

Eigenvalues, Eigenvectors, and Applications

We will look at things related to eigenvalues and eigenvectors in this set of notes, which should cover most of the material in Chapter 7. The non-examinable sections are marked with an asterisk.

Contents

1	Introduction to eigenvalues	1
2	Matrix functions*	3
3	Eigenspace and multiplicities	4
4	Geometric point of view	5
5	Complex eigenvalues	6
6	Trace and determinant	8
7	Examples	9
8	Non-diagonalizable matrices*	11
9	Fast eigenvalue algorithms*	14

1 Introduction to eigenvalues

Example. Consider the *Fibonacci numbers* defined by

$$F_0 = 0, F_1 = 1, F_{n+1} = F_n + F_{n-1}$$

We will try to find a general formula for F_n .

(*Step 1*). This can be reformulated as computing A^n , where

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

In fact, observe that $A \begin{bmatrix} x \\ y \end{bmatrix} = \begin{bmatrix} y \\ x+y \end{bmatrix}$, so it's easy to check that

$$A^n \begin{bmatrix} 0 \\ 1 \end{bmatrix} = \begin{bmatrix} F_n \\ F_{n+1} \end{bmatrix}$$

(*Step 2*). If A were equal to a diagonal matrix $\begin{bmatrix} d_1 & 0 \\ 0 & d_2 \end{bmatrix}$, then $A^n = \begin{bmatrix} d_1^n & 0 \\ 0 & d_2^n \end{bmatrix}$. But we are not so lucky. The next best thing is $A = SDS^{-1}$, where D is diagonal, since then

$$A^n = (SDS^{-1})(SDS^{-1}) \cdots (SDS^{-1}) = SD(S^{-1}S)D(S^{-1}S) \cdots DS^{-1} = SD^n S^{-1}$$

We observed this in the change of basis section. A physicist would say that A is **decoupled** in the basis defined by S , because A acts on the two vectors independently of each other. A mathematician is more straightforward and says A is **diagonalized** by S .

(*Step 3*). How do we find S and D ? Let $S = [\vec{v}_1 \ \vec{v}_2]$, and let $D = \begin{bmatrix} d_1 & 0 \\ 0 & d_2 \end{bmatrix}$, then

$$AS = [A\vec{v}_1 \ A\vec{v}_2], \quad SD = [d_1\vec{v}_1 \ d_2\vec{v}_2]$$

So we need $A\vec{v}_i = d_i\vec{v}_i$ for $i = 1, 2$. Moreover, we need \vec{v}_1 and \vec{v}_2 to be linearly independent, so in particular, they cannot be zero. In general, if we have $A\vec{v} = \lambda\vec{v}$, where λ is a scalar and $\vec{v} \neq 0$, then λ is an **eigenvalue** of A , and \vec{v} is an associated **eigenvector**.

(Step 4). Here comes the trick: $A\vec{v} = \lambda\vec{v}$ is equivalent to $(A - \lambda I)\vec{v} = 0$. If $\vec{v} \neq 0$, then it is in $\ker(A - \lambda I)$, so $A - \lambda I$ is *not invertible*. We can tell when this happens by computing a determinant.

$$p(\lambda) := \det(A - \lambda I) = \det \begin{bmatrix} -\lambda & 1 \\ 1 & 1 - \lambda \end{bmatrix} = \lambda^2 - \lambda - 1$$

This is the **characteristic polynomial** of A . Its roots are exactly the eigenvalues of A . In our case, they are $\lambda_{\pm} = \frac{1}{2}(1 \pm \sqrt{5})$.

(Step 5). Now, finding eigenvectors just comes down to finding non-trivial elements in a kernel. For λ_+ , the relevant matrix is

$$A - \lambda_+ I = \begin{bmatrix} \frac{1}{2}(-1 - \sqrt{5}) & 1 \\ 1 & \frac{1}{2}(1 - \sqrt{5}) \end{bmatrix}$$

Observe that the (column) vector $\vec{v}_+ = [1 \ \frac{1}{2}(1 + \sqrt{5})]^T$ satisfies the equation given by the first row, so it *must* be in the kernel, because the kernel is non-trivial by construction. You can easily verify it by computation as well. We have found our first eigenvector. Similarly, we can find that an eigenvector for λ_- is $\vec{v}_- = [1 \ \frac{1}{2}(1 - \sqrt{5})]^T$.

(Step 6). Now let $S = [\vec{v}_+ \ \vec{v}_-]$, then $\det S = -\sqrt{5} \neq 0$, so S is invertible. By construction, we have

$$AS = S \begin{bmatrix} \lambda_+ & 0 \\ 0 & \lambda_- \end{bmatrix}$$

so we finally get

$$A^n = S \begin{bmatrix} \lambda_+^n & 0 \\ 0 & \lambda_-^n \end{bmatrix} S^{-1} = -\frac{1}{\sqrt{5}} \begin{bmatrix} 1 & 1 \\ \frac{1}{2}(1 + \sqrt{5}) & \frac{1}{2}(1 - \sqrt{5}) \end{bmatrix} \begin{bmatrix} \lambda_+^n & 0 \\ 0 & \lambda_-^n \end{bmatrix} \begin{bmatrix} \frac{1}{2}(1 - \sqrt{5}) & -1 \\ -\frac{1}{2}(1 + \sqrt{5}) & 1 \end{bmatrix}$$

Therefore,

$$F_n = [1 \ 0]A^n \begin{bmatrix} 0 \\ 1 \end{bmatrix} = -\frac{1}{\sqrt{5}} [1 \ 1] \begin{bmatrix} \lambda_+^n & 0 \\ 0 & \lambda_-^n \end{bmatrix} \begin{bmatrix} -1 \\ 1 \end{bmatrix} = \frac{1}{\sqrt{5}} \left(\left(\frac{1 + \sqrt{5}}{2} \right)^n - \left(\frac{1 - \sqrt{5}}{2} \right)^n \right)$$

It may not be clear a priori that this is always an integer, but see remark (1).

We now make three tangential remarks:

1. Observe that we got from $A\vec{v}_+ = \lambda_+\vec{v}_+$ to $A\vec{v}_- = \lambda_-\vec{v}_-$ by simply replacing the signs of all occurrences of $\sqrt{5}$. This turns out to be an extremely profound observation: there is a symmetry between $\sqrt{5}$ and $-\sqrt{5}$, and an equation using only integers which holds for one also holds for the other.

It also holds if $\sqrt{5}$ is replaced by \sqrt{r} for any rational number r which is not a square. If instead we have $\sqrt{2} + \sqrt{3}$, we have four symmetries: we can switch the signs of $\sqrt{2}$ and $\sqrt{3}$ independently. But what happens with $\sqrt{2} + \sqrt{3} + \sqrt{6}$? What about $\sqrt[3]{2}$? This line of thought led to *Galois theory*, and one of the first achievements is the insolubility of the general quintic using radicals. I can go on. Please let me know if you want to learn more.

