

Scientific Computing: An Introductory Survey

Chapter 4 – Eigenvalue Problems

Prof. Michael T. Heath

Department of Computer Science
University of Illinois at Urbana-Champaign

Copyright © 2002. Reproduction permitted
for noncommercial, educational use only.



Eigenvalue Problems

- Eigenvalue problems occur in many areas of science and engineering, such as structural analysis
- Eigenvalues are also important in analyzing numerical methods
- Theory and algorithms apply to complex matrices as well as real matrices
- With complex matrices, we use conjugate transpose, A^H , instead of usual transpose, A^T



Geometric Interpretation

- Matrix expands or shrinks any vector lying in direction of eigenvector by scalar factor
- Expansion or contraction factor is given by corresponding eigenvalue λ
- Eigenvalues and eigenvectors decompose complicated behavior of general linear transformation into simpler actions



Characteristic Polynomial

- Equation $Ax = \lambda x$ is equivalent to

$$(A - \lambda I)x = 0$$
 which has nonzero solution x if, and only if, its matrix is singular
- Eigenvalues of A are roots λ_i of *characteristic polynomial*

$$\det(A - \lambda I) = 0$$
 in λ of degree n
- *Fundamental Theorem of Algebra* implies that $n \times n$ matrix A always has n eigenvalues, but they may not be real nor distinct
- Complex eigenvalues of real matrix occur in complex conjugate pairs: if $\alpha + i\beta$ is eigenvalue of real matrix, then so is $\alpha - i\beta$, where $i = \sqrt{-1}$



Outline

- 1 Eigenvalue Problems
- 2 Existence, Uniqueness, and Conditioning
- 3 Computing Eigenvalues and Eigenvectors



Eigenvalues and Eigenvectors

- Standard *eigenvalue problem*: Given $n \times n$ matrix A , find scalar λ and nonzero vector x such that

$$Ax = \lambda x$$

- λ is *eigenvalue*, and x is corresponding *eigenvector*
- λ may be complex even if A is real
- *Spectrum* = $\lambda(A)$ = set of eigenvalues of A
- *Spectral radius* = $\rho(A) = \max\{|\lambda| : \lambda \in \lambda(A)\}$



Examples: Eigenvalues and Eigenvectors

- $A = \begin{bmatrix} 1 & 0 \\ 0 & 2 \end{bmatrix}$: $\lambda_1 = 1, x_1 = \begin{bmatrix} 1 \\ 0 \end{bmatrix}, \lambda_2 = 2, x_2 = \begin{bmatrix} 0 \\ 1 \end{bmatrix}$
- $A = \begin{bmatrix} 1 & 1 \\ 0 & 2 \end{bmatrix}$: $\lambda_1 = 1, x_1 = \begin{bmatrix} 1 \\ 0 \end{bmatrix}, \lambda_2 = 2, x_2 = \begin{bmatrix} 1 \\ 1 \end{bmatrix}$
- $A = \begin{bmatrix} 3 & -1 \\ -1 & 3 \end{bmatrix}$: $\lambda_1 = 2, x_1 = \begin{bmatrix} 1 \\ 1 \end{bmatrix}, \lambda_2 = 4, x_2 = \begin{bmatrix} 1 \\ -1 \end{bmatrix}$
- $A = \begin{bmatrix} 1.5 & 0.5 \\ 0.5 & 1.5 \end{bmatrix}$: $\lambda_1 = 2, x_1 = \begin{bmatrix} 1 \\ 1 \end{bmatrix}, \lambda_2 = 1, x_2 = \begin{bmatrix} -1 \\ 1 \end{bmatrix}$
- $A = \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix}$: $\lambda_1 = i, x_1 = \begin{bmatrix} 1 \\ i \end{bmatrix}, \lambda_2 = -i, x_2 = \begin{bmatrix} i \\ 1 \end{bmatrix}$
where $i = \sqrt{-1}$



Example: Characteristic Polynomial

- Characteristic polynomial of previous example matrix is

$$\det \left(\begin{bmatrix} 3 & -1 \\ -1 & 3 \end{bmatrix} - \lambda \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \right) =$$

$$\det \left(\begin{bmatrix} 3 - \lambda & -1 \\ -1 & 3 - \lambda \end{bmatrix} \right) =$$

$$(3 - \lambda)(3 - \lambda) - (-1)(-1) = \lambda^2 - 6\lambda + 8 = 0$$

so eigenvalues are given by

$$\lambda = \frac{6 \pm \sqrt{36 - 32}}{2}, \text{ or } \lambda_1 = 2, \lambda_2 = 4$$



Companion Matrix

- Monic polynomial

$$p(\lambda) = c_0 + c_1\lambda + \dots + c_{n-1}\lambda^{n-1} + \lambda^n$$

is characteristic polynomial of *companion matrix*

$$C_n = \begin{bmatrix} 0 & 0 & \dots & 0 & -c_0 \\ 1 & 0 & \dots & 0 & -c_1 \\ 0 & 1 & \dots & 0 & -c_2 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \dots & 1 & -c_{n-1} \end{bmatrix}$$

- Roots of polynomial of degree > 4 cannot always be computed in finite number of steps
- So in general, computation of eigenvalues of matrices of order > 4 requires (theoretically infinite) iterative process



Example: Characteristic Polynomial

- Consider

$$A = \begin{bmatrix} 1 & \epsilon \\ \epsilon & 1 \end{bmatrix}$$

where ϵ is positive number slightly smaller than $\sqrt{\epsilon_{\text{mach}}}$

- Exact eigenvalues of A are $1 + \epsilon$ and $1 - \epsilon$
- Computing characteristic polynomial in floating-point arithmetic, we obtain

$$\det(A - \lambda I) = \lambda^2 - 2\lambda + (1 - \epsilon^2) = \lambda^2 - 2\lambda + 1$$

which has 1 as double root

- Thus, eigenvalues cannot be resolved by this method even though they are distinct in working precision



