Note 15: Upper Triangulation, Schur Decomposition

1 Overview and Motivation

In the previous notes – Note 9, Note 10, Note 11, and Note 12, along with others previously – we have relied a lot on the structure of diagonalizable matrices, namely the ability to separate a vector problem into independent scalar problems. Unfortunately, this does not cover all cases; not all matrices are diagonalizable. In this note, we will find an equivalent characterization for matrices which are not diagonalizable, develop a decomposition which preserves many key properties of diagonalization and which works for all matrices, and discuss some critical implications of this decomposition.

Key Idea 1 (Upper Triangularization)
Any square matrix can be transformed via an orthonormal change of basis to an upper triangular matrix. This decomposition allows us to write vector problems in terms of scalar problems that we can solve or have already solved, similar to diagonalization.

2 Non-Diagonalizable Matrices

To discuss matrices which are not diagonalizable, we will first introduce the idea of *eigenvalue multiplicity*.

Definition 2 (Multiplicities of an Eigenvalue)
Let $A \in \mathbb{R}^{n\times n}$ be a square matrix with characteristic polynomial $p_A(\lambda)$ and distinct eigenvalues $\lambda_1, \ldots, \lambda_d$.

- The *algebraic multiplicity* m_A^a of λ_i is the multiplicity of λ_i as a root of the characteristic polynomial $p_A(\lambda)$. In other words, factoring the characteristic polynomial into linear factors:

$$p_A(\lambda) = \prod_{i=1}^{d}(\lambda - \lambda_i)^{m_i} \tag{1}$$

  we have $m_A^a(\lambda_i) = m_i$.

- The *geometric multiplicity* $m_A^g(\lambda_i)$ of λ_i is the number of linearly independent eigenvectors of A with eigenvalue λ_i. In other words, $m_A^g(\lambda_i) = \dim(\text{Null}(A - \lambda_i I))$.

And, we also have the following theorem which summarizes everything we need to know about multiplicities for now.

Theorem 3 (Results on Multiplicities)
Let $A \in \mathbb{R}^{n\times n}$ with distinct eigenvalues $\lambda_1, \ldots, \lambda_d$.

(i) We have $\sum_{i=1}^{d} m_A^a(\lambda_i) = n$.

(ii) We have $m_A^a(\lambda_i) \geq m_A^g(\lambda_i)$ for every i.

(iii) The following are equivalent:

(a) A is diagonalizable;

(b) for all eigenvalues λ of A, we have $m_A^a(\lambda) = m_A^g(\lambda)$;

(c) $\sum_{i=1}^{d} m_A^g(\lambda_i) = n$.

Not all matrices have equal algebraic and geometric multiplicity for all eigenvalues; such matrices are called *defective*. An example of a defective matrix is

$$A = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix} \tag{2}$$

which has one eigenvalue $\lambda_1 = 0$. This eigenvalue has algebraic multiplicity $m_A^a(0) = 2$ since the characteristic polynomial is $p_A(\lambda) = \lambda^2 = (\lambda - 0)^2$. The eigenvalue also has corresponding one-dimensional eigenspace $\mathrm{Null}(A - \lambda_1 I) = \mathrm{Span}\left(\begin{bmatrix} 1 \\ 0 \end{bmatrix}\right)$, so $m_A^g(0) = 1$.

Defective matrices are exactly those matrices which cannot be diagonalized.

3 Upper Triangular Matrices

We now examine properties of upper triangular matrices which make upper triangularization a useful substitute to diagonalization.

Definition 4 (Upper Triangular Matrix)
A square matrix $T \in \mathbb{C}^{n\times n}$ is *upper triangular* if $t_{ij} = 0$ for $i > j$. That is, it has the form

$$T := \begin{bmatrix} t_{11} & t_{12} & \cdots & t_{1n} \\ 0 & t_{22} & \cdots & t_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & t_{nn} \end{bmatrix}. \tag{3}$$

Upper triangular matrices have a lot of useful properties, but the most important one is that we can read off the eigenvalues from the diagonal of the matrix.

Theorem 5 (Eigenvalues of Upper Triangular Matrices)
If $T \in \mathbb{C}^{n\times n}$ is an upper triangular matrix, i.e.,

$$T := \begin{bmatrix} t_{11} & t_{12} & \cdots & t_{1n} \\ 0 & t_{22} & \cdots & t_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & t_{nn} \end{bmatrix}. \tag{4}$$

then $t_{11}, t_{22}, \ldots, t_{nn}$ are eigenvalues of T.