2. Since $|\lambda_-| < 1$, the second term is exponentially smaller than the first term, so F_n is the integer closest to $\frac{1}{\sqrt{5}} \left(\frac{1 + \sqrt{5}}{2} \right)^n$.
3. In practice, this is not the best way to compute the Fibonacci numbers, since it involves a lot of floating point arithmetic. It is much better to compute A^n directly by repeated squaring.

In summary, we did the following to diagonalize a general $n \times n$ matrix A .

(Step 1). Compute the **characteristic polynomial**

$$p_A(\lambda) = \det(A - \lambda I_n)$$

This is a polynomial of degree n in λ with leading coefficient $(-1)^n$.

(Step 2). A *generic* polynomial of degree n has n distinct roots. Let's pretend that is the case and let $\lambda_1, \dots, \lambda_n$ be the roots of p_A . These are the **eigenvalues** of A .

(Step 3). For each λ_i , by construction, $\ker(A - \lambda_i I_n)$ is non-trivial, so we can pick a non-zero vector \vec{v}_i in it. This is an **eigenvector** associated to λ_i .

(Step 4). If $S = [\vec{v}_1 \ \dots \ \vec{v}_n]$ is invertible, then we are done: $A = SDS^{-1}$, where D is the diagonal matrix with diagonal entries $\lambda_1, \dots, \lambda_n$. We will see later that if the eigenvalues are distinct, then S is automatically invertible.

In the case when the characteristic polynomial has repeated roots, the situation is a lot more subtle, and we will spend some time on that. For now, consider the difference between the matrices

$$A = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, \quad B = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix}$$

and perhaps try to follow the above steps.

2 Matrix functions*

Let A be an $n \times n$ matrix. Let $p(\lambda) = c_0 + c_1\lambda + \dots + c_n\lambda^n$ be a polynomial in λ , then we can define

$$p(A) := I_n + c_1A + c_2A^2 + \dots + c_nA^n$$

So I_n takes the role of 1 here. Observe that while two matrices need not commute in general, we always have $p(A)q(A) = q(A)p(A)$ for any polynomials p and q .

But we don't need to stop there. We can consider a power series instead. First, let's take

$$\exp(\lambda) = 1 + \lambda + \frac{1}{2!}\lambda^2 + \dots + \frac{1}{n!}\lambda^n + \dots$$

Recall from calculus that this converges for all λ , so one can ask if the following series

$$e^A = \exp(A) \stackrel{?}{=} I_n + A + \frac{1}{2}A^2 + \dots + \frac{1}{n!}A^n + \dots$$

makes sense. One way to make this question precise is to ask if each entry in the truncated series tends to a limit. It is actually possible to show that this is true by directly estimating the entries, but inspired by the previous section, we can first try to diagonalize A . First, suppose D is diagonal with diagonal entries $\lambda_1, \dots, \lambda_n$, then it is clear that

$$e^D = \begin{bmatrix} e^{\lambda_1} & & \\ & \ddots & \\ & & e^{\lambda_n} \end{bmatrix}$$

Now, if $A = SDS^{-1}$, then the same cancellation argument applied to each term gives $e^A = Se^D S^{-1}$, so we have convergence, and moreover a practical way of computing e^A .

Now consider the other important power series

$$\log(1 + \lambda) = \lambda - \frac{\lambda^2}{2} + \dots + (-1)^{n+1} \frac{\lambda^n}{n} + \dots$$

The same reduction to the diagonal case works, but in the diagonal case, we have to contend with the radius of convergence. If A has an eigenvalue greater than 1, then the series does not converge for that eigenvalue, so $\log(1 + A)$ cannot *a priori* be defined. We see that an important measurement of the "size" of A is its maximal eigenvalue,¹ which is sometimes called the **spectral radius** of A . The above discussion can be formulated as: a power series converges at a matrix A if the spectral radius of A is less than the radius of convergence of the power series.

¹We will see other ways of measuring the size of a matrix, for example, the maximum of all entries is another good way. Later, the maximum singular value is another way. Each is useful in some circumstances, and it is interesting to study how they are related to each other.

Theorem. If the spectral radius of A is less than 1, then $I - A$ is invertible.

Proof. The function $\frac{1}{1-x}$ has a power series expansion

$$\frac{1}{1-x} = 1 + x + x^2 + \dots$$

which converges when $|x| < 1$. Therefore, if the spectral radius of A is at most 1, then $I - A$ is invertible with inverse $I + A + A^2 + \dots$. \square

Applying functions to matrices has many applications. Suppose we have a system of differential equations

$$\begin{cases} \dot{x}(t) = y(t) \\ \dot{y}(t) = x(t) + y(t) \end{cases}, \quad x(0) = y(0) = 1$$

We can put it in matrix form by letting $\vec{x}(t) = \begin{bmatrix} x(t) \\ y(t) \end{bmatrix}$, $A = \begin{bmatrix} 0 & 1 \\ 1 & 1 \end{bmatrix}$, then

$$\frac{d\vec{x}(t)}{dt} = A\vec{x}(t), \quad \vec{x}(0) = \vec{x}_0 := \begin{bmatrix} 1 \\ 1 \end{bmatrix}$$

The solution has the same form as if \vec{x} were just a regular function.

$$\vec{x}(t) = e^{At}\vec{x}_0$$

From this, one sees that the eigenvalues of A should control the long-term behaviour of the solution. There are also far-reaching generalizations of this to infinite dimensional spaces, allowing us to evaluate functions at general operators.

3 Eigenspace and multiplicities

Before the interlude, the following matrices were considered

$$A = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, \quad B = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix}$$

In both cases, the characteristic polynomial is $p(\lambda) = (\lambda - 1)^2$, so there is a unique eigenvalue. We can compute that $\dim \ker(A - I_2) = 2$ and $\dim \ker(B - I_2) = 1$. In the second case, we *cannot* find enough eigenvectors to form a basis of \mathbb{R}^2 , so B is *not* diagonalizable. On the other hand, A is clearly diagonalizable. This section will introduce some further definitions, enough to state the main theorem on diagonalizability.

Definition. Let A be an $n \times n$ matrix. Let λ be a scalar. The λ -**eigenspace** of A is

$$E_\lambda^{(A)} = \ker(A - \lambda I_n)$$

Its dimension is the **geometric multiplicity** of λ , denoted by $m_{\text{geom}}^{(A)}(\lambda)$.