Eigenspaces and Invariant Subspaces

- Eigenvectors can be scaled arbitrarily: if $Ax = \lambda x$, then $A(\gamma x) = \lambda(\gamma x)$ for any scalar γ , so γx is also eigenvector corresponding to λ
- Eigenvectors are usually *normalized* by requiring some norm of eigenvector to be 1
- *Eigenspace* = $S_\lambda = \{x : Ax = \lambda x\}$
- Subspace S of \mathbb{R}^n (or \mathbb{C}^n) is *invariant* if $AS \subseteq S$
- For eigenvectors $x_1 \dots x_p$, $\text{span}(\{x_1 \dots x_p\})$ is invariant subspace



Examples: Matrix Properties

- Transpose: $\begin{bmatrix} 1 & 2 \\ 3 & 4 \end{bmatrix}^T = \begin{bmatrix} 1 & 3 \\ 2 & 4 \end{bmatrix}$
- Conjugate transpose: $\begin{bmatrix} 1+i & 1+2i \\ 2-i & 2-2i \end{bmatrix}^H = \begin{bmatrix} 1-i & 2+i \\ 1-2i & 2+2i \end{bmatrix}$
- Symmetric: $\begin{bmatrix} 1 & 2 \\ 2 & 3 \end{bmatrix}$
- Nonsymmetric: $\begin{bmatrix} 1 & 3 \\ 2 & 4 \end{bmatrix}$
- Hermitian: $\begin{bmatrix} 1 & 1+i \\ 1-i & 2 \end{bmatrix}$
- NonHermitian: $\begin{bmatrix} 1 & 1+i \\ 1+i & 2 \end{bmatrix}$



Characteristic Polynomial, continued

- Computing eigenvalues using characteristic polynomial is *not* recommended because of
 - work in computing coefficients of characteristic polynomial
 - sensitivity of coefficients of characteristic polynomial
 - work in solving for roots of characteristic polynomial
- Characteristic polynomial is powerful theoretical tool but usually not useful computationally



Multiplicity and Diagonalizability

- *Multiplicity* is number of times root appears when polynomial is written as product of linear factors
- Eigenvalue of multiplicity 1 is *simple*
- *Defective* matrix has eigenvalue of multiplicity $k > 1$ with fewer than k linearly independent corresponding eigenvectors
- Nondefective matrix A has n linearly independent eigenvectors, so it is *diagonalizable*

$$X^{-1}AX = D$$

where X is nonsingular matrix of eigenvectors



Relevant Properties of Matrices

- Properties of matrix A relevant to eigenvalue problems

Property	Definition
diagonal	$a_{ij} = 0$ for $i \neq j$
tridiagonal	$a_{ij} = 0$ for $ i - j > 1$
triangular	$a_{ij} = 0$ for $i > j$ (upper) $a_{ij} = 0$ for $i < j$ (lower)
Hessenberg	$a_{ij} = 0$ for $i > j + 1$ (upper) $a_{ij} = 0$ for $i < j - 1$ (lower)
orthogonal	$A^T A = A A^T = I$
unitary	$A^H A = A A^H = I$
symmetric	$A = A^T$
Hermitian	$A = A^H$
normal	$A^H A = A A^H$



Examples, continued

- Orthogonal: $\begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$, $\begin{bmatrix} -1 & 0 \\ 0 & -1 \end{bmatrix}$, $\begin{bmatrix} \sqrt{2}/2 & \sqrt{2}/2 \\ -\sqrt{2}/2 & \sqrt{2}/2 \end{bmatrix}$
- Unitary: $\begin{bmatrix} i\sqrt{2}/2 & \sqrt{2}/2 \\ -\sqrt{2}/2 & -i\sqrt{2}/2 \end{bmatrix}$
- Nonorthogonal: $\begin{bmatrix} 1 & 1 \\ 1 & 2 \end{bmatrix}$
- Normal: $\begin{bmatrix} 1 & 2 & 0 \\ 0 & 1 & 2 \\ 2 & 0 & 1 \end{bmatrix}$
- Nonnormal: $\begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix}$



Properties of Eigenvalue Problems

Properties of eigenvalue problem affecting choice of algorithm and software

- Are all eigenvalues needed, or only a few?
- Are only eigenvalues needed, or are corresponding eigenvectors also needed?
- Is matrix real or complex?
- Is matrix relatively small and dense, or large and sparse?
- Does matrix have any special properties, such as symmetry, or is it general matrix?



Conditioning of Eigenvalues

- If μ is eigenvalue of perturbation $A + E$ of nondefective matrix A , then

$$|\mu - \lambda_k| \leq \text{cond}_2(X) \|E\|_2$$

where λ_k is closest eigenvalue of A to μ and X is nonsingular matrix of eigenvectors of A

- Absolute condition number of eigenvalues is condition number of matrix of eigenvectors with respect to solving linear equations
- Eigenvalues may be sensitive if eigenvectors are nearly linearly dependent (i.e., matrix is nearly defective)
- For *normal* matrix ($A^H A = A A^H$), eigenvectors are orthogonal, so eigenvalues are well-conditioned



Problem Transformations

- *Shift*: If $Ax = \lambda x$ and σ is any scalar, then $(A - \sigma I)x = (\lambda - \sigma)x$, so eigenvalues of shifted matrix are shifted eigenvalues of original matrix
- *Inversion*: If A is nonsingular and $Ax = \lambda x$ with $x \neq 0$, then $\lambda \neq 0$ and $A^{-1}x = (1/\lambda)x$, so eigenvalues of inverse are reciprocals of eigenvalues of original matrix
- *Powers*: If $Ax = \lambda x$, then $A^k x = \lambda^k x$, so eigenvalues of power of matrix are same power of eigenvalues of original matrix
- *Polynomial*: If $Ax = \lambda x$ and $p(t)$ is polynomial, then $p(A)x = p(\lambda)x$, so eigenvalues of polynomial in matrix are values of polynomial evaluated at eigenvalues of original matrix



Example: Similarity Transformation

- From eigenvalues and eigenvectors for previous example,

$$\begin{bmatrix} 3 & -1 \\ -1 & 3 \end{bmatrix} \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix} = \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix} \begin{bmatrix} 2 & 0 \\ 0 & 4 \end{bmatrix}$$

and hence

$$\begin{bmatrix} 0.5 & 0.5 \\ 0.5 & -0.5 \end{bmatrix} \begin{bmatrix} 3 & -1 \\ -1 & 3 \end{bmatrix} \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix} = \begin{bmatrix} 2 & 0 \\ 0 & 4 \end{bmatrix}$$

