Applied Numerical Linear Algebra. Lecture 12

Regular Matrix Pencils and Weierstrass Canonical Form

The standard eigenvalue problem asks for which scalars z the matrix $A - zI$ is singular; these scalars are the eigenvalues. This notion generalizes in several important ways.

DEFINITION. $A - \lambda B$, where A and B are m-by-n matrices, is called a *matrix pencil*, or just a *pencil*. Here is λ an indeterminate, not a particular, numerical value.

DEFINITION. If A and B are square and $\det(A - \lambda B)$ is not identically zero (or when there exists at least one $\lambda : \det(A - \lambda B) \neq 0$), the pencil $A - \lambda B$ is called *regular*. Otherwise it is called *singular*. When $A - \lambda B$ is regular, $p(\lambda) \equiv \det(A - \lambda B)$ is called the *characteristic polynomial* of $A - \lambda B$ and the eigenvalues of $A - \lambda B$ are defined to be
(1) the roots of $p(\lambda)$,
(2) ∞ (with multiplicity $n - \deg(p)$) if $\deg(p) < n$.

Example

Let

$$A - \lambda B = \begin{bmatrix} 1 & & \\ & 1 & \\ & & 0 \end{bmatrix} - \lambda \begin{bmatrix} 2 & & \\ & 0 & \\ & & 1 \end{bmatrix}.$$

Then $p(\lambda) = \det(A - \lambda B) = (1 - 2\lambda) \cdot (1 - 0\lambda) \cdot (0 - \lambda) = (1 - 2\lambda)(-\lambda)$, so the eigenvalues are $\lambda = \frac{1}{2}$, 0 and ∞. $\diamond$

PROPOSITION. Let $A - \lambda B$ be regular.

- If B is nonsingular, all eigenvalues of $A - \lambda B$ are finite and the same as the eigenvalues of AB^{-1} or $B^{-1}A$.
- If B is singular, $A - \lambda B$ has eigenvalue ∞ with multiplicity $n - \mathrm{rank}(B)$.
- If A is nonsingular, the eigenvalues of $A - \lambda B$ are the same as the reciprocals of the eigenvalues (or $\frac{1}{\lambda_i}$) of $A^{-1}B$ or BA^{-1}, where a zero eigenvalue of $A^{-1}B$ corresponds to an infinite eigenvalue of $A - \lambda B$.

Proof.

- If B is nonsingular and λ' is an eigenvalue, then $0 = \det(A - \lambda' B) = \det(AB^{-1} - \lambda' I) = \det(B^{-1}A - \lambda' I)$ so λ' is also an eigenvalue of AB^{-1} and $B^{-1}A$.
- If B is singular, then take $p(\lambda) = \det(A - \lambda B)$, write the SVD of B as $B = U\Sigma V^T$, and substitute to get

$$p(\lambda) = \det(A - \lambda U\Sigma V^T) = \det(U(U^T AV - \lambda\Sigma)V^T)$$

$$= \pm \det(U^T AV - \lambda\Sigma).$$

Since $\mathrm{rank}(B) = \mathrm{rank}(\Sigma)$, only $\mathrm{rank}(B)$ λ's appear in $U^T AV - \lambda\Sigma$, so the degree of the polynomial $\det(U^T AV - \lambda\Sigma)$ is $\mathrm{rank}(B)$.

- If A is nonsingular, $\det(A - \lambda B) = 0$ and $\det(A(I - \lambda A^{-1}B)) = 0$ if and only if $\det(I - \lambda A^{-1}B) = 0$ or $\det(I - \lambda BA^{-1}) = 0$.

$$\det(I - \lambda A^{-1}B) = 0 \rightarrow \det(\frac{1}{\lambda}I - A^{-1}B) = 0$$

This equality can hold only if $\lambda \neq 0$ and $1/\lambda$ is an eigenvalue of $A^{-1}B$ and BA^{-1}. □

DEFINITION. Let P_L and P_R be nonsingular matrices. Then pencils $A - \lambda B$ and $P_L A P_R - \lambda P_L B P_R$ are called *equivalent*.

PROPOSITION. The equivalent regular pencils $A - \lambda B$ and $P_L A P_R - \lambda P_L B P_R$ have the same eigenvalues. The vector x is a right eigenvector of $A - \lambda B$ if and only if $P_R^{-1} x$ is a right eigenvector of $P_L A P_R - \lambda P_L B P_R$. The vector y is a left eigenvector of $A - \lambda B$ if and only if $(P_L^*)^{-1} y$ is a left eigenvector of $P_L A P_R - \lambda P_L B P_R$.

Proof.