Definition. Let A be an $n \times n$ matrix. Let λ be a scalar. The **algebraic multiplicity** of λ is the unique non-negative integer $m_{\text{alg}}^{(A)}(\lambda)$ such that

$$p_A(x) = (x - \lambda)^{m_{\text{alg}}^{(A)}(\lambda)} g(x)$$

where p_A is the characteristic polynomial of A , and $g(\lambda) \neq 0$.

When we say roots of a polynomial counted with multiplicity, we mean each root contributes its algebraic multiplicity. For example, $p(\lambda) = (\lambda - 1)^2(\lambda + 1)(\lambda - 2)$ has three positive roots, counted with multiplicity.

It follows from definition that the following statements are equivalent

1. λ is an eigenvalue of A .

2. $E_\lambda^{(A)} \neq \{0\}$.
3. $m_{\text{geom}}^{(A)}(\lambda) > 0$.
4. $m_{\text{alg}}^{(A)}(\lambda) > 0$.

Also observe that $E_0^{(A)} = \ker(A)$.

For our examples, we have $m_{\text{geom}}^{(A)}(1) = 2$ and $m_{\text{geom}}^{(B)}(1) = 1$. But $m_{\text{alg}}^{(A)}(1) = m_{\text{alg}}^{(B)}(1) = 2$. The fact that A is diagonalizable but B is not is explained by the following theorem, which is the main examinable theorem of this chapter.

Theorem. Let A be an $n \times n$ matrix.

1. For any λ , we have $m_{\text{geom}}^{(A)}(\lambda) \leq m_{\text{alg}}^{(A)}(\lambda)$.
2. The matrix A is diagonalizable over \mathbb{C} if and only if equality holds for all $\lambda \in \mathbb{C}$.
3. The matrix A is diagonalizable over \mathbb{R} if and only if it is diagonalizable over \mathbb{C} and $p_A(\lambda)$ has n real roots, counted with multiplicity.

This is quite an elaborate theorem, and the proof will occupy much of the next few sections.

4 Geometric point of view

Recall that A is diagonalizable if there exists a diagonal matrix D and an invertible matrix S such that $A = SDS^{-1}$. We now reinterpret everything using linear transformations.

Definition. Let $T : \mathbb{R}^n \rightarrow \mathbb{R}^n$ be a linear transformation. An **eigenbasis** for T is a basis of \mathbb{R}^n consisting of eigenvectors of T . We say T is **diagonalizable** if it has an eigenbasis.

It is clear that T is diagonalizable if and only if its matrix with respect to any basis is diagonalizable.

This is now a coordinate-independent definition. Here is the advantage: **diagonalizing a matrix corresponds exactly to finding an eigenbasis of the underlying linear transformation.** By starting with a linear transformation, we are fixing the space and varying the coordinate system on the space. The problem becomes one of choosing a good basis.

We should check that most of the things we introduced are in fact geometric, in the sense that they do not change when we change coordinates. For the characteristic polynomial, this is a property of the determinant

$$\det(SAS^{-1} - \lambda I) = \det(S(A - \lambda I)S^{-1}) = \det(A - \lambda I)$$

so we can speak of the characteristic polynomial of a linear transformation. The algebraic multiplicity is therefore also a property of a linear transformation. The eigenspace, eigenvector, and geometric multiplicities are already geometric: they are expressed in terms of kernels and dimensions.

Warning. Similar matrices have the same eigenvalues. They do not necessarily have the same eigenvectors. Instead, the eigenvectors are related by a change of basis.

The main theorem is therefore entirely geometric. We use this reformulation to prove its first part.

Proof of Part (1) of Main Theorem. Let $V = \ker(T - \lambda I)$. Fix a basis $\{\vec{v}_1, \dots, \vec{v}_m\}$ for V , then $m = m_{\text{geom}}^{(T)}(\lambda)$. Extend it to a basis of \mathbb{R}^n by adjoining vectors $\vec{v}_{m+1}, \dots, \vec{v}_n$. With respect to this basis, the matrix of T has the form

$$\begin{bmatrix} \lambda I_m & * \\ 0 & * \end{bmatrix}$$

Therefore, $p_T(x) = (x - \lambda)^m g(x)$ for some polynomial $g(x)$. In particular, $m \leq m_{\text{alg}}^{(T)}(\lambda)$. \square

Before proving the other parts, we show that different eigenspaces are disjoint. This also applies to generalized eigenspaces, to be introduced later, with essentially the same proof.

Theorem. Let $\lambda_1, \dots, \lambda_m$ be distinct eigenvalues of T . Suppose we have $\vec{v}_1 + \dots + \vec{v}_m = 0$, where $\vec{v}_i \in E_{\lambda_i}^{(T)}$ for each i , then $\vec{v}_1 = \dots = \vec{v}_m = 0$.

Proof. Suppose the conclusion is false. Let $\vec{v}_1 + \dots + \vec{v}_m = 0$ be a non-trivial relation where the number of non-zero \vec{v}_i is minimal. Without loss of generality, we may suppose $\vec{v}_1 \neq 0$. Observe that $T\vec{v}_i = \lambda_i\vec{v}_i$ by hypothesis, so applying $T - \lambda_1$ gives

$$(\lambda_1 - \lambda_1)\vec{v}_1 + (\lambda_2 - \lambda_1)\vec{v}_2 + \dots + (\lambda_m - \lambda_1)\vec{v}_m = 0$$

The term from $E_{\lambda_1}^{(T)}$ is now zero, so by minimality, we must have $(\lambda_i - \lambda_1)\vec{v}_i = 0$ for all $i > 1$. But the eigenvalues are distinct, so $\vec{v}_i = 0$ if $i > 1$, which also implies $\vec{v}_1 = 0$, which is a contradiction. \square

Proof of Part (3) of Main Theorem. First suppose T is diagonalizable, then using an eigenbasis, we can check the equality of the two multiplicities on diagonal matrices, which makes the claim clear.

In general, let $\lambda_1, \dots, \lambda_p$ be eigenvalues of T , then by concatenating a choice of bases for $E_{\lambda_1}^{(T)}, \dots, E_{\lambda_p}^{(T)}$, we get a set of $\sum_{i=1}^p m_{\text{geom}}^{(T)}(\lambda_i)$ linearly independent eigenvectors in \mathbb{R}^n of size, by the previous theorem. If T has n real eigenvalues counted with multiplicities, and $m_{\text{alg}}^{(T)}(\lambda) = m_{\text{geom}}^{(T)}(\lambda)$ for all λ , then this procedure gives us an independent set of size n , so it is an eigenbasis for T . \square

It follows from the proof that there are two obstructions to diagonalizability:

- $m_{\text{alg}}^{(T)}(\lambda) > m_{\text{geom}}^{(T)}(\lambda)$: the eigenspaces are not big enough. In a non-examinable section, we will see how to fix this using generalized eigenspaces.
- $\sum_{i=1}^p m_{\text{alg}}^{(T)}(\lambda_i) < n$: there are not enough eigenvalues. This obstruction no longer exists if we work with \mathbb{C} , the field of complex numbers.

