by fixed point iteration on the Schur factorization $AQ = QR$.

$$Q_0 = I, \ \textit{iterate}\ AQ_n = Q_{n+1}R_{n+1},$$

where $Q_{n+1}R_{n+1}$ is a QR factorization of AQ_n.
Note: Even if A is Hessenberg and Q is Hessenberg, AQ is not!

The basic QR iteration

Basic QR iteration:

Start: $A = Q_1 R_1$, iterate $Q_{k+1} R_{k+1} = R_k Q_k$

Note: RQ is Hessenberg!

Less clear! However,

$$\begin{array}{lll} A = Q_1 R_1, & R_1 = Q_1^* A, & R_1 Q_1 = Q_1^* A Q_1 \\ Q_2 R_2 = Q_1^* A Q_1, & R_2 = Q_2^* Q_1^* A Q_1, & R_2 Q_2 = Q_2^* Q_1^* A Q_1 Q_2 \\ \ldots & & \end{array}$$

$$A Q_1 Q_2 \cdots Q_n = Q_1 Q_2 \cdots Q_{n+1} R_{n+1}$$

Basic QR iteration

1: Input $A \in \mathbb{C}^{n \times n}$
2: Output $R = A_n$, $U = U_n$ such that $A = URU^*$
3: Set $A_0 = A$, $U_0 = I$
4: **for** $k = 1, \cdots\cdots$ **do**
5: Compute the QR factorization $A_{k-1} = Q_k R_k$
6: Set $A_k = R_k Q_k$
7: Set $U_k = U_{k-1} Q_k$
8: **end for**

Demo: `/home/maya/matlab/Eigenvalues/QR_iteration/Basic_QR_iter.m`

Basic QR iteration - convergence

Theorem: Let A be real invertible $n \times n$ matrix, such that the moduli of the eigenvalues $\lambda_1, \cdots, \lambda_n$ are distinkt,

$$|\lambda_1| > |\lambda_2| > \cdots |\lambda_n| >> 0.$$

Let $A = Q^{-1} D Q$, where $D = diag(\lambda_1, \lambda_2, \cdots, \lambda_n)$. Suppose $Q = LU$, where L is unit lower-triangular. and U is upper-triangular.

- The strictly lower triangular part of A_k converges to zero in $O(t^k)$.
- The diagonal part of A converges to D in $O(t^k)$, where

$$t = \max\left\{\left|\frac{\lambda_2}{\lambda_1}\right|, \cdots, \left|\frac{\lambda_n}{\lambda_{n-1}}\right|\right\} < 1.$$

Note: The upper triangular entries may fail to converge (to the Schur form).

Basic QR iteration with shifts (to accelerate convergence)

1: Input $A \in \mathbb{C}^{n \times n}$
2: Output $R = A_n$, $U = U_n$ such that $AU = UR$
3: Set $A_0 = A$, $U_0 = I$
4: **for** $k = 1, \cdots\cdots$ **do**
5: Select a shift σ_k
6: Compute the QR factorization $A_{k-1} - \sigma_k I = Q_k R_k$
7: Set $A_k = R_k Q_k + \sigma_k I$
8: Set $U_k = U_{k-1} Q_k$
9: **end for**

Computational aspects

If A is upper-Hessenberg, so is Q_n. But AQ_n is NOT upper-Hessenberg. Thus, computationally it is always advisable to convert AQ_n to upper-Hessenberg or tridiagonal in the symmetric case.

Recall, that the Arnoldi factorization leads to $AV = VH$, where H is upper-Hessenberg.

Single vector iterations
Power method - relation to QR iteration

Convergence of the QR iteration

The convergence depends on the shift.

Example of a reference: Tai-Lin Wang, William B. Gragg Convergence of the shifted QR algorithm for unitary Hessenberg matrices Mathematics of Computation, 71, 1473-1496

Shifts:

- The Rayleigh shift: $\sigma = a_{nn}$, a_{nn} - the last diagonal element of A_k.
- the Wilkinson shift: σ is taken as the absolute smallest eigenvalue of the trailing 2×2 submatrix of A_k, closer to a_{nn}; $\begin{bmatrix} a_{n-1,n-1} & a_{n-1,n} \\ \beta n - 1 & a_{nn} \end{bmatrix}$;
- Mixed shift: R-shift if $\theta\beta_{n-2} \geq \beta_{n-1}$ and W-shift otherwise.