Proof. We have

$$T - t_{ii}I_n = \begin{bmatrix} t_{11} - t_{ii} & t_{12} & \cdots & t_{1n} \\ 0 & t_{22} - t_{ii} & \cdots & t_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & t_{nn} - t_{ii} \end{bmatrix}. \tag{5}$$

Since $t_{ii} - t_{ii} = 0$, the i^{th} column pivot is 0. Thus $T - t_{ii}I_n$ must have a null space, and so t_{ii} is an eigenvalue of T. □

In fact, it can be shown that a slightly stronger statement holds (though since we do not know, and will not worry about, taking determinants of large matrices, *the proof is out of scope*; the most direct proof is by the Laplace expansion method to compute the determinant, which, again, is out of scope).

Theorem 6 (Characteristic Polynomials of Upper Triangular Matrices)
If $T \in \mathbb{C}^{n \times n}$ is an upper triangular matrix, i.e.,

$$T := \begin{bmatrix} t_{11} & t_{12} & \cdots & t_{1n} \\ 0 & t_{22} & \cdots & t_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & t_{nn} \end{bmatrix}. \tag{6}$$

then the characteristic polynomial of T is

$$p_T(\lambda) = \prod_{i=1}^{n} (\lambda - t_{ii}). \tag{7}$$

This says that the *only* eigenvalues of T are the t_{ii} (not guaranteed by the earlier theorem, in the case of duplicates), and the number of times each eigenvalue is on the diagonal is equal to its algebraic multiplicity.

Suppose we have some upper triangular matrix T, and some vector $\vec{y}$, and we want to solve $T\vec{x} = \vec{y}$. (Since a lot of the problems we deal with in this course are analogous to this problem, this isn't as contrived as it might seem). Written out, we have

$$\begin{bmatrix} t_{11} & t_{12} & \cdots & t_{1n} \\ 0 & t_{22} & \cdots & t_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & t_{nn} \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{bmatrix} = \begin{bmatrix} y_1 \\ y_2 \\ \vdots \\ y_n \end{bmatrix}. \tag{8}$$

More explicitly,

$$\begin{bmatrix} y_1 \\ y_2 \\ \vdots \\ y_n \end{bmatrix} = \begin{bmatrix} t_{11}x_1 + t_{12}x_2 + \cdots + t_{1n}x_n \\ t_{22}x_2 + \cdots + t_{2n}x_n \\ \vdots \\ t_{nn}x_n \end{bmatrix}. \tag{9}$$

More formally, the k^{th} row is

$$y_k = \sum_{j=k}^{n} t_{kj}x_j. \tag{10}$$

If none of the t_{ii} are zero, then we can solve the bottom row first to get $x_n = \frac{y_n}{t_{nn}}$. Then we can plug that into the penultimate row, getting

$$x_{n-1} = \frac{y_{n-1} - t_{n-1,n}x_n}{t_{n-1,n-1}} = \frac{y_{n-1} - y_n \frac{t_{n-1,n}}{t_{nn}}}{t_{n-1,n-1}}. \tag{11}$$

And plug this into the third-to-last row, and so on. A general formula is

$$x_k = \frac{1}{t_{kk}}\left(y_k - \sum_{j=k+1}^{n} t_{kj}x_j\right) \tag{12}$$

which we can recursively apply until we solve for all the x_i. This method of solving is called *back-substitution*.

Key Idea 7 (Solving an Upper Triangular System)
In an upper triangular system, we can reduce a matrix equation into a bunch of scalar equations, which we can then solve in a specific order to get a solution for all the scalar equations, and thus the original matrix equation.

4 Schur Decomposition

Now we will learn how to get any matrix into upper triangular form via a change of basis; this matrix decomposition is called the *Schur decomposition*.

Theorem 8 (Existence of Schur Decomposition)
Let $A \in \mathbb{C}^{n\times n}$ be a square matrix. Then there is a change-of-basis matrix $U \in \mathbb{C}^{n\times n}$ and an upper-triangular matrix $T \in \mathbb{C}^{n\times n}$ such that $A = UTU^{-1}$. Moreover, the eigenvalues of A are on the diagonal of T according to their multiplicities.

Unfortunately, we are not able to prove this result in full generality until we learn about complex inner products in Note 2j. Once there, we will be able to prove Theorem 8 (and in fact show that we can always find a U which is orthonormal with respect to the complex inner product) using a very similar method as will be used to prove the following theorem, which is is tractable from our point of view.

Theorem 9 (Existence of Real Schur Decomposition)
Let $A \in \mathbb{R}^{n\times n}$ be a square matrix with real eigenvalues. Then there is an orthonormal change-of-basis matrix $U \in \mathbb{R}^{n\times n}$ and an upper-triangular matrix $T \in \mathbb{R}^{n\times n}$ such that $A = UTU^{\top}$. Moreover, the eigenvalues of A are on the diagonal of T according to their multiplicities.

The proof of Theorem 9 is on the longer side and may distract from the overall flow of this note, so it is left to Appendix A. We fully expect you to read the proof and understand it. It is completely in-scope for the course.

The method of proof is constructive, so it also doubles as a proof of correctness for the following algorithm.

