Lecture Notes for Math 104: Fall 2010 (Week 9)

Jacob Bernstein

CHAPTER 1

Twenty-Sixth Lectures

In the this lecture we further discuss properties of the Eigenvalues and Eigenvectors. In particular, we derive some consequences of the Schur factorization discussed last lecture. (Note Lecture Twenty-Five was accidentally overwritten).

1. Applications of the Schur Factorization

Recall, last time we showed that:

THEOREM 1.1. *Every square matrix $A \in \mathbb{C}^{m\times m}$ has a Schur factorization. That is*

$$A = QTQ^*$$

where $Q \in \mathbb{C}^{m\times m}$ is unitary and $T \in \mathbb{C}^{m\times m}$ is upper triangular.

REMARK 1.2. If T is diagonal then A is diagonizable and is indeed is *unitarly* diagonizable.

One nice thing about upper triangular matrices is that the entries on their diagonal are the eigenvalues:

THEOREM 1.3. *Let*

$$X = \begin{bmatrix} x_{11} & x_{12} & \cdots \\ 0 & x_{22} & \cdots \\ \vdots & & \ddots \end{bmatrix}$$

be upper triangular. Then $\Lambda(X) = \{x_{11}, \ldots, x_{mm}\}$.

PROOF. Expanding out the determinant one can compute the characteristic polynomial of X to be

$$P_X(z) = (z - x_{11}) \cdots (z - x_{mm}).$$

One readily sees that the roots are then the elements on the diagonal of X. □

As a consequence, if we can find a Schur factorization of a matrix A we can find the eigenvalues of a matrix. In order to make this precise idea of a *similarity transformation*. This is just another word for changing the basis that one uses to represent the matrix.

DEFINITION 1.4. We say two matrices $A, B \in \mathbb{C}^{m\times m}$ are *similar* if there is a non-singular matrix X so that the matrix $B = X^{-1}AX$.

As we've seen B is the matrix A in the basis given by the columns of X. An important fact which follows from properties of the determinant is that if A and B are similar matrices then $P_A(z) = P_B(z)$, that is A and B have the same characteristic polynomial. In particular, A and B have the same eigenvalues with

the same algebraic multiplicities. In fact, as A and B are similar there is a non-singular $X \in \mathbb{C}^{m\times m}$ so that $B = X^{-1}AX$. Clearly, if $\mathbf{v}$ is an eigenvector of A corresponding to $\lambda \in \Lambda(A) = \Lambda(B)$, then $X^{-1}\mathbf{v}$ is an eigenvector of B corresponding to λ. In particular, the geometric multiplicity of λ with respect to A and B are the same. As mentioned, the proof uses from properties of the determinant. We refer to Theorem 24.3 of Trefethen and Bau.

A consequence of the Schur factorization is that any matrix $A \in \mathbb{C}^{m\times m}$ is similar to an upper triangular matrix $T \in \mathbb{C}^{m\times m}$ and hence the eigenvalues of A can be determined from the diagonal of T.

We can also use the Schur factorization to prove things:

THEOREM 1.5. *Let $A \in \mathbb{C}^{m\times m}$ be hermitian. Then A has real eigenvalues, is non-defective and there is an orthonormal set of eigenvectors of A.*

PROOF. Let $A = QTQ^*$ be a Schur factorization of A. One has $A^* = A$ so $QT^*Q^* = QTQ^*$ that is $T^* = T$. Since T is upper triangular this means that T is diagonal and all the entries on the diagonal are real. This implies that the eigenvalues of A are all real as desired. Finally, as T is diagonal, the columns of Q are the eigenvectors of A. □

REMARK 1.6. Another way to say this is that when A is hermitan it is *unitarily diagonizable*.

2. Applications of Eigenvalues

Once one knows the eigenvalues and eigenvectors of a matrix A one can tell a number of useful facts about A right away. However, one can also tell many useful things about iterates of A that is matrices of the form A^n (i..e the matrix obtained by multiplying A by itself n-times). In particular, it is relatively painless to compute A^n given such information. Indeed, suppose A is non-defective and so is diagonizable, i.e. we can write

$$A = X\Lambda X^{-1}$$

for some non-singular X and diagonal Λ here

$$\Lambda = \begin{bmatrix} \lambda_1 & 0 & \cdots & 0 \\ 0 & \ddots & 0 & \\ 0 & \cdots & 0 & \lambda_m \end{bmatrix}$$

Then it is simple to see that

$$A^2 = X\Lambda X^{-1}X\Lambda X^{-1} = X\Lambda^2 X^{-1}$$

and so by induction

$$A^n = X\Lambda X^{-1}X\Lambda X^{-1} = X\Lambda^n X^{-1}$$

Notice that knowing the SVD does not allow for such a nice formula. In practice, the SVD gives alot of information about the matrix A, but tells one little about iterates of A.