W-shift - the convergence is quadratic for $A \neq A^*$ and cubic for Hermitian A.

Partial QR iteration I

Theorem: Consider the basic QR iteration. The first j columns of Q_{k+1} can be computed from partial iteration
Start with $Q_0^{(j)} = I(:, 1:j)$, iterate $AQ_k^{(j)} = Q_{k+1}^{(j)} R_{k+1}^{(j)}$.
Then $R_{k+1}^{(j)}$ is the top-left part of R_{k+1}.
Question: If the partial QR iteration converges, to what it converges? We will have

$$AQ^{(j)} = Q^{(j)} R^{(j)}$$

with $Q^{(j)}(n \times j)$ orthogonal and $R^{(j)}(j \times j)$ - upper-triangular. The subspace V spanned by the columns of $Q^{(j)}$ is invariant subspace for A and the eigenvalues of $R^{(j)}$ must e eigenvalues of A, but which?

Partial QR iteration II

Take now $j = 1$:
Start with $\boldsymbol{v}_0$, iterate $A\boldsymbol{v}_k = \boldsymbol{v}_{k+1}\boldsymbol{r}_{k+1}$.
We know that it converges to the maximum eigenvalue of A.

Real Schur decomposition

If A is real, non-Hermitian with complex eigenvalues, it is clear that the QR iteration cannot converge to a Schur decomposition.
As long as real shifts are used, all matrices involved will be real!
In that case convergence takes place to the so-called Real Schur Decomposition.
Theorem: Any real matrix A of size n can be factorized as $A = QRQ^*$, where Q is real orthogonal and R is real quasi upper-triangular. The latter denotes that $R = R_0 + D$, where R_0 is upper-triangular and D is block-diagonal with blocks of size at most 2.
Remark: The eigenvalues of the 2×2 quasi upper-triangular matrix are exactly the complex eigenpairs of A.
$A(y|z) = (y|z)\begin{bmatrix} \mu & \nu \\ -\nu & \mu \end{bmatrix}$, $\lambda = \mu \pm i\nu$.

Numerical Linear Algebra

Numerical methods for eigenvalue problems
Subspace methods

General framework – projection methods

Want to solve $\boxed{\boldsymbol{b} - A\boldsymbol{x} = \boldsymbol{0}}$, $\boldsymbol{b}, \boldsymbol{x} \in \mathbb{R}^n, A \in \mathbb{R}^{n \times n}$.

Instead, choose two subspaces $L \subset \mathbb{R}^n$ and $K \subset \mathbb{R}^n$ and

$\boxed{* \text{ find } \widetilde{\boldsymbol{x}} \in \boldsymbol{x}^{(0)} + \delta,\ \delta \in K, \text{ such that } \boldsymbol{b} - A\widetilde{\boldsymbol{x}} \perp L}$

K - search space
L - subspace of constraints

$*$ - basic projection step

The framework is known as Petrov-Galerkin conditions.

There are two major classes of projection methods:

- orthogonal - if $K \equiv L$,
- oblique - if $K \neq L$.

Notations:
$\widetilde{x} = x^0 + \delta$ - (δ - correction)
$r^0 = b - Ax^0$ (r^0 - residual)

$$* \text{ find } \delta \in K, \text{ such that } r^0 - A\delta \perp L$$

Repetition from projection methods

Choose a basis in K and L: $V = \{v_1, v_2, \cdots, v_m\}$ and $W = \{w_1, w_2, \cdots, w_m\}$.
Then, $\widetilde{x} = x^0 + \delta = x^0 + Vy$ for some $y \in R^m$.

The orthogonality condition can be written as

$$(**) \qquad \boxed{W^T(r^0 - AVy)}$$

which is exactly the Petrov-Galerkin condition.
From $(**)$ we get

$$W^T r^0 = W^T AVy$$
$$y = (W^T AV)^{-1} W^T r^0$$
$$\widetilde{x} = x^0 + V(W^T AV)^{-1} W^T r^0$$

In practice, $m < n$, even $m \ll n$, for instance, $m = 1$.

Matrix formulation, cont.

$$\widetilde{x} = x^0 + V(W^T AV)^{-1} W^T r^0$$

The matrix $W^T AV$ will be small and, hopefully, with a nice structure.

