Jim Lambers
CME 335
Spring Quarter 2010-11
Lecture 4 Supplemental Notes

The Unsymmetric Eigenvalue Problem

Properties and Decompositions

Let A be an $n \times n$ matrix. A *nonzero* vector $\mathbf{x}$ is called an *eigenvector* of A if there exists a scalar λ such that

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

The scalar λ is called an *eigenvalue* of A, and we say that $\mathbf{x}$ is an eigenvector of A *corresponding* to λ. We see that an eigenvector of A is a vector for which matrix-vector multiplication with A is equivalent to scalar multiplication by λ.

We say that a nonzero vector $\mathbf{y}$ is a *left eigenvector* of A if there exists a scalar λ such that

$$\lambda\mathbf{y}^H = \mathbf{y}^H A.$$

The superscript H refers to the *Hermitian transpose*, which includes transposition and complex conjugation. That is, for any matrix A, $A^H = \overline{A^T}$. An eigenvector of A, as defined above, is sometimes called a *right eigenvector* of A, to distinguish from a left eigenvector. It can be seen that if $\mathbf{y}$ is a left eigenvector of A with eigenvalue λ, then $\mathbf{y}$ is also a right eigenvector of A^H, with eigenvalue $\overline{\lambda}$.

Because $\mathbf{x}$ is nonzero, it follows that if $\mathbf{x}$ is an eigenvector of A, then the matrix $A - \lambda I$ is singular, where λ is the corresponding eigenvalue. Therefore, λ satisfies the equation

$$\det(A - \lambda I) = 0.$$

The expression $\det(A-\lambda I)$ is a polynomial of degree n in λ, and therefore is called the *characteristic polynomial* of A (eigenvalues are sometimes called *characteristic values*). It follows from the fact that the eigenvalues of A are the roots of the characteristic polynomial that A has n eigenvalues, which can repeat, and can also be complex, even if A is real. However, if A is real, any complex eigenvalues must occur in complex-conjugate pairs.

The set of eigenvalues of A is called the *spectrum* of A, and denoted by $\lambda(A)$. This terminology explains why the magnitude of the largest eigenvalues is called the *spectral radius* of A. The *trace* of A, denoted by $\mathrm{tr}(A)$, is the sum of the diagonal elements of A. It is also equal to the sum of the eigenvalues of A. Furthermore, $\det(A)$ is equal to the *product* of the eigenvalues of A.

Example A 2×2 matrix

$$A = \begin{bmatrix} a & b \\ c & d \end{bmatrix}$$

has trace $\text{tr}(A) = a + d$ and determinant $\det(A) = ad - bc$. Its characteristic polynomial is

$$\det(A-\lambda I) = \begin{vmatrix} a-\lambda & b \\ c & d-\lambda \end{vmatrix} = (a-\lambda)(d-\lambda)-bc = \lambda^2-(a+d)\lambda+(ad-bc) = \lambda^2-\text{tr}(A)\lambda+\det(A).$$

From the quadratic formula, the eigenvalues are

$$\lambda_1 = \frac{a+d}{2} + \frac{\sqrt{(a-d)^2+4bc}}{2}, \quad \lambda_2 = \frac{a+d}{2} - \frac{\sqrt{(a-d)^2+4bc}}{2}.$$

It can be verified directly that the sum of these eigenvalues is equal to $\text{tr}(A)$, and that their product is equal to $\det(A)$. □

A subspace W of $\mathbb{R}^n$ is called an *invariant subspace* of A if, for any vector $\mathbf{x} \in W$, $A\mathbf{x} \in W$. Suppose that $\dim(W) = k$, and let X be an $n \times k$ matrix such that $\text{range}(X) = W$. Then, because each column of X is a vector in W, each column of AX is also a vector in W, and therefore is a linear combination of the columns of X. It follows that $AX = XB$, where B is a $k \times k$ matrix.

Now, suppose that $\mathbf{y}$ is an eigenvector of B, with eigenvalue λ. It follows from $B\mathbf{y} = \lambda\mathbf{y}$ that

$$XB\mathbf{y} = X(B\mathbf{y}) = X(\lambda\mathbf{y}) = \lambda X\mathbf{y},$$

but we also have

$$XB\mathbf{y} = (XB)\mathbf{y} = AX\mathbf{y}.$$

Therefore, we have

$$A(X\mathbf{y}) = \lambda(X\mathbf{y}),$$

which implies that λ is also an eigenvalue of A, with corresponding eigenvector $X\mathbf{y}$. We conclude that $\lambda(B) \subseteq \lambda(A)$.

If $k = n$, then X is an $n \times n$ invertible matrix, and it follows that A and B have the same eigenvalues. Furthermore, from $AX = XB$, we now have $B = X^{-1}AX$. We say that A and B are *similar matrices*, and that B is a *similarity transformation* of A.

Similarity transformations are essential tools in algorithms for computing the eigenvalues of a matrix A, since the basic idea is to apply a sequence of similarity transformations to A in order to obtain a new matrix B whose eigenvalues are easily obtained. For example, suppose that B has a 2×2 block structure

$$B = \begin{bmatrix} B_{11} & B_{12} \\ 0 & B_{22} \end{bmatrix},$$

where B_{11} is $p \times p$ and B_{22} is $q \times q$.