Algorithm 10 Real Schur Decomposition

Input: A square matrix $A \in \mathbb{R}^{n \times n}$ with real eigenvalues.

Output: An orthonormal matrix $U \in \mathbb{R}^{n \times n}$ and an upper-triangular matrix $T \in \mathbb{R}^{n \times n}$ such that $A = UTU^\top$.

function $\textsc{RealSchurDecomposition}(A)$
  if A is 1×1 **then**
    return $\begin{bmatrix} 1 \end{bmatrix}, A$
  end if
  $(\vec{q}_1, \lambda_1) := \textsc{FindEigenvectorEigenvalue}(A)$
  $Q := \textsc{ExtendBasis}(\{\vec{q}_1\}, \mathbb{R}^n)$ ▷ Extend $\{\vec{q}_1\}$ to a basis of $\mathbb{R}^n$ using Gram-Schmidt; see Note 13
  Unpack $Q := \begin{bmatrix} \vec{q}_1 & \widetilde{Q} \end{bmatrix}$
  Compute and unpack $Q^\top A Q = \begin{bmatrix} \lambda_1 & \vec{\widetilde{a}}_{12}^\top \\ \vec{0}_{n-1} & \widetilde{A}_{22} \end{bmatrix}$
  $(P, \widetilde{T}) := \textsc{RealSchurDecomposition}(\widetilde{A}_{22})$
  $U := \begin{bmatrix} \vec{q}_1 & \widetilde{Q}P \end{bmatrix}$
  $T := \begin{bmatrix} \lambda_1 & \vec{\widetilde{a}}_{12}^\top P \\ \vec{0}_{n-1} & \widetilde{T} \end{bmatrix}$
  return (U, T)
end function

Concept Check: Once you read Note 2j, prove Theorem 8 and come up with a complex analogue to Algorithm 10.

We have now just developed a method for, and validated our use of, upper triangularization, which we used in Note 11 for example.

5 Spectral Theorem

Now that we have shown the existence of the Schur decomposition, we can now use it to prove one of the most important and fundamental theorems in linear algebra. This is the spectral theorem (for real symmetric matrices). Spectral theorems (for different classes of symmetric linear maps) are useful in mathematics and engineering, as they reveal useful decompositions of symmetric linear maps.

Theorem 11 (Spectral Theorem for Real Symmetric Matrices)
Let $A \in \mathbb{R}^{n \times n}$ be real and symmetric. Then:

(i) The eigenvalues of A are real.

(ii) A is diagonalizable.

(iii) There is an orthonormal basis of $\mathbb{R}^n$ consisting of eigenvectors of A.

In short, A may be *orthonormally diagonalized*: $A = V \Lambda V^\top$ where $V \in \mathbb{R}^{n \times n}$ is an orthonormal matrix of eigenvectors of A, and $\Lambda \in \mathbb{R}^{n \times n}$ is a real diagonal matrix of eigenvalues.

The proof of Theorem 11 is on the longer side and may distract from the overall flow of this note, so it is left to Appendix B. We fully expect you to read the proof and understand it. It is completely in-scope for the course.

Concept Check: Once you read Note 2j, come up with a complex analogue to Theorem 11 (i.e., a spectral theorem for complex "Hermitian" matrices).

In fact, this gives symmetric matrices a huge amount of useful structure. Namely, orthonormal diagonalization is a key ingredient in proofs and algorithms involving symmetric matrices.

6 Example

We run the upper triangularization algorithm on the matrix

$$A := \begin{bmatrix} 0 & 1 \\ -2 & -3 \end{bmatrix}. \tag{13}$$

The eigenvalues of A and corresponding eigenvectors are

$$\vec{q}_1 = \begin{bmatrix} -\frac{1}{\sqrt{5}} \\ \frac{2}{\sqrt{5}} \end{bmatrix} \quad \lambda_1 = -2 \qquad \vec{q}_2 := \begin{bmatrix} -\frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} \end{bmatrix} \quad \lambda_2 = -1. \tag{14}$$

Applying the algorithm with $\vec{q}_1$ initially, we extend $\vec{q}_1$ to a basis Q of $\mathbb{R}^2$, by letting

$$Q = \begin{bmatrix} \vec{q}_1 & \widetilde{Q} \end{bmatrix} \qquad \text{where} \qquad \widetilde{Q} = \begin{bmatrix} \frac{2}{\sqrt{5}} \\ \frac{1}{\sqrt{5}} \end{bmatrix}. \tag{15}$$

Thus

$$Q = \begin{bmatrix} -\frac{1}{\sqrt{5}} & \frac{2}{\sqrt{5}} \\ \frac{2}{\sqrt{5}} & \frac{1}{\sqrt{5}} \end{bmatrix}. \tag{16}$$