- So original matrix is similar to diagonal matrix, and eigenvectors form columns of similarity transformation matrix



Conditioning of Eigenvalue Problems

- Condition of eigenvalue problem is sensitivity of eigenvalues and eigenvectors to changes in matrix
- Conditioning of eigenvalue problem is *not* same as conditioning of solution to linear system for same matrix
- Different eigenvalues and eigenvectors are not necessarily equally sensitive to perturbations in matrix



Conditioning of Eigenvalues

- If $(A + E)(x + \Delta x) = (\lambda + \Delta\lambda)(x + \Delta x)$, where λ is simple eigenvalue of A , then

$$|\Delta\lambda| \lesssim \frac{\|y\|_2 \cdot \|x\|_2}{|y^H x|} \|E\|_2 = \frac{1}{\cos(\theta)} \|E\|_2$$

where x and y are corresponding right and left eigenvectors and θ is angle between them

- For symmetric or Hermitian matrix, right and left eigenvectors are same, so $\cos(\theta) = 1$ and eigenvalues are inherently well-conditioned
- Eigenvalues of nonnormal matrices may be sensitive
- For multiple or closely clustered eigenvalues, corresponding eigenvectors may be sensitive



Similarity Transformation

- B is *similar* to A if there is nonsingular matrix T such that

$$B = T^{-1} A T$$

- Then

$$B y = \lambda y \Rightarrow T^{-1} A T y = \lambda y \Rightarrow A(T y) = \lambda(T y)$$

so A and B have same eigenvalues, and if y is eigenvector of B , then $x = T y$ is eigenvector of A

- Similarity transformations preserve eigenvalues and eigenvectors are easily recovered



Diagonal Form

- Eigenvalues of diagonal matrix are diagonal entries, and eigenvectors are columns of identity matrix
- Diagonal form is desirable in simplifying eigenvalue problems for general matrices by similarity transformations
- But not all matrices are diagonalizable by similarity transformation
- Closest one can get, in general, is *Jordan form*, which is nearly diagonal but may have some nonzero entries on first superdiagonal, corresponding to one or more multiple eigenvalues



Triangular Form

- Any matrix can be transformed into triangular (*Schur*) form by similarity, and eigenvalues of triangular matrix are diagonal entries
- Eigenvectors of triangular matrix less obvious, but still straightforward to compute
- If

$$A - \lambda I = \begin{bmatrix} U_{11} & u & U_{13} \\ \mathbf{0} & 0 & v^T \\ O & 0 & U_{33} \end{bmatrix}$$

is triangular, then $U_{11}y = u$ can be solved for y , so that

$$x = \begin{bmatrix} y \\ -1 \\ 0 \end{bmatrix}$$

is corresponding eigenvector



Forms Attainable by Similarity

A	T	B
distinct eigenvalues	nonsingular	diagonal
real symmetric	orthogonal	real diagonal
complex Hermitian	unitary	real diagonal
normal	unitary	diagonal
arbitrary real	orthogonal	real block triangular (real Schur)
arbitrary	unitary	upper triangular (Schur)
arbitrary	nonsingular	almost diagonal (Jordan)

- Given matrix A with indicated property, matrices B and T exist with indicated properties such that $B = T^{-1}AT$
- If B is diagonal or triangular, eigenvalues are its diagonal entries
- If B is diagonal, eigenvectors are columns of T



Convergence of Power Iteration

- To see why power iteration converges to dominant eigenvector, express starting vector x_0 as linear combination

$$x_0 = \sum_{i=1}^n \alpha_i v_i$$

where v_i are eigenvectors of A

- Then

$$x_k = Ax_{k-1} = A^2x_{k-2} = \dots = A^k x_0 =$$

$$\sum_{i=1}^n \lambda_i^k \alpha_i v_i = \lambda_1^k \left(\alpha_1 v_1 + \sum_{i=2}^n (\lambda_i/\lambda_1)^k \alpha_i v_i \right)$$

- Since $|\lambda_i/\lambda_1| < 1$ for $i > 1$, successively higher powers go to zero, leaving only component corresponding to v_1



Limitations of Power Iteration

Power iteration can fail for various reasons

- Starting vector may have *no* component in dominant eigenvector v_1 (i.e., $\alpha_1 = 0$) — not problem in practice because rounding error usually introduces such component in any case
- There may be more than one eigenvalue having same (maximum) modulus, in which case iteration may converge to linear combination of corresponding eigenvectors
- For real matrix and starting vector, iteration can never converge to complex vector



Block Triangular Form

- If

$$A = \begin{bmatrix} A_{11} & A_{12} & \dots & A_{1p} \\ & A_{22} & \dots & A_{2p} \\ & & \ddots & \vdots \\ & & & A_{pp} \end{bmatrix}$$

with *square* diagonal blocks, then

$$\lambda(A) = \bigcup_{j=1}^p \lambda(A_{jj})$$

so eigenvalue problem breaks into p smaller eigenvalue problems

- Real* Schur form has 1×1 diagonal blocks corresponding to real eigenvalues and 2×2 diagonal blocks corresponding to pairs of complex conjugate eigenvalues



Power Iteration

- Simplest method for computing one eigenvalue-eigenvector pair is *power iteration*, which repeatedly multiplies matrix times initial starting vector
- Assume A has unique eigenvalue of maximum modulus, say λ_1 , with corresponding eigenvector v_1
- Then, starting from nonzero vector x_0 , iteration scheme

$$x_k = Ax_{k-1}$$

converges to multiple of eigenvector v_1 corresponding to *dominant* eigenvalue λ_1