Another thing we can do is take square-roots of (some) matrices. Consider first a diagonal matix

$$A = \begin{bmatrix} a_1 & 0 & \cdots & 0 \\ 0 & \ddots & 0 & \\ \vdots & & & \\ 0 & \cdots & 0 & a_m \end{bmatrix}$$

We want to find a B so that $B^2 = A$. For simplicity, we assume $a_i \geq 0$. Then we can take

$$B = \begin{bmatrix} \sqrt{a_1} & 0 & \cdots & 0 \\ 0 & \ddots & 0 & \\ \vdots & & & \\ 0 & \cdots & 0 & \sqrt{a_m} \end{bmatrix}$$

and $B^2 = A$ and so we write $B = \sqrt{A}$ in this case. More generally, we say a hermitian matrix A is *positive semi-definite* if all the eigenvalues of A are non-negative. This is equivalent to $\langle A\mathbf{x}, \mathbf{x}\rangle \geq 0$ for all $\mathbf{x} \in \mathbb{C}^m$ (A still hermitian). Then we can check that there is a hermitian matrix B with $B^2 = A$. Indeed, as A is hermitan it is (unitarily) diagonizable, so

$$A = Q\Lambda Q^*$$

As all the eigenvalues of A are non-negative all the entries of Λ are non-negative, so we just set

$$B = Q\sqrt{\Lambda}Q^*$$

CHAPTER 2

Twenty-Seventh Lecture

We discuss here iterative methods of determining eigenvalues and eigenvectors.

As previously mentioned finding eigenvalues is not an easy procedure. This is because finding the roots of the characteristic polynomial is a non-linear problem, and is computationally difficult. Thus, we will instead discuss some other approaches. The methods we discuss are not the ones used in practice, but are related and will give some insight into how one would numerically find eigenvalues.

It turns out to be the case that the discussion is vastly simplified if we restrict attention to real symmetric matrices. I.e. $A \in \mathbb{R}^{m\times m}$ and $A^* = A^\top = A$. Notice, the eigenvalues (and hence also eigenvectors) are real. To fix notation for this lecture we let $\lambda_1, \ldots, \lambda_m$ be the eigenvectors of A and $\mathbf{q}_1, \ldots, \mathbf{q}_m$ the associated eigenvalues normalize so $||\mathbf{q}_j||_2 = 1$ (so they form an orthonormal basis of $\mathbb{R}^m$. We also order the eigenvalus so that $|\lambda_1| \geq |\lambda_2| \geq \cdots \geq |\lambda_m|$.

1. Rayleigh Quotient

One way to think about an eigenvalues is as follows: Fix a vector $\mathbf{x} \in \mathbb{R}^m$. We seek the scalar $\alpha \in \mathbb{R}$ that makes $\mathbf{x}$ an as close as possible to beign an eigenvector. I.e. we want to minimize

$$||A\mathbf{x} - \alpha\mathbf{x}||_2$$

Of course if $\mathbf{x}$ is actually an eigenvector this is minimized when α is actually the associated eigenvalue as the value is zero.

We can re-formulate this question as follows: We are trying to solve the overdetermined system of equations :

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

in the one unknown α in the sense of least squares. To make this easier to parse, let us think of $\mathbf{x}$ as an $m \times 1$ matrix and write it as X. Then we are solving

$$X\left[\alpha\right] = A\mathbf{x}$$

in the sense of least squares.

To do this, we need to find P_X the projector onto $R(X) = span(\mathbf{x})$. This is given by

$$P = \frac{\mathbf{x}\mathbf{x}^*}{||\mathbf{x}||_2^2} = \frac{\mathbf{x}\mathbf{x}^\top}{||\mathbf{x}||_2^2}$$

Hence we may take

$$\alpha\mathbf{x} = PA\mathbf{x} = \frac{\mathbf{x}^\top A\mathbf{x}}{||\mathbf{x}||_2}\mathbf{x} = \frac{\langle A\mathbf{x}, \mathbf{x}\rangle}{||\mathbf{x}||_2^2}\mathbf{x}$$

This value α is called the *Rayleigh Quotient* and denote it by $r(\mathbf{x})$. So

$$r(\mathbf{x}) = \frac{\langle A\mathbf{x}, \mathbf{x}\rangle}{||\mathbf{x}||_2^2}$$

Notice that when $\mathbf{x} = \mathbf{q}_j$ is an eigenvector, $r(\mathbf{x}) = \lambda_j$ the associated eigenvalue.

