Lectures - Week 8
Eigenvalues & Numerical Methods for Finding a Single Eigenvalue and the Singular Value Decomposition Theorem

Review

First we will summarize some of the facts we saw last week for the algebraic eigenvalue problem of finding a nonzero vector $\vec{x}$ and a scalar λ such that $A\vec{x} = \lambda\vec{x}$

- The characteristic equation $\det(A - \lambda I) = 0$ is a polynomial of degree n in λ and the eigenvalues of A are the roots of this equation.
- If we know an eigenvector $\vec{x}$ of A, then the corresponding eigenvalue can be found by the Rayleigh Quotient $\lambda = (\vec{x}TA\vec{x})/(\vec{x}TA\vec{x})$.
- If we know an eigenvalue λ of A, then the corresponding eigenvector is in the null space of $A - \lambda I$.
- A is not invertible if and only if at least one of its eigenvalues is zero; in that case the corresponding eigenvector is a vector in $\mathcal{N}(A)$.
- If $A\vec{x} = \lambda\vec{x}$ then $(\lambda^k, \vec{x})$ is an eigenpair of A^k and $(1/\lambda, \vec{x})$ is an eigenpair of A^{-1}, if it exists.
- The algebraic multiplicity (a.m.) is the number of times an eigenvalue is repeated and the geometric multiplicity (g.m.) is the number of linearly independent eigenvector corresponding to the eigenvalue; g.m. $\leq$ a.m.; if g.m. ¡ a.m. for any eigenvalue we say the matrix is *defective*; if an $n \times n$ matrix has n linearly independent eigenvectors then we say it is *nondefective* or equivalently has a complete set of linearly independent eigenvectors; if an eigenvalue is distinct then a.m.=g.m.=1
- If A is symmetric then its eigenvalues are real.
- If A is symmetric and positive definite then its eigenvalues are real and greater than zero.
- If A is orthogonal then its eigenvalues have magnitude one.
- Eigenvectors corresponding to distinct eigenvalues are linearly independent.
- Similar matrices (i.e., $B = M^{-1}AM$) have the same eigenvalues.
- Schur's Theorem. Any matrix can be made orthogonally similar to a triangular matrix and any symmetric matrix can be made orthogonally similar to a diagonal matrix.
- If $M^{-1}AM = \Lambda$ is a diagonal matrix, then the diagonal entries of Λ are the eigenvalues of A; we say A is diagonalizable. A is diagonalizable if and only if it is nondefective; the matrix that diagonalizes A is the matrix whose columns are the eigenvectors of A.
- What matrices are guaranteed to be similar to a diagonal matrix? A matrix with distinct eigenvalues is diagonalizable because it has n linearly independent eigenvectors; a symmetric matrix.

The matrix norm induced by the Euclidean vector norm

When we studied norms, we defined $\|A\|_2$ to be the matrix norm of A induced by the ℓ_2 vector norm (the Euclidean norm)

$$\|A\|_2 = \max_{x \neq 0} \frac{\|A\vec{x}\|_2}{\|\vec{x}\|_2}$$

but we did not discover an easy way to compute it as we did for the norms $\|A\|_1, \|A\|_\infty$. It turns out that $\|A\|_2$ is related to the eigenvalues of A.
The set of eigenvalues of a matrix is sometimes called its *spectrum* and the largest eigenvalue in magnitude called its *spectral radius* and denoted $\rho(A)$. In many problems the spectral radius tells us a lot about a matrix. For example, in iterative methods for linear systems the spectral radius of the iteration matrix (which depends on the method used) must be ≤ 1 to guarantee convergence.