Computing $Q^\top AQ$, we get

$$Q^\top AQ = \begin{bmatrix} -\frac{1}{\sqrt{5}} & \frac{2}{\sqrt{5}} \\ \frac{2}{\sqrt{5}} & \frac{1}{\sqrt{5}} \end{bmatrix}^\top \begin{bmatrix} 0 & 1 \\ -2 & -3 \end{bmatrix} \begin{bmatrix} -\frac{1}{\sqrt{5}} & \frac{2}{\sqrt{5}} \\ \frac{2}{\sqrt{5}} & \frac{1}{\sqrt{5}} \end{bmatrix} = \begin{bmatrix} -2 & -3 \\ 0 & -1 \end{bmatrix}. \tag{17}$$

We see that $\lambda_1 = -2$, which is indeed the eigenvalue corresponding to λ_1. Also unpacking, $\widetilde{a}_{12}^\top = -3$. And $\widetilde{A}_{22} = -1$, so its Schur decomposition is computed recursively to get $P = 1$ and $\widetilde{T} = -1$. Then

$$U := \begin{bmatrix} \vec{q}_1 & \widetilde{Q}P \end{bmatrix} = \begin{bmatrix} \vec{q}_1 & \widetilde{Q} \end{bmatrix} = Q = \begin{bmatrix} -\frac{1}{\sqrt{5}} & \frac{2}{\sqrt{5}} \\ \frac{2}{\sqrt{5}} & \frac{1}{\sqrt{5}} \end{bmatrix} \tag{18}$$

which is orthonormal, and

$$T := \begin{bmatrix} \lambda_1 & \widetilde{a}_{12}^\top P \\ 0 & \widetilde{T} \end{bmatrix} = \begin{bmatrix} -2 & -3 \\ 0 & -1 \end{bmatrix}. \tag{19}$$

Note that T is upper triangular with the eigenvalues of A on the diagonal. And, $A = UTU^\top$.

7 (OPTIONAL) Numerical Implications

Sometimes when we are solving problems or designing algorithms (say for stability or controllability), we want to do coordinate changes to bases in which we get easily solvable systems.

Previously, we used eigenvector bases for this. Now, we are allowed to use upper-triangularization bases. It turns out that a lot of the time, using the upper-triangularization basis is better, in the sense that our computation is more numerically stable. Let's try to briefly unpack why that is.

For notation's sake, let's say that we're working with a matrix $A \in \mathbb{R}^{n \times n}$.

Right off the bat, we re-emphasize that sometimes A is not diagonalizable. In this case, upper-triangularization is the only method we have developed so far that actually works for A, and so it is the best method to use by default.

Now let's suppose that A is diagonalizable. Let $A = V\Lambda V^{-1}$ be the representation of A in the eigenvector basis V – the diagonalization of A – and let $A = UTU^{\top}$ be the representation of A in the upper triangular basis – the Schur decomposition of A.

One key difference we can see in these formulas is that in the diagonalization representation, we need to compute V and V^{-1}, while in the upper triangularization representation, we only need to compute U and $U^{\top}$ (which is easy to compute given U). This is the difference we are looking for.

We will now try to justify why computing V^{-1} is numerically unstable. We do this by considering a range of scenarios (say configurations of A).

One extreme is when A is symmetric. Then the eigenvectors of A are orthonormal. If they are normalized, then V is a matrix with orthonormal columns and rows, so that $V^{-1} = V^{\top}$. In this case, we have shown earlier in the note that the Schur decomposition is exactly equal to the diagonalization, so there is no advantage to be gained by either side in terms of numeric stability.

The other extreme is when A is not diagonalizable. Then multiple eigenvectors of A are not distinct – in particular, they align perfectly. In this case V is singular, and thus non-invertible, since two of the columns of V are identical. Since we can only use the Schur decomposition, this is the better method.

We can make our point by considering matrices A that are "close to non-diagonalizable" – that is, A with eigenvector matrix V which is "nearly singular". That is, two eigenvectors of A are almost completely aligned. We saw such matrices in the critically damped case for RLC circuits, for example.

In this case, inverting the V matrix and separating the aligned eigenvectors is difficult, and indeed numerically unstable. One can make the argument precise using the notion of *condition number*, which is out of scope for the class. But, heuristically, here is what happens.

When we find a matrix inverse, conceptually it's similar to finding a solution $\vec{x}$ to $V\vec{x} = \vec{y}$. Since there are almost-aligned eigenvectors in V, there is at least one direction in $\mathbb{R}^n$ for which all of the columns of V have really small components in that direction (because there are n directions and effectively $n-1$ vectors to use). If $\vec{y}$ points into that problematic direction, then the coordinates of $\vec{x}$ (i.e., the coefficients of the linear combination of the columns of V) will, more often than not, have to be very large to push the vectors in V to reach $\vec{y}$, while cancelling out in all other directions to perfectly equal $\vec{y}$ – even for benign, generic $\vec{y}$ such as unit vectors! Moreover, *very similar values of $\vec{y}$ (such as a "true" value of $\vec{y}$ compared to a computer representation of $\vec{y}$), lead to very different values for $\vec{x}$!* Since $\vec{x}$ has crazy behavior and $\vec{x} = V^{-1}\vec{y}$, it is reasonable to think of V^{-1} as being numerically unstable.

