

Handout 3 for MATH 220

Notes on Diagonalization and Change of Basis

Laurence Barker, Mathematics Department, Bilkent University,
version: 1st January 2012.

Warning: These notes are intended as a reference for an introductory course on linear algebra. Their purpose is to summarize the rationale behind the method of diagonalization. Illustrative numerical examples can be found in textbooks.

Toy Problem: Let x_n and y_n be the number of female tribbles and male tribbles, respectively, on day n . Every night, each female tribble gives birth to 2 tribbles, both of them male, and each male tribble gives birth to 2 tribbles, both of them female. Tribbles never die. It is given that $x_0 = 3$ and $y_0 = 5$. Give a formula for the number of tribbles on day n . (Tribbles are small cute alien creatures which look like fluffy tennis-balls and which reproduce very fast. See the episode *The Trouble with Tribbles* of the original *Star Trek* television series.)

Answer: We have $x_n = 4 \cdot 3^n - (-1)^n$ and $y_n = 4 \cdot 3^n + (-1)^n$.

Proof 1: We argue by induction on n . The case $n = 0$ is trivial. Now

$$x_{n+1} = x_n + 2y_n, \quad y_{n+1} = 2x_n + y_n$$

for all natural numbers n . Assume, inductively, that the assertion holds for x_n and y_n . Then

$$x_{n+1} = (4 \cdot 3^n - (-1)^n) + 2(4 \cdot 3^n + (-1)^n) = 4 \cdot 3^{n+1} - (-1)^{n+1}$$

and similarly for y_{n+1} , as required. \square

Proof 2: Let $A = \begin{bmatrix} 1 & 2 \\ 2 & 1 \end{bmatrix}$ and $f_1 = \begin{bmatrix} 1 \\ 1 \end{bmatrix}$ and $f_2 = \begin{bmatrix} -1 \\ 1 \end{bmatrix}$. Then $Af_1 = 3f_1$ and $Af_2 = -f_2$, hence

$$\begin{aligned} \begin{bmatrix} x_n \\ y_n \end{bmatrix} &= A \begin{bmatrix} x_{n-1} \\ y_{n-1} \end{bmatrix} = A^n \begin{bmatrix} x_0 \\ y_0 \end{bmatrix} = A^n \begin{bmatrix} 3 \\ 5 \end{bmatrix} = A^n(4f_1 + f_2) = 4A^n f_1 + A^n f_2 \\ &= 4 \cdot 3^n \begin{bmatrix} 1 \\ 1 \end{bmatrix} + (-1)^n \begin{bmatrix} -1 \\ 1 \end{bmatrix} = \begin{bmatrix} 4 \cdot 3^n - (-1)^n \\ 4 \cdot 3^n + (-1)^n \end{bmatrix}. \quad \square \end{aligned}$$

The second proof contains the seeds of a systematic method, which we shall explain below.

Change of basis:

Let us begin by recalling how linear maps can be represented by matrices. Let V be finite-dimensional vector space over a field F and let $\alpha : V \rightarrow V$ be a linear map. Writing $n = \dim(V)$, let us choose a basis $\mathcal{E} = \{e_1, \dots, e_n\}$ for V . The matrix A representing α with respect to \mathcal{E} is the $n \times n$ matrix A such that, given vectors $x = x_1e_1 + \dots + x_n e_n$ and $y = y_1e_1 + \dots + y_n e_n$ with $\alpha(x) = y$, then $A\underline{x} = \underline{y}$ where $\underline{x} = (x_1, \dots, x_n)$ and $\underline{y} = (y_1, \dots, y_n)$. In other words, letting $a_{i,j}$ be the (i, j) entry of A , then $y_i = \sum_j a_{i,j} x_j$.

It is important to note that the matrix A depends not only on α but also on the choice of basis \mathcal{E} . In many contexts of application, it is helpful to change the basis so that α is represented by a different matrix which is easier to work with.