Example: Power Iteration

- Ratio of values of given component of x_k from one iteration to next converges to dominant eigenvalue λ_1
- For example, if $A = \begin{bmatrix} 1.5 & 0.5 \\ 0.5 & 1.5 \end{bmatrix}$ and $x_0 = \begin{bmatrix} 0 \\ 1 \end{bmatrix}$, we obtain

k	x_k^T		ratio
0	0.0	1.0	
1	0.5	1.5	1.500
2	1.5	2.5	1.667
3	3.5	4.5	1.800
4	7.5	8.5	1.889
5	15.5	16.5	1.941
6	31.5	32.5	1.970
7	63.5	64.5	1.985
8	127.5	128.5	1.992

- Ratio is converging to dominant eigenvalue, which is 2



Normalized Power Iteration

- Geometric growth of components at each iteration risks eventual overflow (or underflow if $\lambda_1 < 1$)
- Approximate eigenvector should be normalized at each iteration, say, by requiring its largest component to be 1 in modulus, giving iteration scheme

$$y_k = Ax_{k-1}$$

$$x_k = y_k / \|y_k\|_\infty$$

- With normalization, $\|y_k\|_\infty \rightarrow |\lambda_1|$, and $x_k \rightarrow v_1 / \|v_1\|_\infty$



Example: Normalized Power Iteration

- Repeating previous example with normalized scheme,

k	x_k^T	$\ y_k\ _\infty$
0	0.000 1.0	
1	0.333 1.0	1.500
2	0.600 1.0	1.667
3	0.778 1.0	1.800
4	0.882 1.0	1.889
5	0.939 1.0	1.941
6	0.969 1.0	1.970
7	0.984 1.0	1.985
8	0.992 1.0	1.992

< interactive example >



Power Iteration with Shift

- Convergence rate of power iteration depends on ratio $|\lambda_2/\lambda_1|$, where λ_2 is eigenvalue having second largest modulus
- May be possible to choose shift, $A - \sigma I$, such that

$$\left| \frac{\lambda_2 - \sigma}{\lambda_1 - \sigma} \right| < \left| \frac{\lambda_2}{\lambda_1} \right|$$

so convergence is accelerated

- Shift must then be added to result to obtain eigenvalue of original matrix



Inverse Iteration

- If smallest eigenvalue of matrix required rather than largest, can make use of fact that eigenvalues of A^{-1} are reciprocals of those of A , so smallest eigenvalue of A is reciprocal of largest eigenvalue of A^{-1}
- This leads to *inverse iteration* scheme

$$\begin{aligned} Ay_k &= x_{k-1} \\ x_k &= y_k / \|y_k\|_\infty \end{aligned}$$

which is equivalent to power iteration applied to A^{-1}

- Inverse of A not computed explicitly, but factorization of A used to solve system of linear equations at each iteration



Example: Inverse Iteration

- Applying inverse iteration to previous example to compute smallest eigenvalue yields sequence

k	x_k^T	$\ y_k\ _\infty$
0	0.000 1.0	
1	-0.333 1.0	0.750
2	-0.600 1.0	0.833
3	-0.778 1.0	0.900
4	-0.882 1.0	0.944
5	-0.939 1.0	0.971
6	-0.969 1.0	0.985


which is indeed converging to 1 (which is its own reciprocal in this case)

< interactive example >



Geometric Interpretation

- Behavior of power iteration depicted geometrically



- Initial vector $x_0 = v_1 + v_2$ contains equal components in eigenvectors v_1 and v_2 (dashed arrows)
- Repeated multiplication by A causes component in v_1 (corresponding to larger eigenvalue, 2) to dominate, so sequence of vectors x_k converges to v_1



Example: Power Iteration with Shift

- In earlier example, for instance, if we pick shift of $\sigma = 1$, (which is equal to other eigenvalue) then ratio becomes zero and method converges in one iteration
- In general, we would not be able to make such fortuitous choice, but shifts can still be extremely useful in some contexts, as we will see later



Inverse Iteration, continued

- Inverse iteration converges to eigenvector corresponding to *smallest* eigenvalue of A
- Eigenvalue obtained is dominant eigenvalue of A^{-1} , and hence its reciprocal is smallest eigenvalue of A in modulus



Inverse Iteration with Shift

- As before, shifting strategy, working with $A - \sigma I$ for some scalar σ , can greatly improve convergence
- Inverse iteration is particularly useful for computing eigenvector corresponding to approximate eigenvalue, since it converges rapidly when applied to shifted matrix $A - \lambda I$, where λ is approximate eigenvalue
- Inverse iteration is also useful for computing eigenvalue closest to given value β , since if β is used as shift, then desired eigenvalue corresponds to smallest eigenvalue of shifted matrix



Rayleigh Quotient

- Given approximate eigenvector x for real matrix A , determining best estimate for corresponding eigenvalue λ can be considered as $n \times 1$ linear least squares approximation problem

$$x\lambda \cong Ax$$

- From normal equation $x^T x \lambda = x^T Ax$, least squares solution is given by

$$\lambda = \frac{x^T Ax}{x^T x}$$

- This quantity, known as **Rayleigh quotient**, has many useful properties



Rayleigh Quotient Iteration

- Given approximate eigenvector, Rayleigh quotient yields good estimate for corresponding eigenvalue
- Conversely, inverse iteration converges rapidly to eigenvector if approximate eigenvalue is used as shift, with one iteration often sufficing
- These two ideas combined in **Rayleigh quotient iteration**

$$\sigma_k = x_k^T Ax_k / x_k^T x_k$$

$$(A - \sigma_k I)y_{k+1} = x_k$$

$$x_{k+1} = y_{k+1} / \|y_{k+1}\|_\infty$$

starting from given nonzero vector x_0



Example: Rayleigh Quotient Iteration

- Using same matrix as previous examples and randomly chosen starting vector x_0 , Rayleigh quotient iteration converges in two iterations

k	x_k^T	σ_k
0	0.807 0.397	1.896
1	0.924 1.000	1.998
2	1.000 1.000	2.000



Deflation, continued

- Thus, we can work with B to compute next eigenvalue λ_2
- Moreover, if y_2 is eigenvector of B corresponding to λ_2 , then

$$x_2 = H^{-1} \begin{bmatrix} \alpha \\ y_2 \end{bmatrix}, \quad \text{where } \alpha = \frac{b^T y_2}{\lambda_2 - \lambda_1}$$

is eigenvector corresponding to λ_2 for original matrix A , provided $\lambda_1 \neq \lambda_2$