- $\det(A - \lambda B) = 0$ if and only if $\det(P_L(A - \lambda B)P_R) = 0$.
- $(A - \lambda B)x = 0$ if and only if $P_L(A - \lambda B)P_R P_R^{-1} x = 0$.
- $(A - \lambda B)^* y = 0$ if and only if $P_R^*(A - \lambda B)^* P_L^* (P_L^*)^{-1} y = 0$. $\square$

THEOREM. *Weierstrass canonical form*. Let $A - \lambda B$ be regular. Then there are nonsingular P_L and P_R such that

$$P_L(A - \lambda B)P_R = \mathrm{diag}(J_{n_1}(\lambda_1) - \lambda I_{n_1}, \ldots, J_{n_k}(\lambda_{n_k}) - \lambda I_{n_k}, N_{m_1}, \ldots, N_{m_r}),$$

where $J_{n_i}(\lambda_i)$ is an n_i-by-n_i Jordan block with eigenvalue λ_i,

$$J_{n_i}(\lambda_i) = \begin{bmatrix} \lambda_i & 1 & & \\ & \ddots & \ddots & \\ & & \ddots & 1 \\ & & & \lambda_i \end{bmatrix},$$

and N_{m_i} is a "Jordan block for $\lambda = \infty$ with multiplicity m_i,"

$$N_{m_i} = \begin{bmatrix} 1 & \lambda & & \\ & 1 & \ddots & \\ & & \ddots & \lambda \\ & & & 1 \end{bmatrix} = I_{m_i} + \lambda J_{m_i}(0).$$

For a proof, see [F. Gantmacher. The Theory of Matrices, vol. II (translation). Chelsea, New York, 1959].

Application of Jordan and Weierstrass Forms to Differential Equations

Consider the linear differential equation

$$\dot{x}(t) = Ax(t) + f(t), \tag{1}$$
$$x(0) = x_0. \tag{2}$$

An explicit solution of (1)-(2) is given by

$$x(t) = e^{At}x_0 + \int_0^t e^{A(t-\tau)} f(\tau) d\tau.$$

.

If we know the Jordan form $A = SJS^{-1}$, we may change variables in the differential equation (1) to $y(t) = S^{-1}x(t)$. Then $Sy(t) = x(t)$ and $\dot{x}(t) = S\dot{y}(t)$. We substitute $\dot{x}(t) = S\dot{y}(t)$ into (1) to get
$S\dot{y}(t) = ASy(t) + f(t)$
or $\dot{y}(t) = Jy(t) + S^{-1}f(t)$,
with solution $y(t) = e^{Jt}y_0 + \int_0^t e^{J(t-\tau)} S^{-1} f(\tau) d\tau$.

There is an explicit formula for $f(J) = e^{Jt}$ or any other function $f(J)$.

Suppose that f is given by its Taylor series $f(z) = \sum_{i=0}^{\infty} \frac{f^{(i)}(0)z^i}{i!}$ and J is a single Jordan block $J = \lambda I + N$, where N has ones on the first superdiagonal and zeros elsewhere. Then writing Taylor series for $f(J)$ with $\left(\begin{array}{c} i \\ j \end{array}\right) = \frac{i!}{j!(i-j)!}$ we get:

$$
\begin{aligned}
f(J) &= \sum_{i=0}^{\infty} \frac{f^{(i)}(0)(\lambda I + N)^i}{i!} \\
&= \sum_{i=0}^{\infty} \frac{f^{(i)}(0)}{i!} \sum_{j=0}^{i} \left(\begin{array}{c} i \\ j \end{array}\right) \lambda^{i-j} N^j \text{ by the binomial theorem} \\
&= \sum_{j=0}^{\infty} \sum_{i=j}^{\infty} \frac{f^{(i)}(0)}{i!} \left(\begin{array}{c} i \\ j \end{array}\right) \lambda^{i-j} N^j \text{ reversing the order of summation} \\
&= \sum_{j=0}^{n-1} N^j \sum_{i=j}^{\infty} \frac{f^{(i)}(0)}{i!} \left(\begin{array}{c} i \\ j \end{array}\right) \lambda^{i-j},
\end{aligned}
$$

where in the last equality we used the fact that $N^j = 0$ for $j > n-1$. Note that N^j has ones on the jth superdiagonal and zeros elsewhere.

Finally, note that $\sum_{i=j}^{\infty} \frac{f^{(i)}(0)}{i!} \binom{i}{j} \lambda^{i-j}$ is the Taylor expansion for $f^{(j)}(\lambda)/j!$. Thus

$$
\begin{aligned}
f(J) &= f\left(\begin{bmatrix} \lambda & 1 & & \\ & \ddots & \ddots & \\ & & \ddots & 1 \\ & & & \lambda \end{bmatrix}^{n\times n} \right) = \sum_{j=0}^{n-1} \frac{N^j f^{(j)}(\lambda)}{j!} \\
&= \begin{bmatrix} f(\lambda) & f'(\lambda) & \frac{f''(\lambda)}{2!} & \cdots & \frac{f^{(n-1)}(\lambda)}{(n-1)!} \\ & \ddots & \ddots & \ddots & \vdots \\ & & \ddots & \ddots & \frac{f''(\lambda)}{2!} \\ & & & \ddots & f'(\lambda) \\ & & & & f(\lambda) \end{bmatrix}
\end{aligned}
$$

so that $f(J)$ is upper triangular with $f^{(j)}(\lambda)/j!$ on the jth superdiagonal.