We are in fact sufficiently careful with the wording that the above proof works word-by-word to give a proof of part (2), once we allow scalars to be complex numbers. We will review some facts about complex numbers in the next section and talk about complex eigenvalues.

5 Complex eigenvalues

The fundamental theorem of algebra states that a polynomial of degree n has exactly n complex roots, counted with multiplicities, so the second obstruction no longer occurs. This is why the statement of part (2) of the Main Theorem is much cleaner. We will assume the basics of complex numbers as covered in the textbook and instead highlight some key features.

In this section, let i be a square root of -1 . There is a *choice* involved, and you actually cannot distinguish between $+\sqrt{-1}$ and $-\sqrt{-1}$. This is very profound, so profound that we make the following definition.

Definition. If $z = a + bi$, where $a, b \in \mathbb{R}$, then its **(complex) conjugate** is $\bar{z} := a - bi$.

The fact that i and $-i$ are indistinguishable is the philosophical reason why conjugation preserves all arithmetic operations, so for example

$$\overline{z + w} = \bar{z} + \bar{w}, \quad \overline{z\bar{w}} = \bar{z} \cdot \bar{w}, \quad \overline{z^{-1}} = \bar{z}^{-1}$$

This is related to the earlier remark that $\sqrt{5}$ and $-\sqrt{5}$ are indistinguishable if you only have access to the rational numbers. Another consequence is the following trivial and yet extremely profound fact: if an expression involving z and \bar{z} is supposed to have a real value, then it must be invariant under complex conjugation, so for example

$$\text{Re}(z) = \frac{1}{2}(z + \bar{z}), \quad \text{Im}(z) = \frac{1}{2i}(z - \bar{z})$$

both return real numbers, because they are invariant under conjugation. Once again, please talk to me if you want to learn more.

Since we can perform arithmetic operations, and there is a good notion of size ($|z| := (z\bar{z})^{1/2}$), we can evaluate power series at complex numbers.

Definition. The exponential function is defined by the usual power series

$$e^z = \exp(z) := \sum_{n=0}^{\infty} \frac{z^n}{n!}$$

This satisfies the usual properties, the key one being

$$\exp(z + w) = \exp(z) \exp(w)$$

By substituting in $z = i\theta$, we get

$$e^{i\theta} = \sum_{n=0}^{\infty} i^n \frac{\theta^n}{n!} = \sum_{k=0}^{\infty} (-1)^k \frac{\theta^{2k}}{(2k)!} + i \sum_{k=0}^{\infty} (-1)^k \frac{\theta^{2k+1}}{(2k+1)!} = \cos \theta + i \sin \theta$$

which you may know as Euler's formula, even though it is a complete triviality from definition.

Now, by allowing complex numbers as scalars, the discussions in the previous sections go through, and part (2) of the Main Theorem follows from the proof for part (3) and the fundamental theorem of algebra. We will now look at an instructive example.

Example. Let R be the rotation matrix

$$R = \begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix}$$

Its characteristic polynomial is $p_R(\lambda) = \lambda^2 - (2 \cos \theta)\lambda + 1$, so the eigenvalues are $\cos \theta \pm i \sin \theta = e^{\pm i\theta}$. Now,

$$R - (\cos \theta + i \sin \theta)I_2 = \begin{bmatrix} -i \sin \theta & \sin \theta \\ -\sin \theta & -i \sin \theta \end{bmatrix}$$


so an eigenvector for $e^{i\theta}$ is $\begin{bmatrix} i \\ 1 \end{bmatrix}$. By direct computation or the symmetry heuristics, an eigenvector for $e^{-i\theta}$ is $\begin{bmatrix} -i \\ 1 \end{bmatrix}$, so we have

$$R = \begin{bmatrix} i & -i \\ 1 & 1 \end{bmatrix} \begin{bmatrix} e^{i\theta} & 0 \\ 0 & e^{-i\theta} \end{bmatrix} \begin{bmatrix} i & -i \\ 1 & 1 \end{bmatrix}^{-1}$$

Note that the S -matrix does not depend on the angle, which fundamentally happens because all 2D rotations commute with each other. We use this observation to find two rotation matrices whose sum is also a rotation matrix: observe the complex number identity

$$1 + z = w, \quad z = e^{\frac{2\pi i}{3}}, w = e^{\frac{\pi i}{3}}$$

See the following picture:



It follows that

$$\begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} + \begin{bmatrix} z & 0 \\ 0 & \bar{z} \end{bmatrix} = \begin{bmatrix} w & 0 \\ 0 & \bar{w} \end{bmatrix}$$

so after a change of basis, we get $I_2 + R_{\frac{2\pi}{3}} = R_{\frac{\pi}{3}}$.

In general, by the principle of symmetry under complex conjugation, the complex roots of a real polynomial must appear in conjugate pairs: if z is a root, then \bar{z} is also a root. Let T be a real linear operator with eigenvalues z and \bar{z} . If \vec{v} is an eigenvector for z , then $\bar{\vec{v}}$ is an eigenvector for \bar{z} . Let $z = a + bi$ where $a, b \in \mathbb{R}$, then $\bar{z} = a - bi$. In the diagonal form, we have a block of the form $\begin{bmatrix} a+bi & 0 \\ 0 & a-bi \end{bmatrix}$. Write $\vec{v} = \vec{u} + i\vec{w}$, where \vec{u}, \vec{w} are real vectors, then we have

$$T(\vec{u} + i\vec{w}) = (a + bi)(\vec{u} + i\vec{w})$$

Comparing real and imaginary parts gives

$$T\vec{u} = a\vec{u} - b\vec{w}, \quad T\vec{w} = b\vec{u} + a\vec{w}$$

Therefore, if we replace $\{\vec{v}, \bar{\vec{v}}\}$ by $\{\vec{u}, \vec{w}\}$, then the matrix of T with respect to this new basis is $\begin{bmatrix} a & b \\ -b & a \end{bmatrix}$. We summarize this discussion in the following theorem.

Theorem. Suppose A is a real $n \times n$ matrix which is diagonalizable over \mathbb{C} , then A is similar to a block diagonal matrix, where each block is either a single number or a 2×2 matrix of the form $\begin{bmatrix} a & b \\ -b & a \end{bmatrix}$.