- Process can be repeated to find additional eigenvalues and eigenvectors



Example: Rayleigh Quotient

- Rayleigh quotient can accelerate convergence of iterative methods such as power iteration, since Rayleigh quotient $x_k^T Ax_k / x_k^T x_k$ gives better approximation to eigenvalue at iteration k than does basic method alone
- For previous example using power iteration, value of Rayleigh quotient at each iteration is shown below

k	x_k^T	$\ y_k\ _\infty$	$x_k^T Ax_k / x_k^T x_k$
0	0.000 1.0		
1	0.333 1.0	1.500	1.500
2	0.600 1.0	1.667	1.800
3	0.778 1.0	1.800	1.941
4	0.882 1.0	1.889	1.985
5	0.939 1.0	1.941	1.996
6	0.969 1.0	1.970	1.999



Rayleigh Quotient Iteration, continued

- Rayleigh quotient iteration is especially effective for symmetric matrices and usually converges very rapidly
- Using different shift at each iteration means matrix must be refactored each time to solve linear system, so cost per iteration is high unless matrix has special form that makes factorization easy
- Same idea also works for complex matrices, for which transpose is replaced by conjugate transpose, so Rayleigh quotient becomes $x^H Ax / x^H x$



Deflation

- After eigenvalue λ_1 and corresponding eigenvector x_1 have been computed, then additional eigenvalues $\lambda_2, \dots, \lambda_n$ of A can be computed by **deflation**, which effectively removes known eigenvalue
- Let H be any nonsingular matrix such that $Hx_1 = \alpha e_1$, scalar multiple of first column of identity matrix (Householder transformation is good choice for H)
- Then similarity transformation determined by H transforms A into form

$$HAH^{-1} = \begin{bmatrix} \lambda_1 & b^T \\ 0 & B \end{bmatrix}$$

where B is matrix of order $n - 1$ having eigenvalues $\lambda_2, \dots, \lambda_n$



Deflation, continued

- Alternative approach lets u_1 be any vector such that $u_1^T x_1 = \lambda_1$
- Then $A - x_1 u_1^T$ has eigenvalues $0, \lambda_2, \dots, \lambda_n$
- Possible choices for u_1 include
 - $u_1 = \lambda_1 x_1$, if A is symmetric and x_1 is normalized so that $\|x_1\|_2 = 1$
 - $u_1 = \lambda_1 y_1$, where y_1 is corresponding left eigenvector (i.e., $A^T y_1 = \lambda_1 y_1$) normalized so that $y_1^T x_1 = 1$
 - $u_1 = A^T e_k$, if x_1 is normalized so that $\|x_1\|_\infty = 1$ and k th component of x_1 is 1



Simultaneous Iteration

- Simplest method for computing many eigenvalue-eigenvector pairs is *simultaneous iteration*, which repeatedly multiplies matrix times matrix of initial starting vectors
- Starting from $n \times p$ matrix X_0 of rank p , iteration scheme is

$$X_k = AX_{k-1}$$

- $\text{span}(X_k)$ converges to invariant subspace determined by p largest eigenvalues of A , provided $|\lambda_p| > |\lambda_{p+1}|$
- Also called *subspace iteration*



QR Iteration

- For $p = n$ and $X_0 = I$, matrices

$$A_k = \hat{Q}_k^H A \hat{Q}_k$$

generated by orthogonal iteration converge to triangular or block triangular form, yielding all eigenvalues of A

- QR iteration* computes successive matrices A_k without forming above product explicitly
- Starting with $A_0 = A$, at iteration k compute QR factorization

$$Q_k R_k = A_{k-1}$$

and form reverse product

$$A_k = R_k Q_k$$



Example: QR Iteration

- Let $A_0 = \begin{bmatrix} 7 & 2 \\ 2 & 4 \end{bmatrix}$
- Compute QR factorization

$$A_0 = Q_1 R_1 = \begin{bmatrix} .962 & -.275 \\ .275 & .962 \end{bmatrix} \begin{bmatrix} 7.28 & 3.02 \\ 0 & 3.30 \end{bmatrix}$$

and form reverse product

$$A_1 = R_1 Q_1 = \begin{bmatrix} 7.83 & .906 \\ .906 & 3.17 \end{bmatrix}$$

- Off-diagonal entries are now smaller, and diagonal entries closer to eigenvalues, 8 and 3
- Process continues until matrix is within tolerance of being diagonal, and diagonal entries then closely approximate eigenvalues



Example: QR Iteration with Shifts

- Repeat previous example, but with shift of $\sigma_1 = 4$, which is lower right corner entry of matrix
- We compute QR factorization

$$A_0 - \sigma_1 I = Q_1 R_1 = \begin{bmatrix} .832 & .555 \\ .555 & -.832 \end{bmatrix} \begin{bmatrix} 3.61 & 1.66 \\ 0 & 1.11 \end{bmatrix}$$

and form reverse product, adding back shift to obtain

$$A_1 = R_1 Q_1 + \sigma_1 I = \begin{bmatrix} 7.92 & .615 \\ .615 & 3.08 \end{bmatrix}$$

- After one iteration, off-diagonal entries smaller compared with unshifted algorithm, and eigenvalues closer approximations to eigenvalues

< interactive example >



Orthogonal Iteration

- As with power iteration, normalization is needed with simultaneous iteration
- Each column of X_k converges to dominant eigenvector, so columns of X_k become increasingly ill-conditioned basis for $\text{span}(X_k)$
- Both issues can be addressed by computing QR factorization at each iteration

$$\begin{aligned} \hat{Q}_k R_k &= X_{k-1} \\ X_k &= A \hat{Q}_k \end{aligned}$$

where $\hat{Q}_k R_k$ is *reduced* QR factorization of X_{k-1}

- This *orthogonal iteration* converges to block triangular form, and leading block is triangular if moduli of consecutive eigenvalues are distinct



