

Outline

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



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



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* $= \rho(A) = \max\{|\lambda| : \lambda \in \lambda(A)\}$



Classic Eigenvalue Problem

- Consider the coupled pair of differential equations:

$$\begin{aligned}\frac{dv}{dt} &= 4v - 5w, & v &= 8 \text{ at } t = 0, \\ \frac{dw}{dt} &= 2v - 3w, & w &= 5 \text{ at } t = 0.\end{aligned}$$

- This is an *initial-value problem*.
- With the coefficient matrix,

$$A = \begin{bmatrix} 4 & -5 \\ 2 & -3 \end{bmatrix},$$

we can write this as,

$$\frac{d}{dt} \begin{pmatrix} v(t) \\ w(t) \end{pmatrix} = \begin{bmatrix} 4 & -5 \\ 2 & -3 \end{bmatrix} \begin{pmatrix} v(t) \\ w(t) \end{pmatrix}.$$

- Introducing the *vector unknown*, $\mathbf{u}(t) := [v(t) \ w(t)]^T$ with $\mathbf{u}(0) = [8 \ 5]^T$, we can write the system in vector form,

$$\frac{d\mathbf{u}}{dt} = A\mathbf{u}, \text{ with } \mathbf{u} = \mathbf{u}(0) \text{ at } t = 0.$$

- How do we find $\mathbf{u}(t)$?

- If we had a 1×1 matrix $A = a$, we would have a scalar equation:

$$\frac{du}{dt} = a u \text{ with } u = u(0) \text{ at } t = 0.$$

The solution to this equation is a pure exponential:

$$u(t) = e^{at} u(0),$$

which satisfies the initial condition because $e^0 = 1$.

- The derivative with respect to t is $ae^{at}u(0) = au$, so it satisfies the scalar initial value problem.
- The constant a is critical to how this system behaves.
 - If $a > 0$ then the solution grows in time.
 - If $a < 0$ then the solution decays.
 - If $a \in Im$ then the solution is oscillatory.
(More on this later...)

- Coming back to our system, suppose we again look for solutions that are pure exponentials in time, e.g.,

$$\begin{aligned}v(t) &= e^{\lambda t}y \\w(t) &= e^{\lambda t}z.\end{aligned}$$

- If this is to be a solution to our initial value problem, we require

$$\begin{aligned}\frac{dv}{dt} &= \lambda e^{\lambda t}y = 4e^{\lambda t}y - 5e^{\lambda t}z \\ \frac{dw}{dt} &= \lambda e^{\lambda t}z = 2e^{\lambda t}y - 3e^{\lambda t}z.\end{aligned}$$

- The $e^{\lambda t}$ cancels out from each side, leaving:

$$\begin{aligned}\lambda y &= 4y - 5z \\ \lambda z &= 2y - 3z,\end{aligned}$$

which is the eigenvalue problem.

- In vector form, $\mathbf{u}(t) = e^{\lambda t} \mathbf{x}$, yields

$$\frac{d\mathbf{u}}{dt} = A\mathbf{u} \iff \lambda e^{\lambda t} \mathbf{x} = A(e^{\lambda t} \mathbf{x})$$

which gives the eigenvalue problem in matrix form:

$$\lambda \mathbf{x} = A\mathbf{x} \text{ or}$$

$$A\mathbf{x} = \lambda \mathbf{x}.$$

- As in the scalar case, the solution behavior depends on whether λ has
 - positive real part \longrightarrow a growing solution,
 - negative real part \longrightarrow a decaying solution,
 - an imaginary part \longrightarrow an oscillating solution.
- Note that here we have two unknowns: λ and \mathbf{x} .
- We refer to (λ, \mathbf{x}) as an eigenpair, with *eigenvalue* λ and *eigenvector* \mathbf{x} .

Solving the Eigenvalue Problem

- The eigenpair satisfies

$$(A\mathbf{x} - \lambda I)\mathbf{x} = 0,$$

which is to say,

- \mathbf{x} is in the null-space of $A - \lambda I$
- λ is chosen so that $A - \lambda I$ has a null-space.
- We thus seek λ such that $A - \lambda I$ is singular.
- Singularity implies $\det(A - \lambda I)=0$.
- For our example:

$$0 = \begin{vmatrix} 4 - \lambda & -5 \\ 2 & -3 - \lambda \end{vmatrix} = (4 - \lambda)(-3 - \lambda) - (-5)(2),$$

or

$$\lambda^2 - \lambda - 2 = 0,$$

which has roots $\lambda = -1$ or $\lambda = 2$.

Finding the Eigenvectors

- For the case $\lambda = \lambda_1 = -1$, $(A - \lambda_1 I)\mathbf{x}_1$ satisfies,

$$\begin{bmatrix} 5 & -5 \\ 2 & -2 \end{bmatrix} \begin{pmatrix} y \\ z \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \end{pmatrix},$$

which gives us the eigenvector \mathbf{x}_1

$$\mathbf{x}_1 = \begin{pmatrix} y \\ z \end{pmatrix} = \begin{pmatrix} 1 \\ 1 \end{pmatrix}.$$

- **Note that any nonzero multiple of \mathbf{x}_1 is also an eigenvector.**
- Thus, \mathbf{x}_1 defines a *subspace* that is invariant under multiplication by A .

- For the case $\lambda = \lambda_2 = 2$, $(A - \lambda_2 I)\mathbf{x}_2$ satisfies,

$$\begin{bmatrix} 2 & -5 \\ 2 & -5 \end{bmatrix} \begin{pmatrix} y \\ z \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \end{pmatrix},$$

which gives us the second eigenvector as any multiple of

$$\mathbf{x}_2 = \begin{pmatrix} y \\ z \end{pmatrix} = \begin{pmatrix} 5 \\ 2 \end{pmatrix}.$$

Return to Model Problem

- Note that our model problem $\frac{d\mathbf{u}}{dt} = A\mathbf{u}$, is *linear* in the unknown \mathbf{u} .
- Thus, if we have two solutions $\mathbf{u}_1(t)$ and $\mathbf{u}_2(t)$ satisfying the differential equation, their sum $\mathbf{u} := \mathbf{u}_1 + \mathbf{u}_2$ also satisfies the equation:

$$\begin{array}{r} \frac{d\mathbf{u}_1}{dt} = A\mathbf{u}_1 \\ + \frac{d\mathbf{u}_2}{dt} = A\mathbf{u}_2 \\ \hline \frac{d}{dt}(\mathbf{u}_1 + \mathbf{u}_2) = A(\mathbf{u}_1 + \mathbf{u}_2) \\ \frac{d\mathbf{u}}{dt} = A\mathbf{u} \end{array}$$

- Take $\mathbf{u}_1 = c_1 e^{\lambda_1 t} \mathbf{x}_1$:

$$\begin{aligned} \frac{d\mathbf{u}_1}{dt} &= c_1 \lambda_1 e^{\lambda_1 t} \mathbf{x}_1 \\ A\mathbf{u}_1 &= A(c_1 e^{\lambda_1 t} \mathbf{x}_1) \\ &= c_1 e^{\lambda_1 t} A\mathbf{x}_1 \\ &= c_1 e^{\lambda_1 t} \lambda_1 \mathbf{x}_1 \\ &= \frac{d\mathbf{u}_1}{dt}. \end{aligned}$$

- Similarly, for $\mathbf{u}_2 = c_2 e^{\lambda_2 t} \mathbf{x}_2$:

$$\frac{d\mathbf{u}_2}{dt} = A\mathbf{u}_2.$$

- Thus,

$$\frac{d\mathbf{u}}{dt} = \frac{d}{dt}(\mathbf{u}_1 + \mathbf{u}_2) = A(\mathbf{u}_1 + \mathbf{u}_2)$$

$$\mathbf{u} = c_1 e^{\lambda_1 t} \mathbf{x}_1 + c_2 e^{\lambda_2 t} \mathbf{x}_2.$$

- The only remaining part is to find the coefficients c_1 and c_2 such that $\mathbf{u} = \mathbf{u}(0)$ at time $t = 0$.
- This initial condition yields a 2×2 system,

$$\begin{bmatrix} \mathbf{x}_1 & \mathbf{x}_2 \end{bmatrix} \begin{pmatrix} c_1 \\ c_2 \end{pmatrix} = \begin{pmatrix} 8 \\ 5 \end{pmatrix}.$$

- Solving for c_1 and c_2 via Gaussian elimination:

$$\begin{bmatrix} 1 & 5 \\ 1 & 2 \end{bmatrix} \begin{pmatrix} c_1 \\ c_2 \end{pmatrix} = \begin{pmatrix} 8 \\ 5 \end{pmatrix}$$

$$\begin{bmatrix} 1 & 5 \\ 0 & -3 \end{bmatrix} \begin{pmatrix} c_1 \\ c_2 \end{pmatrix} = \begin{pmatrix} 8 \\ -3 \end{pmatrix}$$

$$c_2 = 1$$

$$c_1 = 8 - 5c_2 = 3.$$

- So, our solution is $\mathbf{u}(t) = \mathbf{x}_1 c_1 e^{\lambda_1 t} + \mathbf{x}_2 c_2 e^{\lambda_2 t}$

$$= \begin{pmatrix} 1 \\ 1 \end{pmatrix} 3e^{-t} + \begin{pmatrix} 5 \\ 2 \end{pmatrix} e^{2t}.$$

- Clearly, after a long time, the solution is going to look like a multiple of $\mathbf{x}_2 = [5 \ 2]^T$ because the component of the solution parallel to \mathbf{x}_1 will decay.
- (More precisely, the component parallel to \mathbf{x}_1 will not grow as fast as the component parallel to \mathbf{x}_2 .)