As a recap, to get to this form, first diagonalize as usual over \mathbb{C} , then replace a pair of conjugate eigenvectors with their real and imaginary parts. In that plane, instead of a scaling, A acts as a scaled rotation, which is responsible for the spirals in the corresponding discrete dynamic systems.

Finally, we remark that this operation of going from $\begin{bmatrix} z & 0 \\ 0 & \bar{z} \end{bmatrix}$ to $\begin{bmatrix} a & b \\ -b & a \end{bmatrix}$ preserves all arithmetic operations, since it is just a change of basis, and the basis does not depend on z and \bar{z} . So we can view complex numbers as 2×2 matrices. The unit circle is sent to rotation matrices.

6 Trace and determinant

Recall that the trace of an $n \times n$ matrix A is the sum of its diagonal entries. In the expansion of $\det(A - \lambda I_n)$, the leading coefficient is $(-1)^n$. The term for λ^{n-1} also comes entirely from the diagonal, and it is not hard to see that it must be $(-1)^{n-1} \text{Tr}(A)$. Finally, if $\lambda = 0$, then we recover on the one hand the constant term of $p_A(\lambda)$, but on the other hand $\det(A)$. Therefore,

Theorem. Let A be an $n \times n$ matrix, then its characteristic polynomial has the form

$$p_A(\lambda) = (-1)^n \lambda^n + (-1)^{n-1} \text{Tr}(A) \lambda^{n-1} + \dots + \det(A)$$

In particular, the characteristic polynomial of a 2×2 matrix is $\lambda^2 - \text{Tr}(A)\lambda + \det(A)$.

By the relation between roots of a polynomial and its coefficients, we get that

Theorem. The product of all eigenvalues is $\det(A)$. The sum of all eigenvalues is $\text{Tr}(A)$.

We now compute the eigenvalues of a TST matrix. Recall that they are matrices of the form

$$D_{a,b}^{(n)} = \begin{bmatrix} a & b & 0 & \dots & 0 \\ b & a & b & \dots & 0 \\ \vdots & \ddots & \ddots & \ddots & \vdots \\ 0 & \dots & b & a & b \\ 0 & \dots & 0 & b & a \end{bmatrix}$$

So we know that $D_{a,b}^{(n)} - \lambda I_n = D_{a-\lambda,b}^{(n)}$. Since we also had a formula for its determinant, it would appear reasonable to compute its eigenvalues directly. This is doable but tricky. Instead, we guess an eigenbasis and compute eigenvalues from there.

The guess is based on the identity

$$\sin \frac{(k-1)\pi}{n} + \sin \frac{(k+1)\pi}{n} = 2 \cos \frac{\pi}{n} \sin \frac{k\pi}{n}$$

Let $\vec{v}_k = [\sin \frac{k\pi}{n+1}, \sin \frac{2k\pi}{n+1}, \dots, \sin \frac{nk\pi}{n+1}]$, then for each i ,

$$(D_{0,1}^{(n)} \vec{v}_k)_i = \sin \frac{(i-1)k\pi}{n+1} + \sin \frac{(i+1)k\pi}{n+1} = 2 \cos \frac{k\pi}{n+1} \sin \frac{ik\pi}{n+1} = 2 \cos \frac{k\pi}{n+1} (\vec{v}_k)_i$$

Therefore, $D_{0,1}^{(n)} \vec{v}_k = 2 \cos \frac{k\pi}{n+1} \vec{v}_k$. For $k = 1, 2, \dots, n$, these give n eigenvectors with n different eigenvalues, so they form an eigenbasis. Therefore, the eigenvalues of $D_{0,1}^{(n)}$ are $2 \cos \frac{k\pi}{n+1}$ for $k = 1, \dots, n$.

Finally, $D_{(a,b)}^{(n)} = I_n + bD_{(0,1)}^{(n)}$, so the eigenvalues of $D_{(a,b)}^{(n)}$ are $a + 2b \cos \frac{k\pi}{n+1}$. In particular, take $a = 2$ and $b = 1$. The half-angle formula gives

$$2 + 2 \cos \frac{k\pi}{n+1} = 4 \cos^2 \frac{k\pi}{2(n+1)}$$

So $\det D_{(2,1)}^{(n)} = \prod_{k=1}^{n+1} (4 \cos^2 \frac{k\pi}{2(n+1)})$. But we have computed that $\det D_{(2,1)}^{(n)} = n + 1$. Their equality gives the formula promised in the last chapter, namely

$$\cos \frac{\pi}{2n} \cos \frac{2\pi}{2n} \dots \cos \frac{(n-1)\pi}{2n} = \frac{\sqrt{n}}{2^{2n-2}}$$

7 Examples

Example 1 We compute the eigenvalues of the following $n \times n$ matrix

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

This requires computing its characteristic polynomial

$$p_A(\lambda) = \det \begin{bmatrix} 1 - \lambda & 1 & 1 & 1 & \cdot & \cdot & \cdot & 1 \\ 1 & -\lambda & 0 & 0 & \cdot & \cdot & \cdot & 0 \\ 1 & 0 & -\lambda & 0 & \cdot & \cdot & \cdot & 0 \\ 1 & 0 & 0 & -\lambda & \cdot & \cdot & \cdot & 0 \\ \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \\ 1 & 0 & 0 & 0 & \cdot & \cdot & \cdot & -\lambda \end{bmatrix}$$

To compute the determinant, we use column operations. For now, suppose $\lambda \neq 0$, then we can add λ^{-1} times column i to column 1 for $i = 2, \dots, n$ to get

$$p_A(\lambda) = \det \begin{bmatrix} 1 - \lambda + (n-1)\lambda^{-1} & 1 & 1 & 1 & \cdot & \cdot & \cdot & 1 \\ 0 & -\lambda & 0 & 0 & \cdot & \cdot & \cdot & 0 \\ 0 & 0 & -\lambda & 0 & \cdot & \cdot & \cdot & 0 \\ 0 & 0 & 0 & -\lambda & \cdot & \cdot & \cdot & 0 \\ \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \\ 0 & 0 & 0 & 0 & \cdot & \cdot & \cdot & -\lambda \end{bmatrix} = (1 - \lambda + (n-1)\lambda^{-1})(-\lambda)^{n-1}$$

We can simplify this expression to $(-\lambda)^{n-2}(\lambda^2 - \lambda - (n-1))$. Therefore, this equals to $p_A(\lambda)$ if $\lambda \neq 0$. If $\lambda = 0$, then we are looking at $\det(A)$. If $n > 2$, then A has repeated columns, so $\det(A) = 0$, and the formula still holds. If $n = 2$, we can check this directly. A more satisfying reason is that the polynomial

$$q(\lambda) = p_A(\lambda) - (-\lambda)^{n-2}(\lambda^2 - \lambda - (n-1))$$

satisfies $q(\lambda) = 0$ if $\lambda \neq 0$. But a non-zero polynomial can have at most finitely many roots, so $q = 0$.