To solve the more general problem

$$B\dot{x} = Ax + f(t),$$

$A - \lambda B$ regular, we use the Weierstrass form: let $P_L(A - \lambda B)P_R$ be in Weierstrass form, and let us rewrite the equation as

$$P_L B P_R P_R^{-1}\dot{x} = P_L A P_R P_R^{-1} x + P_L f(t).$$

Let $P_R^{-1}x = y$ and $P_L f(t) = g(t)$. Then from the above equation we obtain

$$P_L B P_R \dot{y} = P_L A P_R y + g(t),$$

or

$$P_L A P_R y - P_L B P_R \dot{y} + g(t) = 0.$$

Now the problem has been decomposed into subproblems:

$$\begin{bmatrix} I_{n_1} & & & & & \\ & \ddots & & & & \\ & & I_{n_k} & & & \\ & & & J_{m_1}(0) & & \\ & & & & \ddots & \\ & & & & & J_{m_r}(0) \end{bmatrix} \dot{y}$$

$$= \begin{bmatrix} J_{n_1}(\lambda_1) & & & & & \\ & \ddots & & & & \\ & & J_{n_k}(\lambda_k) & & & \\ & & & I_{m_1} & & \\ & & & & \ddots & \\ & & & & & I_{m_r} \end{bmatrix} y + g.$$

We observe that we have to solve 2 subproblems:
(1) $\dot{\tilde{y}} = J_{n_i}(\lambda_i)\tilde{y} + \tilde{g}(t) \equiv J\tilde{y} + \tilde{y}(t)$ and
(2)$J_m(0)\dot{\tilde{y}} = \tilde{y} + \tilde{g}(t)$.
Each subproblem (1) $\dot{\tilde{y}} = J_{n_i}(\lambda_i)\tilde{y} + \tilde{g}(t) \equiv J\tilde{y} + \tilde{y}(t)$ is a standard linear ODE as above with solution

$$\tilde{y}(t) = \tilde{y}(0)e^{Jt} + \int_0^t e^{J(t-\tau)}\tilde{g}(\tau)d\tau.$$

(2) The solution of $J_m(0)\dot{\tilde{y}} = \tilde{y} + \tilde{g}(t)$ is gotten by back substitution starting from the last equation: write $J_m(0)\dot{\tilde{y}} = \tilde{y} + \tilde{g}(t)$ as

$$\begin{bmatrix} 0 & 1 & & \\ & \ddots & \ddots & \\ & & \ddots & 1 \\ & & & 0 \end{bmatrix} \begin{bmatrix} \dot{\tilde{y}}_1 \\ \vdots \\ \vdots \\ \dot{\tilde{y}}_m \end{bmatrix} = \begin{bmatrix} \tilde{y}_1 \\ \vdots \\ \vdots \\ \tilde{y}_m \end{bmatrix} + \begin{bmatrix} \tilde{g}_1 \\ \vdots \\ \vdots \\ \tilde{g}_m \end{bmatrix}.$$

The mth (last) equation says $0 = \tilde{y}_m + \tilde{g}_m$ or $\tilde{y}_m = -\tilde{g}_m$. The ith equation says $\dot{\tilde{y}}_{i+1} = \tilde{y}_i + \tilde{g}_i$, so $\tilde{y}_i = \dot{\tilde{y}}_{i+1} - \tilde{g}_i$ and thus $\tilde{y}_{i+1} = \dot{\tilde{y}}_{i+2} - \tilde{g}_{i+1}, \dot{\tilde{y}}_{i+1} = \ddot{\tilde{y}}_{i+2} - \dot{\tilde{g}}_{i+1}, \tilde{y}_i = \ddot{\tilde{y}}_{i+2} - \dot{\tilde{g}}_{i+1} - \tilde{g}_i, \ldots$ such that we have

$$\tilde{y}_i = \sum_{k=i}^{m} -\frac{d^{k-i}}{dt^{k-i}} \tilde{g}_k(t).$$

Therefore the solution depends on derivatives of $\tilde{g}$, *not* an integral of $\tilde{g}$ as in the usual ODE. Thus a continuous $\tilde{g}$ which is not differentiable can cause a discontinuity in the solution; this is sometimes called an *impulse response* and occurs only if there are infinite eigenvalues. Furthermore, to have a continuous solution $\tilde{y}$ must satisfy certain consistency conditions at $t = 0$:

$$y_i(0) = \sum_{k=m}^{i} -\frac{d^{k-i}}{dt^{k-i}} g_k(0).$$

Numerical methods, based on time-stepping, for solving such differential algebraic equations, or ODEs with algebraic constraints, are described in [K. Brenan, S. Campbell, and L. Petzold. Numerical Solution of Initial-Value Problems in Differential-Algebraic Equations. North Holland, New York, 1989.]

Definite Pencils

A simpler special case that often arises in practice is the pencil $A - \lambda B$, where $A = A^T$, $B = B^T$, and B is positive definite. Such pencils are called *definite pencils*.
THEOREM. Let $A = A^T$, and let $B = B^T$ be positive definite. Then there is a real nonsinqular matrix X so that $X^T AX = \mathrm{diag}(\alpha_1, \ldots, \alpha_n)$ and $X^T BX = \mathrm{diag}(\beta_1, \ldots, \beta_n)$. In particular, all the eigenvalues α_i / β_i are real and finite.

Proof. The proof that we give is actually the algorithm used to solve the problem:
(1) Let $LL^T = B$ be the Cholesky decomposition.
(2) Construct $H = L^{-1}AL^{-T}$; note that H is symmetric.
(3) Let $H = Q\Lambda Q^T$, with Q orthogonal, Λ real and diagonal. If we take $X = L^{-T}Q$ then $X^T A \underbrace{X}_{L^{-T}Q} = Q^T \underbrace{L^{-1}AL^{-T}}_{H} Q = \Lambda$ and

$$X^T BX = \underbrace{Q^T L^{-1}}_{X^T} B \underbrace{L^{-T} Q}_{X} = Q^T L^{-1} \underbrace{LL^T}_{B} L^{-T} Q = I. \ \square$$

Note that the theorem is also true if $\alpha A + \beta B$ is positive definite for some scalars α and β.