Let $\mathbf{x} = \begin{bmatrix} \mathbf{x}_1^T & \mathbf{x}_2^T \end{bmatrix}^T$ be an eigenvector of B, where $\mathbf{x}_1 \in \mathbb{C}^p$ and $\mathbf{x}_2 \in \mathbb{C}^q$. Then, for some scalar $\lambda \in \lambda(B)$, we have

$$\begin{bmatrix} B_{11} & B_{12} \\ 0 & B_{22} \end{bmatrix} \begin{bmatrix} \mathbf{x}_1 \\ \mathbf{x}_2 \end{bmatrix} = \lambda \begin{bmatrix} \mathbf{x}_1 \\ \mathbf{x}_2 \end{bmatrix}.$$

If $\mathbf{x}_2 \neq \mathbf{0}$, then $B_{22}\mathbf{x}_2 = \lambda\mathbf{x}_2$, and $\lambda \in \lambda(B_{22})$. But if $\mathbf{x}_2 = \mathbf{0}$, then $B_{11}\mathbf{x}_1 = \lambda\mathbf{x}_1$, and $\lambda \in \lambda(B_{11})$. It follows that, $\lambda(B) \subseteq \lambda(B_{11}) \cup \lambda(B_{22})$. However, $\lambda(B)$ and $\lambda(B_{11}) \cup \lambda(B_{22})$ have the same number of elements, so the two sets must be equal. Because A and B are similar, we conclude that

$$\lambda(A) = \lambda(B) = \lambda(B_{11}) \cup \lambda(B_{22}).$$

Therefore, if we can use similarity transformations to reduce A to such a block structure, the problem of computing the eigenvalues of A *decouples* into two smaller problems of computing the eigenvalues of B_{ii} for $i = 1, 2$. Using an inductive argument, it can be shown that if A is block upper-triangular, then the eigenvalues of A are equal to the union of the eigenvalues of the diagonal blocks. If each diagonal block is 1×1, then it follows that the eigenvalues of any upper-triangular matrix are the diagonal elements. The same is true of any lower-triangular matrix; in fact, it can be shown that because $\det(A) = \det(A^T)$, the eigenvalues of A^T are the same as the eigenvalues of A.

Example The matrix

$$A = \begin{bmatrix} 1 & -2 & 3 & -3 & 4 \\ 0 & 4 & -5 & 6 & -5 \\ 0 & 0 & 6 & -7 & 8 \\ 0 & 0 & 0 & 7 & 0 \\ 0 & 0 & 0 & -8 & 9 \end{bmatrix}$$

has eigenvalues 1, 4, 6, 7, and 9. This is because A has a block upper-triangular structure

$$A = \begin{bmatrix} A_{11} & A_{12} \\ 0 & A_{22} \end{bmatrix}, \quad A_{11} = \begin{bmatrix} 1 & -2 & 3 \\ 0 & 4 & -5 \\ 0 & 0 & 6 \end{bmatrix}, \quad A_{22} = \begin{bmatrix} 7 & 0 \\ -8 & 9 \end{bmatrix}.$$

Because both of these blocks are themselves triangular, their eigenvalues are equal to their diagonal elements, and the spectrum of A is the union of the spectra of these blocks. □

Suppose that $\mathbf{x}$ is a *normalized* eigenvector of A, with eigenvalue λ. Furthermore, suppose that P is a Householder reflection such that $P\mathbf{x} = \mathbf{e}_1$. Because P is symmetric *and* orthogonal, P is its own inverse, so $P\mathbf{e}_1 = \mathbf{x}$. It follows that the matrix P^TAP, which is a similarity transformation of A, satisfies

$$P^TAP\mathbf{e}_1 = P^TA\mathbf{x} = \lambda P^T\mathbf{x} = \lambda P\mathbf{x} = \lambda\mathbf{e}_1.$$

That is, $\mathbf{e}_1$ is an eigenvector of P^TAP with eigenvalue λ, and therefore P^TAP has the block structure

$$P^TAP = \begin{bmatrix} \lambda & \mathbf{v}^T \\ \mathbf{0} & B \end{bmatrix}.$$

Therefore, $\lambda(A) = \{\lambda\} \cup \lambda(B)$, which means that we can now focus on the $(n-1) \times (n-1)$ matrix B to find the rest of the eigenvalues of A. This process of reducing the eigenvalue problem for A to that of B is called *deflation*.

Continuing this process, we obtain the *Schur Decomposition*

$$A = Q^H T Q$$

where T is an upper-triangular matrix whose diagonal elements are the eigenvalues of A, and Q is a *unitary* matrix, meaning that $Q^H Q = I$. That is, a unitary matrix is the generalization of a real orthogonal matrix to complex matrices. Every square matrix has a Schur decomposition.

The columns of Q are called *Schur vectors*. However, for a general matrix A, there is no relation between Schur vectors of A and eigenvectors of A, as each Schur vector $\mathbf{q}_j$ satisfies $A\mathbf{q}_j = AQ\mathbf{e}_j = QT\mathbf{e}_j$. That is, $A\mathbf{q}_j$ is a linear combination of $\mathbf{q}_1, \ldots, \mathbf{q}_j$. It follows that for $j = 1, 2, \ldots, n$, the first j Schur vectors $\mathbf{q}_1, \mathbf{q}_2, \ldots, \mathbf{q}_j$ span an invariant subspace of A.

The Schur vectors and eigenvectors of A are the same when A is a *normal* matrix, which means that $A^H A = AA^H$. Any symmetric or skew-symmetric matrix, for example, is normal. It can be shown that in this case, the normalized eigenvectors of A form an orthonormal basis for $\mathbb{R}^n$. It follows that if $\lambda_1, \lambda_2, \ldots, \lambda_n$ are the eigenvalues of A, with corresponding (orthonormal) eigenvectors $\mathbf{q}_1, \mathbf{q}_2, \ldots, \mathbf{q}_n$, then we have

$$AQ = QD, \quad Q = \begin{bmatrix} \mathbf{q}_1 & \cdots & \mathbf{q}_n \end{bmatrix}, \quad D = \operatorname{diag}(\lambda_1, \ldots, \lambda_n).$$

Because Q is a unitary matrix, it follows that

$$Q^H A Q = Q^H Q D = D,$$

and A is similar to a diagonal matrix. We say that A is *diagonalizable*. Furthermore, because D can be obtained from A by a similarity transformation involving a unitary matrix, we say that A is *unitarily diagonalizable*.