Therefore, 0 is an eigenvalue with algebraic multiplicity $n - 2$, and there are two other eigenvalues

$$\lambda_{\pm} = \frac{1 \pm \sqrt{4n-3}}{2}$$

Suppose $n > 1$, then the rank of A is 2, since columns 2 to n are equal. By the rank-nullity theorem, $\dim \ker A = n - 2$, so the geometric multiplicity of 0 is $n - 2$. It follows that A is diagonalizable, and all of its eigenvalues are real. We will see in the next chapter that this property holds for all symmetric matrices.

Example 2 Textbook section 7.6, #40. The question studies the following matrix

$$A = \begin{bmatrix} p & -q & -r & -s \\ q & p & s & -r \\ r & -s & p & q \\ s & r & -q & p \end{bmatrix}$$

and relates it to the representation of an integer as a sum of four squares. We will do it with the minimal amount of computation, as is intended.

First observe by direct computation that

$$A^T A = (p^2 + q^2 + r^2 + s^2)I_4$$

It follows that A is invertible if and only if p, q, r, s are not all 0, in which case its inverse is

$$A^{-1} = \frac{1}{p^2 + q^2 + r^2 + s^2} A^T$$

Moreover, $\det(A)^2 = \det(A^T A) = (p^2 + q^2 + r^2 + s^2)^4$, so $\det(A) = \pm(p^2 + q^2 + r^2 + s^2)^2$.

We don't know the sign yet, which a priori depends on p, q, r, s , but we will figure it out from the characteristic polynomial. Observe that $A - \lambda I_4$ can be obtained from A by replacing p with $p - \lambda$, so

$$\det(A - \lambda I_4) = \pm((p - \lambda)^2 + q^2 + r^2 + s^2)^2$$

But the leading coefficient must be positive, so the sign is positive. In particular, setting $\lambda = 0$ gives

$$\det(A) = (p^2 + q^2 + r^2 + s^2)^2$$

This can also be argued by continuity: if λ is large, then the diagonal entries dominate the rest, so the determinant is positive. But the sign cannot jump, so it must always be positive.

It follows easily from the formula for the characteristic polynomial that the eigenvalues of A are

$$\lambda_{\pm} = -p \pm \sqrt{-(q^2 + r^2 + s^2)}$$

each occurring with algebraic multiplicity 2. Since A^T is a multiple of A^{-1} , we have $AA^T = A^T A$, so the spectral theorem for normal operators (see notes for Chapter 8) shows that A is diagonalizable, but we don't need to know this.

Now, observe that

$$\|A\vec{x}\|^2 = \vec{x}^T A^T A \vec{x} = (p^2 + q^2 + r^2 + s^2) \vec{x}^T \vec{x} = (p^2 + q^2 + r^2 + s^2) \|\vec{x}\|^2$$

so $\|A\vec{x}\| = \sqrt{p^2 + q^2 + r^2 + s^2} \|\vec{x}\|$. The squared version is more useful for later: if $\vec{x} = [x_1 \ x_2 \ x_3 \ x_4]^T$ and $A\vec{x} = [y_1 \ y_2 \ y_3 \ y_4]^T$, then it says

$$y_1^2 + y_2^2 + y_3^2 + y_4^2 = (p^2 + q^2 + r^2 + s^2)(x_1^2 + x_2^2 + x_3^2 + x_4^2)$$

Moreover, if all $p, q, r, s, x_1, x_2, x_3, x_4$ are integers, then y_1, y_2, y_3, y_4 are also integers, so we have shown that the product of two sums of four squares is also a sum of four squares. For example, set $[p \ q \ r \ s] = [3 \ 3 \ 4 \ 5]$ and $\vec{x} = [1 \ 2 \ 4 \ 4]^T$, then

$$A\vec{x} = [-39 \ 13 \ 18 \ 13]^T$$

so we must have

$$(-39)^2 + 13^2 + 18^2 + 13^2 = (3^2 + 3^2 + 4^2 + 5^2)(1^2 + 2^2 + 4^2 + 4^2) = 59 \times 37 = 2183$$

Given that all primes are representable as a sum of four squares, this fact shows that all positive integers are representable as a sum of four squares.

We end with two remarks on this problem.

1. One can check that $A_{p,q,r,s}A_{p',q',r',s'}$ has the same form, with parameters given by the coordinates of $A_{p,q,r,s}[p' \ q' \ r' \ s']^T$, up to some sign differences. These matrices form a new system of number-like objects, similar to how a complex number $a + bi$ can be represented as a 2×2 matrix $\begin{bmatrix} a & b \\ -b & a \end{bmatrix}$. They are more commonly known as the *quaternions*, and written as $p + qi + rj + sk$, where i, j, k satisfy the relations $i^2 = j^2 = k^2 = ijk = -1$. This is not commutative: $ij = k = -ji$.

One way to build complex numbers from real numbers is to define a complex number as a pair of two real numbers (a, b) with multiplication $(a, b)(c, d) := (ac - bd, ad + bc)$. Starting instead with the complex numbers and performing this construction gives us the quaternions. Starting with the quaternions gives the Cayley octonions, which are not even associative, but they are sometimes useful in studying certain exotic objects. Continuing this way will lose even more properties, and the resulting objects have not been particularly useful.

2. The number of ways of writing a positive integer n as a sum of four squares is

$$r_4(n) = 8 \sum_{4|d|n} d$$

where the sum goes over all divisors of n which are not multiples of 4. If p is a prime, then the answer is $8(p+1)$. This treats permutations of the terms and changing signs as giving distinct representations, so if $p = 3$, the 32 representations are all of the form $3 = (\pm 1)^2 + (\pm 1)^2 + (\pm 1)^2 + 0^2$. If $p = 37$, there are 3 essentially distinct representations

$$37 = 1^2 + 2^2 + 4^2 + 4^2 = 0^2 + 0^2 + 1^2 + 6^2 = 2^2 + 2^2 + 2^2 + 5^2$$

To count those representations, the problem reduces to studying representation of a number as a sum of three or fewer squares. The two-square case is easily resolved with a similar formula. The three square case is somewhat well-understood, but it lies far deeper and has relation to the generalized Riemann hypothesis.