Lemma The $\|A\|_2$ is determined by

$$\|A\|_2 = \sqrt{\rho(A^T A)}\,.$$

Recall that $A^T A$ is a symmetric matrix which is always at least positive semi-definite and thus has real eigenvalues ≥ 0 so taking the square root makes sense. Also if A has linearly independent columns it is also positive definitWe will not prove this here because it is quite lengthy but we would demonstrate it in a similar way to the proofs for the other norms; i.e., we show $\|A\|_2 \leq \sqrt{\rho(A^T A)}$ and then show the opposite inequality. An easy way to remember this is to consider each term in the numerator and denominator of the definition where we have squared the expression for convenience

$$\|A\|_2^2 = \max_{x \neq 0} \frac{\|A\vec{x}\|_2^2}{\|\vec{x}\|_2^2} = \max_{x \neq 0} \frac{\vec{x}^T A^T A \vec{x}}{\vec{x}^T \vec{x}}$$

Now if $\vec{x}$ is an eigenvector of A then this quotient is just an eigenvalue of $A^T A$. Of course this doesn't prove the result because how do we know that the maximum occurs when $\vec{x}$ is an eigenvector of A but rather it is a way to remember the definition.
Once we have this way of calculating $\|A\|_2$ we can easily determine the $\mathcal{K}_2(A)$. When A is symmetric it has an especially nice form.

Lemma Let A be an $n \times n$ invertible matrix. Then $\mathcal{K}_2(A)$ is determined by

$$\mathcal{K}_2(A) = \|A\|_2 \|A^{-1}\|_2 = \sqrt{\rho(A^T A)}\sqrt{\rho(A^{-T} A^{-1})} = \frac{\sigma_{\max}}{\sigma_{\min}}$$

where $\sigma_{\max}$ is the spectral radius of $A^T A$, i.e., the dominant eigenvalue in magnitude, and $\sigma_{\min}$ is the smallest eigenvalue of $A^T A$ in magnitude.
Moreover, if A is symmetric and has eigenvalues ordered as $\lambda_{\min} = \lambda_1 \leq \lambda_2 \leq \cdots \leq \lambda_n = \lambda_{\max}$ then

$$\mathcal{K}_2(A) = \rho(A)\rho(A^{-1}) = \frac{\lambda_{\max}}{\lambda_{\min}}$$

The expression for $\mathcal{K}_2(A)$ is just an obvious substitution of the definition of the condition number $\mathcal{K}(A)$ and the $\|B\|_2$ for $B = A$ and $B = A^{-1}$. Of course we should point out again that the matrix A^TA is symmetric positive definite because A is invertible and so its eigenvalues σ_i are real and positive so that this expression makes sense. The simplified expression when A is symmetric requires a bit of work. Let $A = A^T$ so that $A^TA = A^2$ and $A^{-T}A^{-1} = (A^{-1})^2$. Using this in the definition of the condition number and the expression for calculating the 2−norm we obtain

$$\mathcal{K}_2(A) = \|A\|_2\|A^{-1}\|_2 = \sqrt{\rho(A^TA)}\sqrt{\rho(A^{-T}A^{-1})} = \sqrt{\rho(A^2)}\sqrt{\rho\big((A^{-1})^2\big)}$$

Recall that if λ is an eigenvalue of A then λ^2 is an eigenvalue of A^2 and $1/\lambda$ is an eigenvalue of A^{-1} and we know $\lambda \neq 0$ because A is invertible. By assumption $\lambda_{\max}$ is the largest eigenvalue of A so $\rho(A) = \lambda_{\max}$ and thus $\rho(A^2) = \lambda^2_{\max}$. Similarly $\lambda_{\min}$ is the smallest eigenvalue of A and thus $1/\lambda_{\min}$ is the largest eigenvalue of A^{-1}; so $\rho(A^{-1}) = 1/\lambda_{\min}$ and thus $\rho((A^{-1})^2) = 1/\lambda^2_{\min}$. Using these results gives us

$$\mathcal{K}_2(A) = \sqrt{\rho(A^2)}\sqrt{\rho\big((A^{-1})^2\big)} = \sqrt{\lambda^2_{\max}}\sqrt{\frac{1}{\lambda^2_{\min}}} = \frac{\lambda_{\max}}{\lambda_{\min}} \quad \text{when } A = A^T.$$

