Math 642:550 — Summer 2001
MTTh 6:15–8:45 PM   Hill 525
Prof. Bumby

Schur's unitary triangularization theorem

Introduction. It may seem that unitary matrices (or real orthogonal matrices) are very special. For a fixed matrix A, the matrix $U^{-1}AU$ expresses the same linear transformation as A using the *orthonormal* basis formed by the columns of U.

Although only a limited family of coordinate changes are allowed here, we shall see that U can be found that makes $U^{-1}AU$ upper triangular. Such a triangular form shows the eigenvalues an allows simple determination of the eigenvectors, so it is not surprising that such information about A will be needed in the construction of U. However, the construction will be done one column at a time, so only a single eigenvector will be needed at each stage. Furthermore, once the first element of such a basis has been chosen, the remaining ones are confined to an $n-1$ dimensional subspace, and it will be possible to determine an action on this space that will be used in later steps. This action will have eigenvectors even if all eigenvectors of A have been used in previous steps. This triangularization result is strong enough to be used to extend proofs of important theoretical results like the Cayley-Hamilton theorem that are usually proved for diagonalizable matrices before approaching the general case.

Although Schur's theorem appears as result 5R in the textbook, the treatment here follows Roger A. Horn & Charles R. Johnson, *Matrix Analysis*, Cambridge, 1985. That book, with its companion, *Topics in Matrix Analysis*, provide a good source for more details on topics introduced in this course, along with some more advanced topics. The treatment is more theoretical than our textbook, but still down-to-earth. The book also has many references, some remarkably recent. Schur's theorem appears in section 2.3, with consequences filling the next several sections. Since U is unitary, $U^{-1} = U^{\mathrm{H}}$, and we use the latter expression in formulating the theorem.

Schur's Theorem. *Given an n by n matrix A with eigenvalues $\lambda_1, \ldots, \lambda_n$ in any prescribed order, there is a unitary n by n matrix U such that $T = U^{\mathrm{H}}AU$ is upper triangular and the diagonal elements $t_{ii} = \lambda_i$. Furthermore, if the entries of A and its eigenvalues are all real, U may be chosen to be real orthogonal.*

Note: There is no claim that the matrices U and T are unique. Indeed, the proof is constructive, and we will see that there will be many choices in following its steps. If the characteristic polynomial of A has multiple roots, the λ_i should contain each zero of the characteristic polynomial as many times as the factor $\lambda - \lambda_i$ appears in the characteristic polynomial. These equal values may appear anywhere in the list of λ_i, but the total number of appearances (called the *algebraic multiplicity* of the eigenvalue is constrained.

The proof will be by induction on n.

The basis for the induction. If $n = 1$, the matrix A looks just like a scalar, which is its only eigenvalue, so it is already in triangular form. Thus, $T = A$ and $U = I$ satisfies the conditions of the theorem. While T is necessarily unique in this case, U could be any complex number of absolute value 1 (or ± 1 in the real case).

Brief digression: Householder matrices. The construction will concentrate on producing U. The given information will lead to identifying the first column of U, which must then be extended to the whole matrix. The approach taken in elementary courses to constructing a unitary matrix whose first column has a given direction is the Gram-Schmidt method: form a matrix consisting of the column you want followed by n more columns forming a basis of $\mathbb{C}^n$ (you can use the standard basis if you like), giving an n by $n+1$ matrix whose first column is the given vector and whose remaining n columns form an identity matrix. Then, construct a QR decomposition of this matrix. In any factorization of our matrix, the first factor will have n rows and the second factor will have $n+1$ columns. In the form of the QR factorization that we will construct, Q will be a square matrix with n columns, so R must have n rows, so R has the same shape as the matrix being factored.