- Solving for c_1 and c_2 via Gaussian elimination:

$$\begin{bmatrix} 1 & 5 \\ 1 & 2 \end{bmatrix} \begin{pmatrix} c_1 \\ c_2 \end{pmatrix} = \begin{pmatrix} 8 \\ 5 \end{pmatrix}$$

$$\begin{bmatrix} 1 & 5 \\ 0 & -3 \end{bmatrix} \begin{pmatrix} c_1 \\ c_2 \end{pmatrix} = \begin{pmatrix} 8 \\ -3 \end{pmatrix}$$

$$c_2 = 1$$

$$c_1 = 8 - 5c_2 = 3.$$

- So, our solution is $\mathbf{u}(t) = \mathbf{x}_1 c_1 e^{\lambda_1 t} + \mathbf{x}_2 c_2 e^{\lambda_2 t}$

$$= \begin{pmatrix} 1 \\ 1 \end{pmatrix} 3e^{-t} + \begin{pmatrix} 5 \\ 2 \end{pmatrix} e^{2t}.$$

- Clearly, after a long time, the solution is going to look like a multiple of $\mathbf{x}_2 = [5 \ 2]^T$ because the component of the solution parallel to \mathbf{x}_1 will decay.
- (More precisely, the component parallel to \mathbf{x}_1 will not grow as fast as the component parallel to \mathbf{x}_2 .)

Growing / Decaying Modes

- Our model problem,

$$\frac{d\mathbf{u}}{dt} = A\mathbf{u} \longrightarrow \mathbf{u}(t) = \mathbf{x}_1 c_1 e^{\lambda_1 t} + \mathbf{x}_2 c_2 e^{\lambda_2 t}$$

leads to growth/decay of components.

- Also get growth/decay through matrix-vector products.
- Consider $\mathbf{u} = c_1 \mathbf{x}_1 + c_2 \mathbf{x}_2$.

$$\begin{aligned} A\mathbf{u} &= c_1 A\mathbf{x}_1 + c_2 A\mathbf{x}_2 \\ &= c_1 \lambda_1 \mathbf{x}_1 + c_2 \lambda_2 \mathbf{x}_2 \end{aligned}$$

$$\begin{aligned} A^k \mathbf{u} &= c_1 \lambda_1^k \mathbf{x}_1 + c_2 \lambda_2^k \mathbf{x}_2 \\ &= \lambda_2^k \left[c_1 \left(\frac{\lambda_1}{\lambda_2} \right)^k \mathbf{x}_1 + c_2 \mathbf{x}_2 \right]. \end{aligned}$$

$$\lim_{k \rightarrow \infty} A^k \mathbf{u} = \lambda_2^k [c_1 \cdot 0 \cdot \mathbf{x}_1 + c_2 \mathbf{x}_2] = c_2 \lambda_2^k \mathbf{x}_2.$$

- So, repeated matrix-vector products lead to emergence of eigenvector associated with the eigenvalue λ that has largest modulus.
- This is the main idea behind the ***power method***, which is a common way to find the eigenvector associated with $\max |\lambda|$.

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



Examples: Eigenvalues and Eigenvectors

- $A = \begin{bmatrix} 1 & 0 \\ 0 & 2 \end{bmatrix}$: $\lambda_1 = 1$, $\mathbf{x}_1 = \begin{bmatrix} 1 \\ 0 \end{bmatrix}$, $\lambda_2 = 2$, $\mathbf{x}_2 = \begin{bmatrix} 0 \\ 1 \end{bmatrix}$
- $A = \begin{bmatrix} 1 & 1 \\ 0 & 2 \end{bmatrix}$: $\lambda_1 = 1$, $\mathbf{x}_1 = \begin{bmatrix} 1 \\ 0 \end{bmatrix}$, $\lambda_2 = 2$, $\mathbf{x}_2 = \begin{bmatrix} 1 \\ 1 \end{bmatrix}$
- $A = \begin{bmatrix} 3 & -1 \\ -1 & 3 \end{bmatrix}$: $\lambda_1 = 2$, $\mathbf{x}_1 = \begin{bmatrix} 1 \\ 1 \end{bmatrix}$, $\lambda_2 = 4$, $\mathbf{x}_2 = \begin{bmatrix} 1 \\ -1 \end{bmatrix}$
- $A = \begin{bmatrix} 1.5 & 0.5 \\ 0.5 & 1.5 \end{bmatrix}$: $\lambda_1 = 2$, $\mathbf{x}_1 = \begin{bmatrix} 1 \\ 1 \end{bmatrix}$, $\lambda_2 = 1$, $\mathbf{x}_2 = \begin{bmatrix} -1 \\ 1 \end{bmatrix}$
- $A = \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix}$: $\lambda_1 = i$, $\mathbf{x}_1 = \begin{bmatrix} 1 \\ i \end{bmatrix}$, $\lambda_2 = -i$, $\mathbf{x}_2 = \begin{bmatrix} i \\ 1 \end{bmatrix}$

where $i = \sqrt{-1}$



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 $i = \sqrt{-1}$



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

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



Companion Matrix

- Monic polynomial

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

is characteristic polynomial of *companion matrix*

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

- 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



Companion Matrix, n=3

□ Look at determinant of $(C - I \lambda)$: $|C - I \lambda| = 0$

$$C := \begin{pmatrix} 0 & 0 & -c_0 \\ 1 & 0 & -c_1 \\ 0 & 1 & -c_2 \end{pmatrix}$$

$$\begin{aligned} \text{Eigenvalues: } |C - I\lambda| &= \begin{vmatrix} -\lambda & 0 & -c_0 \\ 1 & -\lambda & -c_1 \\ 0 & 1 & (-c_2 - \lambda) \end{vmatrix} \\ &= -c_0 \begin{vmatrix} 1 & -\lambda \\ 0 & 1 \end{vmatrix} + c_1 \begin{vmatrix} -\lambda & 1 \\ 0 & 1 \end{vmatrix} - (c_2 + \lambda) \begin{vmatrix} -\lambda & 0 \\ 1 & -\lambda \end{vmatrix} \\ &= -c_0 - c_1\lambda - c_2\lambda^2 - \lambda^3 = 0 \end{aligned}$$

$$\implies p(\lambda) = c_0 + c_1\lambda + c_2\lambda^2 + \lambda^3 = 0$$

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



Example: Characteristic Polynomial

- Consider

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

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

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

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

which has 1 as double root

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



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*

*Algebraic
Multiplicity*

*Geometric
Multiplicity*

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

where X is nonsingular matrix of eigenvectors

Note: Every matrix is ϵ away from being diagonalizable.



Diagonalization

- The real merit of eigenvalue decomposition is that it simplifies powers of a matrix.

- Consider $X^{-1}AX = D, \text{ diagonal}$

$$AX = XD$$

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

$$\begin{aligned} A^2 &= (XDX^{-1})(XDX^{-1}) \\ &= XD^2X^{-1} \end{aligned}$$

$$\begin{aligned} A^k &= (XDX^{-1})(XDX^{-1}) \cdots (XDX^{-1}) \\ &= XD^kX^{-1} \end{aligned}$$

$$= X \begin{bmatrix} \lambda_1^k & & & \\ & \lambda_2^k & & \\ & & \ddots & \\ & & & \lambda_n^k \end{bmatrix} X^{-1}$$

- High powers of A tend to be dominated by largest eigenpair $(\lambda_1, \underline{x}_1)$, assuming $|\lambda_1| \geq |\lambda_2| \geq \cdots \geq |\lambda_n|$.

Matrix Powers Example

- Consider our 1D finite difference example introduced earlier.

$$\frac{d^2u}{dx^2} = f(x) \longrightarrow -\frac{u_{i-1} - 2u_i + u_{i+1}}{\Delta x^2} \approx f(x_i).$$

where $u(0) = u(1) = 0$ and $\Delta x = 1/(n+1)$.

- In matrix form,

$$A\mathbf{u} = \frac{1}{\Delta x^2} \begin{pmatrix} 2 & -1 & & & \\ -1 & 2 & -1 & & \\ & -1 & \ddots & \ddots & \\ & & \ddots & \ddots & -1 \\ & & & -1 & 2 \end{pmatrix} \begin{pmatrix} u_1 \\ u_2 \\ \vdots \\ \vdots \\ u_m \end{pmatrix} = \begin{pmatrix} f_1 \\ f_2 \\ \vdots \\ \vdots \\ f_m \end{pmatrix}$$

- Eigenvectors and eigenvalues have closed-form expression:

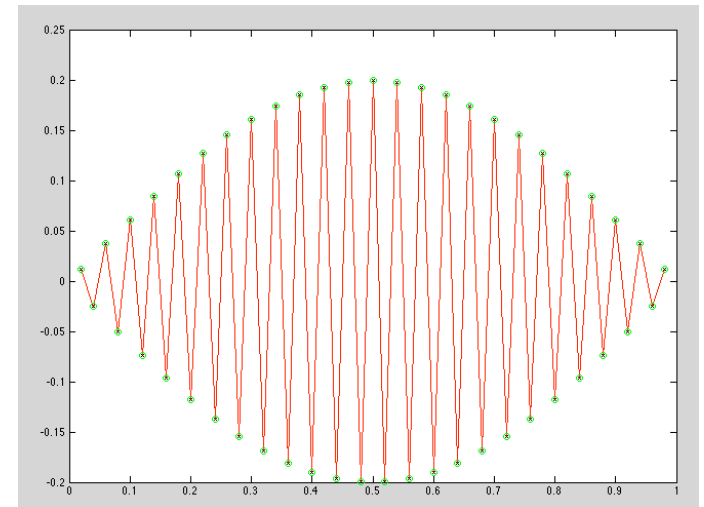
$$(\mathbf{z}_k)_i = \sin k\pi x_i = \sin k\pi i\Delta x \qquad \lambda_k = \frac{2}{\Delta x^2} (1 - \cos k\pi\Delta x)$$

- Eigenvalues are in the interval $\sim [\pi^2, 4(n+1)^2]$.