!!! $W^T AV$ should be invertible.

A prototype projection-based iterative method:

Given $x^{(0)}$; $x = x^{(0)}$
Until convergence do:
Choose K and L
Choose basis V in K and W in L
Compute $r = b - Ax$
$y = (W^T AV)^{-1} W^T r$
$x = x + Vy$

Degrees of freedom: m, K, L, V, W.
Clearly, if $K \equiv L$, then $V = W$.

(1) Consider two important cases: $L = K$ and $L = AK$
(2) Make a special choice of
$K = \mathcal{K}_{k-1} = \{\boldsymbol{v}, A\boldsymbol{v}, A^2\boldsymbol{v}, \cdots, A^{k-1}\boldsymbol{v}\}$.

Back to 'Eig' problems: subspace methods

Let A - square non-singular of order n.
Let $\mathcal{V} \in \mathbb{C}^n$ be a subspace with $dim(\mathcal{V}) = k < n$, $\mathcal{V}$ - full rank, V - basis of $\mathcal{V}$.
Let $\mathcal{W} = A\mathcal{V}$, $W = AV$.
Task: Find approximate eigenpairs of the eigenvalue problem $A\boldsymbol{z} = \boldsymbol{z}\lambda$ based on information from V and W.

$\mathcal{V}$-orthogonal residuals, Ritz values and Ritz vectors I

Strategy 1 ($V - orth$)**:** Find approximate pair $\boldsymbol{v} \in \mathcal{V}$ and $\mu \in \mathbb{C}$, s.t. the eigenvalue residual $\boldsymbol{r} = A\boldsymbol{v} - \boldsymbol{v}\mu$ is orthogonal to $\mathcal{V}$.
$\boldsymbol{v} \in \mathcal{V} \Rightarrow \boldsymbol{v} = V\boldsymbol{y} = \sum_{1 \le i \le k} y_i \boldsymbol{v}_i, \boldsymbol{y} \in \mathbb{C}^n$.
If $\boldsymbol{r} \perp \mathcal{V}$ then $V^*\boldsymbol{r} = \boldsymbol{0}$, thus

$$V^*(AV\boldsymbol{y} - \boldsymbol{y}\mu) = \boldsymbol{0} \;\equiv\; V^*W\boldsymbol{y} = V^*V\boldsymbol{y}\mu.$$

The pair $(\mu, V\boldsymbol{y})$ can be considered as an approximate eigenpair of A.
Def. (Ritz data): The approximate eigenpairs $(\mu, V\boldsymbol{y})$ are called Ritz pairs, Ritz values and Ritz vectors.

$\mathcal{V}$-orthogonal residuals, Ritz values and Ritz vectors II

Strategy 2 ($W - orth$)**:** Provided that A is invertible, $A\boldsymbol{z} = \boldsymbol{z}\lambda$ can be reformulated as $A^{-1}\boldsymbol{z} = \boldsymbol{z}\lambda^{-1}$. Since $V = A^{-1}W$ we look for $\boldsymbol{w} \in \mathcal{W}, \nu \in \mathbb{C}$, s.t. the eigenvalue residual $\boldsymbol{r} = A\boldsymbol{v} - \boldsymbol{v}\nu$ is orthogonal to $\mathcal{V}$.

$$W^*V\boldsymbol{y} = W^*W\boldsymbol{y}\nu.$$

The pairs $(\nu, W\boldsymbol{y})$ are Ritz pairs of A^{-1} in W.
The following interpretations are equivalent:

1. replace$A, \lambda, W = AV$ from the previous section by $A^{-1}, \lambda^{-1}1, V = A^{-1}W$
2. instead of finding $\boldsymbol{v} \in \mathcal{V}$ with $\boldsymbol{r} \perp \mathcal{V}$, we look for $\boldsymbol{v} \in \mathcal{V}$ with $\boldsymbol{r} \perp \mathcal{W}$.

$\mathcal{V}$-orthogonal residuals, Ritz values and Ritz vectors III

Def. (Harmonic Ritz data): The approximate eigenpairs $(\nu^{-1}, V\boldsymbol{y})$ are called Harmonic Ritz pairs, Harmonic Ritz values and Harmonic Ritz vectors.
The above does not depend on the basis, as long as $\mathcal{W} = A\mathcal{V}$.
However, due to $\mathcal{W} = A\mathcal{V}$ we cannot have simultaneously $V^*V = I$ and $W^*W = I$, i.e. V and W simultaneously orthogonal.