Example Find $\mathcal{K}_2(A)$ for each A

$$A_1 = \begin{pmatrix} 2 & 1 \\ 1 & 3 \end{pmatrix} \qquad A_2 = \begin{pmatrix} 2 & 1 \\ 0 & 3 \end{pmatrix}$$

A_1 is a symmetric matrix so we can simply calculate the eigenvalues and take the ratio of the maximum over the minimum. The eigenvalues of A are the roots of the characteristic polynomial

$$\det\begin{pmatrix} 2-\lambda & 1 \\ 1 & 3-\lambda \end{pmatrix} = (2-\lambda)(3-\lambda) - 1 = 0 \Rightarrow \lambda^2 - 5\lambda + 5 = 0 \Rightarrow \lambda = \frac{5 \pm \sqrt{5}}{2}$$

so $\mathcal{K}_2(A_2) = \frac{5+\sqrt{5}}{5-\sqrt{5}} \approx 17.94$.

The matrix A_2 is not symmetric so we have to form $A_2^TA_2$ and look at its eigenvalues; we have

$$A_2^TA_2 = \begin{pmatrix} 2 & 0 \\ 1 & 3 \end{pmatrix}\begin{pmatrix} 2 & 1 \\ 0 & 3 \end{pmatrix} = \begin{pmatrix} 4 & 2 \\ 2 & 10 \end{pmatrix}$$

and its eigenvalues are the roots of the characterisitic polynomial

$$(4-\lambda)(10-\lambda) - 4 = 0 \Rightarrow \lambda^2 - 14\lambda + 36 = 0 \Rightarrow \lambda = 7 \pm \sqrt{13}$$

and thus $\mathcal{K}_2(A_2) \approx 3.12$.

Gerschgorin Circle Theorem

Before we look at methods for finding a specific eigenvalue of a matrix, we consider a localization theorem which gives us an estimate for a region that the eigenvalues are in; this could give us a "ball park" estimate on the size of the eigenvalues. This theorem is just one of many localization theorems available for eigenvalues.

Theorem Gerschgorin Circle Theorem Let A be an $n \times n$ real matrix and define the circles or disks

$$\mathcal{D}_i = \{z \text{ such that } |z - a_{ii}| \leq \sum_{\substack{i=1 \\ i \neq j}}^{n} |a_{ij}|\}$$

Then the eigenvalues of A lie in the union of these disks, $\cup_{i=1}^{n} \mathcal{D}_i$. Moreover if p of the disks are disjoint then exactly p eigenvalues lie in the union of these p disks.

Lets look at what this theorem says. First of all, we did not assume that A was symmetric so its eigenvalues can be complex and so the disks are regions in the complex plane and z here represents a complex number. If A happens to be symmetric then the disks are really just intervals on the real line. The $i$ith disk is centered at the diagonal entry a_{ii} of the matrix with radius determined by the sum of the absolute values of the off-diagonal entries in that row.

Example Use the Gerschgorin Circle Theorem to obtain an estimate for each eigenvalue

$$A_1 = \begin{pmatrix} 2 & 2 & -2 \\ 2 & 4 & 1 \\ 1 & 1 & 10 \end{pmatrix}$$

For A_1 we first create the disks in the complex plane

$$\mathcal{D}_1 = \{z \text{ such that } |z - 2| \leq 4\}$$

$$\mathcal{D}_2 = \{z \text{ such that } |z - 4| \leq 3\}$$

$$\mathcal{D}_3 = \{z \text{ such that } |z - 10| \leq 2\}$$

When we draw these disks we see that $\mathcal{D}_3$ is disjoint from the other two so one eigenvalue lies in $\mathcal{D}_3$ and two in $\mathcal{D}_1 \cup \mathcal{D}_2$. Note that this does NOT say that one eigenvalue is in each disk; in this case all we know is that two eigenvalues are in the union of $\mathcal{D}_1$ and $\mathcal{D}_2$ so, for example, both could be in $\mathcal{D}_1$ or both in $\mathcal{D}_2$.