QR Iteration, continued

- Successive matrices A_k are unitarily similar to each other

$$A_k = R_k Q_k = Q_k^H A_{k-1} Q_k$$

- Diagonal entries (or eigenvalues of diagonal blocks) of A_k converge to eigenvalues of A
- Product of orthogonal matrices Q_k converges to matrix of corresponding eigenvectors
- If A is symmetric, then symmetry is preserved by QR iteration, so A_k converge to matrix that is both triangular and symmetric, hence diagonal



QR Iteration with Shifts

- Convergence rate of QR iteration can be accelerated by incorporating *shifts*

$$\begin{aligned} Q_k R_k &= A_{k-1} - \sigma_k I \\ A_k &= R_k Q_k + \sigma_k I \end{aligned}$$

where σ_k is rough approximation to eigenvalue

- Good shift can be determined by computing eigenvalues of 2×2 submatrix in lower right corner of matrix



Preliminary Reduction

- Efficiency of QR iteration can be enhanced by first transforming matrix as close to triangular form as possible before beginning iterations
- Hessenberg matrix* is triangular except for one additional nonzero diagonal immediately adjacent to main diagonal
- Any matrix can be reduced to Hessenberg form in finite number of steps by orthogonal similarity transformation, for example using Householder transformations
- Symmetric Hessenberg matrix is tridiagonal
- Hessenberg or tridiagonal form is preserved during successive QR iterations



Preliminary Reduction, continued

Advantages of initial reduction to upper Hessenberg or tridiagonal form

- Work per QR iteration is reduced from $\mathcal{O}(n^3)$ to $\mathcal{O}(n^2)$ for general matrix or $\mathcal{O}(n)$ for symmetric matrix
- Fewer QR iterations are required because matrix nearly triangular (or diagonal) already
- If any zero entries on first subdiagonal, then matrix is block triangular and problem can be broken into two or more smaller subproblems



Cost of QR Iteration

Approximate overall cost of preliminary reduction and QR iteration, counting both additions and multiplications

- Symmetric matrices
 - $\frac{4}{3}n^3$ for eigenvalues only
 - $9n^3$ for eigenvalues and eigenvectors
- General matrices
 - $10n^3$ for eigenvalues only
 - $25n^3$ for eigenvalues and eigenvectors



Krylov Subspace Methods

- Equating k th columns on each side of equation $AQ_n = Q_n H$ yields recurrence

$$Aq_k = h_{1k}q_1 + \dots + h_{kk}q_k + h_{k+1,k}q_{k+1}$$

relating q_{k+1} to preceding vectors q_1, \dots, q_k

- Premultiplying by q_j^H and using orthonormality,

$$h_{jk} = q_j^H Aq_k, \quad j = 1, \dots, k$$

- These relationships yield **Arnoldi iteration**, which produces unitary matrix Q_n and upper Hessenberg matrix H_n column by column using only matrix-vector multiplication by A and inner products of vectors



Arnoldi Iteration, continued

- If

$$Q_k = [q_1 \ \dots \ q_k]$$

then

$$H_k = Q_k^H A Q_k$$

is upper Hessenberg matrix

- Eigenvalues of H_k , called **Ritz values**, are approximate eigenvalues of A , and **Ritz vectors** given by $Q_k y$, where y is eigenvector of H_k , are corresponding approximate eigenvectors of A
- Eigenvalues of H_k must be computed by another method, such as QR iteration, but this is easier problem if $k \ll n$



Preliminary Reduction, continued

- QR iteration is implemented in two-stages

symmetric \rightarrow tridiagonal \rightarrow diagonal

or

general \rightarrow Hessenberg \rightarrow triangular

- Preliminary reduction requires definite number of steps, whereas subsequent iterative stage continues until convergence
- In practice only modest number of iterations usually required, so much of work is in preliminary reduction
- Cost of accumulating eigenvectors, if needed, dominates total cost



Krylov Subspace Methods

- Krylov subspace methods reduce matrix to Hessenberg (or tridiagonal) form using only matrix-vector multiplication
- For arbitrary starting vector x_0 , if

$$K_k = [x_0 \ Ax_0 \ \dots \ A^{k-1}x_0]$$

then

$$K_n^{-1} A K_n = C_n$$

where C_n is upper Hessenberg (in fact, companion matrix)

- To obtain better conditioned basis for $\text{span}(K_n)$, compute QR factorization

$$Q_n R_n = K_n$$

so that

$$Q_n^H A Q_n = R_n C_n R_n^{-1} \equiv H$$

with H upper Hessenberg



Arnoldi Iteration

$x_0 =$ arbitrary nonzero starting vector

$$q_1 = x_0 / \|x_0\|_2$$

for $k = 1, 2, \dots$

$$u_k = Aq_k$$

for $j = 1$ to k

$$h_{jk} = q_j^H u_k$$

$$u_k = u_k - h_{jk} q_j$$

end

$$h_{k+1,k} = \|u_k\|_2$$

if $h_{k+1,k} = 0$ then stop

$$q_{k+1} = u_k / h_{k+1,k}$$

end



Arnoldi Iteration, continued

- Arnoldi iteration fairly expensive in work and storage because each new vector q_k must be orthogonalized against all previous columns of Q_k , and all must be stored for that purpose.
- So Arnoldi process usually restarted periodically with carefully chosen starting vector
- Ritz values and vectors produced are often good approximations to eigenvalues and eigenvectors of A after relatively few iterations



Lanczos Iteration

- Work and storage costs drop dramatically if matrix is symmetric or Hermitian, since recurrence then has only three terms and H_k is tridiagonal (so usually denoted T_k)

$$q_0 = 0$$

$$\beta_0 = 0$$

$x_0 =$ arbitrary nonzero starting vector

$$q_1 = x_0 / \|x_0\|_2$$

for $k = 1, 2, \dots$

$$u_k = Aq_k$$

$$\alpha_k = q_k^H u_k$$

$$u_k = u_k - \beta_{k-1}q_{k-1} - \alpha_k q_k$$

$$\beta_k = \|u_k\|_2$$

if $\beta_k = 0$ then stop

$$q_{k+1} = u_k / \beta_k$$

end



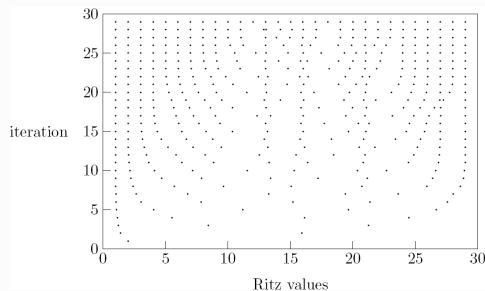
Lanczos Iteration, continued

- In principle, if Lanczos algorithm were run until $k = n$, resulting tridiagonal matrix would be orthogonally similar to A
- In practice, rounding error causes loss of orthogonality, invalidating this expectation
- Problem can be overcome by reorthogonalizing vectors as needed, but expense can be substantial
- Alternatively, can ignore problem, in which case algorithm still produces good eigenvalue approximations, but multiple copies of some eigenvalues may be generated