Let $\mathcal{F} = \{f_1, \dots, f_n\}$ be another basis for V . Since \mathcal{E} and \mathcal{F} are bases, there exist unique scalars $t_{i,j}$ and $s_{j,i}$ such that

$$f_j = \sum_i t_{i,j} e_i, \quad e_i = \sum_j s_{j,i} f_j$$

Writing $x = x_1 e_1 + \dots + x_n e_n = \sum_i x_i e_i$ and $x = \sum_j x'_j f_j$, then $x = \sum_{i,j} t_{i,j} x'_j e_i$. By the uniqueness of coordinates with respect to a given basis,

$$x_i = \sum_j t_{i,j} x'_j, \quad x'_j = \sum_i s_{j,i} x_i.$$

The coordinate vectors $\underline{x} = (x_1, \dots, x_n)$ and $\underline{x}' = (x'_1, \dots, x'_n)$ represent x with respect to \mathcal{E} and \mathcal{F} , respectively. We have

$$\underline{x} = T \underline{x}'$$

where T is the matrix with (i, j) -entry $t_{i,j}$. Similarly, $\underline{x}' = S \underline{x}$, where S is the matrix with (j, i) -entry $s_{j,i}$. Plainly, $ST = I = TS$ where I denotes the identity $n \times n$ matrix. In other words, $S = T^{-1}$ and

$$\underline{x}' = T^{-1} \underline{x}.$$

We call T the **coordinate transformation matrix** from \mathcal{F} -coordinates to \mathcal{E} -coordinates. Evidently, T^{-1} is the coordinate transformation matrix in the other direction, \mathcal{E} -coordinates to \mathcal{F} -coordinates.

Still letting A be the matrix representing α with respect to \mathcal{E} , now let B be the matrix representing α with respect to \mathcal{F} . The equation $y = \alpha(x)$ can be expressed in coordinate form as $\underline{y} = A \underline{x}$ and as $\underline{y}' = B \underline{x}'$. But $\underline{y}' = T^{-1} \underline{y} = T^{-1} A \underline{x} = T^{-1} A T \underline{x}'$. It follows that

$$B = T^{-1} A T, \quad A = T B T^{-1}.$$

These observations motivate the following definition. Given $n \times n$ matrices A and B , then A is said to be **similar** to B provided there exists an invertible $n \times n$ matrix T such that $A = T B T^{-1}$. The following remark is easy to check.

Remark: Let A, B, C be $n \times n$ matrices. Then A is similar to A . If A is similar to B , then B is similar to A . If A is similar to B and if B is similar to C , then A is similar to C .

In other words, similarity of $n \times n$ matrices is an equivalence relation.

We can now clear up a loose-end from earlier on in the course. We define the **determinant** of the linear map α to be the scalar

$$\det(\alpha) = \det(A)$$

where A is a matrix representing α with respect to some basis. The determinant of α is well-defined, independently of the choice of basis, thanks to the following remark.

Remark: Given similar $n \times n$ matrices A and B , then $\det(A) = \det(B)$.

Proof: Write $A = T B T^{-1}$. Then $\det(A) = \det(T) \det(B) \det(T)^{-1} = \det(B)$. \square

Exercise: Given an $n \times n$ matrix A , writing $a_{i,j}$ for the (i, j) -entry of A , we define the **trace** of A to be $\text{tr}(A) = \sum_{i=1}^n a_{i,i}$. Let B be an $n \times n$ matrix similar to A . Show that $\text{tr}(A) = \text{tr}(B)$.

In view of the latest exercise, we can define the **trace** of α to be

$$\operatorname{tr}(\alpha) = \operatorname{tr}(A)$$

where A is a matrix representing α . Indeed, the exercise implies that $\operatorname{tr}(\alpha)$ is well-defined.