For A_2 the matrix is symmetric so instead of disks we have intervals on the real line.

The Power Method and its Variants

We now look at a algorithms for determining a specific eigenvalue; for example the spectral radius of a matrix, the minimum eigenvalue (in magnitude) or the eigenvalue nearest a specific value. We will not look at the QR method or its variants for finding all eigenvalues and eigenvectors of a matrix.

The Power Method is an iterative method for find $\rho(A)$ which simply takes powers of the original matrix A. The assumption for the method to work is that A must have a complete set of linearly independent eigenvectors. We will assume that A has real eigenvalues and a complete set of linearly independent eigenvectors; for example, if A is symmetric so that we know its eigenpairs are real and it is guaranteed to have a complete set of linearly independent eigenvectors. The method will give us an approximation to the eigenvector corresponding to $\rho(A)$ and we will use the Rayleigh Quotient to obtain an approximation to the corresponding eigenvalue.

Power Method - Case 1

Let A have real eigenvalues and a complete set of linearly independent eigenvectors. Assume that the dominant eigenvalue is unique (not repeated) and the eigenvalues are ordered as follows

$$|\lambda_1| > |\lambda_2| \geq |\lambda_3| \geq \cdots \geq |\lambda_n|$$

and the associated eigenvectors are denoted $\vec{v}_i$, $i = 1, \ldots, n$.

Given $\vec{x}^0$ then we generate a sequence of iterates $\vec{x}^k$ from

$$x^{k+1} = \sigma_k A x^k$$

where σ_k is a suitably chosen scaling factor. Then $x^k \to C\vec{v}_1$ as $k \to \infty$ and the Rayleigh Quotient $((\vec{x}^k)^T A\vec{x}^k)/(\vec{x}^k)^T\vec{x}^k \to \lambda_1$.

In this algorithm we are simply multiplying the previous iterate by the matrix A and scaling the result. We have to scale because otherwise the entries in the resulting vectors could become unbounded or approach zero. We now want to see why this method converges to the eigenvector corresponding to the dominant eigenvalue and at what rate it converges.

From the algorithm we choose an arbitrary starting vector $\vec{x}^0$ so

$$\vec{x}^1 = \sigma_1 A\vec{x}^0, \quad \vec{x}^2 = \sigma_2 A\vec{x}^1 = \sigma_1\sigma_2 A^2\vec{x}^0$$

Continuing in this manner we see that

$$\vec{x}^k = \sigma_k A\vec{x}^{k-1} = \sigma_k\sigma_{k-1}A^2\vec{x}^{k-2} = \cdots = \left(\Pi_{i=1}^k \sigma_i\right) A^k\vec{x}^0 \,.$$

Because we have assumed that A has a complete set of linearly independent eigenvectors these vectors can be used as a basis for $\mathbf{R}^n$. Consequently there are constants c_i such that

$$x^0 = \sum_{i=1}^n c_i\vec{v}_i$$

Using this expression in our formula for x^k and the fact that $A^k\vec{v}_i = \lambda_i^k\vec{v}_i$ gives

$$\vec{x}^k = \left(\Pi_{i=1}^k \sigma_i\right) A^k\vec{x}^0 = \epsilon_k A^k\Big[\sum_{i=1}^n c_i\vec{v}_i\Big] = \epsilon_k\Big[\sum_{i=1}^n c_i A\vec{v}_i\Big] = \epsilon_k\Big[\sum_{i=1}^n c_i\lambda_i\vec{v}_i\Big] \,.$$

where we have written the product of the constants $\Pi_{i=1}^{k}\sigma_i$ as ϵ_k for ease of exposition. We now factor out the dominant eigenvalue term λ_1^k to get

$$\vec{x}^k = \epsilon_k \lambda_1^k \Big[c_1\vec{v}_1 + c_2\Big(\frac{\lambda_2}{\lambda_1}\Big)^k \vec{v}_2 + c_3\Big(\frac{\lambda_3}{\lambda_1}\Big)^k \vec{v}_3 + \cdots + c_n\Big(\frac{\lambda_n}{\lambda_1}\Big)^k \vec{v}_n \Big] .$$