Matlab Example: heat_demo.m

- ❑ Repeatedly applying A to a random input vector reveals the eigenvalue of maximum modulus.
- ❑ This idea leads to one of the most common (but not most efficient) ways of finding an eigenvalue/vector pair, called the power method.

```
h = 1/(n+1);  
e = ones(n,1);  
A = spdiags([-e 2*e -e],[-1:1, n,n)/(h*h);  
  
x=1:n; x=h*x';  
  
z=rand(n,1); hold off;  
for k=1:2000;  
    z=A*z; z=z/norm(z);  
    if mod(k,100)==0; plot(x,z,'r-'); pause(.1); end;  
end;
```



Diagonalization

- Note that if we define $A^0 = I$, we have any polynomial of A defined as

$$p_k(A)\underline{x} = X \begin{bmatrix} p_k(\lambda_1) & & & \\ & p_k(\lambda_2) & & \\ & & \ddots & \\ & & & p_k(\lambda_n) \end{bmatrix} X^{-1}\underline{x}.$$

- We can further extend this to other functions,

$$f(A)\underline{x} = X \begin{bmatrix} f(\lambda_1) & & & \\ & f(\lambda_2) & & \\ & & \ddots & \\ & & & f(\lambda_n) \end{bmatrix} X^{-1}\underline{x}.$$

- For example, the solution to $f(A)\underline{x} = \underline{b}$ is would be

$$\underline{x} = X [f(D)]^{-1} X^{-1}\underline{b}.$$

- The diagonalization concept is very powerful because it transforms *systems* of equations into scalar equations.

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* $= \mathcal{S}_\lambda = \{x : Ax = \lambda x\}$
- Subspace \mathcal{S} of \mathbb{R}^n (or \mathbb{C}^n) is *invariant* if $A\mathcal{S} \subseteq \mathcal{S}$
- For eigenvectors $x_1 \cdots x_p$, $\text{span}([x_1 \cdots x_p])$ is invariant subspace

In theory...



Relevant Properties of Matrices

- Properties of matrix \mathbf{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	$\mathbf{A}^T \mathbf{A} = \mathbf{A} \mathbf{A}^T = \mathbf{I}$
unitary	$\mathbf{A}^H \mathbf{A} = \mathbf{A} \mathbf{A}^H = \mathbf{I}$
symmetric	$\mathbf{A} = \mathbf{A}^T$ Skew symmetric: $\mathbf{A} = -\mathbf{A}^T$
Hermitian	$\mathbf{A} = \mathbf{A}^H$
normal	$\mathbf{A}^H \mathbf{A} = \mathbf{A} \mathbf{A}^H$



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: $\begin{bmatrix} 1 & 2 \\ 2 & 3 \end{bmatrix}$ Skew-Symmetric: $\begin{bmatrix} 0 & -2 \\ 2 & 0 \end{bmatrix} = -\begin{bmatrix} 0 & -2 \\ 2 & 0 \end{bmatrix}^T$
- Nonsymmetric: $\begin{bmatrix} 1 & 3 \\ 2 & 4 \end{bmatrix}$
- Hermitian: $\begin{bmatrix} 1 & 1+i \\ 1-i & 2 \end{bmatrix}$
- NonHermitian: $\begin{bmatrix} 1 & 1+i \\ 1+i & 2 \end{bmatrix}$



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: $\begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix}$



Normal Matrices

Normal matrices have orthogonal eigenvectors, so $\underline{x}_i^H \underline{x}_j = \delta_{ij}$

$$X^T = X^{-1}$$

$$A = XDX^H$$

Normal matrices include

- symmetric ($A = A^T$)
- skew-symmetric ($A = -A^T$)
- unitary ($U^H U = I$)
- circulant (periodic+Toeplitz)
- others ...

$$\begin{bmatrix} 1 & 2 & 0 \\ 0 & 1 & 2 \\ 2 & 0 & 1 \end{bmatrix}$$

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? *Sparsity*



Sparsity

- ❑ Sparsity, either direct or implied, is a big driver in choice of eigenvalue solvers.
- ❑ Typically, only $O(n)$ entries in entire matrix, where $n \sim 10^9$ — 10^{18} might be anticipated.
- ❑ Examples include Big Data (e.g., google page rank) and physics simulations (fluid, heat transfer, electromagnetics, fusion, etc.).
- ❑ Usually, need only a few eigenvectors / eigenvalues, $k \ll n$.
- ❑ Often, there are special properties of A that make it difficult to create A . Instead, work strictly with matrix-vector products

$$\mathbf{y} = \mathbf{A} \mathbf{x}$$

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 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



Conditioning of Eigenvalues

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

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

where \mathbf{x} and \mathbf{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



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



Similarity Transformation

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

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

- Then

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

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

- Similarity transformations preserve eigenvalues and eigenvectors are easily recovered



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



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



Simple non-diagonalizable example, 2×2 Jordan block:

$$\begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} = \lambda \begin{pmatrix} x_1 \\ x_2 \end{pmatrix}$$

$$\begin{vmatrix} 1 - \lambda & 1 \\ 0 & 1 - \lambda \end{vmatrix} = (1 - \lambda)^2 = 0$$

Only one eigenvector: $\underline{x} = \begin{pmatrix} 1 \\ 0 \end{pmatrix}$

$$\begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix} \begin{pmatrix} 1 \\ 0 \end{pmatrix} = \begin{pmatrix} 1 \\ 0 \end{pmatrix}$$

3×3 Non-Diagonalizable Example

$$A = \begin{bmatrix} 2 & & \\ & 2 & \\ & & 2 \end{bmatrix}, \quad B = \begin{bmatrix} 2 & 1 & \\ & 2 & 1 \\ & & 2 \end{bmatrix}.$$

- Characteristic polynomial is $(\lambda - 2)^3$ for both A and B .
- Algebraic multiplicity is 3.
- For A , three eigenvectors. Say, \mathbf{e}_1 , \mathbf{e}_2 , and \mathbf{e}_3 .
- For B , only one eigenvector ($\alpha\mathbf{e}_1$), so geometric multiplicity of B is 1.

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

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

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



Eigenvectors / Eigenvalues of Upper Triangular Matrix

Suppose A is upper triangular

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

Then

$$0 = (A - \lambda I)\underline{x} = \begin{bmatrix} U_{11} & \underline{u} & U_{13} \\ 0 & 0 & \underline{v}^T \\ O & 0 & U_{33} \end{bmatrix} \begin{pmatrix} \underline{y} \\ -1 \\ 0 \end{pmatrix} = \begin{pmatrix} U_{11}\underline{y} - \underline{u} \\ 0 \\ 0 \end{pmatrix}$$

$(A - \lambda I) \quad \underline{x} \quad 0$

Because U_{11} is nonsingular, can solve $U_{11}\underline{y} = \underline{u}$ to find eigenvector \underline{x} .

Block Triangular Form

- If

$$\mathbf{A} = \begin{bmatrix} \mathbf{A}_{11} & \mathbf{A}_{12} & \cdots & \mathbf{A}_{1p} \\ & \mathbf{A}_{22} & \cdots & \mathbf{A}_{2p} \\ & & \ddots & \vdots \\ & & & \mathbf{A}_{pp} \end{bmatrix}$$

with *square* diagonal blocks, then

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

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



Eigenvalue-Revealing Factorizations

- Diagonalization: $A = X\Lambda X^{-1}$ if A is nondefective.
- Unitary diagonalization: $A = Q\Lambda Q^*$ if A is normal.
- Unitary triangularization: $A = QTQ^*$ always exists.
(T upper triangular.)

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



Similarity Transformations

- Given

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

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

- If A is normal ($AA^H = A^H A$),

$$A = Q \Lambda Q^H$$

B is diagonal, T is unitary ($T^{-1} = T^H$).

- If A is symmetric real,

$$A = Q \Lambda Q^T$$

B is diagonal, T is orthogonal ($T^{-1} = T^T$).

- If B is diagonal, T is the matrix of eigenvectors.

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



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

$$\begin{aligned} x_k &= Ax_{k-1} = A^2 x_{k-2} = \cdots = A^k x_0 = \\ &= \sum_{i=1}^n \lambda_i^k \alpha_i v_i = \lambda_1^k \left(\alpha_1 v_1 + \sum_{i=2}^n (\lambda_i / \lambda_1)^k \alpha_i v_i \right) \end{aligned}$$

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



2 x 2 Example

$$A = \begin{bmatrix} 1.5 & 0.5 \\ 0.5 & 1.5 \end{bmatrix}$$

$$D = \begin{bmatrix} 1 & \\ & 2 \end{bmatrix} \quad X = \begin{bmatrix} -1/\sqrt{2} & 1/\sqrt{2} \\ 1/\sqrt{2} & 1/\sqrt{2} \end{bmatrix}$$

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

$$\underline{x}_k = A \underline{x}_{k-1}$$

k	\underline{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



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
This issue corrected by simultaneous iteration.
- For real matrix and starting vector, iteration can never converge to complex vector



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

$$\begin{aligned} \mathbf{y}_k &= \mathbf{A}\mathbf{x}_{k-1} \\ \mathbf{x}_k &= \mathbf{y}_k / \|\mathbf{y}_k\|_\infty \end{aligned}$$

- With normalization, $\|\mathbf{y}_k\|_\infty \rightarrow |\lambda_1|$, and $\mathbf{x}_k \rightarrow \mathbf{v}_1 / \|\mathbf{v}_1\|_\infty$



Example: Normalized Power Iteration

- Repeating previous example with normalized scheme,

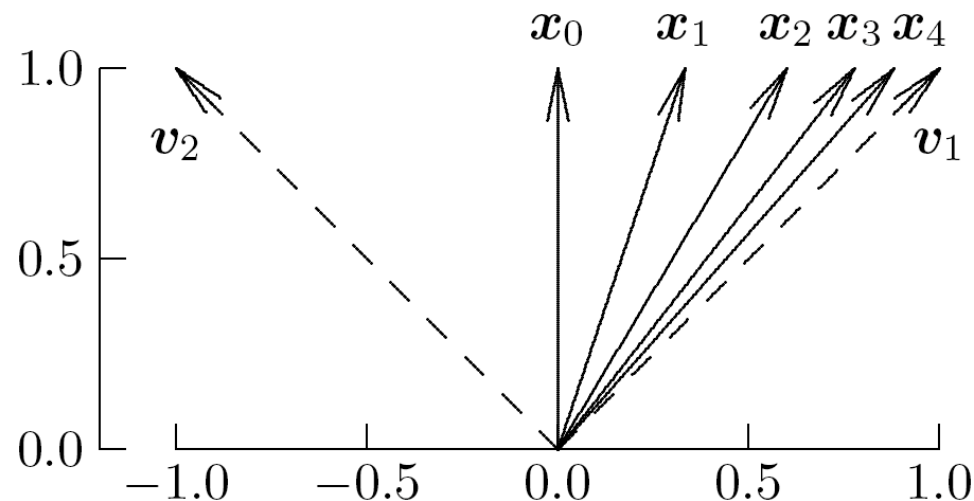
k	\mathbf{x}_k^T		$\ \mathbf{y}_k\ _\infty$
0	0.000	1.0	
1	0.333	1.0	1.500
2	0.600	1.0	1.667
3	0.778	1.0	1.800
4	0.882	1.0	1.889
5	0.939	1.0	1.941
6	0.969	1.0	1.970
7	0.984	1.0	1.985
8	0.992	1.0	1.992

< interactive example >



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



Convergence Rate of Power Iteration

Convergence rate of power iteration depends on relative separation of λ_1 and λ_2 .