Expansion of the subspaces $V_+ = (V|\boldsymbol{v})$, $W_+ = (W|\boldsymbol{w})$ I

Say, we want to add an arbitrary vector $\boldsymbol{v}$ to V, $V_+ = (V|\boldsymbol{v})$. Then we need to compute $\boldsymbol{w} = A\boldsymbol{v}$ and expand $W_+ = (W|\boldsymbol{w})$.
The matrices we need to extract eigenvalue approximations are:

$$V_+^*V_+ = (V|\boldsymbol{v})^*(V|\boldsymbol{v}) = \left(\begin{array}{c|c} V^*V & V^*\boldsymbol{v} \\ \hline \boldsymbol{v}^*V & \boldsymbol{v}^*\boldsymbol{v} \end{array}\right)$$

respectively

$$V_+^*W_+ = (V|\boldsymbol{v})^*(W|\boldsymbol{w}) = \left(\begin{array}{c|c} V^*W & V^*\boldsymbol{w} \\ \hline \boldsymbol{v}^*W & \boldsymbol{v}^*\boldsymbol{w} \end{array}\right)$$

Expansion of the subspaces $V_+ = (V|\boldsymbol{v})$, $W_+ = (W|\boldsymbol{w})$ II

Theorem

Let $A = A^*$ be positive definite with eigenvalues $\lambda_n \leq \cdots \leq \lambda_1$, $\mu_k \leq \cdots \leq \mu_1$ be the $\mathcal{V}$-orthogonal residual approximations of eigenvalues of A and $\mu_{k+1}^+ \leq \cdots \leq \mu_2^+ \leq \mu_1^+$ be the ones in $\mathcal{V}^+$. Then

$$\mu_{k+1}^+ \leq \mu_k \leq \mu_k^+ \cdots \leq \mu_2 \leq \mu_2^+ \leq \mu_1 \leq \mu_1^+$$

Expansion of the subspaces $V_+ = (V|\boldsymbol{v})$, $W_+ = (W|\boldsymbol{w})$ III

Theorem

Let $A = A^*$ be pos.def. and $\mathcal{V}_1 \subset \mathcal{V}_2 \subset \cdots \subset \mathcal{V}_n$ be a nested sequence of subspaces of dimension $dim(\mathcal{V}_j) = j$. Related to $\mathcal{V}_k$ we denote the corresponding eigenvalues $\mu_k \leq \mu_{k-1} \leq \cdots \leq \mu_1$. Then we have

$$\mu_j^{(j)} \leq \mu_j^{(j+1)} \leq \cdots \leq \mu_j^{(n)} = \lambda_j \; for \; any \; j$$

The sequence of jth largest approximate eigenvalues converges monotonically to the jth largest eigenvalue of A
Conversely,

$$\lambda_j = \mu_j^{(n)} \leq \mu_{j-1}^{(n-1)} \leq \cdots \leq \mu_1^{(n-j+1)} \; for \; any \; j,$$

the sequence of the jth smallest approximate eigenvalue converges monotonically to the jth smallest exact eigenvalue.

Expansion of the subspaces $V_+ = (V|\boldsymbol{v})$, $W_+ = (W|\boldsymbol{w})$ IV

$$\begin{array}{ccccccccc}
 & & & & \mu_1^{(1)} & & & & \\
 & & & \mu_2^{(2)} & & \mu_1^{(2)} & & & \\
 & & \mu_3^{(3)} & & \mu_2^{(3)} & & \mu_1^{(3)} & & \\
 & \mu_4^{(4)} & & \mu_3^{(4)} & & \mu_2^{(4)} & & \mu_1^{(4)} & \\
\mu_5^{(5)} & & \mu_4^{(5)} & & \mu_3^{(5)} & & \mu_2^{(5)} & & \mu_1^{(5)}
\end{array}$$

The jth left diagonal with positive slope decreases monotonically to λ_{n+1-j}
The jth right diagonal with negative slope increases monotonically to λ_j