8 Non-diagonalizable matrices*

Not all hopes are lost when the matrix is not diagonalizable. We will talk about three fixes. But first, we remark that a *generic* matrix is diagonalizable, since a generic polynomial has no repeated roots. One way to make “generic” precise is to say that there is a polynomial Δ in n^2 variables such that if an $n \times n$ matrix A is not diagonalizable, then Δ evaluated at the entries of A is 0. In other words, if $\Delta \neq 0$, then A is diagonalizable. For $n = 2$, we can take $\Delta = (\text{Tr } A)^2 - 4 \det A = (a + c)^2 - 4(ad - bc)$. Therefore, in many results, one can assume the matrix is diagonalizable and then do some sort of continuity argument.

Theorem. Every $n \times n$ matrix is similar to an upper triangular matrix (over \mathbb{C}).

Proof. We prove more: if T is a linear transformation, then there exists an orthonormal basis such that T is upper triangular in that basis (we need to change the definitions slightly to accommodate complex numbers). The proof proceeds by induction.

By the fundamental theorem of algebra, T has a complex eigenvector \vec{u}_1 , which we may assume has unit length by scaling. Let $V = \text{span}(\vec{u}_1)^\perp$, and let P be the projection onto V . We have a linear operator PT on V , which has dimension $n - 1$, so by induction hypothesis, it has an orthonormal basis with respect to which T is upper triangular. Appending \vec{u}_1 to it gives the required basis. \square

This is sufficient to prove the Cayley–Hamilton theorem for general matrices.

Theorem (Cayley–Hamilton). Let A be an $n \times n$ matrix with characteristic polynomial p_A , then $p_A(A) = 0$.

Proof. By the previous theorem, we may assume A is upper triangular, in which case we have

$$p_A(\lambda) = \prod_{i=1}^n (\lambda - a_{ii})$$

so $p_A(A) = \prod_{i=1}^n (A - a_{ii}I_n)$. It is now easy to prove by induction that the first k columns of $\prod_{i=1}^k (A - a_{ii}I_n)$ are zero, so $p_A(A) = 0$. \square

This is usually not good enough for the computations we want to do, so there is a more powerful theorem.

Theorem (Jordan normal form). For each $\lambda \in \mathbb{C}$ and positive integer n , define a Jordan block

$$J_{\lambda,n} = \begin{bmatrix} \lambda & 1 & 0 & \cdots & 0 \\ 0 & \lambda & 1 & \cdots & 0 \\ \vdots & \vdots & \ddots & \ddots & \vdots \\ \vdots & \vdots & \cdots & \lambda & 1 \\ 0 & 0 & \cdots & 0 & \lambda \end{bmatrix}$$

Then every matrix A is similar to a block diagonal matrix, where each block is a Jordan block. Moreover, this is unique up to re-arranging the blocks.

Therefore, the similarity classes of $n \times n$ complex matrices are indexed by a multi-set of distinct eigenvalues with a partition for each multiplicity.

We now indicate some of the inputs to the proof. The first is a generalized eigenspace

$$\tilde{E}_\lambda^{(T)} := \bigcup_{k \geq 1} \ker((T - \lambda)^k)$$

So if $A = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}$, then its 0-eigenspace is $\text{span}(\vec{e}_1)$, but its generalized 0-eigenspace is \mathbb{R}^2 . In general, the dimension of the generalized eigenspace for λ is equal to its algebraic multiplicity. The proof consists of first reducing to the case where $\lambda = 0$ is the unique eigenvalue occurring with algebraic multiplicity n , and then applying the previous theorem on upper triangular form to conclude that $T^n = 0$.

Similarly to the proof of the Main Theorem, we get a decomposition of the entire space into generalized eigenspaces. To obtain the Jordan normal form, we need to study the structure of a generalized eigenspace. This requires looking at the following sequence of subspaces

$$E_\lambda^{(T)} = \ker(T - \lambda) \subseteq \ker((T - \lambda)^2) \subseteq \ker((T - \lambda)^3) \cdots$$

The amount by which the dimension jumps in each step determines the sizes of the Jordan blocks. It can be helpful to think about this sequence when T is already a Jordan block.

As an example, we prove a simple case of the theorem. It follows very easily by reducing to the upper triangular case, but we give a much longer proof to illustrate the general idea.

Theorem. If A is a 2×2 matrix with an eigenvalue λ of algebraic multiplicity 2, then either A is equal to λI_2 or A is similar to $J = \begin{bmatrix} \lambda & 1 \\ 0 & \lambda \end{bmatrix}$.

Proof. If λ has geometric multiplicity 2, then A is diagonalizable, so A is similar to λI_2 , which implies $A = \lambda I_2$. We now suppose that $m_{\text{geom}}^{(\lambda)}(A) = 1$, so $\dim \ker(A - \lambda) = 1$.

Consider the usual equation

$$A[\vec{v}_1, \vec{v}_2] = [\vec{v}_1, \vec{v}_2]J = [\vec{v}_1, \vec{v}_1 + \vec{v}_2]$$

So we need to solve two equations

$$(A - \lambda I)\vec{v}_1 = 0, \quad (A - \lambda I)\vec{v}_2 = \vec{v}_1$$

Since A has an eigenvalue equal to λ , \vec{v}_1 can be taken to be an eigenvector. Now the question is if $\vec{v}_1 \in \text{Im}(A - \lambda I)$. By assumption, $m_{\text{geom}}^{(\lambda)}(A) = 1$, so $\dim \text{Im}(A - \lambda I) = 1$. The characteristic polynomial of A is $p_A(X) = (X - \lambda)^2$, so by the Cayley–Hamilton theorem $(A - \lambda I)^2 = 0$, so $\text{Im}(A - \lambda I)$ is contained in $\ker(A - \lambda I)$. But both spaces have dimension 1, so they are equal. By construction, $\vec{v}_1 \in \ker(A - \lambda I)$, so it is also in $\text{Im}(A - \lambda I)$. It follows that \vec{v}_2 exists.

Finally, \vec{v}_2 and \vec{v}_1 are linearly independent: suppose not, then $\vec{v}_2 \in \ker(A - \lambda I)$. But by definition, $(A - \lambda I)\vec{v}_2 = \vec{v}_1 \neq 0$, which is a contradiction. \square

Finally, if the entries of the starting matrix are all rational numbers, we might not want to introduce irrational numbers into the picture at all. One such possibility is the rational canonical form, given in the next theorem. Its proof is actually not hard, but it involves a long digression into rational polynomials.

Theorem. Given a polynomial $p(\lambda) = a_0 + a_1\lambda + \cdots + a_n\lambda^n$, where $a_n = 1$, define

$$B_p = \begin{bmatrix} 0 & 0 & \cdots & 0 & -a_0 \\ 1 & 0 & \cdots & 0 & -a_1 \\ \vdots & \ddots & \ddots & \vdots & \vdots \\ 0 & & \ddots & 0 & -a_{n-2} \\ 0 & 0 & \cdots & 1 & -a_{n-1} \end{bmatrix}$$

Then every matrix A is similar to a block diagonal matrix, where each block is isomorphic to B_p , where p is a factor of the characteristic polynomial of A . This representation is unique up to permuting the blocks.