As $k \to \infty$ all the terms in the expression except the first approach zero because we have assumed that $\lambda_1 > \lambda_i$ for all $i \neq 1$. Now if we don't scale, then the first term either approaches ∞ or 0 depending on whether $\lambda_1 > 1$ or $\lambda_1 < 1$. One can demonstrate that if the scaling is chosen as, e.g., $\|\cdot\|_\infty$ then the first term goes to a a scalar times $\vec{v}_1$, the eigenvector corresponding to λ_1; this calculation is straightforward but tedious and we don't include it here. As $k \to \infty$ the largest term in the expression (not counting the first) is a constant times $(\lambda_2/\lambda_1)^k$ because we have assumed the ordering $|\lambda_2| \geq |\lambda_3| \geq \cdots \geq |\lambda_n|$. Consequently the rate at which we have convergence is governed by $(|\lambda_2|/|\lambda_1|)^k$. This means, e.g., if $\lambda_2 = .5\lambda_1$ at the tenth iteration we would have $(\lambda_2/\lambda_1)^{10} = (0.5)^10 \approx .000976$ but if the eigenvalues are clustered, e.g., $\lambda_2 = 0.95\lambda_1$ at the tenth iteration we would have $(\lambda_2/\lambda_1)^{10} = (.95)^{10} \approx 0.5987$ and the convergence would be very slow.

In practice we must scale the iterates at each step. One way to do this is to normalize the iterates by some norm. For example, we could compute our iterates as

$$\vec{y}^{k+1} = A\vec{x}^k \quad \Rightarrow \quad \vec{x}^{k+1} = \frac{\vec{y}^{k+1}}{\|\vec{y}^{k+1}\|}$$

where we could use the Euclidean norm, the infinity norm or any of our choosing. This way our iterates have length one in some norm. Also note that our analysis works for any starting vector $\vec{x}^0$; it only relies on the fact that A is non-defective.

Example Calculate four iterates of the power method for

$$A = \begin{pmatrix} 2 & -1 \\ -1 & 2 \end{pmatrix} \qquad \vec{x}_0 = \begin{pmatrix} 1 \\ 0 \end{pmatrix}$$

This matrix has eigenvalues of 3 and 1. As a scaling use $\|\cdot\|_\infty$ and calculate the corresponding approximate eigenvalue by using the Rayleigh quotient. Tabulate your results and include columns for the error in the eigenvalue and a column giving $(\lambda_2/\lambda_1)^k = (1/3)^k$.

k	$y^k = A\vec{x}^{k-1}$	$x^{k+1} = \vec{y}^k/\|y^k\|_\infty$	λ_1^k	error	$(\lambda_2/\lambda_1)^k$
0		$(1,0)^T$			
1	$(2,-1)^T$	$(1,-1/2)^T$	2.8	0.2	0.33333
2	$(5/2,-2)^T$	$(1,-4/5)^T$	2.9761	0.02439	0.11111
3	$(14/5,-13/15)^T$	$(1,-13/14)^T$	2.9973	0.00274	0.037037
4	$(41/14,-20/7)^T$	$(1,-40/41))^T$	2.99997	0.00003387	0.004115

It turns out that if the dominant eigenvalue of A is repeated r times but A still has a complete set of linearly independent eigenvectors then $\vec{x}^k$ approaches a linear combination of the r eigenvectors corresponding to λ_1 which is itself an eigenvector corresponding to λ_1 (see homework for proof). If the root is not unique, e.g., the spectral radius is one but both ± 1 are eigenvalues then the algorithm can be modified to work.