Algorithm $(V - orth)$

Recall that if $\boldsymbol{r} \perp \mathcal{V}$ then $V^*\boldsymbol{r} = \boldsymbol{0}$, we have the relation

$$V^*(AV\boldsymbol{y} - \boldsymbol{y}\mu) = \boldsymbol{0} \;\equiv\; V^*W\boldsymbol{y} = V^*V\boldsymbol{y}\mu.$$

and we need to compute the eigenvalues of $M = V^*AV = V^*W$ to find approximated eigenvalues of A. The pair $(\mu, V\boldsymbol{y})$ can be considered as an approximate eigenpair of A.
We have also seen that when adding a new vector to $\mathcal{V}$, the matrix M_+ can be easily constructed.
Note: By the definition of $\mathcal{V}$-orthogonal methods, each residual is orthogonal to $\mathcal{V}$. This, if we expand in the direction of $\boldsymbol{r}$, we need not to orthogonalize.
The resulting algorithm is referred to as the Arnoldi method for general matrices and its version for Hermitian matrices, the Lanczos method.

Recall: how to construct a basis for a Krylov subspace

Arnoldi's method for general matrices
Consider $\mathcal{K}^m(A, \boldsymbol{v}) = \{\boldsymbol{v}, A\boldsymbol{v}, A^2\boldsymbol{v}, \cdots, A^{m-1}\boldsymbol{v}\}$, generated by some matrix A and vector $\boldsymbol{v}$.

1: Choose a vector $\boldsymbol{v}_1$ such that $\|\boldsymbol{v}_1\| = 1$
2: **for** $j = 1, 2, \cdots, m$ **do**
3: **for** $i = 1, 2, \cdots, j$ **do**
4: $h_{ij} = (A\boldsymbol{v}_j, \boldsymbol{v}_i)$
5: **end for**
6: $\boldsymbol{w}_j = A\boldsymbol{v}_j - \sum_{i=1}^{j} h_{ij}\boldsymbol{v}_i$
7: $h_{j+1,j} = \|\boldsymbol{w}_j\|$
8: If $h_{j+1,j} = 0$, stop
9: $\boldsymbol{v}_{j+1} = \boldsymbol{w}_j / h_{j+1,j}$
10: **end for**

The result of Arnoldi's process

- $V^m = \{\boldsymbol{v}_1, \boldsymbol{v}_2, \cdots, \boldsymbol{v}_m\}$ is an orthonormal basis in $\mathcal{K}^m(A, \boldsymbol{v})$
- $AV^m = V^m H^m + \boldsymbol{w}_{m+1}\boldsymbol{e}_m^T$

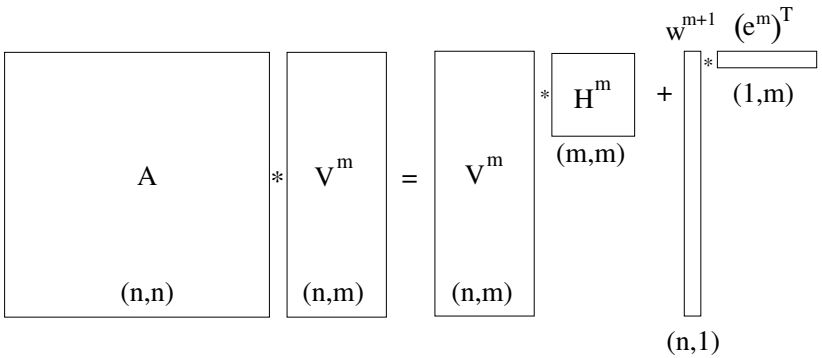

The result of Arnoldi's process

- $V^m = \{\boldsymbol{v}_1, \boldsymbol{v}_2, \cdots, \boldsymbol{v}_m\}$ is an orthonormal basis in $\mathcal{K}^m(A, \boldsymbol{v})$
- $AV^m = V^m H^m + \boldsymbol{w}_{m+1} \boldsymbol{e}_m^T$