Singular Matrix Pencils and the Kronecker Canonical Form

The pencil $A - \lambda B$ is singular if either A and B are nonsquare, or they are square and $\det(A - \lambda B) = 0$ for all λ.

Example

Let $A = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}$ and $B = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}$ are such that the pencil $A - \lambda B$ is singular. Then by making arbitrarily small changes to get $A' = \begin{bmatrix} 1 & \epsilon_1 \\ \epsilon_2 & 0 \end{bmatrix}$ and $B' = \begin{bmatrix} 1 & \epsilon_3 \\ \epsilon_4 & 0 \end{bmatrix}$, we have $A' - \lambda B' = \begin{bmatrix} 1 & \epsilon_1 \\ \epsilon_2 & 0 \end{bmatrix} - \lambda \begin{bmatrix} 1 & \epsilon_3 \\ \epsilon_4 & 0 \end{bmatrix} = \begin{bmatrix} 1-\lambda & \epsilon_1 - \lambda\epsilon_3 \\ \epsilon_2 - \lambda\epsilon_4 & 0 \end{bmatrix}$ and the eigenvalues become ϵ_1/ϵ_3 and ϵ_2/ϵ_4, which can be arbitrary complex numbers. So the eigenvalues are *infinitely* sensitive.

$\diamond$

THEOREM. *Kronecker canonical form.* Let A and B be arbitrary rectangular m-by-n matrices. Then there are square nonsingular matrices P_L and P_R so that $P_L A P_R - \lambda P_L B P_R$ is block diagonal with four kinds of blocks:

$$J_m(\lambda') - \lambda I = \begin{bmatrix} \lambda' - \lambda & 1 & & \\ & \ddots & \ddots & \\ & & \ddots & 1 \\ & & & \lambda' - \lambda \end{bmatrix}, \quad \text{\textit{m-by-m Jordan block}};$$

$$N_m = \begin{bmatrix} 1 & \lambda & & \\ & 1 & \ddots & \\ & & \ddots & \lambda \\ & & & 1 \end{bmatrix}, \qquad \text{\textit{m-by-m Jordan block for} } \lambda = \infty;$$

$$L_m = \begin{bmatrix} 1 & \lambda & & \\ & \ddots & \ddots & \\ & & 1 & \lambda \end{bmatrix}, \qquad \text{\textit{m-by-(m+1) right singular block}};$$

$$L_m^T = \begin{bmatrix} 1 & & \\ \lambda & \ddots & \\ & \ddots & 1 \\ & & \lambda \end{bmatrix}, \qquad \textit{(m+1)-by-m left singular block}.$$

We call L_m a right singular block since it has a right null vector $[\lambda^m, -\lambda^{m-1}, \ldots, \pm 1]$ for all λ. L_m^T has an analogous left null vector. For a proof, see [F. Gantmacher. The Theory of Matrices, vol. II (translation). Chelsea, New York, 1959.].

Application of Kronecker Form to Differential Equations

Suppose that we want to solve $B\dot{x} = Ax + f(t)$, where $A - \lambda B$ is a singular pencil. Write

$$P_L B P_R P_R^{-1} \dot{x} = P_L A P_R P_R^{-1} x + P_L f(t).$$

Let $P_R^{-1} x = y$ and $P_L f(t) = g(t)$. Then from the above equation we obtain

$$P_L B P_R \dot{y} = P_L A P_R y + g(t),$$

or

$$P_L A P_R y - P_L B P_R \dot{y} + g(t) = 0$$

to decompose the problem into independent blocks.

There are four kinds, one for each kind in the Kronecker form. We have already dealt with $J_m(\lambda') - \lambda I$, and N_m blocks when we considered regular pencils and Weierstrass form, so we have to consider only L_m and L_m^T blocks. From the L_m blocks we get

$$\begin{bmatrix} 0 & 1 & & \\ & \ddots & \ddots & \\ & & 0 & 1 \end{bmatrix} \begin{bmatrix} \dot{y}_1 \\ \vdots \\ \dot{y}_{m+1} \end{bmatrix} = \begin{bmatrix} 1 & 0 & & \\ & \ddots & \ddots & \\ & & 1 & 0 \end{bmatrix} \begin{bmatrix} y_1 \\ \vdots \\ y_{m+1} \end{bmatrix} + \begin{bmatrix} g_1 \\ \vdots \\ g_m \end{bmatrix}$$

or

$$\begin{array}{rclcrcl} \dot{y}_2 & = & y_1 + g_1 & or & y_2(t) & = & y_2(0) + \int_0^t (y_1(\tau) + g_1(\tau)) d\tau, \\ \dot{y}_3 & = & y_2 + g_2 & or & y_3(t) & = & y_3(0) + \int_0^t (y_2(\tau) + g_2(\tau)) d\tau, \\ & \vdots & & & & & \\ \dot{y}_{m+1} & = & y_m + g_m & or & y_{m+1}(t) & = & y_{m+1}(0) + \int_0^t (y_m(\tau) + g_m(\tau)) d\tau. \end{array}$$