We may think of

$$r : \mathbb{R}^m \to \mathbb{R}$$

as a function of several variables. It is not hard to see that away from $\mathbf{x} = 0$ this is a smooth function (i.e. all partial derivatives exist and are continuous). A straight forward computation gives

$$\nabla r(\mathbf{x}) = \frac{2}{||\mathbf{x}||_2^2} \left(A\mathbf{x} - r(\mathbf{x})\mathbf{x}\right)$$

(here we have taken the gradient of r). In particular, the *critical values* of r are precisely the eigenvalues of A. While the *critical points* are the eigenvectors. To make this clearer one usually restricts r to the sphere $||\mathbf{x}||_2 = 1$.

How does this help us? Well we always know that the maximum of r and the minimum of r on the sphere $||\mathbf{x}||_2 = 1$ are critical points of r. In particular, these give us eigenvalues. This should remind you of the SVD. This approach (i.e. looking for the maximum and minimum) won't help us directly as it is not very computational. However, it does give us some useful information. Namely, notice that when $\mathbf{x}$ is an eigenvector, $r(\mathbf{x})$ gives the associated eigenvalue. Moreover, if we have reason to believe that $\mathbf{x}$ is near to an eigenvector $\mathbf{q}_j$ then $r(\mathbf{x})$ is near λ_j. Indeed, we Taylor's theorem gives that the following estimate holds:

$$r(\mathbf{x}) - r(\mathbf{q}_j) = O(||\mathbf{x} - \mathbf{q}_j||_2^2), \mathbf{x} \to \mathbf{q}_j \tag{1.1}$$

Here we used that ∇r vanishes at $\mathbf{x} = \mathbf{q}_j$. Notice that for $\epsilon > 0$ small that ϵ^2 is much smaller. In other words, if $\mathbf{x}$ is close to $\mathbf{q}_j$ then $r(\mathbf{x})$ is very close to λ_j.

2. Power Iteration

We introduce now a method called *power iteration* that finds the largest eigenvalue and eigenvector of matrices A under certain conditions on A. The basic idea is that repeated multplication by A tends to amplify the eigenvector corresponding to the largest eigenvalue more than the other eigenvectors. That is if we start with an appropriate $\mathbf{v}$ then consider $\mathbf{v}^{(k)} = A^k\mathbf{v}$, if one expresses $\mathbf{v}$ in the basis of eigenvectors $\{\mathbf{q}_1, \ldots, \mathbf{q}_m\}$ the coefficient in front of $\mathbf{q}_1$ should be much larger than all the other coefficients.

More precisely, start with a (randomly choosen) vector $\mathbf{v}^{(0)}$ with $||\mathbf{v}^{(0)}||_2 = 1$. And consider the iterative construction:

$$\mathbf{v}^{(k+1)} = \frac{A\mathbf{v}^{(k)}}{||A\mathbf{v}^{(k)}||_2}, \lambda^{(k+1)} = r(\mathbf{v}^{(k+1)})$$

Then in good circumstances one has that $\mathbf{v}^{(k)} \to \mathbf{q}_1$ and $\lambda^{(k)} \to \lambda_1$.

THEOREM 2.1. *Suppose that* $|\lambda_1| > |\lambda_2| \geq \ldots |\lambda_m| \geq 0$ *and* $\langle \mathbf{q}_1, \mathbf{v}^{(0)} \rangle \neq 0$ *then the iterates above satisfy*

$$||\mathbf{v}^{(k)} - (\pm\mathbf{q}_1)||_2 = O(\left|\frac{\lambda_2}{\lambda_1}\right|^k), |\lambda^{(k)} - \lambda_1| = O(\left|\frac{\lambda_2}{\lambda_1}\right|^{2k})$$

REMARK 2.2. The signs in front of the $\mathbf{q}_1$ are an unfortunate technical annoyance. If $\lambda_1 > 0$ they may always be taken to be positive, while if $\lambda_1 < 0$ they alternate in k.

PROOF. We note that as the $\mathbf{q}_i$ form an orthonormal basis we can write $\mathbf{v}^{(0)}$ as

$$\mathbf{v}^{(0)} = a_1\mathbf{q}_1 + a_2\mathbf{q}_2 + \cdots + a_m\mathbf{q}_m.$$

Notice that $a_1 = \langle \mathbf{v}^{(0)}, \mathbf{q}_1 \rangle \neq 0$. Then (here c_k is a normalizing term):

$$\begin{aligned}\mathbf{v}^{(k)} &= c_k A^k \mathbf{v}^(0) \\ &= c_k(a_1\lambda_1^k\mathbf{q}_1 + \ldots + a_m\lambda_m^k\mathbf{q}_m) \\ &= c_k\lambda_1^k\left(a_1\mathbf{q}_1 + a_2(\lambda_2/\lambda_1)^k\mathbf{q}_2 + \ldots + a_m(\lambda_m/\lambda_1)^k\mathbf{q}_m\right)\end{aligned}$$