Diagonal representation of linear maps

An $n \times n$ matrix A is said to be **diagonal** provided the (i, j) -entry is zero whenever $i \neq j$. In that case, letting $a_{i,i}$ denote the (i, i) -entry, we write $A = \operatorname{diag}(a_{1,1}, \dots, a_{n,n})$. Diagonal matrices tend to be very easy to work with. For instance, if A is diagonal, then $A^m = \operatorname{diag}(a_{1,1}^m, \dots, a_{n,n}^m)$ for any positive integer m , moreover, A is invertible if and only if each $a_{i,i} \neq 0$ and, in that case, $A^{-1} = \operatorname{diag}(a_{1,1}^{-1}, \dots, a_{n,n}^{-1})$.

We say that A is **diagonalizable** provided A is similar to a diagonal matrix. In other words, A is diagonalizable provided there exists a diagonal matrix B and an invertible matrix T such that $A = TBT^{-1}$. In that case, one way of calculating A^m is to make use of the equality $A^m = TB^mT^{-1}$. Also, if A is invertible, then B is invertible, and $A^{-1} = TB^{-1}T^{-1}$.

Again, let V be an n -dimensional vector space over a field F , and let $\alpha : V \rightarrow V$ be a linear map. We say that α is **diagonal** provided α is represented by a diagonal matrix with respect to some basis. Below, we shall describe a method for finding a diagonal matrix representing a given diagonal linear map.

First, we need a definition. Given a non-zero vector $f \in V$ and a scalar $\lambda \in F$ such that $\alpha(f) = \lambda f$, we call f an **eigenvector** of α with **eigenvalue** λ .

Remark: A scalar $\lambda \in F$ is an eigenvalue of α if and only if $\det(\alpha - \lambda I) = 0$, where I denotes the identity map on V .

Proof: Both of the specified conditions are plainly equivalent to the condition that the equation $(\alpha - \lambda I)x = 0$ has a non-zero solution $x \in V$. \square

The equation $\det(\alpha - \lambda I) = 0$ is called the **characteristic** equation of the linear map α .

Choosing a basis \mathcal{E} of V and letting A be the matrix representing α with respect to \mathcal{E} , the characteristic equation of α can be rewritten as $\det(A - \lambda I) = 0$, where I now denotes the identity matrix. Sometimes, we call this equation the **characteristic equation** of the matrix A . In an evident way, we can also speak of the **eigenvectors** and **eigenvalues** of A .

At last, we can explain the idea behind the second proof pertaining to the toy problem above. If we can find a basis $\mathcal{F} = \{f_1, \dots, f_n\}$ of V such that each f_j is an eigenvector for α , say $\alpha(f_j) = \lambda_j f_j$, then the matrix B representing α with respect to \mathcal{F} is the diagonal matrix $B = \operatorname{diag}(\lambda_1, \dots, \lambda_n)$, where λ_i is the eigenvalue associated with the eigenvector f_j .

In applications, two kinds of scenario often arise. In one of them, a diagonal linear map α is given, the matrix A representing α with respect to some basis \mathcal{E} has been determined, and the task is to find another basis \mathcal{F} such that α is represented by a diagonal matrix B with respect to \mathcal{F} . Letting T be the coordinate transformation matrix from \mathcal{F} -coordinates to \mathcal{E} -coordinates, then $A = TBT^{-1}$. In the other kind of scenario, a diagonalizable matrix A is given, and we seek a diagonal matrix B and an invertible matrix T such that $A = TBT^{-1}$. This is really the same problem as before, and we can understand α to be the linear map on F^n such that α is represented by A with respect to the standard basis of F^n .

The procedure is as follows:

Step 1: Find the eigenvalues $\lambda_1, \dots, \lambda_n$, which are the solutions to the polynomial equation $\det(A - \lambda I) = 0$ (possibly with repeated solutions). Then $B = \operatorname{diag}(\lambda_1, \dots, \lambda_n)$.

Step 2: Find the corresponding eigenvectors f_i by solving the equation $(A - \lambda I)f_j = 0$ (taking care to find m linearly independent eigenvectors associated with an eigenvalue that has multiplicity m as a repeated solution). The matrix T is the matrix whose j -th column is the coordinate vector representing f_j with respect to \mathcal{E} .