This means that we can choose y_1 as an arbitrary integrable function and use the above recurrence relations to get a solution. This is because we have one more unknown than equation, so the ODE is *underdetermined*. From the L_m^T blocks we get

$$\begin{bmatrix} 0 & & \\ 1 & \ddots & \\ & \ddots & 0 \\ & & 1 \end{bmatrix} \begin{bmatrix} \dot{y}_1 \\ \vdots \\ \dot{y}_m \end{bmatrix} = \begin{bmatrix} 1 & & \\ 0 & \ddots & \\ & \ddots & 1 \\ & & 0 \end{bmatrix} \begin{bmatrix} y_1 \\ \vdots \\ y_m \end{bmatrix} + \begin{bmatrix} g_1 \\ \vdots \\ g_{m+1} \end{bmatrix}$$

or

$$\begin{aligned} 0 &= y_1 + g_1, \\ \dot{y}_1 &= y_2 + g_2, \\ &\vdots \\ \dot{y}_{m-1} &= y_m + g_m, \\ \dot{y}_m &= g_{m+1}. \end{aligned}$$

Starting with the first equation, we solve to get

$$\begin{aligned} y_1 &= -g_1, \\ y_2 &= -g_2 - \dot{g}_1, \\ &\vdots \\ y_m &= -g_m - \dot{g}_{m-1} - \cdots - \frac{d^{m-1}}{dt^{m-1}} g_1 \end{aligned}$$

and the *consistency condition* $g_{m+1} = -\dot{g}_m - \cdots - \frac{d^m}{dt^m} g_1$. So unless the g_i satisfy this equation, there is no solution. Here we have one more equation than unknown, and the subproblem is *overdetermined*.

Nonlinear Eigenvalue Problems

Finally, we consider the *nonlinear eigenvalue problem* or *matrix polynomial*

$$\sum_{i=0}^{d} \lambda^i A_i = \lambda^d A_d + \lambda^{d-1} A_{d-1} + \cdots + \lambda A_1 + A_0.$$

Suppose for simplicity that the A_i are n-by-n matrices and A_d is nonsingular.

DEFINITION. The *characteristic polynomial* of the matrix polynomial (upper) is $p(\lambda) = \det(\sum_{i=0}^{d} \lambda^i A_i)$. The roots of $p(\lambda) = 0$ are defined to be the eigenvalues. One can confirm that $p(\lambda)$ has degree $d \cdot n$, so there are $d \cdot n$ eigenvalues. Suppose that γ is an eigenvalue. A nonzero vector x satisfying $\sum_{i=0}^{d} \gamma^i A_i x = 0$ is a right eigenvector for γ. A left eigenvector y is defined analogously by $\sum_{i=0}^{d} \gamma^i y^* A_i = 0$.

Example

The ODE arising in equation $M\ddot{x}(t) = -B\dot{x}(t) - Kx(t)$ is $M\ddot{x}(t) + B\dot{x}(t) + Kx(t) = 0$. If we seek solutions of the form $x(t) = e^{\lambda_i t} x_i(0)$, we get

$$e^{\lambda_i t}(\lambda_i^2 M x_i(0) + \lambda_i B x_i(0) + K x_i(0)) = 0,$$

or $\lambda_i^2 M x_i(0) + \lambda_i B x_i(0) + K x_i(0) = 0$. Thus λ_i is an eigenvalue and $x_i(0)$ is an eigenvector of the matrix polynomial $\lambda^2 M + \lambda B + K$. $\diamond$

The Symmetric Eigenproblems: perturbation theory

Suppose that A is symmetric, with eigenvalues $\alpha_1 \geq \cdots \geq \alpha_n$ and corresponding unit eigenvectors $q_1, \ldots, q_n$. Suppose E is also symmetric, and let $\hat{A} = A + E$ have perturbed eigenvalues $\hat{\alpha}_1 \geq \cdots \geq \hat{\alpha}_n$ and corresponding perturbed eigenvectors $\hat{q}_1, \ldots, \hat{q}_n$. We want to bound the differences between the eigenvalues α_i and $\hat{\alpha}_i$ and between the eigenvectors q_i and $\hat{q}_i$ in terms of the "size" of E.

THEOREM. *Weyl.* Let A and E be n-by-n symmetric matrices. Let $\alpha_1 \geq \cdots \geq \alpha_n$ be the eigenvalues of A and $\hat{\alpha}_1 \geq \cdots \geq \hat{\alpha}_n$ be the eigenvalues of $\hat{A} = A + E$. Then $|\alpha_i - \hat{\alpha}_i| \leq ||E||_2$.

COROLLARY. Let G and F be arbitrary matrices (of the same size) where $\sigma_1 \geq \cdots \geq \sigma_n$ are the singular values of G and $\sigma'_1 \geq \cdots \geq \sigma'_n$ are the singular values of $G + F$. Then $|\sigma_i - \sigma'_i| \leq ||F||_2$.

DEFINITION. The *Rayleigh quotient* of a symmetric matrix A and nonzero vector u is $\rho(u, A) \equiv (u^T Au)/(u^T u)$.