Assuming $c_1 \neq 0$ and $|\lambda_1| > |\lambda_j|$, $j > 1$, we have

$$\begin{aligned} A^k \underline{x} &= \sum_{j=1}^n \underline{x}_j \lambda_j^k c_j \\ &= \lambda_1^k c_1 \left[\underline{x}_1 + \sum_{j=2}^n \underline{x}_j \frac{\lambda_j^k}{\lambda_1^k} \frac{c_j}{c_1} \right] \\ &\sim \underline{x}_1 \lambda_1^k c_1 \text{ as } k \longrightarrow \infty \\ &\sim \lambda_1^k c_1 \left[\underline{x}_1 + \underline{x}_2 \left(\frac{\lambda_2}{\lambda_1} \right)^k \frac{c_2}{c_1} \right] \end{aligned}$$

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

Note: in general, need $\max_j |\lambda_j - \sigma| / |\lambda_1 - \sigma| < |\lambda_2| / |\lambda_1|$.



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

Idea: pick shifts at each iteration as estimate of spectrum emerges.



Power Iteration with Shift

$$(A - \sigma I)\underline{x} = \lambda \underline{x} - \sigma \underline{x} = (\lambda - \sigma)\underline{x} = \mu \underline{x}$$

If $\lambda_k \in \{1.9 \dots .1\}$, then

$$\frac{\lambda_2}{\lambda_1} = 0.9$$

If $\sigma = 0.4$, then $\mu_k \in \{.5 \ .4 \dots - .4\}$ and

$$\frac{\mu_2}{\mu_1} = 0.8,$$

so about twice the convergence rate.

Shifted power iteration, however, is somewhat limited.

The real power derives from *inverse* power iterations with shifts.

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



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



Example: Inverse Iteration

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

k	\mathbf{x}_k^T		$\ \mathbf{y}_k\ _\infty$
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 >



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



- Power Iteration: $\mathbf{x} = A^k \mathbf{x} \longrightarrow c\mathbf{x}_1$
- Normalized Power Iteration:
$$\left. \begin{array}{l} \mathbf{y} = A\mathbf{x} \\ \mathbf{x} = \mathbf{y}/\|\mathbf{y}\| \end{array} \right\} \begin{array}{l} \|\mathbf{y}\| \longrightarrow |\lambda_1| \\ \mathbf{x} \longrightarrow \mathbf{x}_1 \end{array}$$
- Inverse Iteration:
$$\left. \begin{array}{l} \mathbf{y} = A^{-1}\mathbf{x} \\ \mathbf{x} = \mathbf{y}/\|\mathbf{y}\| \end{array} \right\} \begin{array}{l} \|\mathbf{y}\| \longrightarrow |\lambda_n|^{-1} \\ \mathbf{x} \longrightarrow \mathbf{x}_n \end{array}$$
- Inverse Iteration with shift:
$$\left. \begin{array}{l} M = A - \sigma I \\ \mathbf{y} = M^{-1}\mathbf{x} \\ \mathbf{x} = \mathbf{y}/\|\mathbf{y}\| \end{array} \right\} \begin{array}{l} \|\mathbf{y}\| \longrightarrow |\mu_k| = \max |\lambda_k - \sigma|^{-1} \\ \mathbf{x} \longrightarrow \mathbf{x}_k \end{array}$$

Inverse iteration with shift can be arbitrarily fast since separation ratio can be 0.

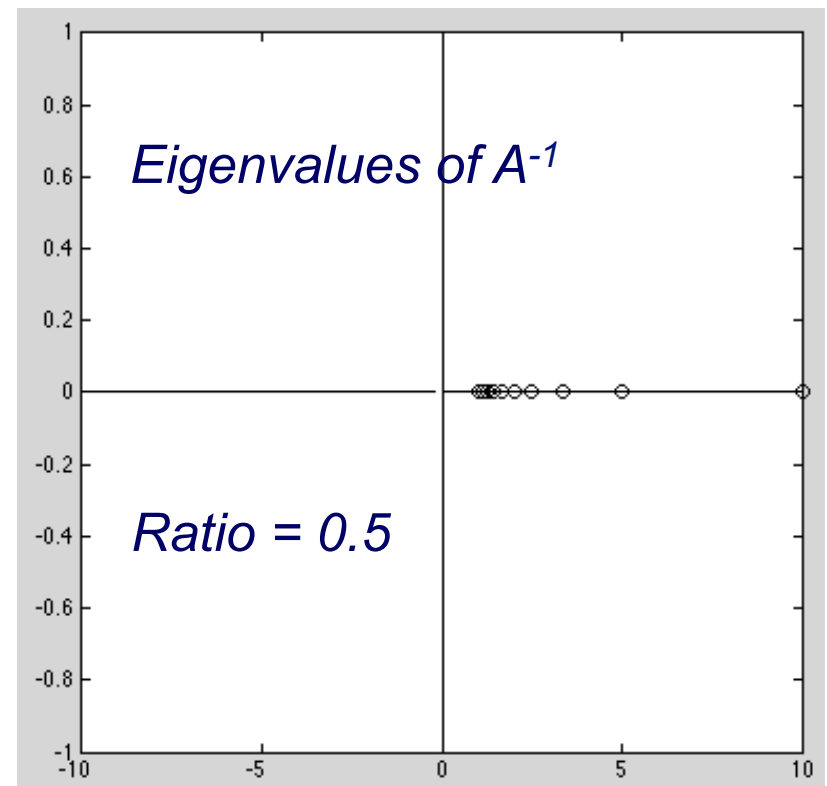
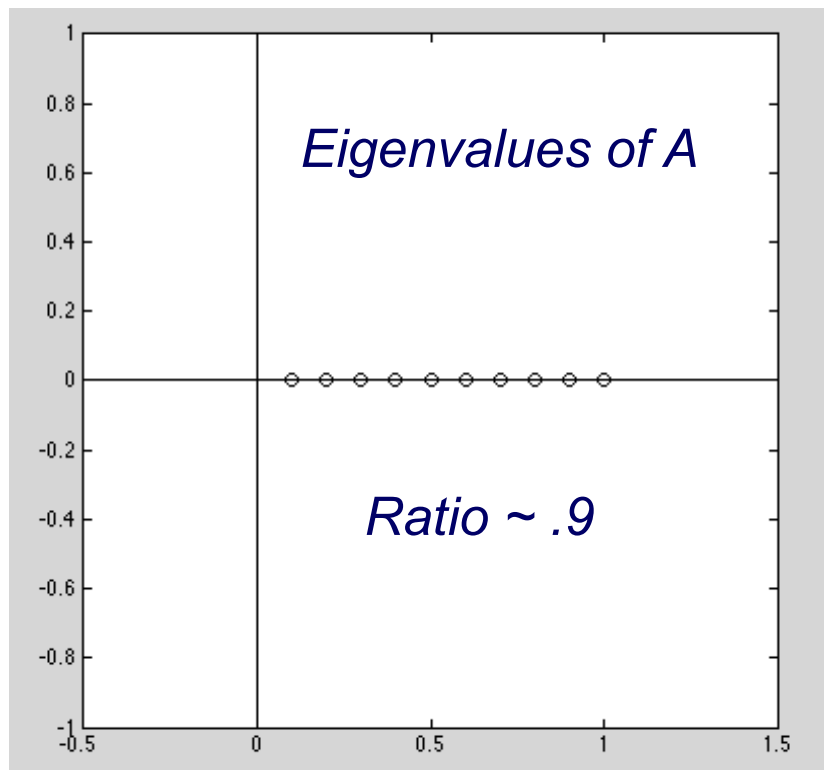
- Inverse Iteration:

$$\begin{aligned}\mathbf{x}^k &= A^{-k} \mathbf{x}^0 \\ &= \sum_{j=1}^n \mathbf{x}_j \left(\frac{1}{\lambda_j} \right)^k c_j \\ &= \left(\frac{1}{\lambda_n} \right)^k c_n \left[\mathbf{x}_n + \sum_{j=1}^{n-1} \mathbf{x}_j \left(\frac{\lambda_n}{\lambda_j} \right)^k c_j \right].\end{aligned}$$

eig_inv_demo.m

Inverse Iteration Illustration

- With shift and invert, can get significant ratios of dominant eigenvalue



- Inverse Iteration with Shift: Let

$$M := A - \sigma I$$

$$\mu_j := \lambda_j - \sigma, \text{ and}$$

$$l \text{ such that } |\sigma - \mu_l| < |\sigma - \mu_j|, j \neq l.$$

Then,

$$\begin{aligned} \mathbf{x}^k &= M^{-k} \mathbf{x}^0 \\ &= \sum_{j=1}^n \mathbf{x}_j \left(\frac{1}{\mu_j} \right)^k c_j \\ &= \left(\frac{1}{\mu_l} \right)^k c_l \left[\mathbf{x}_l + \sum_{j \neq l} \mathbf{x}_j \left(\frac{\mu_l}{\mu_j} \right)^k c_j \right]. \end{aligned}$$

- Using current approximation to λ_l can select $|\sigma - \lambda_l| = |\mu_l|$ to be small.
(Cannot do the same with shifted power iteration.)
- Blow-up is contained by normalizing after each iteration.

eig_shift_invert.m

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 Convergence

- Assume A is a symmetric matrix and \mathbf{x} is close to an eigenvector,

$$\mathbf{x} = c_1 \mathbf{x}_1 + \sum_{j=2}^n c_j \mathbf{x}_j, \quad |c_j| = \epsilon \text{ for } j \geq 2$$