Return of the toy problem

As a first little example, let us deal with the above toy problem systematically, using the method that we have described. The eigenvalues of the matrix $A = \begin{bmatrix} 1 & 2 \\ 2 & 1 \end{bmatrix}$ are the solutions to the equation

$$0 = \det(A - \lambda I) = \det \begin{bmatrix} 1 - \lambda & 2 \\ 2 & 1 - \lambda \end{bmatrix} = (1 - \lambda)^2 - 2^2 = \lambda^2 - 2\lambda - 3.$$

The solutions are $\lambda_1 = 3$ and $\lambda_2 = -1$. Write $f_1 = (u, v)$ as a coordinate vector with respect to the standard basis $\mathcal{E} = \{(1, 0), (0, 1)\}$ of \mathbb{R}^2 . To find f_1 , we solve

$$0 = \begin{bmatrix} 1 - \lambda_1 & 2 \\ 2 & 1 - \lambda_1 \end{bmatrix} \begin{bmatrix} u \\ v \end{bmatrix} = \begin{bmatrix} -2 & 2 \\ 2 & -2 \end{bmatrix} \begin{bmatrix} u \\ v \end{bmatrix}$$

which yields $u = v$. So we can put $f_1 = (1, 1)$. A similar calculation with λ_2 in place of λ_1 yields a solution $f_2 = (-1, 1)$ as a coordinate vector with respect to \mathcal{E} . Taking the columns of T to be the \mathcal{E} -coordinates of the eigenvectors, $T = \begin{bmatrix} 1 & -1 \\ 1 & 1 \end{bmatrix}$, whence $T^{-1} = \frac{1}{2} \begin{bmatrix} 1 & 1 \\ -1 & 1 \end{bmatrix}$. Then $A = TBT^{-1}$ and

$$\begin{aligned} A^n &= TB^nT^{-1} = \frac{1}{2} \begin{bmatrix} 1 & -1 \\ 1 & 1 \end{bmatrix} \begin{bmatrix} 3 & 0 \\ 0 & -1 \end{bmatrix}^n \begin{bmatrix} 1 & 1 \\ -1 & 1 \end{bmatrix} \\ &= \frac{1}{2} \begin{bmatrix} 1 & -1 \\ 1 & 1 \end{bmatrix} \begin{bmatrix} 3^n & 3^n \\ -(-1)^n & (-1)^n \end{bmatrix} = \frac{1}{2} \begin{bmatrix} 3^n + (-1)^n & 3^n - (-1)^n \\ 3^n - (-1)^n & 3^n + (-1)^n \end{bmatrix}. \end{aligned}$$

As a check, we note that, putting $n = 1$, the latest equality reduces to the definition of A . Finally, we recover the answer

$$\begin{bmatrix} x_n \\ y_n \end{bmatrix} = A^n \begin{bmatrix} x_0 \\ y_0 \end{bmatrix} = \frac{1}{2} \begin{bmatrix} 3^n + (-1)^n & 3^n - (-1)^n \\ 3^n - (-1)^n & 3^n + (-1)^n \end{bmatrix} \begin{bmatrix} 3 \\ 5 \end{bmatrix} = \begin{bmatrix} 4 \cdot 3^n - (-1)^n \\ 4 \cdot 3^n + (-1)^n \end{bmatrix}.$$

Actually, it was not really necessary to calculate T^{-1} . We could, instead, have argued more along the lines that we presented earlier.

Exercise: Let A be an $n \times n$ matrix with eigenvalues $\lambda_1, \dots, \lambda_n$, up to multiplicity. Thus $\det(A - \lambda I) = (\lambda_1 - \lambda) \dots (\lambda_n - \lambda)$. Show that $\text{tr}(A) = \lambda_1 + \dots + \lambda_n$.

A more difficult problem

The following problem is somewhat similar to the one discussed above, but it is of interest because repeated eigenvalues appear.

Problem: A machine has three possible states, labelled 1, 2, 3. For distinct states i and j , if the machine is in state i at time $t = n$, then the probability of the machine being in state j

at time $t = n + 1$ is $1/4$. Suppose that the machine is in state 1 at time $t = 0$. What is the probability of the machine being in state 1 at time $t = n$?