Here are some simple but important properties of $\rho(u, A)$. First, $\rho(\gamma u, A) = \rho(u, A)$ for any nonzero scalar γ. Second, if $Aq_i = \alpha_i q_i$, then $\rho(q_i, A) = \alpha_i$. More generally, suppose $Q^T AQ = \Lambda = \mathrm{diag}(\alpha_i)$ is the eigendecomposition of A, with $Q = [q_1, \ldots, q_n]$. Expand u in the basis of eigenvectors q_i as follows: $u = Q(Q^T u) \equiv Q\xi = \sum_i q_i \xi_i$. Then we can write

$$\rho(u, A) = \frac{\underbrace{\xi^T Q^T}_{u^T} A \underbrace{Q\xi}_{u}}{\xi^T Q^T Q\xi} = \frac{\xi^T \Lambda \xi}{\xi^T \xi} = \frac{\sum_i \alpha_i \xi_i^2}{\sum_i \xi_i^2}.$$

$$\rho(u, A) = \frac{\xi^T Q^T A Q \xi}{\xi^T Q^T Q \xi} = \frac{\xi^T \Lambda \xi}{\xi^T \xi} = \frac{\sum_i \alpha_i \xi_i^2}{\sum_i \xi_i^2}.$$

In other words, $\rho(u, A)$ is a weighted average of the eigenvalues of A. Its largest value, $\max_{u \neq 0} \rho(u, A)$, occurs for $u = q_1$ ($\xi = e_1$) and equals $\rho(q_1, A) = \alpha_1$.
Its smallest value, $\min_{u \neq 0} \rho(u, A)$, occurs for $u = q_n$ ($\xi = e_n$) and equals $\rho(q_n, A) = \alpha_n$. Together, these facts imply

$$\max_{u \neq 0} |\rho(u, A)| = \max(|\alpha_1|, |\alpha_n|) = ||A||_2.$$

DEFINITION. The *inertia* of a symmetric matrix A is the triple of integers $\mathrm{Inertia}(A) \equiv (\nu, \zeta, \pi)$ where ν is the number of negative eigenvalues of A, ζ is the number of zero eigenvalues of A, and π is the number of positive eigenvalues of A.

If X is orthogonal, then $X^T AX$ and A are similar and so have the same eigenvalues. When X is only nonsingular, we say $X^T AX$ and A are *congruent*.

THEOREM. *Sylvester's inertia theorem.* Let A be symmetric and X be nonsingular. Then A and X^TAX have the same inertia.

DEFINITION. Let A have eigenvalues $\alpha_1 \geq \cdots \geq \alpha_n$. Then the gap between an eigenvalue α_i and the rest of the spectrum is defined to be $\mathrm{gap}(i, A) = \min_{j \neq i} |\alpha_j - \alpha_i|$. We will also write $\mathrm{gap}(i)$ if A is understood from the context.

THEOREM. Let A be symmetric, x be a unit vector, and β be a scalar. Then A has an eigenpair $Aq_i = \alpha_i q_i$, satisfying $|\alpha_i - \beta| \leq ||Ax - \beta x||_2$. Given x, the choice $\beta = \rho(x, A)$ minimizes $||Ax - \beta x||_2$.

Proof. If β is an eigenvalue of A, then it is clear that this β is the solution of eigenvalue problem $Ax = \beta x$. Assume instead that $det(A - \beta I) \neq 0$ (nonsingular). Then $x = (A - \beta I)^{-1}(A - \beta I)x$ and

$$1 = ||x||_2 \leq ||(A - \beta I)^{-1}||_2 \cdot ||(A - \beta I)x||_2. \tag{3}$$

Writing A's eigendecomposition as $A = Q\Lambda Q^T = Q\mathrm{diag}(\alpha_1, \ldots, \alpha_n)Q^T$, we get

$$||(A - \beta I)^{-1}||_2 = ||Q(\Lambda - \beta I)^{-1}Q^T||_2 = ||(\Lambda - \beta I)^{-1}||_2 = 1/\min_i |\alpha_i - \beta|.$$

Inserting the above estimate in (3) we have

$$1 \leq \frac{1}{\min_i |\alpha_i - \beta|} \cdot ||(A - \beta I)x||_2$$

or $\min_i |\alpha_i - \beta| \leq ||(A - \beta I)x||_2$ as desired.

To show that $\beta = \rho(x, A)$ minimizes $||Ax - \beta x||_2$ we will show that x is orthogonal to $Ax - \rho(x, A)x$. Applying the Pythagorean theorem to the sum of orthogonal vectors

$$Ax - \beta x = [Ax - \underbrace{\rho(x, A)}_{\beta} x] + [(\underbrace{\rho(x, A)}_{\beta} - \beta)x]$$

yields

$$\begin{aligned} ||Ax - \beta x||_2^2 &= ||Ax - \rho(x, A)x||_2^2 + ||(\rho(x, A) - \beta)x||_2^2 \\ &\geq ||Ax - \rho(x, A)x||_2^2 \end{aligned}$$

with equality only when $\beta = \rho(x, A)$.
To confirm orthogonality of x and $Ax - \rho(x, A)x$ we need to verify that

$$x^T(Ax - \rho(x, A)x) = x^T \left(Ax - \frac{(x^T Ax)}{x^T x} x \right) = x^T Ax - x^T Ax \frac{(x^T x)}{x^T x} = 0$$

as desired. □

Relative Perturbation Theory

THEOREM. *"Relative" Weyl*. Let A have eiqenvalues α_i and $\hat{A} = X^T A X$ have eigenvalues $\hat{\alpha}_i$. Let $\epsilon \equiv \|X^T X - I\|_2$. Then $|\hat{\alpha}_i - \alpha_i| \le |\alpha_i|\epsilon$. If $\alpha_i \neq 0$, then we can also write