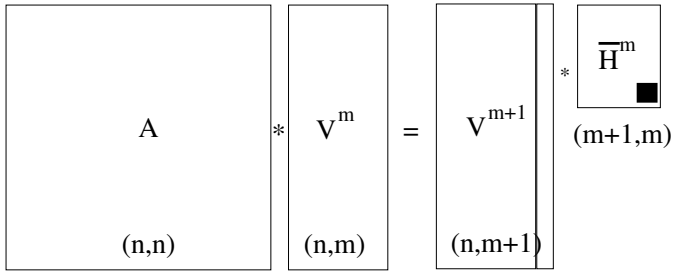

Arnoldi in detail I

Step 1: Choose $\boldsymbol{v}_1$, compute $\boldsymbol{w}_1 = A\boldsymbol{v}_1$
Compute μ such that $\boldsymbol{r}_1 = A\boldsymbol{v}_1 - \boldsymbol{v}_1\mu_1 \perp \boldsymbol{v}_1$
This gives $\mu_1^{(1)} = \boldsymbol{v}_1^* \boldsymbol{w}_1$.
$(\mu_1^{(1)}, \boldsymbol{v}_1)$ - first approximate eigenpair

Arnoldi in detail II

Step 2: Expand $\mathcal{V}_1$ in direction of $\boldsymbol{r}_1$ ($\boldsymbol{r}_1 \perp \mathcal{V}_1$). Note: $\|\boldsymbol{r}_1\|\boldsymbol{v}_2 = \boldsymbol{r}_1 = A\boldsymbol{v}_1 - \boldsymbol{v}_1\mu_1^{(1)}$
$\boldsymbol{v}_2 = \boldsymbol{r}_1/\|\boldsymbol{r}_1\|$, $V2 = (V_1|\boldsymbol{v}_2)$
Compute $\boldsymbol{w}_2 = A\boldsymbol{v}_2$
Solve $V_2^* W_2 \boldsymbol{y} = \boldsymbol{y}\mu$; $V_2^* W_2$ is 2x2, two eigenvalues, $\mu_1^{(2)}, \mu_2^{(2)}$.
Corresponding residuals:

$$R_2 = W_2(\boldsymbol{y}_1|\boldsymbol{y}_2) - V_2(\boldsymbol{y}_1|\boldsymbol{y}_2) \begin{bmatrix} \mu_1^{(2)} & 0 \\ 0 & \mu_2^{(2)} \end{bmatrix}$$

Both residuals are linear combinations of the columns of V_2 and W_2, so it suffices to take one of them, $\boldsymbol{r}$.

Arnoldi in detail III

Step 3: Again, $\|\boldsymbol{r}\|\boldsymbol{v}_3 = \boldsymbol{r} = A\boldsymbol{v}_2 - V_2 V_1^* A\boldsymbol{v}_2$, thus $\boldsymbol{w}_2 = A\boldsymbol{v}_2$ is a linear combination of $\boldsymbol{v}_1, \boldsymbol{v}_2, \boldsymbol{v}_3$ and

$$W_2 = AV_2 = V_3 \begin{bmatrix} * & * \\ * & * \\ 0 & * \end{bmatrix}$$

.
Thus, the relation is

$$W_k = AV_k = V_{k+1} H_{k+1,k}$$

with $H_{k+1,k}$ - upper Hessenberg.
Remark: The Arnoldi factorization is uniquely defined by the initial vector $\boldsymbol{v}_1$.

1: $V = \boldsymbol{v}$, $\|\boldsymbol{v}\| = 1$, $H = []$
2: **while** $\gamma > \varepsilon$ **do**
3: $\quad \boldsymbol{v} = A\boldsymbol{v}$
4: $\quad \boldsymbol{h} = V^*\boldsymbol{v}$
5: $\quad v = V\boldsymbol{h}$
6: $\quad \gamma = \sqrt{\boldsymbol{v}^*\boldsymbol{v}}$
7: $\quad H = \begin{bmatrix} H & \boldsymbol{h} \\ 0 & \gamma \end{bmatrix}$
8: $\quad \boldsymbol{v}_{j+1} = \boldsymbol{w}_j / h_{j+1,j}$
9: $\quad V = (V|\boldsymbol{v}/\gamma)$
10: **end while**

We start expand with $A\boldsymbol{v}_1$, orthogonalization is needed; no need to compute the residual!