Solution: Consider the matrix $A = \begin{bmatrix} 2 & 1 & 1 \\ 1 & 2 & 1 \\ 1 & 1 & 2 \end{bmatrix}$. Let $p_i(n)$ denote the probability of the

machine being in state i at time $t = n$. Then, writing column vectors as rows for convenience, $4(p_{n+1}(1), p_{n+1}(2), p_{n+1}(3)) = A(p_n(1), p_n(2), p_n(3))$ and $(p_0(1), p_0(2), p_0(3)) = (1, 0, 0)$.

The matrix $A - I$ has three identical non-zero rows and hence has nullity 2. So 1 appears twice as an eigenvalue of A . The matrix $A - 4I$ is non-invertible because the sum of its rows is zero, hence 4 is an eigenvalue of A . Therefore, the eigenvalues of A are $\lambda_1 = 1$ and $\lambda_2 = 1$ and $\lambda_3 = 4$. We can put $f_1 = (1, -1, 0)$ and $f_2 = (1, 0, -1)$ because these two vectors are eigenvectors with associated eigenvalue 1 and the set $\{f_1, f_2\}$ is linearly independent. We can put $f_3 = (1, 1, 1)$ as an eigenvector with associated eigenvalue 4. Thus $Af_1 = f_1$ and $Af_2 = f_2$ and $Af_3 = 4f_3$. We have $(p_0(1), p_0(2), p_0(3)) = (f_1 + f_2 + f_3)/3$, hence

$$\begin{aligned} (p_n(1), p_n(2), p_n(3)) &= (A/4)^n(f_1 + f_2 + f_3)/3 = (f_1 + f_2 + 4^n f_3)/3 \cdot 4^n \\ &= (4^n + 2, 4^n - 1, 4^n - 1)/3 \cdot 4^n. \end{aligned}$$

In conclusion, $p_n(1) = (4^n + 2)/3 \cdot 4^n = (1 + 1/4^{n-1})/3$.

Let us mention that, as a variant of the proof that the eigenvalues of A are 1, 1, 4, we could have observed that the matrices $A - I$ and $A - 4I$ are non-invertible, hence the eigenvalues are $\lambda, 1, 4$ for some scalar λ . Then, using an exercise above, we could have noted that $\lambda + 1 + 4 = \text{tr}(A) = 2 + 2 + 2 = 6$, hence $\lambda = 1$. As another alternative, more routine (rather boring, in fact), we could have made the calculation

$$\begin{aligned} \det(A - \lambda I) &= \begin{vmatrix} 2 - \lambda & 1 & 1 \\ 1 & 2 - \lambda & 1 \\ 1 & 1 & 2 - \lambda \end{vmatrix} = (2 - \lambda) \begin{vmatrix} 2 - \lambda & 1 \\ 1 & 2 - \lambda \end{vmatrix} - \begin{vmatrix} 1 & 1 \\ 1 & 2 - \lambda \end{vmatrix} + \begin{vmatrix} 1 & 2 - \lambda \\ 1 & 1 \end{vmatrix} \\ &= (2 - \lambda)(\lambda^2 - 4\lambda + 3) - (1 - \lambda) + (-1 + \lambda) = -\lambda^3 + 6\lambda^2 - 9\lambda + 4 = (1 - \lambda)(1 - \lambda)(4 - \lambda). \end{aligned}$$

When does the diagonalization procedure work?

Again, we let α be a linear map $V \rightarrow V$, where V is an n -dimensional vector space over a field F . We choose a basis $\mathcal{E} = \{e_1, \dots, e_n\}$, and we let A be the matrix representing α with respect to \mathcal{E} . The next remark is obvious.

Remark: The following three conditions are equivalent:

- (a) The linear map α is diagonal.
- (b) The matrix A is diagonalizable.
- (c) There exists a bases $\mathcal{F} = \{f_1, \dots, f_n\}$ of V such that each f_i is an eigenvector of α .