$$\frac{|\hat{\alpha}_i - \alpha_i|}{|\alpha_i|} \le \epsilon.$$

COROLLARY. Let G be an arbitrary matrix with singular values σ_i, and let $\hat{G} = Y^T G X$ have singular values $\hat{\sigma}_i$. Let $\epsilon \equiv \max(\|X^T X - I\|_2, ||Y^T Y - I||_2)$. Then $|\hat{\sigma}_i - \sigma_i| \leq \epsilon \sigma_i$. If $\sigma_i \neq 0$, then we can write

$$\frac{|\hat{\sigma}_i - \sigma_i|}{|\sigma_i|} \leq \epsilon.$$

DEFINITION. The *relative gap* between an eigenvalue α_i of A and the rest of the spectrum is defined to be $\mathrm{rel\_gap}(i, A) = \min_{j \neq i} \frac{|\alpha_j - \alpha_i|}{|\alpha_i|}$.

Algorithms for the Symmetric Eigenproblem

- 1. *Tridiagonal QR iteration*. This algorithm finds all the eigenvalues, and optionally all the eigenvectors, of a symmetric tridiagonal matrix. Implemented efficiently, it is currently the fastest practical method to find all the eigenvalues of a symmetric tridiagonal matrix, taking $O(n^2)$ flops. But for finding all the eigenvectors as well, QR iteration takes a little over $6n^3$ flops on average and is only the fastest algorithm for small matrices, up to about $n = 25$. This is the algorithm underlying the `Matlab" command eig".`
- `2.` *`Rayleigh quotient iteration.`* `This algorithm underlies QR iteration, but we present it separately in order to more easily analyze its extremely rapid convergence and because it may be used as an algorithm by itself. In fact, it generally converges cubically (as does QR iteration).`

- 3. *Divide-and-conquer*. This is currently the fastest method to find all the eigenvalues and eigenvectors of symmetric tridiagonal matrices larger than $n = 25$. In the worst case, divide-and-conquer requires $O(n^3)$ flops, but in practice the constant is quite small.
- 4. *Bisection and inverse iteration*. Bisection may be used to find just a subset of the eigenvalues of a symmetric tridiagonal matrix, say, those in an interval $[a, b]$ or $[\alpha_i, \alpha_{i-j}]$. It needs only $O(nk)$ flops, where k is the number of eigenvalues desired. Thus Bisection can be much faster than QR iteration when $k \ll n$, since QR iteration requires $O(n^2)$ flops. Algorithm of Inverse iteration can then be used to find the corresponding eigenvectors.

There is current research on inverse iteration addressing the problem of close eigenvalues, which may make it the fastest method to find all the eigenvectors (besides, theoretically, divide-and-conquer with the FMM). However, software implementing this improved version of inverse iteration is not yet available.:

- V. Fernando, B. Parlett, and I. Dhillon. A way to find the most redundant equation in a tridiagonal system. Berkeley Mathematics Dept. Preprint, 1995.
- B. N. Parlett and I. S. Dhillon. Fernando's solution to Wilkinson's problem: An application of double factorization. Linear Algebra AppL, 1997.
- I. S. Dhillon. A New O(n2) Algorithm for the Symmetric Tridiagonal Eigenvalue/Eigenvector Problem. Ph.D. thesis, Computer Science Division, University of California, Berkeley, May 1997.
- B. Parlett. The construction of orthogonal eigenvectors for tight clusters by use of submatrices. Center for Pure and Applied Mathematics PAM-664, University of California, Berkeley, CA, January 1996. Submitted to SIAM J. Matrix Anal. Appl.
- T.-Y. Li, H. Zhang, and X.-H. Sun. Parallel homotopy algorithm for symmetric tridiagonal eigenvalue problem. SIAM J. Sci. Statist. Comput., 12:469-487, 1991.
- T.-Y. Li and Z. Zeng. Homotopy-determinant algorithm for solving nonsymmetric eigenvalue problems. Math. Comp., 59:483-502, 1992.
- T.-Y. Li, Z. Zeng, and L. Cong. Solving eigenvalue problems of nonsymmetric matrices with real homotopies. SIAM J. Numer. Anal., 29:229-248, 1992.
- Z. Zeng. Homotopy-Determinant Algorithm for Solving Matrix Eigenvalue Problems and Its Parallelizations. Ph.D. thesis, Michigan State University, East Lansing, MI, 1991.

5. *Jacobi's method*. This method is historically the oldest method for the eigenproblem, dating to 1846. It is usually much slower than any of the above methods, taking $O(n^3)$ flops with a large constant. But the method remains interesting, because it is sometimes much more accurate than the above methods. This is because Jacobi's method can sometimes compute tiny eigenvalues much more accurately than the previous methods [J. Demmel and K. Veselic. Jacobi's method is more accurate than QR. SIAM J. Matrix Anal. AppL, 13:1204-1246, 1992.].