Tthe first estimate follows from this by noting that for $j > 1$, $\left(\frac{\lambda_j}{\lambda_1}\right)^k \to 0$ as $k \to \infty$ at the desired rate. The second follows from this and the quadratic estimate (1.1). When $\lambda_1 > 0$ the signs are all positive if $\lambda_1 < 0$ they alternate. □

Notice that as long as the largest (in magnituded) two eigenvalues have distinct magnitudes (not something one know a priori–a serious drawback) then the power iterates converge to the largest eigenvalue at a rate determined by the ratio between the two eigenvalues. This illustrates some of the drawbacks of this method.

Again the main idea of the method is that successive multiplications by A tends to amplify the part of $\mathbf{v}^{(0)}$ that corresponds the eigenvector $\mathbf{q}_1$ (i.e. the eigenvector associated to λ_1) much more than an other part of $\mathbf{v}^{(0)}$. In particular, after many iterations "most" of $A^k\mathbf{v}^{(0)}$ is in the direction of $\mathbf{q}_1$.

3. Inverse Iteration:NIC

As we saw above there are two major drawbacks to power iteration. First it only finds the largest eigenvalue. Second if there is not a large amount of separation between the largest two eigenvalues the convergence is slow. A way to overcome the first issue is to consider the matrix

$$A - \mu I$$

for some $\mu \in \mathbb{R}$ to be specified. The important point is that $\Lambda(A - \mu I) = \Lambda(A) - \mu$. I.e. the eigenvalues are $\lambda_i - \mu$. If μ is not an eigenvalue, then it is straightforward to see that $A - \mu I$ is invertible and the eigenvalues of

$$(A - \mu I)^{-1}$$

are $(\lambda_i - \mu)^{-1}$. That is if μ is near to the eigenvalue λ_{i_0} then $(A - \mu I)^{-1}$ has a very large eigenvalue given by $(\lambda_{i_0} - \mu)^{-1}$. If we then use this with power iteration it converges to λ_{i_0} and the associated eigenvector $\mathbf{q}_{i_0}$. That is we can find all the eigenvalues, at least as long as we start near enougth.

This procedure is known as inverse iteration. The basic idea is the start with a vector $\mathbf{v}^{(0)}$ with $||\mathbf{v}^{(0)}||_2 = 1$. Now iteratively do the following procedure:

(1) Solve for $(A - \mu I)\mathbf{w} = \mathbf{v}^{(k)}$.
(2) Set $\mathbf{v}^{(k+1)} = \frac{\mathbf{w}}{||\mathbf{w}||_2}$
(3) Set $\lambda^{(k+1)} = r(\mathbf{v}^{(k+1)})$.

In ideal situations one then gets convergence to an eigenvalue.

THEOREM 3.1. *Suppose that λ_{i_0} is the closes eigenvalue of A to μ and λ_{i_1} is the second closest and that $|\mu - \lambda_{i_0}| < |\mu - \lambda_{i_1}| \leq |\mu - \lambda_j|$ for $j \neq i_0$. Further suppose that $\langle \mathbf{v}^{(0)}, \mathbf{q}_{i_0} \rangle \neq 0$. Then the iterates of inverse iteration satisfy*

$$||\mathbf{v}^{(k)} - (\pm \mathbf{q}_{i_0})||_2 = O\left(\left|\frac{\mu - \lambda_{i_0}}{\mu - \lambda_{i_1}}\right|^k\right), |\lambda^{(k)} - \lambda_{i_0}| = O\left(\left|\frac{\mu - \lambda_{i_0}}{\mu - \lambda_{i_1}}\right|^{2k}\right)$$

as $k \to \infty$.

Notice this is essentially the same rate of growth as before, but it does allow one to find different eigenvalues. There is still an issue when one has nearby eigenvalues.

One way to think of the inverse iteration is that it is a way to transform an eigenvalue estimate to an eigenvector estimate. I.e. if we you have a pretty good idea of what one of the eigenvalues is, application of inverse iteration gives you a much better idea, as well as giving an associated eigenvector. This contrasts with (1.1) where having a vector $\mathbf{x}$ that is near to an eigenvector means that $r(\mathbf{x})$ is quite close to an eigenvalue. By a clever combination of inverse iteration with the Rayleigh quotient one obtains a algorithm that converges quite rapidly to an eigenvalue, provided one starts near enough.