On the other hand, the Schur decomposition is like a boon, in the sense that when we compute it, we never have to take matrix inverses. All we have to do is take matrix transposes, use Gram-Schmidt[1], and we're in business – we get a fully orthonormal basis that turns our system into one that's easily solvable. This process is more numerically stable, since our change-of-basis is orthonormal and we never have to take an inverse anywhere.

This is the crux of why the Schur decomposition is more numerically stable than diagonalization. This type of analysis can be explored more in e.g., EE 127, and Math 128.

8 Final Comments

In this note, we first discussed why some matrices are not diagonalizable. Then, we discussed upper-triangular matrices, and in particular properties of upper-triangular matrices that make upper-triangularization a good alternative to diagonalization for solving problems, more specifically the property of this basis to split vector equations into scalar equations that can be solved one at a time. Then, we discussed the Schur decomposition, as a way to upper-triangularize arbitrary matrices. Finally, we stated and proved the spectral theorem for real symmetric matrices, using the extra structure that the Schur decomposition gave us.

This validates the upper-triangularization decompositions we used earlier, for example in Note 11.

[1] Gram-Schmidt is also not great as a numerical linear algebra tool, because it suffers from a phenomenon called *catastrophic cancellation*. The gist of it is that we end up subtracting a lot of vectors, end up with vectors that should be – but are not quite, on our computer, due to computer arithmetic limitations – zero, and then we normalize it and get a unit vector in an essentially random direction. Using this vector in our computation can lead to crazy results. There are other methods to do orthonormalization, such as one of many algorithms for the QR decomposition, although they are more technically complex. *All of this footnote is out of scope for the course.*

A Proof of Theorem 9

To show this theorem, we will need the following two lemmas.

The ability to do this decomposition relies on the following fundamental claim.

Lemma 12 (Existence of an Eigenvalue/Eigenvector Pair)
Let $A \in \mathbb{C}^{n\times n}$ be a square matrix. Then A has an eigenvalue and corresponding eigenvector.

Proof. Since $p_A(\lambda)$, i.e., the characteristic polynomial of A, has an eigenvalue, the Fundamental Theorem of Algebra asserts that it has at least one distinct root, which is our eigenvalue λ_0. Then $\det(A - \lambda_0 I_n) = p_A(\lambda_0) = 0$, so $A - \lambda_0 I_n$ has a null space. Thus there is at least one eigenvector $\vec{v}_0$ such that $(\vec{v}_0, \lambda_0)$ is an eigenvector/eigenvalue pair for A. □

In order to ensure the eigenvalues of all matrices we deal with are real, we require the following lemma, which follows from the same calculation as Theorem 6.

Lemma 13 (Characteristic Polynomial of Block Upper Triangular Matrices)
If $T \in \mathbb{C}^{(m+n)\times(m+n)}$ is a block upper triangular matrix, i.e.,

$$T := \begin{bmatrix} T_{11} & T_{12} \\ 0_{n\times m} & T_{22} \end{bmatrix}. \tag{20}$$

where $T_{11} \in \mathbb{C}^{m\times m}$, $T_{12} \in \mathbb{C}^{m\times n}$, $T_{22} \in \mathbb{C}^{n\times n}$, then the characteristic polynomial of T is

$$p_T(\lambda) = p_{T_{11}}(\lambda) p_{T_{22}}(\lambda). \tag{21}$$

Proof of Theorem 9. We use a recursive approach.

The recursive base case is $n = 1$. If $A \in \mathbb{R}^{1\times 1}$ then A is a scalar, so the orthonormal change-of-basis matrix U can just be $U = 1$, and $T = A$. Thus $A = UTU^{-1}$ is a Schur decomposition of A.

Now consider the general recursive case. Suppose $A \in \mathbb{R}^{n\times n}$ is a matrix with n (not necessarily distinct) real eigenvalues $\lambda_1, \ldots, \lambda_n$. Pick a normalized eigenvector $\vec{q}$ of A which corresponds to λ_1; this exists by Lemma 12. Use Gram-Schmidt to extend $\vec{q}$ to an orthonormal basis $Q = \begin{bmatrix} \vec{q} & \widetilde{Q} \end{bmatrix}$ of $\mathbb{R}^n$. Then

$$\begin{aligned} Q^\top A Q &= \begin{bmatrix} \vec{q}^\top \\ \widetilde{Q}^\top \end{bmatrix} A \begin{bmatrix} \vec{q} & \widetilde{Q} \end{bmatrix} && (22) \\ &= \begin{bmatrix} \vec{q}^\top A \vec{q} & \vec{q}^\top A \widetilde{Q} \\ \widetilde{Q}^\top A \vec{q} & \widetilde{Q}^\top A \widetilde{Q} \end{bmatrix} && (23) \\ &= \begin{bmatrix} \lambda_1 \vec{q}^\top \vec{q} & \vec{q}^\top A \widetilde{Q} \\ \lambda_1 \widetilde{Q}^\top \vec{q} & \widetilde{Q}^\top A \widetilde{Q} \end{bmatrix}. && (24) \end{aligned}$$