Arnoldi: the residual is available

Question: for which k the matrix $H_{k,k}$ has eigenvalues with small residual?
We have $\boldsymbol{r} = AV_k\boldsymbol{y}_k - V_k\boldsymbol{y}_k\mu = h_{k+1,k}\boldsymbol{v}_{k+1}\boldsymbol{e}_k^*\boldsymbol{y}_k$. Using $\|\boldsymbol{v}_{k+1}\| = \|\boldsymbol{e}_k\| = 1$ we obtain $\|\boldsymbol{r}_k\| = |h_{k+1,k}be_k^*\boldsymbol{y}_k| \leq |h_{k+1,k}|$.

Implicitly restarted Arnoldi (IRA) method I

Observations:

- When increasing the dimension of $\mathcal{V}_k$ the computational work becomes prohibitive.
- Roundoff errors creep in.
- The computed spectrum of $H_{k,k}$ contains approximations of uninteresting eigenvalues.
- Perhaps we wish to have started with a different initial vector.

Question: Can we 'reverse' the Arnoldi factorization without heavy recomputations?
Answer: Yes!

Implicitly restarted Arnoldi (IRA) method II

Assume that we have performed $m = \ell + p$ Arnoldi steps. We split the approximate eigenvalues into two groups: $\mu_1, \cdots, \mu_\ell$ wanted and $\theta_1, \cdots, \theta_p$ unwanted, roots of the polynomial $P_p(t) = (t-\theta_1)(t-\theta_2)\cdots,(t-\theta_p)$, We have

$$\boxed{AV_m = V_mH_m + \beta_m\boldsymbol{v}_{m+1}\boldsymbol{e}_m^T.}$$

Apply the first factor $A - \theta_1 I$ to all basis vectors V:

$$\begin{array}{lrcl} & (A-\theta_1 I)V_m & = & V_m(H_m - \theta_1 I) + \beta_m\boldsymbol{v}_{m+1}\boldsymbol{e}_m^T \\ & \text{QR-factorize} & & H_m - \theta_1 I = Q_1R_1 \\ (1) & (A-\theta_1 I)V_m & = & V_mQ_1R_1 + \beta_m\boldsymbol{v}_{m+1}\boldsymbol{e}_m^T \\ (2) & (A-\theta_1 I)(V_mQ_1) & = & (V_mQ_1)R_1Q_1 + \beta_m\boldsymbol{v}_{m+1}\boldsymbol{e}_m^TQ_1 \\ (3) & A(V_mQ_1) & = & (V_mQ_1)(R_1Q_1 + \theta_1 I) + \beta_m\boldsymbol{v}_{m+1}\boldsymbol{e}_m^TQ_1 \end{array}$$

Implicitly restarted Arnoldi (IRA) method III

Denote:

$$\begin{aligned} H_m^{(1)} &= R_1 R_1 + \theta_1 I \\ (\boldsymbol{b}_{m+1}^{(1)})^T &= \boldsymbol{e}_m Q_1 \\ V_m^{(1)} &= V_m Q_1 \end{aligned}$$

Then (3) becomes

$$\boxed{AV_m^{(1)} = V_m^{(1)} H_m^{(1)} + \boldsymbol{v}_{m+1}(\boldsymbol{b}_{m+1}^{(1)})^T.}$$

- Resembles ordinary Arnoldi
- The first vector in $V_m^{(1)}$ is a multiple of $(A - \theta_1 I)\boldsymbol{v}_1$.
  (Multiply (1) by $\boldsymbol{e}_1$) and use the fact the R_1 is upper-triangular
- The columns of $V_m^{(1)}$ are orthonormal.

Implicitly restarted Arnoldi (IRA) method IV

In the same way we can apply $A - \theta_2$ to $V_m^{(1)}$:

$$\begin{aligned} & (A-\theta_2 I)V_m^{(1)} &&= V_m^{(1)}(H_m^{(1)} - \theta_2 I) + \boldsymbol{v}_{m+1}(\boldsymbol{b}_{m+1}^{(1)})^T \\ & \text{QR-factorize} && H_m^{(1)} - \theta_2 I = Q_2 R_2 \\ (4)\quad & (A-\theta_2 I)V_m^{(1)} &&= V_m^{(1)} Q_2 R_2 + \boldsymbol{v}_{m+1}(\boldsymbol{b}_{m+1}^{(1)})^T \\ (5)\quad & (A-\theta_2 I)(V_m^{(1)} Q_2) &&= (V_m^{(1)} Q_2) R_2 Q_2 + \boldsymbol{v}_{m+1}(\boldsymbol{b}_{m+1}^{(1)})^T Q_2 \\ (6)\quad & A(V_m Q_2) &&= (V_m Q_2)(R_2 Q_2 + \theta_2 I) + \boldsymbol{v}_{m+1}(\boldsymbol{b}_{m+1}^{(1)})^T Q_2 \end{aligned}$$