Even if A is not a normal matrix, it may be diagonalizable, meaning that there exists an invertible matrix P such that $P^{-1}AP = D$, where D is a diagonal matrix. If this is the case, then, because $AP = PD$, the columns of P are eigenvectors of A, and the rows of P^{-1} are eigenvectors of A^T (as well as the left eigenvectors of A, if P is real).

By definition, an eigenvalue of A corresponds to at least one eigenvector. Because any nonzero scalar multiple of an eigenvector is also an eigenvector, corresponding to the same eigenvalue, an eigenvalue actually corresponds to an *eigenspace*, which is the span of any set of eigenvectors corresponding to the same eigenvalue, and this eigenspace must have a dimension of at least one. Any invariant subspace of a diagonalizable matrix A is a union of eigenspaces.

Now, suppose that λ_1 and λ_2 are *distinct* eigenvalues, with corresponding eigenvectors $\mathbf{x}_1$ and $\mathbf{x}_2$, respectively. Furthermore, suppose that $\mathbf{x}_1$ and $\mathbf{x}_2$ are linearly *dependent*. This means that they must be parallel; that is, there exists a nonzero constant c such that $\mathbf{x}_2 = c\mathbf{x}_1$. However, this implies that $A\mathbf{x}_2 = \lambda_2\mathbf{x}_2$ and $A\mathbf{x}_2 = cA\mathbf{x}_1 = c\lambda_1\mathbf{x}_1 = \lambda_1\mathbf{x}_2$. However, because $\lambda_1 \neq \lambda_2$, this is a contradiction. Therefore, $\mathbf{x}_1$ and $\mathbf{x}_2$ must be linearly independent.

More generally, it can be shown, using an inductive argument, that a set of k eigenvectors corresponding to k distinct eigenvalues must be linearly independent. Suppose that $\mathbf{x}_1, \ldots, \mathbf{x}_k$ are

eigenvectors of A, with distinct eigenvalues $\lambda_1, \ldots, \lambda_k$. Trivially, $\mathbf{x}_1$ is linearly independent. Using induction, we assume that we have shown that $\mathbf{x}_1, \ldots, \mathbf{x}_{k-1}$ are linearly independent, and show that $\mathbf{x}_1, \ldots, \mathbf{x}_k$ must be linearly independent as well. If they are not, then there must be constants $c_1, \ldots, c_{k-1}$, not all zero, such that

$$\mathbf{x}_k = c_1\mathbf{x}_1 + c_2\mathbf{x}_2 + \cdots + c_{k-1}\mathbf{x}_{k-1}.$$

Multiplying both sides by A yields

$$A\mathbf{x}_k = c_1\lambda_1\mathbf{x}_1 + c_2\lambda_2\mathbf{x}_2 + \cdots + c_{k-1}\lambda_{k-1}\mathbf{x}_{k-1},$$

because $A\mathbf{x}_i = \lambda_i\mathbf{x}_i$ for $i = 1, 2, \ldots, k-1$. However, because both sides are equal to $\mathbf{x}_k$, and $A\mathbf{x}_k = \lambda_k\mathbf{x}_k$, we also have

$$A\mathbf{x}_k = c_1\lambda_k\mathbf{x}_1 + c_2\lambda_k\mathbf{x}_2 + \cdots + c_{k-1}\lambda_k\mathbf{x}_{k-1}.$$

It follows that

$$c_1(\lambda_k - \lambda_1)\mathbf{x}_1 + c_2(\lambda_k - \lambda_2)\mathbf{x}_2 + \cdots + c_{k-1}(\lambda_k - \lambda_{k-1})\mathbf{x}_{k-1} = \mathbf{0}.$$

However, because the eigenvalues $\lambda_1, \ldots, \lambda_k$ are distinct, and not all of the coefficients $c_1, \ldots, c_{k-1}$ are zero, this means that we have a nontrivial linear combination of linearly independent vectors being equal to the zero vector, which is a contradiction. We conclude that eigenvectors corresponding to distinct eigenvalues are linearly independent.

It follows that if A has n distinct eigenvalues, then it has a set of n linearly independent eigenvectors. If X is a matrix whose columns are these eigenvectors, then $AX = XD$, where D is a diagonal matrix of the eigenvectors, and because the columns of X are linearly independent, X is invertible, and therefore $X^{-1}AX = D$, and A is diagonalizable.

Now, suppose that the eigenvalues of A are not distinct; that is, the characteristic polynomial has repeated roots. Then an eigenvalue with multiplicity m does not necessarily correspond to m linearly independent eigenvectors. The *algebraic multiplicity* of an eigenvalue λ is the number of times that λ occurs as a root of the characteristic polynomial. The *geometric multiplicity* of λ is the dimension of the eigenspace corresponding to λ, which is equal to the maximal size of a set of linearly independent eigenvectors corresponding to λ. The geometric multiplicity of an eigenvalue λ is always less than or equal to the algebraic multiplicity. When it is strictly less, then we say that the eigenvalue is *defective*. When both multiplicities are equal to one, then we say that the eigenvalue is *simple*.