Now since Q is orthonormal, we have

$$I_n = Q^\top Q = \begin{bmatrix} \vec{q}^\top \\ \widetilde{Q}^\top \end{bmatrix} \begin{bmatrix} \vec{q} & \widetilde{Q} \end{bmatrix} = \begin{bmatrix} \vec{q}^\top \vec{q} & \vec{q}^\top \widetilde{Q} \\ \widetilde{Q}^\top \vec{q} & \widetilde{Q}^\top \widetilde{Q} \end{bmatrix}. \tag{25}$$

But we also know

$$I_n = \begin{bmatrix} 1 & \vec{0}_{n-1}^\top \\ \vec{0}_{n-1} & I_{n-1} \end{bmatrix} \tag{26}$$

so $\widetilde{Q}^\top \vec{q} = \vec{0}_{n-1}$, and also $\vec{q}^\top \vec{q} = 1$. Thus

$$Q^\top A Q = \begin{bmatrix} \lambda_1 \vec{q}^\top \vec{q} & \vec{q}^\top A \widetilde{Q} \\ \lambda_1 \widetilde{Q}^\top \vec{q} & \widetilde{Q}^\top A \widetilde{Q} \end{bmatrix} \tag{27}$$

$$= \begin{bmatrix} \lambda_1 & \vec{q}^\top A \widetilde{Q} \\ \vec{0}_{n-1} & \widetilde{Q}^\top A \widetilde{Q} \end{bmatrix}. \tag{28}$$

To clean up a little, we introduce the notation

$$\widetilde{\vec{a}}_{12}^\top := \vec{q}^\top A \widetilde{Q} \qquad \text{and} \qquad \widetilde{A}_{22} := \widetilde{Q}^\top A \widetilde{Q}. \tag{29}$$

Thus we have

$$Q^\top A Q = \begin{bmatrix} \lambda_1 & \vec{q}^\top A \widetilde{Q} \\ \vec{0}_{n-1} & \widetilde{Q}^\top A \widetilde{Q} \end{bmatrix} \tag{30}$$

$$= \begin{bmatrix} \lambda_1 & \widetilde{\vec{a}}_{12}^\top \\ \vec{0}_{n-1} & \widetilde{A}_{22} \end{bmatrix}. \tag{31}$$

This is where we set up for the recursive call, where we will try to recursively upper triangularize $\widetilde{A}_{22}$. We first need to show that $\widetilde{A}_{22}$ is a smaller subproblem of A, which fulfills all assumptions of the theorem, i.e., is a real square matrix with real eigenvalues.

Since $\widetilde{Q} \in \mathbb{R}^{n \times (n-1)}$, we have $\widetilde{A}_{22} \in \mathbb{R}^{(n-1)\times(n-1)}$. Thus $\widetilde{A}_{22}$ is smaller than our original matrix A and also a square matrix.

We can write the characteristic polynomial of A as

$$p_A(\lambda) = \prod_{i=1}^{n} (\lambda - \lambda_i). \tag{32}$$

Since $Q^\top A Q$ is a change of basis from A, we know from the invariance of polynomials under change of basis (proved in Note 12) that

$$p_A(\lambda) = p_{Q^\top A Q}(\lambda). \tag{33}$$

Since $Q^\top A Q$ is block upper triangular, by Lemma 13, we have

$$p_A(\lambda) = p_{Q^\top A Q}(\lambda) = (\lambda - \lambda_1) p_{\widetilde{A}_{22}}(\lambda). \tag{34}$$

Hence

$$p_{\widetilde{A}_{22}}(\lambda) = \frac{p_A(\lambda)}{\lambda - \lambda_1} = \prod_{i=2}^{n} (\lambda - \lambda_i). \tag{35}$$

Thus $\widetilde{A}_{22}$ has all real eigenvalues, in particular $\lambda_2, \dots, \lambda_n$.