We move through the given matrix finding the unique expression of each column as a linear combination of previous columns plus a vector orthogonal (in the Hermitian sense) to the previous columns. If this orthogonal vector is not zero, write it as the next column of Q, and form a column of R such that the product QR gives the expression for the column of A as a linear combination of the columns of Q that have already been written, with zeros in the lower part of the column to force later columns of Q to be ignored. If the orthogonal vector is the zero vector, do not change Q. However, you must still form a column of R expressing the current column of A in terms of the columns of Q that you already have, since R has the same shape as the original matrix. Since the columns of Q will be a basis for the column space of A at the end of the process and A was constructed to have all of $\mathbb{R}^n$ as its column space, we must find n independent vectors to form the columns of Q. If A is a real matrix, Q can be constructed to be a real matrix. So far, we have not been concerned about the lengths of the vectors forming the columns of Q, as long as the length isn't zero. However, to arrive at a unitary matrix, we need to modify Q so that all columns have length 1. To divide each column by its length, we form the matrix QD^{-1}, where D is the matrix whose diagonal entries are the lengths of the columns of Q in order, and whose off-diagonal entries are zero. Then $QR = (QD^{-1})(DR)$ and $(QD^{-1})^{\mathrm{H}}(QD^{-1}) = I$, so that (QD^{-1}) is a unitary matrix. Note that DR is obtained from our original R by multiplying the rows by the length of the vector in corresponding column of the original Q. Also note that the normalization process forces the columns of QD^{-1} to be unit vectors.

Programmers often say, "First make it work, then make it fast". Now that we are sure that there are unitary matrices with any given unit vector as first column, we can try to write one without doing so much work. This can be used both as an alternate to the Gram-Schmidt method for finding the QR factorization and to produce the matrices used in the construction proving Schur's theorem. A geometric interpretation of the desired matrix is that it gives a rigid motion of $\mathbb{C}^n$ (or $\mathbb{R}^n$ in the real case) that takes the first vector in the standard basis into a unit vector that has the same direction as the vector that forms the first column of the given matrix. In the real case, an easy way to do this is to reflect in the $(n-1)$-dimensional space that is the perpendicular bisector of the segment joining these points (the heads of the vectors with tails at the origin if you think of vectors as arrows). The complex case will have an extra complication that we will describe after describing the form of the matrix we seek. With e_1 as the first member of the standard basis and v as the desired vector, let $u = v - e_1$. If $u = 0$, there is noting to do, so we can use the identity as our unitary matrix. Otherwise, let

$$H_u = I - 2\frac{uu^{\mathrm{H}}}{u^{\mathrm{H}}u},$$

where the numerator of the last term is an n by n matrix and the denominator is a scalar which is the square of the length. Direct calculation shows that $H_u^{\mathrm{H}} = H_u$ and $H_u^2 = I$, so H_u is simultaneous unitary and Hermitian. It remains to show that v can be chosen so that $H_u e_1 = v$ (and, since H_u is a reflection, ${}_u v = e_1$).

Since H_u is unitary, we have, for all vectors w,

$$(H_u w)^{\mathrm{H}}(H_u w) = w^{\mathrm{H}} H_u^{\mathrm{H}} H_u w = w^{\mathrm{H}} w.$$

That is, H_u preserves length. Thus, we will need to choose v to be a unit vector. In the real case, this suffices, but there is an extra condition in the complex Hermitian case.

Since H_u is Hermitian, we have, for all vectors w,

$$\begin{aligned}\left(w^{\mathrm{H}}(H_u w)\right)^{\mathrm{H}} &= (w^{\mathrm{H}} H_u^{\mathrm{H}})w \\ &= w^{\mathrm{H}} H_u w \\ &= w^{\mathrm{H}}(H_u w).\end{aligned}$$

That is, $\langle w, H_u w\rangle$ is real. Thus a vector can be taken to the first vector in the standard basis only if it has length 1 and a real first entry.

Conversely, if v is such a vector and $u = v - e_1$,

$$\begin{aligned}u^{\mathrm{H}} u &= v^{\mathrm{H}} v - v^{\mathrm{H}} e_1 - e_1^{\mathrm{H}} v + e_1^{\mathrm{H}} e_1 \\ &= 2 - 2v^{\mathrm{H}} e_1,\end{aligned}$$

since the terms at the end are both 1 and the middle terms are equal.

From this, it is easy to see that $H_u e_1 = v$ and $H_u v = e_1$.

Matrices of the form H_u are called **Householder matrices**.