The next result gives a sufficient criterion for those three equivalent conditions to hold.

Proposition: Suppose that α has n mutually distinct eigenvalues $\lambda_1, \dots, \lambda_n$ in F . Let f_1, \dots, f_n be corresponding eigenvectors, in order. Then $\{f_1, \dots, f_n\}$ is a basis for F^n . In particular, α is diagonal and A is diagonalizable.

Proof: For a contradiction, suppose that $\{f_1, \dots, f_n\}$ is not linearly independent. Write

$$\mu_1 f_1 + \dots + \mu_n f_n = 0$$

where each $\mu_i \in F$, some $\mu_i \neq 0$, and the positive integer $m = |\{i : \mu_i \neq 0\}|$ is as small as possible. Renumbering the f_i if necessary, we may assume that

$$\mu_1 f_1 + \dots + \mu_m f_m = 0$$

and that $\mu_i \neq 0$ for all $1 \leq i \leq m$. Plainly, $m \geq 2$. But

$$\lambda_1 \mu_1 f_1 + \dots + \lambda_m \mu_m f_m = \alpha(\mu_1 f_1 + \dots + \mu_m f_m) = \alpha(0) = 0.$$

Multiplying the first equation by λ_m and then subtracting it from the second equation, we obtain

$$(\lambda_1 - \lambda_m)\mu_1 f_1 + \dots + (\lambda_{m-1} - \lambda_m)\mu_{m-1} f_{m-1} = 0.$$

But $m-1 \geq 1$ and all the coefficients $(\lambda_i - \lambda_m)\mu_i$ are non-zero. This contradicts the minimality of m . \square

It is not hard to see that, if A is diagonalizable, then the above procedure for expressing A in the form $A = TBT^{-1}$ can always be applied successfully.

Sometimes, A may not be diagonalizable over F , yet A may be diagonalizable as a matrix over a larger field. An interesting example of this is the matrix

$$R_\theta = \begin{bmatrix} \cos(\theta) & -\sin(\theta) \\ \sin(\theta) & \cos(\theta) \end{bmatrix}$$

which, of course, represents an anticlockwise rotation of the Euclidian plane \mathbb{R}^2 through an angle of θ . Regarding R_θ as a matrix over \mathbb{R} , then plainly R_θ is not diagonalizable unless θ is an integer multiple of π (in which case, $R_\theta = \pm I$.) But let us now regard R_θ as a matrix over the field of complex numbers \mathbb{C} . Thus, R_θ now represents a linear map $\mathbb{C}^2 \rightarrow \mathbb{C}^2$. The characteristic equation of R_θ is

$$0 = \det(R_\theta - \lambda I) = \begin{vmatrix} c - \lambda & -s \\ s & c - \lambda \end{vmatrix} = (c - \lambda)^2 + s^2 = \lambda^2 - 2c\lambda + 1$$

where $c = \cos(\theta)$ and $s = \sin(\theta)$. The solutions are $\lambda_1 = c + is = e^{i\theta}$ and $\lambda_2 = c - is = e^{-i\theta}$. It is easy to check that the corresponding eigenvectors are $f_1 = (1, -i)$ and $f_2 = (1, i)$. So $T = \begin{bmatrix} 1 & 1 \\ -i & i \end{bmatrix}$. We have $\det(T) = 2i$, hence $T^{-1} = \frac{1}{\det(T)} \begin{bmatrix} i & -1 \\ i & 1 \end{bmatrix} = \frac{1}{2} \begin{bmatrix} 1 & i \\ 1 & -i \end{bmatrix}$. Therefore

$$R_\theta = T \operatorname{diag}(\lambda_1, \lambda_2) T^{-1} = \frac{1}{2} \begin{bmatrix} 1 & 1 \\ -i & i \end{bmatrix} \begin{bmatrix} e^{i\theta} & 0 \\ 0 & e^{-i\theta} \end{bmatrix} \begin{bmatrix} 1 & i \\ 1 & -i \end{bmatrix}.$$