Therefore we can recursively take the Schur decomposition of $\widetilde{A}_{22}$. Write

$$\widetilde{A}_{22} := P \widetilde{T} P^\top \tag{36}$$

where $P \in \mathbb{R}^{(n-1)\times(n-1)}$ is orthonormal and $\widetilde{T} \in \mathbb{R}^{(n-1)\times(n-1)}$ is upper triangular. Then

$$Q^\top A Q = \begin{bmatrix} \lambda_1 & \widetilde{\vec{a}}_{12}^\top \\ \vec{0}_{n-1} & \widetilde{A}_{22} \end{bmatrix} \tag{37}$$

$$= \begin{bmatrix} \lambda_1 & \vec{\tilde{a}}_{12}^\top \\ \vec{0}_{n-1} & P\widetilde{T}P^\top \end{bmatrix} \tag{38}$$

$$= \begin{bmatrix} 1 & \vec{0}_{n-1}^\top \\ \vec{0}_{n-1} & P \end{bmatrix} \begin{bmatrix} \lambda_1 & \vec{\tilde{a}}_{12}^\top P \\ \vec{0}_{n-1} & \widetilde{T} \end{bmatrix} \begin{bmatrix} 1 & \vec{0}_{n-1}^\top \\ \vec{0}_{n-1} & P^\top \end{bmatrix}. \tag{39}$$

where the motivation to reach the last line is that we want to find a matrix factorization that isolates $\widetilde{T}$ in the bottom right corner, making the middle matrix upper-triangular. Again cleaning up notation, let

$$R := \begin{bmatrix} 1 & \vec{0}_{n-1}^\top \\ \vec{0}_{n-1} & P \end{bmatrix}. \tag{40}$$

Then, using the orthonormality of P, we have

$$R^\top R = \begin{bmatrix} 1 & \vec{0}_{n-1}^\top \\ \vec{0}_{n-1} & P \end{bmatrix}^\top \begin{bmatrix} 1 & \vec{0}_{n-1}^\top \\ \vec{0}_{n-1} & P \end{bmatrix} \tag{41}$$

$$= \begin{bmatrix} 1 & \vec{0}_{n-1}^\top \\ \vec{0}_{n-1} & P^\top \end{bmatrix} \begin{bmatrix} 1 & \vec{0}_{n-1}^\top \\ \vec{0}_{n-1} & P \end{bmatrix} \tag{42}$$

$$= \begin{bmatrix} 1 & \vec{0}_{n-1}^\top \\ \vec{0}_{n-1} & P^\top P \end{bmatrix} \tag{43}$$

$$= \begin{bmatrix} 1 & \vec{0}_{n-1}^\top \\ \vec{0}_{n-1} & I_{n-1} \end{bmatrix} \tag{44}$$

$$= I_n. \tag{45}$$

Thus R is orthonormal. Thus we have

$$Q^\top A Q = \begin{bmatrix} 1 & \vec{0}_{n-1}^\top \\ \vec{0}_{n-1} & P \end{bmatrix} \begin{bmatrix} \lambda_1 & \vec{\tilde{a}}_{12}^\top P \\ \vec{0}_{n-1} & \widetilde{T} \end{bmatrix} \begin{bmatrix} 1 & \vec{0}_{n-1}^\top \\ \vec{0}_{n-1} & P^\top \end{bmatrix} \tag{46}$$

$$= R \begin{bmatrix} \lambda_1 & \vec{\tilde{a}}_{12}^\top P \\ \vec{0}_{n-1} & \widetilde{T} \end{bmatrix} R^\top \tag{47}$$

$$\implies A = \underbrace{QR}_{:=U} \underbrace{\begin{bmatrix} \lambda_1 & \vec{\tilde{a}}_{12}^\top P \\ \vec{0}_{n-1} & \widetilde{T} \end{bmatrix}}_{:=T} \underbrace{R^\top Q^\top}_{=U^\top}. \tag{48}$$

Here Q is orthonormal so

$$U^\top U = (QR)^\top (QR) = R^\top Q^\top Q R = R^\top I_n R = R^\top R = I_n. \tag{49}$$

Thus $U \in \mathbb{R}^{n\times n}$ is orthonormal. And $T \in \mathbb{R}^{n\times n}$ is upper triangular. Also, $A = UTU^\top$ by our calculation. Thus we have shown the first claim.

By Theorem 6 and the invariance of polynomials under change of basis (proved in Note 12), we have that

$$p_A(\lambda) = p_{UTU^\top}(\lambda) \tag{50}$$

$$= p_T(\lambda) \tag{51}$$

$$= \prod_{i=1}^{n} (\lambda - t_{ii}) \tag{52}$$

and this shows that the eigenvalues of A are on the diagonal of T according to their algebraic multiplicities, showing the second claim.

In order to build an efficient algorithm, we would now like to build expressions for U and T, so that we can compute them without wasting effort computing intermediate results like R.

$$U = QR \tag{53}$$

$$= \begin{bmatrix} \vec{q} & \widetilde{Q} \end{bmatrix} \begin{bmatrix} 1 & \vec{0}_{n-1}^{\top} \\ \vec{0}_{n-1} & P \end{bmatrix} \tag{54}$$

$$= \begin{bmatrix} \vec{q} & \widetilde{Q}P \end{bmatrix} \tag{55}$$

$$T = \begin{bmatrix} \lambda_1 & \vec{\widetilde{a}}_{12}^{\top} P \\ \vec{0}_{n-1} & \widetilde{T} \end{bmatrix}. \tag{56}$$

We can now implement the algorithm, at the end of Section 4. □

B Proof of Theorem 11

Proof of Theorem 11.