The inductive step. To return to the proof of Schur's theorem, consider λ_1. Since this is an eigenvalue, it must have at least one eigenvector. Let v be an eigenvector for λ_1, normalized to have length 1 and real first entry. Now, we can take U_1 to be a unitary matrix whose first column is v. The first column of AU_1 is $\lambda_1 v$, so if A_1 is defined by $AU_1 = U_1 A_1$, the first column of A_1 gives the expression of $\lambda_1 v$ in terms of a basis formed by the columns of U_1. However, this expression must be λ_1 times v plus 0 times each other vector in the basis. The first column of A_1 thus has λ_1 in the first position and zero everywhere else.

Since we only want to find a *triangular* matrix equivalent to A (or A_1), we can ignore the rest of the first row of A_1. Rows and columns 2 through n describe a mapping from the subspace V_1 orthogonal to v to itself. Restricting the original linear transformation to this subspace gives a linear transformation whose matrix, using the rest of the image of the standard basis under U_1, can be taken to be the result of dropping the first column of A_1. Here, the columns of U_1 except the first are a basis for the domain, but all columns of U_1 are a basis for the codomain. To ignore the first row as well means that we follow this linear transformation by one which sends v to zero and is the identity on V_1, i.e., a projection onto V_1.

The induction hypothesis applies to the composite linear transformation of V_1 which is represented by rows 2 through n of A_1. The $n-1$ by $n-1$ matrices expressing this result are extended to n by n matrices by putting a 1 in the (1, 1) position and zero elsewhere in the first row and column — a process we will call **bordering**. This applies a block structure on n by n matrices in which the first row is separated from all remaining rows and the first column is separated from all remaining columns. If you multiply two such block matrices, the result has a block structure and it is as if you took the rules for multiplying 2 by 2 matrices applied to the submatrices making up the blocks. In particular, the bordering of a product of two matrices is the product of the borderings of the factors. (These properties of block matrices are easily seen from the formulas expressing matrix multiplication.)

The λ_i were defined to be the roots of the characteristic polynomial of A with multiple roots counted according to their multiplicities. In order to apply induction, we note that A and A_1 have the same characteristic polynomial because they are similar. The characteristic polynomial of A_1 can be calculated by expanding $\det(A_1 - \lambda I)$ by its first column. This gives $(\lambda_1 - \lambda)$ times the characteristic polynomial of the matrix in rows and columns 2 through n of A_1. This is independent of the first row of A_1. The eigenvalues of this block are thus the same as those of A with the multiplicity of λ_1 reduced by 1.

The bordered form of $AU = UT$ on the $(n-1)$ dimensional subspace gives a triangular form of a matrix that looks like A_1 except that its upper right block has been replaced by zero. Replacing this matrix by A_1 only changes this part of the product. We can now multiply this result on the left by the bordered U^{H} (which is equal to the bordered U^{-1}) introducing no further changes. This tells us how to modify the bordered T. The result may not be bordered, but it is upper triangular. Multiplying the unitary matrices from these two parts of the process gives a unitary matrix relating the original A to an upper triangular matrix.

If A had real entries and real eigenvalues, we can construct real eigenvectors, allowing this construction to be done using real matrices. A unitary matrix with real entries is called an *orthogonal* matrix.

With the completion of the inductive step, the proof of Schur's theorem is finished.

Multiple eigenvalues. If the eigenvalue taken as λ_1 is a multiple eigenvalue of A, it will also be an eigenvalue of A_1. If we take an eigenvector v^* of A_1 in V_1 corresponding to λ_1, the component of $A - \lambda_1 I$ in V_1 will be zero, which means that it is a multiple of the original eigenvector v. Because v is an eigenvector for *the same* eigenvalue, this multiple is the same for all vectors $v^* + cv$. The vectors found in this way are sometimes called **generalized eigenvectors**. They can be used to produce a space on which some power of $A - \lambda I$ is zero whose dimension is the algebraic multiplicity of λ. As with ordinary eigenvalues, elements of such spaces for different eigenvalues are linearly independent, so there are spaces whose total dimension is n that account for the algebraic multiplicity of each eigenvalue, so there can be no other eigenvectors or generalized eigenvectors. This is one way to show (though there are easier ways) that the geometric multiplicity of each eigenvalue is no larger that its algebraic multiplicity, and to give a geometric interpretation of the algebraic multiplicity.

Normal matrices. Result **5U** of the text characterizes matrices that can be diagonalized by unitary matrices. The treatment is very brief with key results given only as exercises, but the result is important, so it is appropriate to elaborate on it.