The *Jordan canonical form* of an $n \times n$ matrix A is a decomposition that yields information about the eigenspaces of A. It has the form

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

where J has the block diagonal structure

$$J = \begin{bmatrix} J_1 & 0 & \cdots & 0 \\ 0 & J_2 & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ 0 & \cdots & 0 & J_p \end{bmatrix}.$$

Each diagonal block J_p is a *Jordan block* that has the form

$$J_i = \begin{bmatrix} \lambda_i & 1 & & \\ & \lambda_i & \ddots & \\ & & \lambda_i & 1 \\ & & & \lambda_i \end{bmatrix}, \quad i = 1, 2, \ldots, p.$$

The number of Jordan blocks, p, is equal to the number of linearly independent eigenvectors of A. The diagonal element of J_i, λ_i, is an eigenvalue of A. The number of Jordan blocks associated with λ_i is equal to the geometric multiplicity of λ_i. The sum of the sizes of these blocks is equal to the algebraic multiplicity of λ_i. If A is diagonalizable, then each Jordan block is 1×1.

Example Consider a matrix with Jordan canonical form

$$J = \begin{bmatrix} 2 & 1 & 0 & & & \\ 0 & 2 & 1 & & & \\ 0 & 0 & 2 & & & \\ & & & 3 & 1 & \\ & & & 0 & 3 & \\ & & & & & 2 \end{bmatrix}.$$

The eigenvalues of this matrix are 2, with algebraic multiplicity 4, and 3, with algebraic multiplicity 2. The geometric multiplicity of the eigenvalue 2 is 2, because it is associated with two Jordan blocks. The geometric multiplicity of the eigenvalue 3 is 1, because it is associated with only one Jordan block. Therefore, there are a total of three linearly independent eigenvectors, and the matrix is not diagonalizable. □

The Jordan canonical form, while very informative about the eigensystem of A, is not practical to compute using floating-point arithmetic. This is due to the fact that while the eigenvalues of a matrix are continuous functions of its entries, the Jordan canonical form is not. If two computed eigenvalues are nearly equal, and their computed corresponding eigenvectors are nearly parallel, we do not know if they represent two distinct eigenvalues with linearly independent eigenvectors, or a multiple eigenvalue that could be defective.

Perturbation Theory

Just as the problem of solving a system of linear equations $A\mathbf{x} = \mathbf{b}$ can be sensitive to perturbations in the data, the problem of computing the eigenvalues of a matrix can also be sensitive to perturbations in the matrix. We will now obtain some results concerning the extent of this sensitivity.

Suppose that A is obtained by perturbing a diagonal matrix D by a matrix F whose diagonal entries are zero; that is, $A = D + F$. If λ is an eigenvalue of A with corresponding eigenvector $\mathbf{x}$, then we have

$$(D - \lambda I)\mathbf{x} + F\mathbf{x} = \mathbf{0}.$$

If λ is not equal to any of the diagonal entries of A, then $D - \lambda I$ is nonsingular and we have

$$\mathbf{x} = -(D - \lambda I)^{-1}F\mathbf{x}.$$

Taking ∞-norms of both sides, we obtain

$$\|\mathbf{x}\|_\infty = \|(D - \lambda I)^{-1}F\mathbf{x}\|_\infty \leq \|(D - \lambda I)^{-1}F\|_\infty \|\mathbf{x}\|_\infty,$$

which yields

$$\|(D - \lambda I)^{-1}F\|_\infty = \max_{1\leq i\leq n} \sum_{j=1, j\neq i}^{n} \frac{|f_{ij}|}{|d_{ii} - \lambda|} \geq 1.$$

It follows that for some i, $1 \leq i \leq n$, λ satisfies

$$|d_{ii} - \lambda| \leq \sum_{j=1, j\neq i}^{n} |f_{ij}|.$$

That is, λ lies within one of the *Gerschgorin circles* in the complex plane, that has center a_{ii} and radius

$$r_i = \sum_{j=1, j\neq i}^{n} |a_{ij}|.$$

This is result is known as the *Gerschgorin Circle Theorem.*

Example The eigenvalues of the matrix

$$A = \begin{bmatrix} -5 & -1 & 1 \\ -2 & 2 & -1 \\ 1 & -3 & 7 \end{bmatrix}$$

are

$$\lambda(A) = \{6.4971, 2.7930, -5.2902\}.$$

The Gerschgorin disks are

$$D_1 = \{z \in \mathbb{C} \,|\, |z-7| \leq 4\}, \quad D_2 = \{z \in \mathbb{C} \,|\, |z-2| \leq 3\}, \quad D_3 = \{z \in \mathbb{C} \,|\, |z+5| \leq 2\}.$$

We see that each disk contains one eigenvalue. □

It is important to note that while there are n eigenvalues and n Gerschgorin disks, it is not necessarily true that each disk contains an eigenvalue. The Gerschgorin Circle Theorem only states that all of the eigenvalues are contained within the *union* of the disks.

Another useful sensitivity result that applies to diagonalizable matrices is the *Bauer-Fike Theorem*, which states that if $X^{-1}AX = \text{diag}(\lambda_1, \ldots, \lambda_n)$, and μ is an eigenvalue of a perturbed matrix $A + E$, then

$$\min_{\lambda \in \lambda(A)} |\lambda - \mu| \leq \kappa_p(X)\|E\|_p.$$

That is, μ is within $\kappa_p(X)\|E\|_p$ of an eigenvalue of A. It follows that if A is "nearly non-diagonalizable", which can be the case if eigenvectors are nearly linearly dependent, then a small perturbation in A could still cause a large change in the eigenvalues.