- If $\|\mathbf{x}\| = 1$, then the Rayleigh quotient

$$\begin{aligned} r(\mathbf{x}) &= \frac{\mathbf{x}^T A \mathbf{x}}{\mathbf{x}^T \mathbf{x}} = \mathbf{x}^T A \mathbf{x} \\ &= \sum_{i=1}^n \sum_{j=1}^n c_i c_j \lambda_i \mathbf{x}_i^T \mathbf{x}_j \\ &= \sum_{i=1}^n c_i^2 \lambda_i \\ &= (1 - n\epsilon^2) \lambda_1 + \epsilon^2 \sum_{i=2}^n \lambda_i \end{aligned}$$

- ❑ Rayleigh quotient is stationary at eigenvectors.
- ❑ It takes on min/max when $j=n$ or 1 .

```
% Plot Rayleigh Quotient on S1 for 2 x 2 matrix

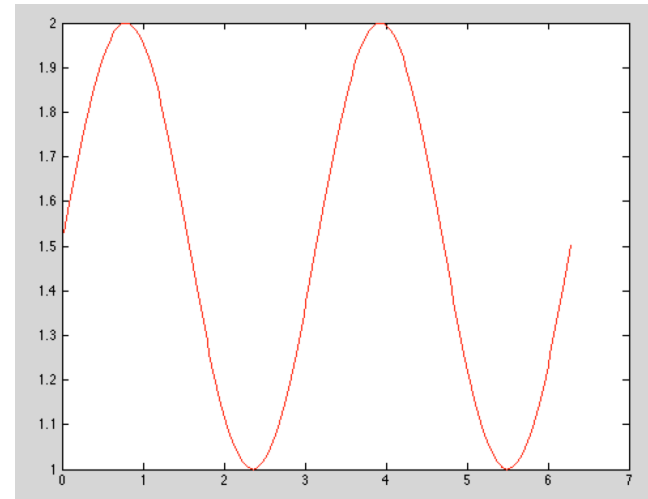
V=[ 1  1 ; -1  1 ];
D=[ 1  0 ;  0  2 ];
for k=1:2; V(:,k)=V(:,k)/norm(V(:,k)); end;

A=V*D*inv(V);

m=200; dt=2*pi/m;

for k=1:m; t=k*dt;
    x=[ cos(t) ; sin(t) ];
    z=A*x;
    rq(k)=z'*x;
    th(k)=t;
end;

plot(th,rq,'r-');
```



```
% Plot Rayleigh Quotient on S2 for 3 x 3 matrix

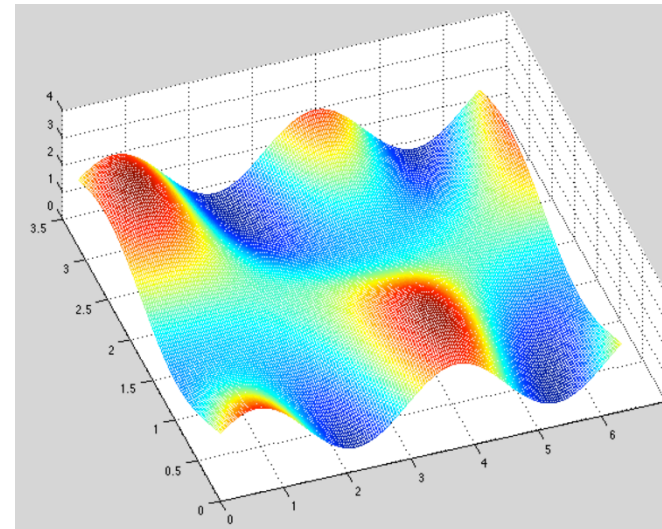
V=[ 1  1  1 ; -1  1  1; 1  1  0];
D=[ 1  0  0 ;  0  2  0;  0  0  3];
A=V*D*inv(V);

m=200; dt=2*pi/m; t=0:dt:2*pi; %% theta
n=100; dp=pi/n; p=0:dp:pi; %% phi
[T,P]=ndgrid(t,p);

r=cos(P); z=sin(P); x=r.*cos(T); y=r.*sin(T);

rq=0*x;
for i=1:m+1; for j=1:n+1;
    v = [ x(i,j); y(i,j); z(i,j)]; w=A*v;
    rq(i,j)= v'*w;
end; end;

mesh(T,P,rq)
```



Example: Rayleigh Quotient

- Rayleigh quotient can accelerate convergence of iterative methods such as power iteration, since Rayleigh quotient $\mathbf{x}_k^T \mathbf{A} \mathbf{x}_k / \mathbf{x}_k^T \mathbf{x}_k$ 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	\mathbf{x}_k^T		$\ \mathbf{y}_k\ _\infty$	$\mathbf{x}_k^T \mathbf{A} \mathbf{x}_k / \mathbf{x}_k^T \mathbf{x}_k$
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
puts more emphasis
on \mathbf{x}_1 and
convergence is
quadratic.*



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*

$$\sigma_k = \mathbf{x}_k^T \mathbf{A} \mathbf{x}_k / \mathbf{x}_k^T \mathbf{x}_k$$

$$(\mathbf{A} - \sigma_k \mathbf{I}) \mathbf{y}_{k+1} = \mathbf{x}_k \quad \text{Solve a system!}$$

$$\mathbf{x}_{k+1} = \mathbf{y}_{k+1} / \|\mathbf{y}_{k+1}\|_\infty$$

starting from given nonzero vector \mathbf{x}_0



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 $x^H A x / x^H x$

Matlab Demo: [eig_shifted_rq.m](#)



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 – Finding Second Eigenpair

Choose H to be elementary Householder matrix such that $H\mathbf{x}_1 = \alpha\mathbf{e}_1$. Consider

$$A\mathbf{x}_1 = \lambda_1 \mathbf{x}_1$$

$$AH^{-1}H\mathbf{x}_1 = \lambda_1 H^{-1}H\mathbf{x}_1$$

$$AH^{-1}\alpha\mathbf{e}_1 = \lambda_1 H^{-1}\alpha\mathbf{e}_1$$

$$HAH^{-1}\mathbf{e}_1 = \lambda_1 \mathbf{e}_1$$

$$A_2 := HAH^{-1} = HAH^{-1}I$$

$$= HAH^{-1} [\mathbf{e}_1 \ \mathbf{e}_2 \ \dots \ \mathbf{e}_n]$$

$$= \begin{bmatrix} \lambda_1 & \mathbf{b}^T \\ 0 & B \\ 0 & \\ 0 & \end{bmatrix}$$

Apply method of choice to B to find λ_2 .

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

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

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



Deflation, continued

- 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

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

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



Matrix-Free Deflation

Deflation via Projection.

for $k = 1, 2, \dots$

$$\mathbf{y} = A\mathbf{x}$$

$$\mathbf{y} = \mathbf{y} - \mathbf{x}_1 \frac{\mathbf{x}_1^T \mathbf{y}}{\mathbf{x}_1^T \mathbf{x}_1} = \mathbf{y} - \mathbf{x}_1 \mathbf{x}_1^T \mathbf{y}$$

$$\mathbf{x} = \mathbf{y} / \|\mathbf{y}\|.$$

- Do not require knowledge of A .
- Only need a routine that provides $\mathbf{y} \longleftarrow A\mathbf{x}$.

Deflation on the Fly – *Subspace Iteration*

Can effect deflation on the fly — *Subspace Iteration*.

- Take two independent vectors $Y = (\mathbf{y}_1 \ \mathbf{y}_2)$.

for $k = 1, 2, \dots$

$$Z = AY$$

$$\mathbf{y}_1 = \mathbf{z}_1 / \|\mathbf{z}_1\|.$$

$$\mathbf{y}_2 = \mathbf{z}_2 - \mathbf{y}_1 \mathbf{y}_1^T \mathbf{z}_2$$

$$\mathbf{y}_2 = \mathbf{y}_2 / \|\mathbf{y}_2\|$$

- $(\mathbf{y}_1 \ \mathbf{y}_2)$ converge to the first two eigenvectors $(\mathbf{x}_1 \ \mathbf{x}_2)$.

Matlab Demo: eig_simultaneous

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*



Subspace Iteration Variants

Let $X_0 \in \mathbb{R}^{n \times p}$ be a matrix of rank p .

- Alg. 1:

for $k = 1, 2, \dots$

$$X_k = A X_{k-1}$$

end

- Alg. 2:

for $k = 1, 2, \dots$

$$Q R = X \quad Q \in \mathbb{R}^{n \times p}, \text{ orthogonal}$$

$$X = A Q$$

end

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

- 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

$$A_k = R_k Q_k$$



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

Recall similarity transformation:

$$A : B = T^{-1}AT \quad \text{same eigenvalues}$$

With $B\underline{z} = \lambda\underline{z}$, we have:

$$\begin{aligned} B\underline{z} = T^{-1}AT\underline{z} &= \lambda\underline{z} \\ AT\underline{z} &= \lambda T\underline{z} \\ A\underline{x} &= \lambda\underline{x}, \quad \text{with } \underline{x} := T\underline{z} \end{aligned}$$

Starting with $A_0 = A$, we consider a sequence of similarity transformations:

$$A_k = Q_k^T A_{k-1} Q_k = Q_k^{-1} A_{k-1} Q_k$$

QR Iteration

Start with $A_0 = A$.

- Alg. QR:

for $k = 1, 2, \dots$

$$\begin{array}{l} Q_k R_k = A_{k-1} \\ A_k = R_k Q_k \end{array} \quad \left| \begin{array}{l} R_k = Q_k^T A_{k-1} \\ A_k = Q_k^T A_{k-1} Q_k \end{array} \right. \quad \left| \begin{array}{l} Q_k = A_{k-1} R_k^{-1} \\ A_k = R_k A_{k-1} R_k^{-1} \end{array} \right.$$

end

- Net result is similarity transformation

$$A_k = (Q_k^T Q_{k-1}^T \cdots Q_1^T) A (Q_1 Q_2 \cdots Q_k)$$

- Eigenvalues and symmetry are preserved.
- Can use a different orthogonal matrix $A_0 := Q_0^T A Q_0$ to start.