One advantage of working over \mathbb{C} rather than \mathbb{R} is that, for any $n \times n$ matrix A over \mathbb{C} , there always exist scalars $\lambda_i \in \mathbb{C}$ such that

$$0 = \det(A - \lambda I) = (\lambda_1 - \lambda) \dots (\lambda_n - \lambda).$$

Indeed, the Fundamental Theorem of Algebra asserts that, given complex numbers a_{n-1}, \dots, a_0 , then there exist complex numbers $\lambda_1, \dots, \lambda_n$ such that, for all complex numbers λ , we have

$$\lambda^n + a_{n-1}\lambda^{n-1} + \dots + a_1\lambda + a_0 = (\lambda - \lambda_1) \dots (\lambda - \lambda_n).$$

It is perhaps rather surprising that, even over \mathbb{C} , non-diagonalizable matrices exist.

Proposition: Let A be a 2×2 matrix over \mathbb{C} . Then A is non-diagonalizable if and only if A is similar to the matrix $\begin{bmatrix} a & b \\ 0 & a \end{bmatrix}$ where $a, b \in \mathbb{C}$ with $b \neq 0$.

Proof: Let B be the specified matrix. The characteristic equation of B is $0 = \det(C - \lambda I) - (a - \lambda)^2$. So a is the unique eigenvalue of C . It is now easy to see that the eigenvectors of C are precisely the vectors having the form $(x, 0)$ where x is a non-zero complex number. These vectors do not span the vector space \mathbb{C}^2 , so C is not diagonalizable, and any matrix similar to C is non-diagonalizable.

Conversely, suppose that A is non-diagonalizable. By the Fundamental Theorem of Algebra, A has an eigenvalue $a \in \mathbb{C}$. On the other hand, by the previous Proposition, A cannot have two distinct eigenvalues. That is to say, a must be the unique eigenvalue of A . Choose an eigenvector f_1 of A , and choose any vector f_2 in \mathbb{C}^2 such that $\{f_1, f_2\}$ is a basis for \mathbb{C}^2 . Letting α be the linear map $\mathbb{C}^2 \rightarrow \mathbb{C}^2$ represented by A with respect to the standard basis of \mathbb{C}^2 , then $\alpha(f_1) = Af_1 = af_1$ and $\alpha(f_2) = bf_1 + df_2$ for some $b, d \in \mathbb{C}$. So, letting B be the matrix representing α with respect to the basis $\{f_1, f_2\}$, we have $B = \begin{bmatrix} a & b \\ 0 & d \end{bmatrix}$. But B has a unique eigenvalue and the characteristic equation of B is $0 = \det(B - \lambda I) = (a - \lambda)(d - \lambda)$, hence $a = d$. Furthermore, B cannot be a diagonal matrix, so $b \neq 0$. \square

For many kinds of matrix that appear frequently in contexts of application — symmetric matrices or unitary matrices, for instance — there are results which guarantee diagonalizability. But the proof of the next theorem illustrates a scenario where failure of diagonalizability arises naturally. Actually, the fastest way to prove the theorem is by induction, nevertheless, the argument we present is an entertaining exercise in the theory developed above.

Incidentally, the following proof is also an illustration of the use of theory as opposed to calculation. We shall be arguing simply by making deductions from conceptual principles, without carrying out any substantial manipulations of written symbols.

Theorem: Let a, b, c be complex numbers with $a \neq 0$. Let x_0, x_1, \dots be an infinite sequence of complex numbers such that $ax_{n+2} + bx_{n+1} + cx_n = 0$ for all natural numbers n . If the quadratic equation $a\lambda^2 + b\lambda + c = 0$ has two distinct solutions λ_1 and λ_2 , then there exist complex numbers u_1 and u_2 such that $x_n = u_1\lambda_1^n + u_2\lambda_2^n$ for all n . If the quadratic equation has a unique non-zero solution λ , then there exist complex numbers μ and ν such that $x_n = (\mu + \nu)\lambda^n$ for all n . If 0 is the unique solution to the quadratic equation, then $x_n = 0$ for all $n \geq 2$.