It would be desirable to have a concrete measure of the sensitivity of an eigenvalue, just as we have the condition number for a system of linear equations. To that end, we assume that λ is a simple eigenvalue of an $n \times n$ matrix A that has Jordan canonical form $J = X^{-1}AX$. Then, $\lambda = J_{ii}$ for some i, and $\mathbf{x}_i$, the ith column of X, is a corresponding right eigenvector.

If we define $Y = X^{-H} = (X^{-1})^H$, then $\mathbf{y}_i$ is a left eigenvector of A corresponding to λ. From $Y^H X = I$, it follows that $\mathbf{y}^H\mathbf{x} = 1$. We now let A, λ, and $\mathbf{x}$ be functions of a parameter ϵ that satisfy

$$A(\epsilon)\mathbf{x}(\epsilon) = \lambda(\epsilon)\mathbf{x}(\epsilon), \quad A(\epsilon) = A + \epsilon F, \quad \|F\|_2 = 1.$$

Differentiating with respect to ϵ, and evaluating at $\epsilon = 0$, yields

$$F\mathbf{x} + A\mathbf{x}'(0) = \lambda\mathbf{x}'(0) + \lambda'(0)\mathbf{x}.$$

Taking the inner product of both sides with $\mathbf{y}$ yields

$$\mathbf{y}^H F\mathbf{x} + \mathbf{y}^H A\mathbf{x}'(0) = \lambda\mathbf{y}^H\mathbf{x}'(0) + \lambda'(0)\mathbf{y}^H\mathbf{x}.$$

Because $\mathbf{y}$ is a left eigenvector corresponding to λ, and $\mathbf{y}^H\mathbf{x} = 1$, we have

$$\mathbf{y}^H F\mathbf{x} + \lambda\mathbf{y}^H\mathbf{x}'(0) = \lambda\mathbf{y}^H\mathbf{x}'(0) + \lambda'(0).$$

We conclude that

$$|\lambda'(0)| = |\mathbf{y}^H F\mathbf{x}| \leq \|\mathbf{y}\|_2\|F\|_2\|\mathbf{x}\|_2 \leq \|\mathbf{y}\|_2\|\mathbf{x}\|_2.$$

However, because θ, the angle between $\mathbf{x}$ and $\mathbf{y}$, is given by

$$\cos\theta = \frac{\mathbf{y}^H\mathbf{x}}{\|\mathbf{y}\|_2\|\mathbf{x}\|_2} = \frac{1}{\|\mathbf{y}\|_2\|\mathbf{x}\|_2},$$

it follows that

$$|\lambda'(0)| \leq \frac{1}{|\cos\theta|}.$$

We define $1/|\cos\theta|$ to be the *condition number* of the simple eigenvalue λ. We require λ to be simple because otherwise, the angle between the left and right eigenvectors is not unique, because the eigenvectors themselves are not unique.

It should be noted that the condition number is also defined by $1/|\mathbf{y}^H\mathbf{x}|$, where $\mathbf{x}$ and $\mathbf{y}$ are normalized so that $\|\mathbf{x}\|_2 = \|\mathbf{y}\|_2 = 1$, but either way, the condition number is equal to $1/|\cos\theta|$. The interpretation of the condition number is that an $O(\epsilon)$ perturbation in A can cause an $O(\epsilon/|\cos\theta|)$ perturbation in the eigenvalue λ. Therefore, if $\mathbf{x}$ and $\mathbf{y}$ are nearly orthogonal, a large change in the eigenvalue can occur. Furthermore, if the condition number is large, then A is close to a matrix with a multiple eigenvalue.

Example The matrix

$$A = \begin{bmatrix} 3.1482 & -0.2017 & -0.5363 \\ 0.4196 & 0.5171 & 1.0888 \\ 0.3658 & -1.7169 & 3.3361 \end{bmatrix}$$

has a simple eigenvalue $\lambda = 1.9833$ with left and right eigenvectors

$$\mathbf{x} = \begin{bmatrix} 0.4150 & 0.6160 & 0.6696 \end{bmatrix}^T, \quad \mathbf{y} = \begin{bmatrix} -7.9435 & 83.0701 & -70.0066 \end{bmatrix}^T$$

such that $\mathbf{y}^H\mathbf{x} = 1$. It follows that the condition number of this eigenvalue is $\|\mathbf{x}\|_2\|\mathbf{y}\|_2 = 108.925$. In fact, the nearby matrix

$$B = \begin{bmatrix} 3.1477 & -0.2023 & -0.5366 \\ 0.4187 & 0.5169 & 1.0883 \\ 0.3654 & -1.7176 & 3.3354 \end{bmatrix}$$

has a double eigenvalue that is equal to 2. □

We now consider the sensitivity of repeated eigenvalues. First, it is important to note that while the eigenvalues of a matrix A are continuous functions of the entries of A, they are not necessarily differentiable functions of the entries. To see this, we consider the matrix

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

where $a > 0$. Computing its characteristic polynomial $\det(A - \lambda I) = \lambda^2 - 2\lambda + 1 - a\epsilon$ and computings its roots yields the eigenvalues $\lambda = 1 \pm \sqrt{a\epsilon}$. Differentiating these eigenvalues with respect to ϵ yields

$$\frac{d\lambda}{d\epsilon} = \pm\sqrt{\frac{a}{\epsilon}},$$

which is undefined at $\epsilon = 0$. In general, an $O(\epsilon)$ perturbation in A causes an $O(\epsilon^{1/p})$ perturbation in an eigenvalue associated with a $p \times p$ Jordan block, meaning that the "more defective" an eigenvalue is, the more sensitive it is.