Comparison of QR and Subspace Iteration

Thus, we have a way of generating $\hat{Q}_k := Q_1 Q_2 \cdots Q_k$ through the following QR iteration:

for $k = 1, 2, \dots$

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

$$A_k = R_k Q_k$$

$$\hat{Q}_k = \hat{Q}_{k-1} Q_k$$

end

If A is normal, columns approach eigenvectors

QR Iteration

Note that $Q_1 R_1 = A_0 = A$.

$$\begin{aligned} A^2 &= Q_1 R_1 Q_1 R_1 \\ &= Q_1 Q_2 R_2 R_1 =: \hat{Q}_2 \hat{R}_2 \end{aligned}$$

$$\begin{aligned} A^3 &= Q_1 R_1 Q_1 R_1 Q_1 R_1 \\ &= Q_1 Q_2 R_2 Q_2 R_2 R_1 =: \hat{Q}_2 \hat{R}_2 \\ &= Q_1 Q_2 Q_3 R_3 R_2 R_1 =: \hat{Q}_3 \hat{R}_3 \end{aligned}$$

$$A^k = \hat{Q}_k \hat{R}_k$$

- Algorithm produces successive powers of A .

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



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



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

eig_qr_w_shifts.m < interactive example >



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



Upper Hessenberg X Upper Triangular is Upper Hessenberg

$$\begin{bmatrix} x & x & x & x & x \\ x & x & x & x & x \\ & x & x & x & x \\ & & x & x & x \\ & & & x & x \end{bmatrix} \begin{bmatrix} x & x & x & x & x \\ & x & x & x & x \\ & & x & x & x \\ & & & x & x \\ & & & & x \end{bmatrix} = \begin{bmatrix} x & x & x & x & x \\ x & x & x & x & x \\ & x & x & x & x \\ & & x & x & x \\ & & & x & x \end{bmatrix}$$

Upper Hessenberg X Upper Triangular is Upper Hessenberg

$$\begin{bmatrix}
 x & x & x & x & x \\
 x & x & x & x & x \\
 \boxed{} & x & x & x & x \\
 & & x & x & x \\
 & & & x & x
 \end{bmatrix}
 \begin{bmatrix}
 \boxed{x} & x & x & x & x \\
 & x & x & x & x \\
 & & x & x & x \\
 & & & x & x \\
 & & & & x
 \end{bmatrix}
 =
 \begin{bmatrix}
 x & x & x & x & x \\
 x & x & x & x & x \\
 \boxed{} & x & x & x & x \\
 & & x & x & x \\
 & & & x & x
 \end{bmatrix}$$

Upper Hessenberg X Upper Triangular is Upper Hessenberg

$$\begin{bmatrix}
 x & x & x & x & x \\
 x & x & x & x & x \\
 \boxed{x} & \boxed{x} & \boxed{x} & \boxed{x} & \boxed{x} \\
 & & x & x & x \\
 & & & x & x
 \end{bmatrix}
 \begin{bmatrix}
 \boxed{x} & x & x & x & x \\
 & x & x & x & x \\
 & & x & x & x \\
 & & & x & x \\
 & & & & x
 \end{bmatrix}
 =
 \begin{bmatrix}
 x & x & x & x & x \\
 x & x & x & x & x \\
 \boxed{} & x & x & x & x \\
 & & x & x & x \\
 & & & x & x
 \end{bmatrix}$$

- Same for upper triangular \times upper Hessenberg.
- Start QR iteration with $A_0 := Q_0^T A Q_0$ such that A_0 is upper Hessenberg.
- Then each A_k is upper Hessenberg, and $Q_k R_k$ can be done with Givens rotations in $O(n^2)$ operations, instead of $O(n^3)$.
- With shifted QR , need only $O(n)$ iterations, so total cost is $O(n^3)$.

QR iteration:

for $k = 1, 2, \dots$

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

$$A_k = R_k Q_k$$

end

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



Preliminary Reduction, continued

- QR iteration is implemented in two-stages

symmetric \longrightarrow tridiagonal \longrightarrow diagonal

or

general \longrightarrow Hessenberg \longrightarrow 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



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

- Assume in the following that A is a symmetric positive definite matrix with eigenvalues $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_n$.
- Suppose we take k iterations of the power method to approximate λ_1 :

Algorithm:

$$\mathbf{y}_k = A^k \mathbf{x}$$

$$\lambda = \frac{\mathbf{y}_k^T A \mathbf{y}_k}{\mathbf{y}_k^T \mathbf{y}_k} \approx \lambda_1$$

Code:

```
 $\mathbf{y} = \mathbf{x}$   
for  $j = 1 : k$   
     $\mathbf{y} = A\mathbf{y}$   
end  
 $\lambda = \frac{\mathbf{y}^T A \mathbf{y}}{\mathbf{y}^T \mathbf{y}}$ 
```

- This approach uses no information from preceding iterates, $\mathbf{y}_{k-1}, \mathbf{y}_{k-2}, \dots, \mathbf{y}_1$.

Krylov Subspace Methods

- Suppose instead, we seek (for $\mathbf{y} \neq 0$),

$$\begin{aligned}\lambda &= \max_{\mathbf{y} \in \mathcal{K}_{k+1}} \frac{\mathbf{y}^T A \mathbf{y}}{\mathbf{y}^T \mathbf{y}} \\ &= \max_{\mathbf{y} \in \mathbb{P}_k(A)\mathbf{x}} \frac{\mathbf{y}^T A \mathbf{y}}{\mathbf{y}^T \mathbf{y}} \\ &\leq \max_{\mathbf{y} \in \mathbb{R}^n} \frac{\mathbf{y}^T A \mathbf{y}}{\mathbf{y}^T \mathbf{y}} = \lambda_1.\end{aligned}$$

- Here, the **Krylov subspace**

$$\mathcal{K}_{k+1} = \mathcal{K}_{k+1}(A, \mathbf{x}) = \text{span}\{\mathbf{x} \ A\mathbf{x} \ A^2\mathbf{x} \ \dots \ A^k\mathbf{x}\}$$

is the space of all polynomials of degree $\leq k$ in A times \mathbf{x} :

$$\mathcal{K}_{k+1}(A, \mathbf{x}) = \mathbb{P}_k(A)\mathbf{x}.$$

Krylov Subspace Methods

- Consider the $n \times (k + 1)$ matrix

$$V_{k+1} = [\mathbf{x} \ A\mathbf{x} \ A^2\mathbf{x} \ \dots \ A^k\mathbf{x}]$$

- Assuming V_{k+1} has full rank ($k+1$ linearly independent columns), then any $\mathbf{y} \in \mathcal{K}_{k+1}$ has the form

$$\mathbf{y} = V_{k+1} \mathbf{z}, \quad \mathbf{z} \in \mathbb{R}^{k+1}.$$

- Our optimal Rayleigh quotient amounts to

$$\lambda = \max_{\mathbf{y} \in \mathcal{K}_{k+1}} \frac{\mathbf{y}^T A \mathbf{y}}{\mathbf{y}^T \mathbf{y}} = \max_{\mathbf{z} \in \mathbb{R}^{k+1}} \frac{\mathbf{z}^T (V_{k+1}^T A V_{k+1}) \mathbf{z}}{\mathbf{z}^T (V_{k+1}^T V_{k+1}) \mathbf{z}}$$

Krylov Subspace Methods

- If we had columns \mathbf{v}_j that were orthonormal ($\mathbf{v}_i^T \mathbf{v}_j = \delta_{ij}$), then we'd have $V_{k+1}^T V_{k+1} = I$ and

$$\begin{aligned}\lambda &= \max_{\mathbf{y} \in \mathcal{R}(V_{k+1})} \frac{\mathbf{y}^T A \mathbf{y}}{\mathbf{y}^T \mathbf{y}} \\ &= \max_{\mathbf{z} \in \mathbb{R}^{k+1}} \frac{\mathbf{z}^T (V_{k+1}^T A V_{k+1}) \mathbf{z}}{\mathbf{z}^T (V_{k+1}^T V_{k+1}) \mathbf{z}} \\ &= \max_{\mathbf{z} \in \mathbb{R}^{k+1}} \frac{\mathbf{z}^T T_{k+1} \mathbf{z}}{\mathbf{z}^T \mathbf{z}} \\ &= \mu_1(T_{k+1}) \leq \lambda_1(A).\end{aligned}$$

- Here, μ_1 is the maximum eigenvalue of $T_{k+1} := V_{k+1}^T A V_{k+1}$.

Krylov Subspace Methods

- This is the idea behind Arnoldi / Lanczos iteration.
- We use information from the entire Krylov subspace to generate optimal eigenpair approximations.
- They require only matrix-vector products (unlike QR iteration, which requires all of A).
- The approximation to λ_1 is given by the eigenvalue μ_1 of the much smaller $(k+1) \times (k+1)$ matrix, T_{k+1} .
- It is the closest approximation to λ_1 out of all possible polynomials of degree k in A and therefore superior (or equal to) the power method.
- Similarly, μ_k is the closest approximation to λ_n .
- The methods produce the best possible approximations (in this subspace) to the extreme eigenvalue/vector pairs.
- Middle eigenpairs are more challenging—must use *shift & invert*.

Krylov Subspace Methods

- Note, for $\|\mathbf{z}\| = 1$,

$$\mu_1 = \max_{\|\mathbf{z}\|=1} \mathbf{z}^T T_{k+1} \mathbf{z}$$

corresponds to $\mathbf{z} = \mathbf{z}_1$, so

$$\mu_1 = \mathbf{z}_1^T T_{k+1} \mathbf{z}_1 = \mathbf{z}_1^T V_{k+1}^T A V_{k+1} \mathbf{z}_1 \approx \lambda_1.$$

- So, corresponding eigenvector approximation for $A\mathbf{x}_1 = \lambda_1\mathbf{x}_1$ is

$$\mathbf{x}_1 \approx V_{k+1} \mathbf{z}_1.$$

Krylov Subspace Methods

- **Remark:** Shifting does not improve Lanczos / Arnoldi. WHY?