Example: Lanczos Iteration

- For 29×29 symmetric matrix with eigenvalues $1, \dots, 29$, behavior of Lanczos iteration is shown below



Jacobi Method, continued

- Transformed matrix diagonal if

$$c^2b + cs(a - d) - s^2b = 0$$

- Dividing both sides by c^2b , we obtain

$$1 + \frac{s}{c} \frac{(a-d)}{b} - \frac{s^2}{c^2} = 0$$

- Making substitution $t = s/c$, we get quadratic equation

$$1 + t \frac{(a-d)}{b} - t^2 = 0$$

for tangent t of angle of rotation, from which we can recover $c = 1/\sqrt{1+t^2}$ and $s = c \cdot t$

- Advantageous numerically to use root of smaller magnitude



Lanczos Iteration, continued

- α_k and β_k are diagonal and subdiagonal entries of symmetric tridiagonal matrix T_k
- As with Arnoldi, Lanczos iteration does not produce eigenvalues and eigenvectors directly, but only tridiagonal matrix T_k , whose eigenvalues and eigenvectors must be computed by another method to obtain Ritz values and vectors
- If $\beta_k = 0$, then algorithm appears to break down, but in that case invariant subspace has already been identified (i.e., Ritz values and vectors are already exact at that point)



Krylov Subspace Methods, continued

- Great virtue of Arnoldi and Lanczos iterations is their ability to produce good approximations to extreme eigenvalues for $k \ll n$
- Moreover, they require only one matrix-vector multiplication by A per step and little auxiliary storage, so are ideally suited to large sparse matrices
- If eigenvalues are needed in middle of spectrum, say near σ , then algorithm can be applied to matrix $(A - \sigma I)^{-1}$, assuming it is practical to solve systems of form $(A - \sigma I)x = y$



Jacobi Method

- One of oldest methods for computing eigenvalues is *Jacobi method*, which uses similarity transformation based on plane rotations
- Sequence of plane rotations chosen to annihilate symmetric pairs of matrix entries, eventually converging to diagonal form
- Choice of plane rotation slightly more complicated than in Givens method for QR factorization
- To annihilate given off-diagonal pair, choose c and s so that

$$\begin{aligned} J^T A J &= \begin{bmatrix} c & -s \\ s & c \end{bmatrix} \begin{bmatrix} a & b \\ b & d \end{bmatrix} \begin{bmatrix} c & s \\ -s & c \end{bmatrix} \\ &= \begin{bmatrix} c^2a - 2csb + s^2d & c^2b + cs(a-d) - s^2b \\ c^2b + cs(a-d) - s^2b & c^2d + 2csb + s^2a \end{bmatrix} \end{aligned}$$

is diagonal



Example: Plane Rotation

- Consider 2×2 matrix

$$A = \begin{bmatrix} 1 & 2 \\ 2 & 1 \end{bmatrix}$$

- Quadratic equation for tangent reduces to $t^2 = 1$, so $t = \pm 1$

- Two roots of same magnitude, so we arbitrarily choose $t = -1$, which yields $c = 1/\sqrt{2}$ and $s = -1/\sqrt{2}$

- Resulting plane rotation J gives

$$J^T A J = \begin{bmatrix} 1/\sqrt{2} & 1/\sqrt{2} \\ -1/\sqrt{2} & 1/\sqrt{2} \end{bmatrix} \begin{bmatrix} 1 & 2 \\ 2 & 1 \end{bmatrix} \begin{bmatrix} 1/\sqrt{2} & -1/\sqrt{2} \\ 1/\sqrt{2} & 1/\sqrt{2} \end{bmatrix} = \begin{bmatrix} 3 & 0 \\ 0 & -1 \end{bmatrix}$$



Jacobi Method, continued

- Starting with symmetric matrix $A_0 = A$, each iteration has form

$$A_{k+1} = J_k^T A_k J_k$$

where J_k is plane rotation chosen to annihilate a symmetric pair of entries in A_k

- Plane rotations repeatedly applied from both sides in systematic sweeps through matrix until off-diagonal mass of matrix is reduced to within some tolerance of zero
- Resulting diagonal matrix orthogonally similar to original matrix, so diagonal entries are eigenvalues, and eigenvectors are given by product of plane rotations



Example: Jacobi Method

- Let $A_0 = \begin{bmatrix} 1 & 0 & 2 \\ 0 & 2 & 1 \\ 2 & 1 & 1 \end{bmatrix}$
- First annihilate (1,3) and (3,1) entries using rotation

$$J_0 = \begin{bmatrix} 0.707 & 0 & -0.707 \\ 0 & 1 & 0 \\ 0.707 & 0 & 0.707 \end{bmatrix}$$

to obtain

$$A_1 = J_0^T A_0 J_0 = \begin{bmatrix} 3 & 0.707 & 0 \\ 0.707 & 2 & 0.707 \\ 0 & 0.707 & -1 \end{bmatrix}$$



Example, continued

- Next annihilate (2,3) and (3,2) entries using rotation

$$J_2 = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 0.975 & -0.221 \\ 0 & 0.221 & 0.975 \end{bmatrix}$$

to obtain

$$A_3 = J_2^T A_2 J_2 = \begin{bmatrix} 3.366 & 0.072 & 0.317 \\ 0.072 & 1.776 & 0 \\ 0.317 & 0 & -1.142 \end{bmatrix}$$