We now consider the sensitivity of eigenvectors, or, more generally, invariant subspaces of a matrix A, such as a subspace spanned by the first k Schur vectors, which are the first k columns in a matrix Q such that $Q^H AQ$ is upper triangular. Suppose that an $n \times n$ matrix A has the Schur decomposition

$$A = QTQ^H, \quad Q = \begin{bmatrix} Q_1 & Q_2 \end{bmatrix}, \quad T = \begin{bmatrix} T_{11} & T_{12} \\ 0 & T_{22} \end{bmatrix},$$

where Q_1 is $n \times r$ and T_{11} is $r \times r$. We define the *separation* between the matrices T_{11} and T_{22} by

$$\text{sep}(T_{11}, T_{22}) = \min_{X \neq 0} \frac{\|T_{11}X - XT_{22}\|_F}{\|X\|_F}.$$

It can be shown that an $O(\epsilon)$ perturbation in A causes a $O(\epsilon/\text{sep}(T_{11}, T_{22}))$ perturbation in the invariant subspace Q_1.

We now consider the case where $r = 1$, meaning that Q_1 is actually a vector $\mathbf{q}_1$, that is also an eigenvector, and T_{11} is the corresponding eigenvalue, λ. Then, we have

$$\begin{aligned}
\text{sep}(\lambda, T_{22}) &= \min_{X \neq 0} \frac{\|\lambda X - XT_{22}\|_F}{\|X\|_F} \\
&= \min_{\|\mathbf{y}\|_2 = 1} \|\mathbf{y}^H (T_2 2 - \lambda I)\|_2 \\
&= \min_{\|\mathbf{y}\|_2 = 1} \|(T_2 2 - \lambda I)^H \mathbf{y}\|_2 \\
&= \sigma_{\min}((T_{22} - \lambda_I)^H) \\
&= \sigma_{\min}(T_{22} - \lambda I),
\end{aligned}$$

since the Frobenius norm of a vector is equivalent to the vector 2-norm. Because the smallest singular value indicates the distance to a singular matrix, $\text{sep}(\lambda, T_{22})$ provides a measure of the separation of λ from the other eigenvalues of A. It follows that eigenvectors are more sensitive to perturbation if the corresponding eigenvalues are clustered near one another. That is, eigenvectors associated with nearby eigenvalues are "wobbly".

It should be emphasized that there is no direct relationship between the sensitivity of an eigenvalue and the sensitivity of its corresponding invariant subspace. The sensitivity of a simple eigenvalue depends on the angle between its left and right eigenvectors, while the sensitivity of an invariant subspace depends on the clustering of the eigenvalues. Therefore, a sensitive eigenvalue, that is nearly defective, can be associated with an insensitive invariant subspace, if it is distant from other eigenvalues, while an insensitive eigenvalue can have a sensitive invariant subspace if it is very close to other eigenvalues.

The Symmetric Eigenvalue Problem

Properties and Decompositions

The eigenvalue problem for a real, symmetric matrix A, or a complex, *Hermitian* matrix A, for which $A = A^H$, is a considerable simplification of the eigenvalue problem for a general matrix. Consider the Schur decomposition $A = QTQ^H$, where T is upper-triangular. Then, if A is Hermitian, it follows that $T = T^H$. But because T is upper-triangular, it follows that T must be diagonal. That is, any symmetric real matrix, or Hermitian complex matrix, is unitarily diagonalizable, as stated previously because A is normal. What's more, because the Hermitian transpose includes complex conjugation, T must equal its complex conjugate, which implies that the eigenvalues of A are real, even if A itself is complex.

Because the eigenvalues are real, we can order them. By convention, we prescribe that if A is an $n \times n$ symmetric matrix, then it has eigenvalues

$$\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_n.$$

Furthermore, by the *Courant-Fischer Minimax Theorem*, each of these eigenvalues has the following characterization:

$$\lambda_k = \max_{\dim(S)=k} \min_{\mathbf{y} \in S, \mathbf{y} \neq \mathbf{0}} \frac{\mathbf{y}^H A \mathbf{y}}{\mathbf{y}^H \mathbf{y}}.$$

That is, the kth largest eigenvalue of A is equal to the maximum, over all k-dimensional subspaces of $\mathbb{C}^n$, of the minimum value of the *Rayleigh quotient*

$$r(\mathbf{y}) = \frac{\mathbf{y}^H A \mathbf{y}}{\mathbf{y}^H \mathbf{y}}, \quad \mathbf{y} \neq \mathbf{0},$$

on each subspace. It follows that λ_1, the largest eigenvalue, is the absolute maximum value of the Rayleigh quotient on all of $\mathbb{C}^n$, while λ_n, the smallest eigenvalue, is the absolute minimum value. In fact, by computing the gradient of $r(\mathbf{y})$, it can be shown that every eigenvector of A is a critical point of $r(\mathbf{y})$, with the corresponding eigenvalue being the value of $r(\mathbf{y})$ at that critical point.

Perturbation Theory

In the symmetric case, the Gerschgorin circles become Gerschgorin intervals, because the eigenvalues of a symmetric matrix are real.

Example The eigenvalues of the 3×3 symmetric matrix

$$A = \begin{bmatrix} -10 & -3 & 2 \\ -3 & 4 & -2 \\ 2 & -2 & 14 \end{bmatrix}$$

are

$$\lambda(A) = \{14.6515, 4.0638, -10.7153\}.$$

The Gerschgorin intervals are

$$D_1 = \{x \in \mathbb{R} \,|\, |x - 14| \leq 4\}, \quad D_2 = \{x \in \mathbb{R} \,|\, |x - 4| \leq 5\}, \quad D_3 = \{x \in \mathbb{R} \,|\, |x + 10| \leq 5\}.$$

We see that each intervals contains one eigenvalue. □

The characterization of the eigenvalues of a symmetric matrix as constrained maxima of the Rayleigh quotient lead to the following results about the eigenvalues of a perturbed symmetric matrix. As the eigenvalues are real, and therefore can be ordered, we denote by $\lambda_i(A)$ the ith largest eigenvalue of A.