Krylov Subspace Methods

- **Remark:** Shifting does not improve Lanczos / Arnoldi. WHY?
 - If $p(x) \in \mathbb{P}_k(x)$, so is $p(x+1)$, $p(ax+b)$, *etc.*
 - So: $\mathcal{K}(A, \mathbf{x}) \equiv \mathcal{K}(A + \alpha I, \mathbf{x})$.
 - The spaces are the same and the Krylov subspace projections will find the same optimal solutions.
- Shifting *may help* with conditioning, however, in certain circumstances.
- The essential steps of the algorithms is to construct, step by step, an orthogonal basis for \mathcal{K}_k .
- We turn to this for the symmetric (Lanczos) case.

Krylov Subspace Methods

- We start with the symmetric case, known as Lanczos iteration.
- The essence of the method is to construct, step by step, an orthogonal basis for \mathcal{K}_k .

```
 $\mathbf{q}_1 = \mathbf{x}/\|\mathbf{x}\|;$   
for  $k = 1, \dots$   
   $\mathbf{u} = A\mathbf{q}_k, \quad u_0 = \|\mathbf{u}\|$   
   $\mathbf{u} = \mathbf{u} - P_k \mathbf{u}$   
   $\beta_k = \|\mathbf{u}\|$   
  if  $\beta_k/u_0 < \epsilon$ , break  
   $\mathbf{q}_{k+1} = \mathbf{u}/\beta_k$   
end
```

- Here, $P_k := Q_k Q_k^T$, is the orthogonal projector onto $\mathcal{R}(Q_k)$, so

$$\mathbf{u} = \mathbf{u} - P_k \mathbf{u} = \mathbf{u} - \sum_{j=1}^k \mathbf{q}_j \mathbf{q}_j^T \mathbf{u},$$

which is implemented as modified Gram-Schmidt orthogonalization.

- Q: Why is u_0 useful?

Krylov Subspace Generation

- Notice how the orthogonal subspace is constructed.

$$\begin{aligned}\mathbf{q}_1 &\in \{\mathbf{x}\} \\ \mathbf{q}_2 &\in \{\mathbf{x}, A\mathbf{x}\} \\ \mathbf{q}_k &\in \{\mathbf{x}, A\mathbf{x}, \dots, A^{k-1}\mathbf{x}\} = \mathcal{K}_k\end{aligned}$$

- In the algorithm, we have

$$\begin{aligned}\mathbf{u} &= A\mathbf{q}_k \\ \mathbf{q}_j^T \mathbf{u} &= \mathbf{q}_j^T A\mathbf{q}_k \\ &= \mathbf{q}_k^T A^T \mathbf{q}_j = \mathbf{q}_k^T A\mathbf{q}_j \\ &= \mathbf{q}_k^T \mathbf{w}_{j+1}, \quad \mathbf{w}_{j+1} := A\mathbf{q}_j \in \mathcal{K}_{j+1}\end{aligned}$$

- However,

$$\mathbf{q}_k \perp \mathcal{K}_{j+1} \quad \forall j+1 < k.$$

- Therefore

$$\begin{aligned}\mathbf{u} &= \mathbf{u} - P_k \mathbf{u} \\ &= \mathbf{u} - \mathbf{q}_k (\mathbf{q}_k^T A\mathbf{q}_k) - \mathbf{q}_{k-1} (\mathbf{q}_{k-1}^T A\mathbf{q}_k) \\ &= \mathbf{u} - \mathbf{q}_k \alpha_k - \mathbf{q}_{k-1} \beta_{k-1}, \\ \alpha_k &:= \mathbf{q}_k^T A\mathbf{q}_k \quad \beta_k := \|\mathbf{u}_{k-1}\|.\end{aligned}$$

Lanczos Iteration (A Symmetric)

- The Lanczos iteration is:

$$\mathbf{q}_1 = \mathbf{x}/\|\mathbf{x}\|;$$

for $k = 1, \dots$

$$\mathbf{u} = A\mathbf{q}_k, \quad u_0 = \|\mathbf{u}\|$$

$$\alpha_k = \mathbf{q}_k^T \mathbf{u}$$

$$\mathbf{u} = \mathbf{u} - \alpha_k \mathbf{q}_k - \beta_{k-1} \mathbf{q}_{k-1}$$

$$\beta_k = \|\mathbf{u}\|$$

if $\beta_k/u_0 < \epsilon$, break

$$\mathbf{q}_{k+1} = \mathbf{u}/\beta_k$$

end

Lanczos Iteration (A Symmetric)

- In matrix form,

$$AQ_k = Q_k T_k + \beta_{k+1} \mathbf{q}_{k+1} \mathbf{e}_k^T$$

- Or,

$$A [\mathbf{q}_1 \ \mathbf{q}_2 \ \cdots \ \mathbf{q}_k] = [\mathbf{q}_1 \ \mathbf{q}_2 \ \cdots \ \mathbf{q}_k] \begin{bmatrix} \alpha_1 & \beta_1 & & \\ \beta_1 & \alpha_1 & \beta_1 & \\ & \ddots & \ddots & \beta_{k-1} \\ & & \beta_{k-1} & \alpha_k \end{bmatrix} + \beta_{k+1} \mathbf{q}_{k+1} \mathbf{e}_k^T.$$

Arnoldi Iteration (*A* Nonsymmetric)

- Arnoldi iteration is essentially the same as Lanczos, save that we do not get the short term recurrence.

$$\mathbf{q}_1 = \mathbf{x}/\|\mathbf{x}\|;$$

for $k = 1, \dots$

$$\mathbf{u} = A\mathbf{q}_k, \quad u_0 = \|\mathbf{u}\|$$

$$\mathbf{u} = \mathbf{u} - P_k \mathbf{u}$$

$$\beta_k = \|\mathbf{u}\|$$

if $\beta_k/u_0 < \epsilon$, break

$$\mathbf{q}_{k+1} = \mathbf{u}/\beta_k$$

end

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

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

then

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

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



Krylov Subspace Methods

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

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

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

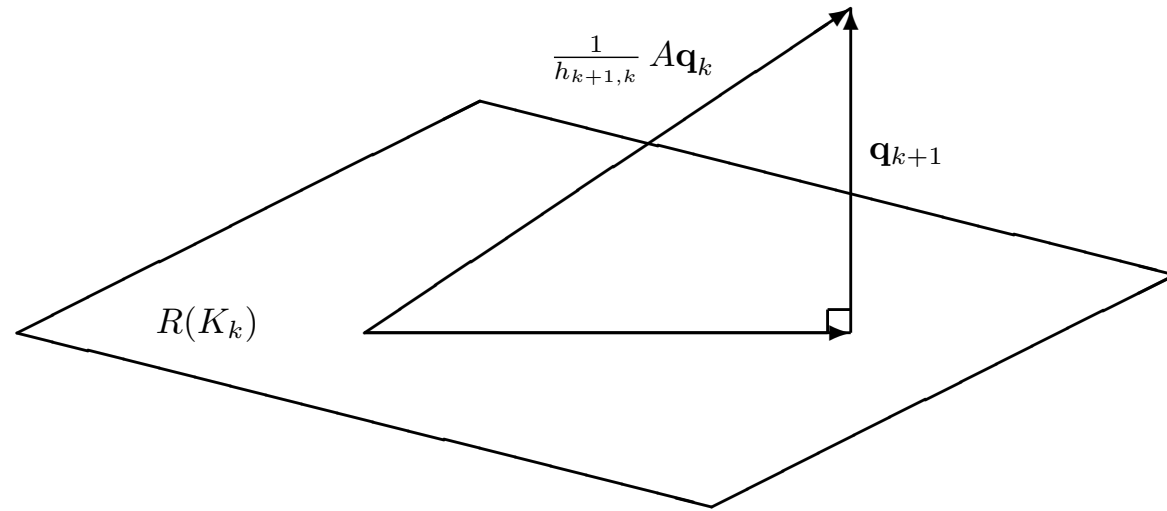
- Premultiplying by q_j^H and using orthonormality,

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

- 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



Krylov Subspace Projections



- Notice that $A\mathbf{q}_j$ will be in $\mathcal{R}(K_k)$ for all $j < k$.
- $\mathbf{q}_{k+1} \perp \mathcal{R}(K_k)$
- $\mathbf{q}_{k+1} \perp \mathcal{R}(AK_{k-1}) \subset \mathcal{R}(K_k)$

Krylov Subspace and Similarity Transformation

- Consider the rank k matrix

$$K_k := (\mathbf{x}_0 \ A\mathbf{x}_0 \ A^2\mathbf{x}_0 \ \cdots \ A^{k-1}\mathbf{x}_0),$$

and associated Krylov subspace $\mathcal{K}_k := R(K_k)$.

- Krylov subspace methods work with the orthogonal vectors $\mathbf{q}_k \in \mathcal{K}_k$, $k=1, 2, \dots$, satisfying $QR = K_k$.
- The similarity transformation

$$Q^{-1}AQ = Q^T A Q = H$$

with entries $h_{ij} = \mathbf{q}_i^T A \mathbf{q}_j$ is upper Hessenberg.

- Proof:
 - $\mathbf{q}_i^T \mathbf{u} = 0$ for any $\mathbf{u} \in K_k$ if $i > k$.
 - $\mathbf{u} = A\mathbf{q}_j \in K_{j+1}$.
 - $\mathbf{q}_i^T \mathbf{u} = \mathbf{q}_i^T A \mathbf{q}_j = 0$ for $i > j + 1$.

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

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

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

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

} *Modified Gram-Schmidt*

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

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

Modified Gram-Schmidt

$$\begin{aligned} \underline{u}_k &= (I - QQ^T) \underline{u}_k \\ &= \underline{u}_k - QQ^T \underline{u}_k \end{aligned}$$

$$= \underline{u}_k - [h_{1k} \underline{q}_1 \ h_{2k} \underline{q}_2 \ \dots \ h_{kk} \underline{q}_k]$$

Q: There are two projectors in the lines above. Where are they ??