(i) Take an arbitrary eigenvector λ of A with corresponding eigenvector $\vec{v}$. Then

$$A\vec{v} = \lambda\vec{v}. \tag{57}$$

Taking the conjugate of both sides and using the fact that A is real so that $A = \overline{A}$, we get

$$A\overline{\vec{v}} = \overline{A}\overline{\vec{v}} = \overline{A\vec{v}} = \overline{\lambda\vec{v}} = \overline{\lambda}\overline{\vec{v}}. \tag{58}$$

Taking advantage of the fact that A is symmetric so that $A = A^{\top}$, we take the transpose of both sides to get

$$A\overline{\vec{v}} = \overline{\lambda}\overline{\vec{v}} \tag{59}$$

$$\overline{\vec{v}}^{\top} A^{\top} = \overline{\lambda}\overline{\vec{v}}^{\top} \tag{60}$$

$$\overline{\vec{v}}^{\top} A = \overline{\lambda}\overline{\vec{v}}^{\top}. \tag{61}$$

Then we multiply by $\vec{v}$ on both sides to get

$$\overline{\vec{v}}^{\top} A\vec{v} = \overline{\lambda}\overline{\vec{v}}^{\top}\vec{v} \tag{62}$$

$$\lambda\overline{\vec{v}}^{\top}\vec{v} = \overline{\lambda}\overline{\vec{v}}^{\top}\vec{v} \tag{63}$$

$$0 = (\lambda - \overline{\lambda})\overline{\vec{v}}^{\top}\vec{v} \tag{64}$$

using the fact that $(\vec{v}, \lambda)$ is an eigenvector-eigenvalue pair of of A. Now

$$\overline{\vec{v}}^{\top}\vec{v} = \sum_{i=1}^{n} \overline{v_i} \cdot v_i = \sum_{i=1}^{n} |v_i|^2. \tag{65}$$

so $\overline{\vec{v}}^{\top}\vec{v}$ is nonzero if and only if $\vec{v}$ is nonzero. Since $\vec{v}$ is an eigenvector and thus nonzero, we know that $\overline{\vec{v}}^{\top}\vec{v}$ is nonzero, and thus that $\lambda - \overline{\lambda} = 0$. Thus $\lambda = \overline{\lambda}$ so λ is real. Since λ is an arbitrary eigenvalue, all eigenvalues of A are real.

(ii) Since A is a square matrix with real eigenvalues, $(U, T) := \textsc{RealSchurDecomposition}(A)$ outputs an orthonormal matrix U and upper triangular matrix T such that $A = UTU^\top$. Since $A = A^\top$, we have

$$\begin{aligned} A &= A^\top && (66) \\ UTU^\top &= (UTU^\top)^\top && (67) \\ UTU^\top &= UT^\top U^\top && (68) \\ T &= T^\top && (69) \end{aligned}$$

which means that $T = T^\top$. Since T is upper triangular, $T^\top$ is upper triangular, so T is diagonal. Since U is orthonormal, $A = UTU^\top = UTU^{-1}$ is a diagonalization of A.

(iii) Since $A = UTU^\top$, right-multiplying by U, we get $AU = UT$. Looking at the columns $\vec{u}_1, \ldots, \vec{u}_n$ of U and the diagonal entries $t_{11}, \ldots, t_{nn}$ of T, we have

$$\begin{aligned} AU &= UT && (70) \\ A \begin{bmatrix} \vec{u}_1 & \cdots & \vec{u}_n \end{bmatrix} &= \begin{bmatrix} \vec{u}_1 & \cdots & \vec{u}_n \end{bmatrix} \begin{bmatrix} t_{11} & & \\ & \ddots & \\ & & t_{nn} \end{bmatrix} && (71) \\ \begin{bmatrix} A\vec{u}_1 & \cdots & A\vec{u}_n \end{bmatrix} &= \begin{bmatrix} t_{11}\vec{u}_1 & \cdots & t_{nn}\vec{u}_n \end{bmatrix}. && (72) \end{aligned}$$

Thus $A\vec{u}_i = t_{ii}\vec{u}_i$ for all i, confirming that $\vec{u}_i$ is an eigenvector of A with eigenvalue t_{ii}. Since U is an orthonormal matrix of eigenvectors, it follows that the vectors in U form an orthonormal basis of $\mathbb{R}^n$ consisting of eigenvectors of A.

□

Contributors:

- Druv Pai.
- Rahul Arya.
- Anant Sahai.