First, **normal matrices** are defined to be complex matrices N that commute with their conjugate-transpose, i.e., $NN^{\mathrm{H}} = N^{\mathrm{H}}N$. This class includes Hermitian, skew-Hermitian and unitary matrices. This definition can be tested by matrix multiplication, and so is computationally easy.

A matrix that can be diagonalized by a unitary matrix is necessarily normal, since $M = UDU^{-1} = UDU^{\mathrm{H}}$ implies

$$\begin{aligned} MM^{\mathrm{H}} &= UDU^{\mathrm{H}}UD^{\mathrm{H}}U^{\mathrm{H}} = UDD^{\mathrm{H}}U^{\mathrm{H}} \\ M^{\mathrm{H}}M &= UD^{\mathrm{H}}U^{\mathrm{H}}UDU^{\mathrm{H}} = UD^{\mathrm{H}}DU^{\mathrm{H}} \end{aligned}$$

and these quantities are equal since diagonal matrices always commute. This allows one to construct normal matrices whose eigenvalues are arbitrary complex numbers, while the special types that we enumerated have strong limitations on their eigenvalues.

Similarly, if N is normal and $T = U^{-1}NU = U^{\mathrm{H}}NU$ with U unitary, then $TT^{\mathrm{H}} = T^{\mathrm{H}}T$. Schur's theorem says that U can be found for which T is upper triangular. If matrices that were both normal and upper triangular could only be diagonal, then Schur's theorem would actually diagonalize normal matrices. This is exactly what happens.

The proof can be organized in the same way as the induction argument used in the proof of Schur's theorem. The result clearly holds for 1×1 matrices, so we turn to the induction step.. We begin by looking at the $(1, 1)$ entry of $TT^{\mathrm{H}} = T^{\mathrm{H}}T$. The left side is the square of the length of the first row of T, while the right side is the square of the length of the first column. Since T is upper triangular, the first column has a simple form: its first entry is t_{11}, and the others are zero. Thus, its length is $|t_{11}|$. The first row is more complicated: its first entry is t_{11}, as before, but nothing is known about the other entries. However, the *length* of a vector has some nice properties: its square is the sum of nonnegative quantities, one of which is $|t_{11}|^2$. This means that the first row and first column can have the same length only if all entries in the first row other than the first are zero. The matrix T is then formed by bordering an $(n-1)\times(n-1)$ matrix T_1, that is easily seen to be normal and triangular if T is.

Real symmetric matrices. Since real symmetric matrices M have real eigenvalues, and the eigenvectors corresponding to real eigenvalues can always be chosen to have real entries, all of the steps of Schur's theorem can be done using real matrices, so there is a real orthogonal matrix such that $S^{-1}MS$ is triangular and symmetric, hence diagonal. Conversely, if M is a real matrix that can be diagonalized by a real orthogonal matrix, the diagonal matrix $S^{-1}MS$ must have all real entries. However, those entries are just the eigenvalues of M.

Other real normal matrices, like the orthogonal matrix

$$\begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}$$

can be diagonalized by a unitary matrix, but that matrix cannot be real. For this matrix, the columns of the diagonalizing matrix U must be proportional to $(1\ i)^{\mathrm{T}}$ and $(1\ \ -i)^{\mathrm{T}}$. For U to be unitary, these columns must be multiplied by complex numbers of norm $\sqrt{2}$. Any such multipliers can be used, and they can be chosen independently for the each column.

Exercises.

1. Find a real Householder matrix whose first column is

$$\begin{pmatrix} 1/2 \\ 1/2 \\ 1/2 \\ 1/2 \end{pmatrix}.$$

2. Find a complex Householder matrix whose first column is

$$\begin{pmatrix} 2/3 \\ 2/3 + 1/3 i \end{pmatrix}.$$

3. Use the fact that $(1\ \ -1)^{\mathrm{T}}$ (here, the use of "transpose" is only to remind you that the vector is a column although it was written as a row for typographical reasons) is an eigenvector of

$$\begin{pmatrix} 1 & -9 \\ -3 & 7 \end{pmatrix}$$

to find a Schur form of this matrix and both eigenvalues. Don't bother finding the other eigenvector, although it will be easy to find at almost any stage of working with this matrix.