Theorem (Wielandt-Hoffman) If A and $A + E$ are $n \times n$ symmetric matrices, then

$$\sum_{i=1}^{n} (\lambda_i(A + E) - \lambda_i(A))^2 \leq \|E\|_F^2.$$

It is also possible to bound the distance between individual eigenvalues of A and $A + E$.

Theorem If A and $A + E$ are $n \times n$ symmetric matrices, then

$$\lambda_n(E) \leq \lambda_k(A + E) - \lambda_k(A) \leq \lambda_1(E).$$

Furthermore,

$$|\lambda_k(A + E) - \lambda_k(A)| \leq \|E\|_2.$$

The second inequality in the above theorem follows directly from the first, as the 2-norm of the symmetric matrix E, being equal to its spectral radius, must be equal to the larger of the absolute value of $\lambda_1(E)$ or $\lambda_n(E)$.

Theorem (Interlacing Property) If A is an $n \times n$ symmetric matrix, and A_r is the $r \times r$ leading principal minor of A, then, for $r = 1, 2, \ldots, n - 1$,

$$\lambda_{r+1}(A_{r+1}) \leq \lambda_r(A_r) \leq \lambda_r(A_{r+1}) \leq \cdots \leq \lambda_2(A_{r+1}) \leq \lambda_1(A_r) \leq \lambda_1(A_{r+1}).$$

For a symmetric matrix, or even a more general normal matrix, the left eigenvectors and right eigenvectors are the same, from which it follows that every simple eigenvalue is "perfectly conditioned"; that is, the condition number $1/|\cos\theta|$ is equal to 1 because $\theta = 0$ in this case. However, the same results concerning the sensitivity of invariant subspaces from the nonsymmetric case apply in the symmetric case as well: such sensitivity increases as the eigenvalues become more clustered,

even though there is no chance of a defective eigenvalue. This is because for a nondefective, repeated eigenvalue, there are infinitely many possible bases of the corresponding invariant subspace. Therefore, as the eigenvalues approach one another, the eigenvectors become more sensitive to small perturbations, for any matrix.

Let Q be an $n \times r$ matrix with orthonormal columns, meaning that $Q_1^T Q_1 = I_r$. If it spans an invariant subspace of an $n \times n$ symmetric matrix A, then $AQ_1 = Q_1 S$, where $S = Q_1^T A Q_1$. On the other hand, if range(Q_1) is *not* an invariant subspace, but the matrix

$$AQ_1 - Q_1 S = E_1$$

is small for any given $r \times r$ symmetric matrix S, then the columns of Q_1 define an *approximate* invariant subspace.

It turns out that $\|E_1\|_F$ is minimized by choosing $S = Q_1^T A Q_1$. Furthermore, we have

$$\|AQ_1 - S_1 S\|_F = \|P_1^{\perp} A Q_1\|_F,$$

where $P_1^{\perp} = I - Q_1 Q_1^T$ is the orthogonal projection into $(\text{range}(Q_1))^{\perp}$, and there exist eigenvalues $\mu_1, \ldots, \mu_r \in \lambda(A)$ such that

$$|\mu_k - \lambda_k(S)| \leq \sqrt{2}\|E_1\|_2, \quad k = 1, \ldots, r.$$

That is, r eigenvalues of A are close to the eigenvalues of S, which are known as *Ritz values*, while the corresponding eigenvectors are called *Ritz vectors*. If $(\theta_k, \mathbf{y}_k)$ is an eigenvalue-eigenvector pair, or an *eigenpair* of S, then, because S is defined by $S = Q_1^T A Q_1$, it is also known as a *Ritz pair*. Furthermore, as θ_k is an approximate eigenvalue of A, $Q_1 \mathbf{y}_k$ is an approximate corresponding eigenvector.

To see this, let σ_k (not to be confused with singular values) be an eigenvalue of S, with eigenvector $\mathbf{y}_k$. We multiply both sides of the equation $S\mathbf{y}_k = \sigma_k \mathbf{y}_k$ by Q_1:

$$Q_1 S \mathbf{y}_k = \sigma_k Q_1 \mathbf{y}_k.$$

Then, we use the relation $AQ_1 - Q_1 S = E_1$ to obtain

$$(AQ_1 - E_1)\mathbf{y}_k = \sigma_k Q_1 \mathbf{y}_k.$$

Rearranging yields

$$A(Q_1 \mathbf{y}_k) = \sigma_k (Q_1 \mathbf{y}_k) + E_1 \mathbf{y}_k.$$

If we let $\mathbf{x}_k = Q_1 \mathbf{y}_k$, then we conclude

$$A\mathbf{x}_k = \sigma_k \mathbf{x}_k + E_1 \mathbf{y}_k.$$

Therefore, $\|E_1\|$ is small in some norm, $Q_1 \mathbf{y}_k$ is nearly an eigenvector.

Power Iterations

The Power Method, when applied to a symmetric matrix to obtain its largest eigenvalue, is more effective than for a general matrix: its rate of convergence $|\lambda_2/\lambda_1|^2$, meaning that it generally converges twice as rapidly.