Among those three, the Jordan normal form is probably the hardest, but also the most useful for theory and applications. We give one application to ODE.

Example. Consider the second order linear ODE

$$\ddot{x}(t) + 2\dot{x}(t) + x(t) = 0$$

In the usual procedure for finding the general solution, one solves the characteristic equation $r^2 + 2r + 1 = 0$, which has a repeated root $r = -1$, so the general solution is $x(t) = e^{-t}(A + Bt)$.

We now do this using matrix exponentiation. Let $y(t) = \dot{x}(t)$, then we have a system of equations

$$\begin{cases} \dot{x}(t) = y(t) \\ \dot{y}(t) = -x(t) - 2y(t) \end{cases}$$

which we have seen can be re-written in matrix form

$$\dot{\vec{x}}(t) = \begin{bmatrix} 0 & 1 \\ -1 & -2 \end{bmatrix} \vec{x}(t)$$

Let A be the matrix, then $\det(A - \lambda I_2) = \lambda^2 + 2\lambda + 1$, which has a repeated root $\lambda = -1$ (think about how this is related to the rational canonical form). The geometric multiplicity of -1 is only 1, so A has Jordan normal form $J = \begin{bmatrix} -1 & 1 \\ 0 & -1 \end{bmatrix}$. In fact,

$$A = SJS^{-1}, \quad S = \begin{bmatrix} 1 & 1 \\ -1 & 0 \end{bmatrix}$$

The general solution is therefore

$$\vec{x}(t) = e^{At}\vec{x}_0 = Se^{Jt}S^{-1}\vec{x}_0$$

It remains to compute e^{Jt} . Write $J = -I + E$, so $E = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}$, then $E^2 = 0$, so by binomial expansion

$$J^2 = (-I + E)^n = (-1)^n I + (-1)^{n-1} nE$$

Therefore,

$$e^{Jt} = \sum_{n=0}^{\infty} \frac{1}{n!} ((-t)^n I + (-t)^{n-1} nE) = \begin{bmatrix} e^{-t} & te^{-t} \\ 0 & e^{-t} \end{bmatrix}$$

Substituting this into the solution gives the required general solution.

9 Fast eigenvalue algorithms*

In practice, people do not compute eigenvalues by finding roots of the characteristic polynomial. Recall that the Gauss–Jordan elimination is $O(n^3)$, which is okay if all entries are numbers. But $\det(A - \lambda I)$ is a polynomial of degree n , so addition is actually $O(n)$ and multiplication is $O(n^2)$ (or $O(n \log n)$ with FFT). Division is even worse, so the end-result is something like $O(n^5)$. On top of that, there are serious numerical stability issues. We also need a root-finding algorithm, but those are well-developed already.

We will briefly discuss two families of eigenvalue-finding strategies, which comes close to scratching the surface of the vast amount of work in this area.

The power method The idea is very simple: if A has a *unique* eigenvalue λ_1 with maximum modulus, then $\frac{A^k \vec{v}}{\|A^k \vec{v}\|}$ should converge to an eigenvector associated with it for a generic initial \vec{v} . This is quite clear if A is diagonal, and in general we can either use Jordan normal form (there is a unique Jordan block of size 1 for λ_1) or we can wiggle the matrix elements a bit to reduce to the diagonal case.

Example. For the Fibonacci example, the above discussion shows that

$$\frac{F_n}{F_{n-1}} \rightarrow \frac{1 + \sqrt{5}}{2}$$

as $n \rightarrow \infty$, which is clear from the formula we wrote down. In fact, the power method comes down to computing the Fibonacci numbers and taking the above ratio every time. The error in the approximation is around $O(0.38^n)$, so we have a linear algorithm (we need to take log because the correct measure of size is the *number* of correct digits).

Example. The power method does not converge exponentially if λ_1 is a multiple root. It may not even converge if there are more than one eigenvalues with modulus equal to the spectral radius. Here are two easy examples:

$$A = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix}, \vec{v} = \begin{bmatrix} 0 \\ 1 \end{bmatrix} \implies \frac{A^n \vec{v}}{\|A^n \vec{v}\|} = \frac{1}{\sqrt{n^2 + 1}} \begin{bmatrix} n \\ 0 \end{bmatrix} \rightarrow \begin{bmatrix} 1 \\ 0 \end{bmatrix} \text{ with error } O(n^{-2})$$
$$A = \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix}, \vec{v} = \begin{bmatrix} 0 \\ 1 \end{bmatrix} \implies \frac{A^n \vec{v}}{\|A^n \vec{v}\|} \text{ is periodic.}$$

In general, the eigenvalue contributing to the spectral radius may not be real, in which case we have at least two eigenvalues on the spectral circle (this does not happen if A is symmetric, as we will see in the next chapter). Moreover, for large matrices, the eigenvalues may be clumped together, so we have a bad convergence rate. This is fixed using shifts.

There are two ideas involved:

1. If we perform the power method on A^{-1} , we can isolate the *minimal* eigenvalue of A .
2. If we have a good estimate μ of an eigenvalue of A , then the minimal eigenvalue of $A - \mu I$ is much smaller than the other eigenvalues.

We use them in conjunction to get the inverse power method with shifts.

1. Find a good estimate μ for an eigenvalue of A .
2. Solve the equation $(A - \mu I)\vec{v}_{k+1} = \vec{v}_k$.
3. Normalize \vec{v}_{k+1} and decide if we are close enough.
4. If not, update the estimate μ and repeat from (2).

There are various schemes for the initial guess and updating estimate part. A good implementation will usually get a cubic convergence rate. To speed up the second step, we can pre-process A to introduce 0s, typically iteratively replacing A by $Q A Q^{-1}$ for some orthogonal matrix Q . This can be done using the same strategy we used for QR-factorization.

QR-based algorithms This is based on the following simple observation: if $A = QR$ is its QR factorization, then RQ is similar to A , since

$$RQ = Q^{-1}(QR)Q = Q^{-1}AQ$$

The basic QR-eigenvalue algorithm consists of repeatedly performing this. At a first glance, there is no reason that this leads anywhere, but in fact the terms below the diagonal tend to converge to 0, so eventually we end with an upper triangular matrix (if all eigenvalues are real) or a block upper triangular matrix with small blocks. The reason is that we are still secretly performing a bunch of power iterations together.

The same shifting idea can still be applied to dramatically speed up convergence. If we pre-process A to have good zeros placements, then the QR-factorization steps are usually $O(n^2)$, so they are not as computationally expensive as they may appear.