Proof: We may assume that $a = 1$, since otherwise we can replace b and c with b/a and c/a , respectively. We may also assume that $c \neq 0$, since otherwise the required conclusion is trivial. We can now understand x_n to be defined for all integers n , with x_{-1}, x_{-2}, \dots and x_2, x_3, \dots recursively determined by x_0 and x_1 via the equality $x_{n+1} + bx_n + cx_{n-1} = 0$. Let us rewrite the equality as $\begin{bmatrix} x_{n+1} \\ x_n \end{bmatrix} = A \begin{bmatrix} x_n \\ x_{n-1} \end{bmatrix} = A^n \begin{bmatrix} x_1 \\ x_0 \end{bmatrix}$ where $A = \begin{bmatrix} -b & -c \\ 1 & 0 \end{bmatrix}$ and n is any integer. This makes sense because $\det(A) = c \neq 0$ and A^n is defined for all n .

Let V be the vector space over \mathbb{C} consisting of the functions $\mathbb{Z} \rightarrow \mathbb{C}$, with the evident addition and scalar multiplication operations. We shall be making use of the observation that the function $n \mapsto x_n$ can be regarded as a vector in V .

The eigenvalues of A are the complex numbers λ such that $\det(A - \lambda I) = 0$, in other words, $\lambda^2 + b\lambda + c = 0$. It is easy to see that, given any eigenvalue λ of A , then the corresponding eigenvectors are precisely the vectors $(\lambda y, y)$ where y is a non-zero complex number.

Suppose that the equation has two distinct solutions λ_1 and λ_2 . Note that λ_1 and λ_2 are both non-zero because $\lambda_1\lambda_2 = c \neq 0$. The matrix A is diagonalizable by a proposition above. (Alternatively, we can argue that A must be diagonalizable because the eigenvectors $(\lambda_1, 1)$ and $(\lambda_2, 1)$ comprise a basis for \mathbb{C}^2 .) Therefore $A = T \operatorname{diag}(\lambda_1, \lambda_2) T^{-1}$ for some invertible matrix T , and $A^n = T \operatorname{diag}(\lambda_1^n, \lambda_2^n) T^{-1}$ for all integers n . Observing that T is independent of n , it is not hard to see that, as a vector in V , the function $n \mapsto x_n$ is a linear combination of the functions $n \mapsto \lambda_1^n$ and $n \mapsto \lambda_2^n$.

It remains to deal with the case where the quadratic equation has a unique solution λ . We have $\lambda \neq 0$ because $\lambda^2 = c \neq 0$. All the eigenvectors of A are the scalar multiples of the vector $(\lambda, 1)$. These vectors do not span \mathbb{C}^2 , so A cannot be diagonalizable. By another proposition above, A is similar to the matrix $B = \begin{bmatrix} \lambda & 1 \\ 0 & \lambda \end{bmatrix}$. A straightforward inductive argument yields

$B^n = \begin{bmatrix} \lambda^n & n\lambda^{n-1} \\ 0 & \lambda^n \end{bmatrix}$ for all integers n . Writing $A = TBT^{-1}$, then $A^n = TB^nT^{-1}$ and

$$\begin{bmatrix} x_{n+1} \\ x_n \end{bmatrix} = T \begin{bmatrix} \lambda_0^n & n\lambda_0^{n-1} \\ 0 & \lambda_0^n \end{bmatrix} T^{-1} \begin{bmatrix} x_1 \\ x_0 \end{bmatrix}.$$

But the functions $n \mapsto \lambda^{n-1}$ and $n \mapsto \lambda^n$ and $n \mapsto \lambda^{n+1}$ are all scalar multiples of each other, and similarly for the functions $n \mapsto n\lambda^{n-1}$ and $n \mapsto n\lambda^n$ and $n \mapsto n\lambda^{n+1}$. So the function $n \mapsto x_n$ is a linear combination of the functions $n \mapsto \lambda^n$ and $n \mapsto n\lambda^n$. \square