Denote:

$$\begin{aligned} H_m^{(2)} &= R_2 R_2 + \theta_2 I \\ (\boldsymbol{b}_{m+1}^{(2)})^T &= (\boldsymbol{b}_{m+1}^{(1)})^T Q_2 \\ V_m^{(2)} &= V_m^{(1)} Q_2 \end{aligned}$$

Implicitly restarted Arnoldi (IRA) method V

Then (6) becomes

$$\boxed{AV_m^{(2)} = V_m^{(2)} H_m^{(2)} + \boldsymbol{v}_{m+1}(\boldsymbol{b}_{m+1}^{(2)})^T.}$$

Algorithm IRA

1: Perform m-step Arnoldi to obtain the factorization
   $AV_m = V_m H_m + \widehat{\boldsymbol{v}}_{m+1}\boldsymbol{e}_m^T$
2: Choose p of the eigenvalues of H_m to be eliminated, $\theta_1, \cdots, \theta_p$
3: Perform a p-step QR with these shifts
4: $[H_m, Q] = qr(H_m, \theta_1, \cdots, \theta_p)$
5: Set $k = m - p$, $H_k = H_m(1:k, 1:k)$, $V_k = V_k Q$
6: Set $\widehat{\boldsymbol{v}}_{k+1} = \widehat{\boldsymbol{v}}_{k+1} + \widehat{\boldsymbol{v}}_{m+1} Q_{m,k}$
7: Continue the Arnoldi factorization from $AV_k = V_k H_k + \widehat{\boldsymbol{v}}_{k+1}\boldsymbol{e}_k^T$

($\widehat{\boldsymbol{v}}_{k+1} = h_{k+1,k}\boldsymbol{v}_k$ is the unscaled Arnoldi vector of the kth step.)

p-step QR with shifts $\theta_1, \cdots, \theta_p$

1: **for** $j = 1, \cdots, p$ **do**
2: $\quad (H - \theta_j I) = QR$
3: $\quad H = RQ + \theta_j I$
4: **end for**

Eliminating a pair of complex eigenvalues

The double shift strategy is recommended to be used if we want to eliminate two complex conjugate eigenvalues θ_1 and θ_2. Then, these two eigenvalues are used as shifts in consecutive iterations to achieve quadratic convergence in the complex case. That is, we compute

1: $(H - \theta_1 I) = Q_1 R_1$
2: $H_1 = R_1 Q_1 + \theta_1 I$
3: $(H_1 - \theta_2 I) = Q_2 R_2$
4: $H_2 = R_2 Q_2 + \theta_2 I$

Here we have to use complex arithmetic.

Avoiding complex arithmetic

Let $\theta_1 = a + ib, \theta_2 = a - ib$. Then $\theta_1 + \theta_2 = 2a$ and $\theta_1\theta_2 = a^2 + b^2$ are real.

$$
\begin{aligned}
Q_1 Q_2 R_2 R_1 &= Q_1 (H_1 - \theta_2 I) R_1 \\
&= Q_1 H_1 R_1 - \theta_2 Q_1 R_1 \\
&= Q_1 (R_1 Q_1 + \theta_1 I) R_1 - \theta_2 (H - \theta_1 I) \\
&= Q_1 R_1 Q_1 R_1 + \theta_1 Q_1 R_1 - \theta_2 (H - \theta_1 I) \\
&= (H - \theta_1 I)^2 + \theta_1 (H - \theta_1 I) - \theta_2 (H - \theta_1 I) \\
&= H^2 - 2\theta_1 H + \theta_1^2 I + \theta_1 H - \theta_1^2 I - \theta_2 H + \theta_1\theta_2 I \\
&= H^2 - (\theta_1 + \theta_2) H + \theta_1\theta_2 I.
\end{aligned}
$$

... Requires H^2 and is not of practical interest.... Francis QR step to circumvent it.

Assignment

Implement the Arnoldi factorization with implicit restart. Eliminate both real and complex conjugate eigenvalues.