Tridiagonal QR iteration

Our first algorithm for the symmetric eigenproblem is completely analogous to this:
1. Given $A = A^T$, use the variation of Algorithm of reduction to upper Hessenberg form, see Lecture 10, to find an orthogonal Q so that $QAQ^T = T$ is tridiagonal.
2. Apply QR iteration to T to get a sequence $T = T_0, T_1, T_2, \ldots$ of tridiagonal matrices converging to diagonal form.

```
ALGORITHM. QR iteration: Given T_0, we iterate


  i=0
  repeat
 Factorize: T_i = Q_i R_i (the QR  decomposition)
  T_i+1 = R_i Q_i
  i=i+1
  until  convergence
```

Since $T_{i+1} = R_i Q_i = Q_i^T (Q_i R_i) Q_i = Q_i^T T_i Q_i$, T_{i+1} and T_i are orthogonally similar.

We must still describe how the shifts are chosen to implement each QR iteration. Denote the ith iterate by

$$T_i = \begin{bmatrix} a_1 & b_1 & & \\ b_1 & \ddots & \ddots & \\ & \ddots & \ddots & b_{n-1} \\ & & b_{n-1} & a_n \end{bmatrix}.$$

The simplest choice of shift would be $\sigma_i = a_n$; this is the single shift QR iteration. It turns out to be cubically convergent for almost all matrices. Unfortunately, examples exist where it does not converge [B. Parlett. The Symmetric Eigenvalue Problem. Prentice Hall, Englewood Cliffs, NJ, 1980, p. 76], so to get global convergence a slightly more complicated shift strategy is needed: We let the shift σ_i be the eigenvalue of $\begin{bmatrix} a_{n-1} & b_{n-1} \\ b_{n-1} & a_n \end{bmatrix}$ that is closest to a_n. This is called *Wilkinson's shift*.

Wilkinson's shift

To get global convergence a slightly more complicated shift strategy is needed: We let the shift σ_i be the eigenvalue of

$$\begin{bmatrix} a_{n-1} & b_{n-1} \\ b_{n-1} & a_n \end{bmatrix}$$

that is closest to a_n. This is called *Wilkinson's shift*.

THEOREM. *Wilkinson*. QR iteration with Wilkinson's shift is globally, and at least linearly, convergent. It is asymptotically cubically convergent for almost all matrices.

A proof of this theorem can be found in [B. Parlett. The Symmetric Eigenvalue Problem. Prentice Hall, Englewood Cliffs, NJ, 1980].

Rayleigh Quotient Iteration (RQI)

ALGORITHM. Rayleigh quotient iteration (RQI): Given x_0 with $||x_0|| = 1$, and a user-supplied stopping tolerance tol, we iterate

$$
\begin{aligned}
&\rho_0 = \rho(x_0, A) = \tfrac{x_0^T A x_0}{x_0^T x_0} \\
&i = 1 \\
&\textit{repeat} \\
&\quad y_i = (A - \rho_{i-1} I)^{-1} x_{i-1} \\
&\quad x_i = y_i / ||y_i||_2 \\
&\quad \rho_i = \rho(x_i, A) \\
&\quad i = i + 1 \\
&\textit{until convergence } (||Ax_i - \rho_i x_i||_2 < \mathrm{tol})
\end{aligned}
$$

When the stopping criterion is satisfied, Theorem tells us that ρ_i is within tol of an eigenvalue of A.

If one uses the shift $\sigma_i = a_{nn}$ in QR iteration and starts Rayleigh quotient iteration with $x_0 = [0, \ldots, 0, 1]^T$, then the connection between QR and inverse iteration can be used to show that the sequence of σ_i and ρ_i from the two algorithms are identical. In this case we will prove that convergence is almost always cubic.

THEOREM. Rayleigh quotient iteration is locally cubically convergent; i.e., the number of correct digits triples at each step once the error is small enough and the eigenvalue is simple.

Proof. We claim that it is enough to analyze the case when A is diagonal. To see why, write $Q^T A Q = \Lambda$, where Q is the orthogonal matrix whose columns are eigenvectors, and $\Lambda = \mathrm{diag}(\alpha_1, \ldots, \alpha_n)$ is the diagonal matrix of eigenvalues. Now change variables in Rayleigh quotient iteration to $\hat{x}_i \equiv Q^T x_i$ and $\hat{y}_i \equiv Q^T y_i$. Then

$$\rho_i = \rho(x_i, A) = \frac{x_i^T A x_i}{x_i^T x_i} = \frac{\hat{x}_i^T Q^T A Q \hat{x}_i}{\hat{x}_i^T Q^T Q \hat{x}_i} = \frac{\hat{x}_i^T \Lambda \hat{x}_i}{\hat{x}_i^T \hat{x}_i} = \rho(\hat{x}_i, \Lambda)$$

and from algorithm RQI it follows that $y_{i+1} = (A - \rho_i I)^{-1} x_i$ and thus $Q\hat{y}_{i+1} = (A - \rho_i I)^{-1} Q \hat{x}_i$, so

$$\hat{y}_{i+1} = Q^T (A - \rho_i I)^{-1} Q \hat{x}_i = (Q^T A Q - \rho_i I)^{-1} \hat{x}_i = (\Lambda - \rho_i I)^{-1} \hat{x}_i.$$

Therefore, running Rayleigh quotient iteration with A and x_0 is equivalent to running Rayleigh quotient iteration with Λ and $\hat{x}_0$. Thus we will assume without loss of generality that $A = \Lambda$ is already diagonal, so the eigenvectors of A are e_i, the columns of the identity matrix.