Let A be an $n \times n$ symmetric matrix. Even more rapid convergence can be obtained if we consider a variation of the Power Method. *Inverse Iteration* is the Power Method applied to $(A - \mu I)^{-1}$. The algorithm is as follows:

Choose $\mathbf{x}_0$ so that $\|\mathbf{x}_0\|_2 = 1$
for $k = 0, 1, 2, \ldots$ **do**
  Solve $(A - \mu I)\mathbf{z}_k = \mathbf{x}_k$ for $\mathbf{z}_k$
  $\mathbf{x}_{k+1} = \mathbf{z}_k / \|\mathbf{z}_k\|_2$
end

Let A have eigenvalues $\lambda_1, \ldots, \lambda_n$. Then, the eigenvalues of $(A - \mu I)^{-1}$ matrix are $1/(\lambda_i - \mu)$, for $i - 1, 2, \ldots, n$. Therefore, this method finds the eigenvalue that is closest to μ.

Now, suppose that we vary μ from iteration to iteration, by setting it equal to the *Rayleigh quotient*

$$r(\mathbf{x}) = \frac{\mathbf{x}^H A \mathbf{x}}{\mathbf{x}^H \mathbf{x}},$$

of which the eigenvalues of A are constrained extrema. We then obtain *Rayleigh Quotient Iteration*:

Choose a vector $\mathbf{x}_0$, $\|\mathbf{x}_0\|_2 = 1$
for $k = 0, 1, 2, \ldots$ **do**
  $\mu_k = \mathbf{x}_k^H A \mathbf{x}_k$
  Solve $(A - \mu_k I)\mathbf{z}_k = \mathbf{x}_k$ for $\mathbf{z}_k$
  $\mathbf{x}_{k+1} = \mathbf{z}_k / \|\mathbf{z}_k\|_2$
end

When this method converges, it converges *cubically* to an eigenvalue-eigenvector pair. To see this, consider the diagonal 2×2 matrix

$$A = \begin{bmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{bmatrix}, \quad \lambda_1 > \lambda_2.$$

This matrix has eigenvalues λ_1 and λ_2, with eigenvectors $\mathbf{e}_1$ and $\mathbf{e}_2$. Suppose that $\mathbf{x}_k = \begin{bmatrix} c_k & s_k \end{bmatrix}^T$, where $c_k^2 + s_k^2 = 1$. Then we have

$$\mu_k = r(\mathbf{x}_k) = \begin{bmatrix} c_k & s_k \end{bmatrix} \begin{bmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{bmatrix} \begin{bmatrix} c_k \\ s_k \end{bmatrix} = \lambda_1 c_k^2 + \lambda_2 s_k^2.$$

From

$$A - \mu_k I = \begin{bmatrix} \lambda_1 - (\lambda_1 c_k^2 + \lambda_2 s_k^2) & 0 \\ 0 & \lambda_2 - (\lambda_1 c_k^2 + \lambda_2 s_k^2) \end{bmatrix} = (\lambda_1 - \lambda_2) \begin{bmatrix} s_k^2 & 0 \\ 0 & -c_k^2 \end{bmatrix},$$

we obtain

$$\mathbf{z}_k = \frac{1}{\lambda_1 - \lambda_2} \begin{bmatrix} c_k/s_k^2 \\ -s_k/c_k^2 \end{bmatrix} = \frac{1}{c_k^2 s_k^2 (\lambda_1 - \lambda_2} \begin{bmatrix} c_k^3 \\ -s_k^3 \end{bmatrix}.$$

Normalizing yields

$$\mathbf{x}_{k+1} = \frac{1}{\sqrt{c_k^6 + s_k^6}} \begin{bmatrix} c_k^3 \\ -s_k^3 \end{bmatrix},$$

which indicates cubic convergence to a vector that is parallel to $\mathbf{e}_1$ or $\mathbf{e}_2$, provided $|c_k| \neq |s_k|$.

It should be noted that Inverse Iteration is also useful for a general (unsymmetric) matrix A, for finding *selected* eigenvectors after computing the Schur decomposition $A = QTQ^H$, which reveals the eigenvalues of A, but not the eigenvectors. Then, a computed eigenvalue can be used as the shift μ, causing rapid convergence to a corresponding eigenvector. In fact, in practice a single iteration is sufficient. However, when no such information about eigenvalues is available, Inverse Iteration is far more practical for a symmetric matrix than an unsymmetric matrix, due to the superior convergence of the Power Method in the symmetric case.

Reduction to Tridiagonal Form

A symmetric Hessenberg matrix is tridiagonal. Therefore, the same kind of Householder reflections that can be used to reduce a general matrix to Hessenberg form can be used to reduce a symmetric matrix A to a tridiagonal matrix T. However, the symmetry of A can be exploited to reduce the number of operations needed to apply each Householder reflection on the left and right of A.

It can be verified by examining the structure of the matrices involved, and the rows and columns influenced by Givens rotations, that if T is a symmetric tridiagonal matrix, and $T = QR$ is its QR factorization, then Q is upper Hessenberg, and R is *upper-bidiagonal* (meaning that it is upper-triangular, with upper bandwidth 1, so that all entries below the main diagonal and above the superdiagonal are zero). Furthermore, $\tilde{T} = RQ$ is also tridiagonal.

Because each Givens rotation only affects $O(1)$ nonzero elements of a tridiagonal matrix T, it follows that it only takes $O(n)$ operations to compute the QR factorization of a tridiagonal matrix, and to multiply the factors in reverse order. However, to compute the eigenvectors of A as well as the eigenvalues, it is necessary to compute the product of all of the Givens rotations, which still takes $O(n^2)$ operations.

The Implicit Q Theorem applies to symmetric matrices as well, meaning that if two orthogonal similarity transformations reduce a matrix A to *unreduced* tridiagonal form, and they have the same first column, then they are essentially equal, as are the tridiagonal matrices that they produce.