Inverse Power Method

Suppose we want to determine the smallest eigenvalue in magnitude. Because we know that if λ_i is an eigenvalue of A then $1/\lambda_i$ is an eigenvalue of A^{-1}; thus one over the smallest eigenvalue of A in magnitude is the largest eigenvalue of A^{-1}. Consequently we can simply use the Power Method with the matrix A^{-1} to determine the eigenvector corresponding to the smallest eigenvalue of A. However, we also know that we never really want to calculate A^{-1} so we want to implement this without actually forming A^{-1}. How do we do this? As usual, we do an LU decomposition of A once and then perform solves at each iteration.

Inverse Power Method - Case 1 Let A have real eigenvalues and a complete set of linearly independent eigenvectors. Assume that the smallest eigenvalue in magnitude is unique (not repeated) and the eigenvalues are ordered as follows

$$|\lambda_1| \geq |\lambda_2| \geq \cdots \geq |\lambda_{n-1}| > |\lambda_n| \Rightarrow \frac{1}{|\lambda_n|} > \frac{1}{|\lambda_{n-1}|} \geq \frac{1}{|\lambda_{n-2}|} \geq \cdots \geq \frac{1}{|\lambda_1|}$$

and the associated eigenvectors are denoted $\vec{v}_i$, $i = 1, \ldots, n$.
Given $\vec{x}^0$ then we generate a sequence of iterates $\vec{x}^k$ from

$$x^{k+1} = \sigma_k A^{-1} x^k$$

where σ_k is a suitably chosen scaling factor. This is equivalent to solving the linear system

$$A x^{k+1} = \sigma_k x^k$$

To implement this algorithm we determine $A = LU$ (or use $PA = LU$) and for each iteration we perform a back and forward solve. Convergence is analogous to the Power Method. Specifically we form $A = LU$ before the iteration starts then for each iteration we must do the following.

- solve $L\vec{y} = \sigma_k \vec{x}^k$ using a forward solve
- solve $U\vec{z} = \vec{y}$ using a back solve
- set $\vec{x}^{k+1} = \vec{z}/\sigma_{k+1}$ for a scaling σ_{k+1} Of course when we use the Rayleigh Quotient we are computing an approximation to an eigenvalue of A^{-1} so we need to take the reciprocal to obtain the eigenvalue of A.

Shifted Inverse Power

Sometimes we might want to find the eigenvalue nearest some number, say μ. In addition, we saw that the convergence for the Power Method (and thus the Inverse Power Method) can be slow if the eigenvalues are clustered so sometimes we can perform a shift in the method to obtain faster convergence.

First note that if $A\vec{x} = \lambda\vec{x}$ then for the shifted matrix $A - \mu I$ we have

$$A\vec{x} = \lambda\vec{x} \Rightarrow A\vec{x} - \mu I\vec{x} = \lambda\vec{x} - \mu I\vec{x} \Rightarrow (A - \mu I)\vec{x} = (\lambda - \mu)\vec{x}$$

so the eigenvalues of $A - \mu I$ are just $\lambda - \mu$ and the eigenvectors are the same as A. Suppose now that we want to find the eigenvalue of A that is nearest to μ. This means that we are looking for the *smallest* eigenvalue of $A - \mu I$ so we just apply the Inverse Power Method using the matrix $A - \mu I$. When we use the Rayleigh Quotient we are finding an approximation to an eigenvalue of $(A - \mu I)^{-1}$ so we first take the reciprocal to get an approximation of an eigenvalue of $A - \mu I$ and then shift the result by μ because if σ is an eigenvalue of $A - \mu I$ then $\sigma = \lambda - \mu \Rightarrow \lambda = \sigma + \mu$.