Arnoldi Iteration, continued

- If

$$Q_k = [q_1 \quad \cdots \quad q_k]$$

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 $k \ll n$



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

A reasonable restart choice: current approximate eigenvector 

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)

$$\mathbf{q}_0 = \mathbf{0}$$

$$\beta_0 = 0$$

$\mathbf{x}_0 =$ arbitrary nonzero starting vector

$$\mathbf{q}_1 = \mathbf{x}_0 / \|\mathbf{x}_0\|_2$$

for $k = 1, 2, \dots$

$$\mathbf{u}_k = \mathbf{A}\mathbf{q}_k$$

$$\alpha_k = \mathbf{q}_k^H \mathbf{u}_k$$

$$\mathbf{u}_k = \mathbf{u}_k - \beta_{k-1} \mathbf{q}_{k-1} - \alpha_k \mathbf{q}_k$$

$$\beta_k = \|\mathbf{u}_k\|_2$$

if $\beta_k = 0$ **then** stop

$$\mathbf{q}_{k+1} = \mathbf{u}_k / \beta_k$$

end



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)



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



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 $(A - \sigma I)x = y$



Optimality of Lanczos, Case A is SPD

- Recall, if A is SPD, then $\mathbf{x}^T A \mathbf{x} > 0 \ \forall \ \mathbf{x} \neq 0$.
- If Q is full rank, then $T := Q^T A Q$ is SPD.

$$(Q\mathbf{y})^T A (Q\mathbf{y}) = \mathbf{y}^T Q^T A Q \mathbf{y} > 0 \ \forall \ \mathbf{y} \neq 0.$$

- If A is SPD then $\|A\|_2 = \lambda_1$ (max eigenvalue). Thus,

$$\lambda_1 = \max_{\mathbf{x} \neq 0} \frac{\mathbf{x}^T A \mathbf{x}}{\mathbf{x}^T \mathbf{x}} = \max_{\|\mathbf{x}\|=1} \mathbf{x}^T A \mathbf{x}.$$

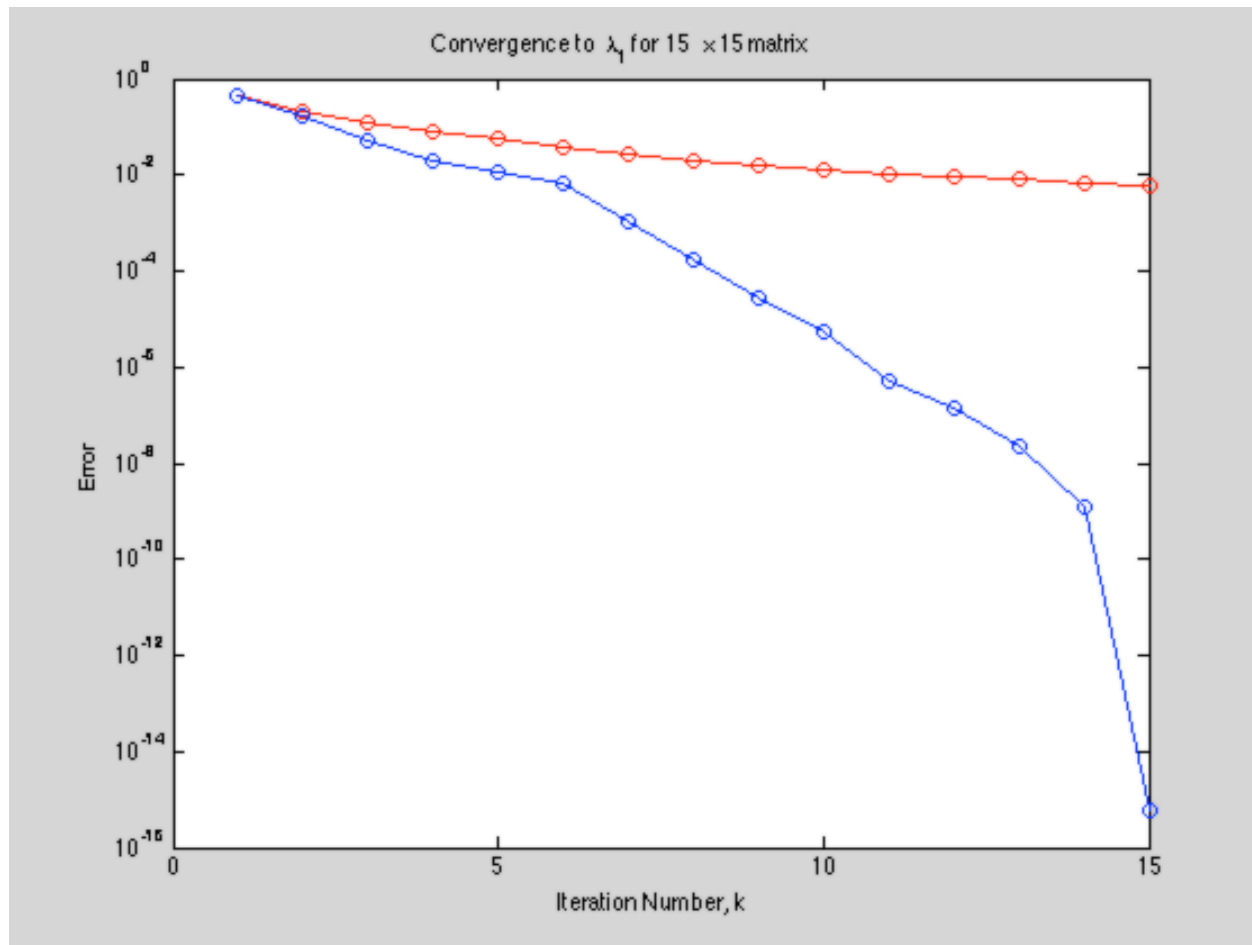
- Let $T\mathbf{y} = \mu\mathbf{y}$. (T is $k \times k$ tridiagonal, $k \ll n$.)

$$\mu_1 = \max_{\|\mathbf{y}\|=1} \mathbf{y}^T T \mathbf{y} = \max_{\|\mathbf{y}\|=1} \mathbf{y}^T Q^T A Q \mathbf{y} \geq \mathbf{x}^T A \mathbf{x} \ \forall \ \mathbf{x} \in \mathcal{K}_k$$

- Therefore, μ_1 is the closest Rayleigh quotient estimate for all $\mathbf{x} \in \mathcal{K}_k$, $\|\mathbf{x}\| = 1$.
- Lanczos is as good as (or much better than) the power method for the same number of matrix-vector products in A .

Matlab Demo

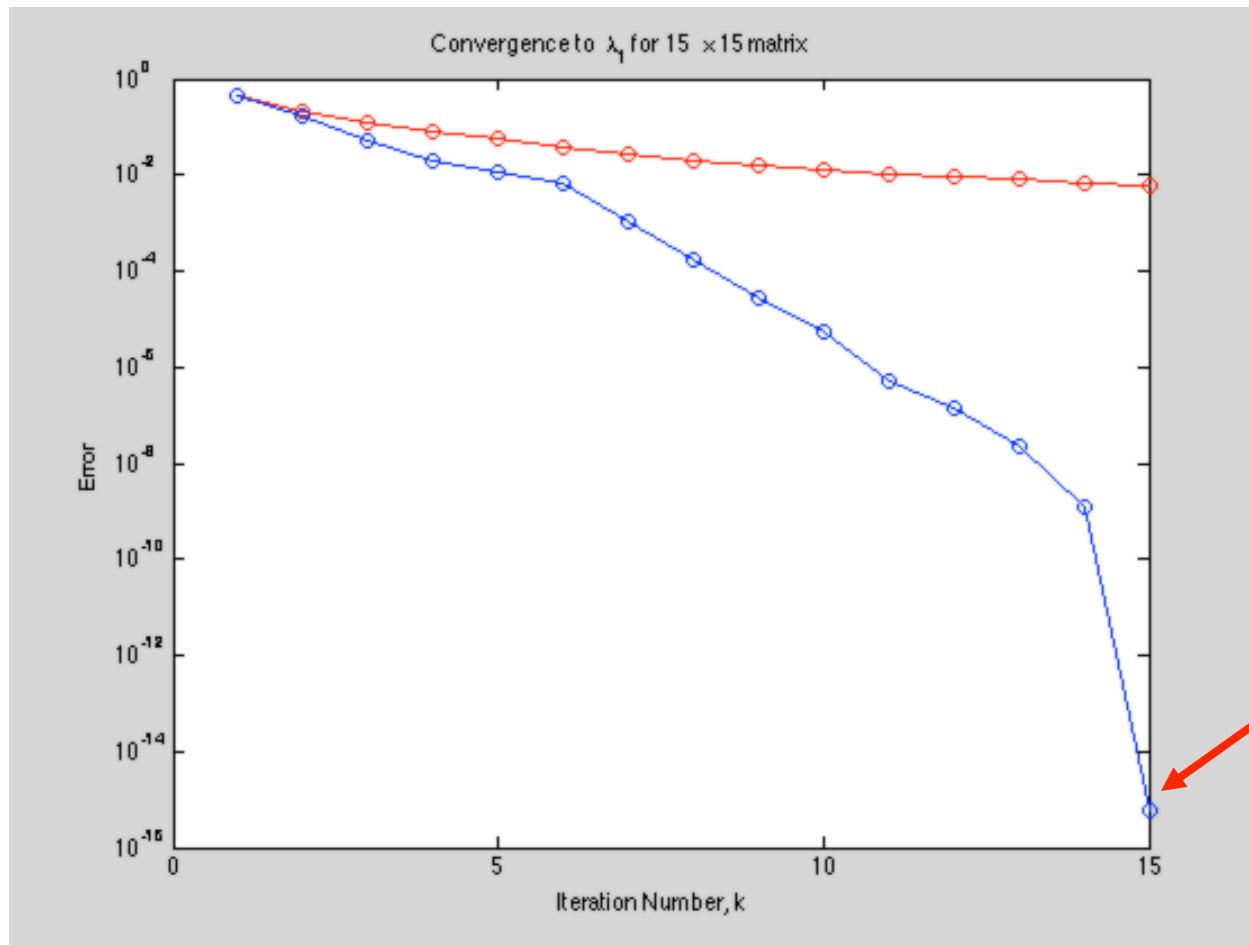
❑ Lanczos vs Power Iteration



❑ Lanczos does a reasonable job of converging to extreme eigenvalues.

Matlab Demo

❑ Lanczos vs Power Iteration



Q: What is happening here?

- ❑ Lanczos does a reasonable job of converging to extreme eigenvalues.

Example: Lanczos Iteration

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