Example, continued

- Process continues until off-diagonal entries reduced to as small as desired
- Result is diagonal matrix orthogonally similar to original matrix, with the orthogonal similarity transformation given by product of plane rotations

< interactive example >



Jacobi Method, continued

- Jacobi method is reliable, simple to program, and capable of high accuracy, but converges rather slowly and difficult to generalize beyond symmetric matrices
- Except for small problems, more modern methods usually require 5 to 10 times less work than Jacobi
- One source of inefficiency is that previously annihilated entries can subsequently become nonzero again, thereby requiring repeated annihilation
- Newer methods such as QR iteration preserve zero entries introduced into matrix



Example, continued

- Next annihilate (1,2) and (2,1) entries using rotation

$$J_1 = \begin{bmatrix} 0.888 & -0.460 & 0 \\ 0.460 & 0.888 & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

to obtain

$$A_2 = J_1^T A_1 J_1 = \begin{bmatrix} 3.366 & 0 & 0.325 \\ 0 & 1.634 & 0.628 \\ 0.325 & 0.628 & -1 \end{bmatrix}$$



Example, continued

- Beginning new sweep, again annihilate (1,3) and (3,1) entries using rotation

$$J_3 = \begin{bmatrix} 0.998 & 0 & -0.070 \\ 0 & 1 & 0 \\ 0.070 & 0 & 0.998 \end{bmatrix}$$

to obtain

$$A_4 = J_3^T A_3 J_3 = \begin{bmatrix} 3.388 & 0.072 & 0 \\ 0.072 & 1.776 & -0.005 \\ 0 & -0.005 & -1.164 \end{bmatrix}$$



Bisection or Spectrum-Slicing

- For real symmetric matrix, can determine how many eigenvalues are less than given real number σ
- By systematically choosing various values for σ (*slicing spectrum* at σ) and monitoring resulting count, any eigenvalue can be isolated as accurately as desired
- For example, symmetric indefinite factorization $A = LDL^T$ makes *inertia* (numbers of positive, negative, and zero eigenvalues) of symmetric matrix A easy to determine
- By applying factorization to matrix $A - \sigma I$ for various values of σ , individual eigenvalues can be isolated as accurately as desired using interval bisection technique



Sturm Sequence

- Another spectrum-slicing method for computing individual eigenvalues is based on *Sturm sequence* property of symmetric matrices
- Let A be symmetric matrix and let $p_r(\sigma)$ denote determinant of leading principal minor of order r of $A - \sigma I$
- Then zeros of $p_r(\sigma)$ strictly separate those of $p_{r-1}(\sigma)$, and number of agreements in sign of successive members of sequence $p_r(\sigma)$, for $r = 1, \dots, n$, equals number of eigenvalues of A strictly greater than σ
- Determinants $p_r(\sigma)$ are easy to compute if A is transformed to tridiagonal form before applying Sturm sequence technique



Relatively Robust Representation

- With conventional methods, cost of computing eigenvalues of symmetric tridiagonal matrix is $\mathcal{O}(n^2)$, but if orthogonal eigenvectors are also computed, then cost rises to $\mathcal{O}(n^3)$
- Another possibility is to compute eigenvalues first at $\mathcal{O}(n^2)$ cost, and then compute corresponding eigenvectors separately using inverse iteration with computed eigenvalues as shifts
- Key to making this idea work is computing eigenvalues and corresponding eigenvectors to very high relative accuracy so that expensive explicit orthogonalization of eigenvectors is not needed
- RRR algorithm exploits this approach to produce eigenvalues *and* orthogonal eigenvectors at $\mathcal{O}(n^2)$ cost



QZ Algorithm

- If A and B are triangular, then eigenvalues are given by $\lambda_i = a_{ii}/b_{ii}$, for $b_{ii} \neq 0$
- *QZ algorithm* reduces A and B simultaneously to upper triangular form by orthogonal transformations
- First, B is reduced to upper triangular form by orthogonal transformation from left, which is also applied to A
- Next, transformed A is reduced to upper Hessenberg form by orthogonal transformation from left, while maintaining triangular form of B , which requires additional transformations from right



Computing SVD

- *Singular values* of A are nonnegative square roots of eigenvalues of $A^T A$, and columns of U and V are orthonormal eigenvectors of AA^T and $A^T A$, respectively
- Algorithms for computing SVD work directly with A , without forming AA^T or $A^T A$, to avoid loss of information associated with forming these matrix products explicitly
- SVD is usually computed by variant of QR iteration, with A first reduced to bidiagonal form by orthogonal transformations, then remaining off-diagonal entries are annihilated iteratively
- SVD can also be computed by variant of Jacobi method, which can be useful on parallel computers or if matrix has special structure



Divide-and-Conquer Method

- Express symmetric tridiagonal matrix T as

$$T = \begin{bmatrix} T_1 & O \\ O & T_2 \end{bmatrix} + \beta uu^T$$

- Can now compute eigenvalues and eigenvectors of smaller matrices T_1 and T_2
- To relate these back to eigenvalues and eigenvectors of original matrix requires solution of *secular equation*, which can be done reliably and efficiently
- Applying this approach recursively yields *divide-and-conquer* algorithm for symmetric tridiagonal eigenproblems



Generalized Eigenvalue Problems

- *Generalized eigenvalue problem* has form

$$Ax = \lambda Bx$$

where A and B are given $n \times n$ matrices

- If either A or B is nonsingular, then generalized eigenvalue problem can be converted to standard eigenvalue problem, either

$$(B^{-1}A)x = \lambda x \quad \text{or} \quad (A^{-1}B)x = (1/\lambda)x$$

- This is not recommended, since it may cause
 - loss of accuracy due to rounding error
 - loss of symmetry if A and B are symmetric
- Better alternative for generalized eigenvalue problems is QZ algorithm



QZ Algorithm, continued

- Finally, analogous to QR iteration, A is reduced to triangular form while still maintaining triangular form of B , which again requires transformations from both sides
- Eigenvalues can now be determined from mutually triangular form, and eigenvectors can be recovered from products of left and right transformations, denoted by Q and Z

