

When Is A Linear Operator Diagonalizable?

Marco Abate

0. Introduction

As it often happens, everything began with a mistake. I was teaching for the third year in a row a linear algebra course to engineering freshmen. One of the highlights of the course was eigenvector theory, and in particular the diagonalization of linear operators on finite-dimensional vector spaces (i.e., of square real or complex matrices). Toward the end of the course I assigned a standard homework: prove that the matrix

$$A = \begin{vmatrix} -1 & -1 & 2 \\ -1 & 0 & 1 \\ 0 & -1 & 1 \end{vmatrix},$$

is diagonalizable. Easy enough, I thought. The characteristic polynomial is

$$p_A(\lambda) = \det(A - \lambda I_3) = -\lambda^3 + \lambda,$$

whose roots are evidently 0, 1, -1 . We have three distinct eigenvalues in a three-dimensional space, and a standard theorem ensures that A is diagonalizable.

To my surprise, the students came complaining that they were unable to solve the exercise. Perplexed (some of the complaining students were very bright), I looked over the exercise again—and I understood. What happened was that, in the homework, I *actually* gave them the matrix

$$B = \begin{vmatrix} -1 & -1 & 2 \\ -1 & 0 & 1 \\ 0 & -1 & -1 \end{vmatrix},$$

whose characteristic polynomial is

$$p_B(\lambda) = -\lambda^3 - 2\lambda^2 - \lambda + 2,$$

which has *no* rational roots. The students were unable to compute the eigenvalues of B , and they got stuck.

This accident started me wondering whether it might be possible to decide when a linear operator T on a finite-dimensional real or complex vector space is diagonalizable *without* computing the eigenvalues. If one is looking for an orthonormal basis of eigenvectors, the answer is well known to be yes: the spectral theorem says that such a basis exists in the complex case if and only if T is normal (i.e., it commutes with its adjoint), and if and only if T is symmetric in the real case. The aim of this note is to give an explicit procedure to decide whether a given linear operator on a finite-dimensional real or complex vector space is diagonalizable. By “explicit” I mean that it can always be worked out with pen and paper; it can be long, it can be tedious, but it can be done. Its ingredients (the minimal polynomial and Sturm’s theorem) are not new; but putting them together yields a result that can be useful as an aside in linear algebra classes.

1. The minimal polynomial

The first main ingredient in our procedure is the minimal polynomial. Let $T: V \rightarrow V$ be a linear operator on a finite-dimensional vector space over the field \mathbb{K} . We shall denote by T^k the composition of T with itself k times, and for any polynomial $p(t) = a_k t^k + \dots + a_0 \in \mathbb{K}[t]$ we put

$$p(T) = a_k T^k + \dots + a_1 T + \text{id}_V,$$

and say that p is *monic* if $a_k = 1$. A *minimal polynomial* $\mu_T \in \mathbb{K}[t]$ of the linear operator T is a monic polynomial of minimal degree such that $\mu(T) = 0$.

The theory of the minimal polynomial is standard. For completeness' sake, I briefly recall the results we shall need. First of all:

Proposition 1.1: *Let $T: V \rightarrow V$ be a linear operator on a finite-dimensional vector space V over the field \mathbb{K} . Then:*

- (i) *the minimal polynomial μ_T of T exists, has degree at most $n = \dim V$, and is unique;*
- (ii) *if $p \in \mathbb{K}[t]$ is such that $p(T) = 0$, then there is some $q \in \mathbb{K}[t]$ such that $p = q\mu_T$.*

For our procedure it is important to show that the minimal polynomial can be explicitly computed. Take $v \in V$, and let d be the minimal non-negative integer such that the vectors $\{v, T(v), \dots, T^d(v)\}$ are linearly dependent. Clearly $d \leq n$ always; $d = 0$ if and only if $v = 0$, and $d = 1$ if and only if v is an eigenvector of T . Choose $a_0, \dots, a_{d-1} \in \mathbb{K}$ such that

$$T^d(v) + a_{d-1}T^{d-1}(v) + \dots + a_1T(v) + a_0v = 0$$

(note that we can assume the coefficient of $T^d(v)$ to be 1 because of the minimality of d), and then set

$$\mu_{T,v}(t) = t^d + a_{d-1}t^{d-1} + \dots + a_1t + a_0 \in \mathbb{K}[t].$$

By definition, $v \in \text{Ker } \mu_{T,v}(T)$; more precisely, $\mu_{T,v}$ is the monic polynomial $p \in \mathbb{K}[t]$ of least degree such that $v \in \text{Ker } p(T)$.

Now, if $p \in \mathbb{K}[t]$ is any common multiple of μ_{T,v_1} and μ_{T,v_2} for any two vectors v_1 and v_2 , then both v_1 and v_2 belong to $\text{Ker } p(T)$. More generally, if $\mathcal{B} = \{v_1, \dots, v_n\}$ is a basis of V , and p is any common multiple of $\mu_{T,v_1}, \dots, \mu_{T,v_n}$, then $\mathcal{B} \subset \text{Ker } p(T)$, and thus $p(T) = 0$. Hence the following result comes as no surprise:

Proposition 1.2: *Let $T: V \rightarrow V$ be a linear operator on a finite-dimensional vector space V over the field \mathbb{K} . Let $\mathcal{B} = \{v_1, \dots, v_n\}$ be a basis of V . Then μ_T is the least common multiple of $\mu_{T,v_1}, \dots, \mu_{T,v_n}$.*

Proof: Let $p \in \mathbb{K}[t]$ be the least common multiple of $\mu_{T,v_1}, \dots, \mu_{T,v_n}$. We have already remarked that $p(T) = 0$, and so μ_T divides p . Conversely, for $j = 1, \dots, n$ write $\mu_T = q_j \mu_{T,v_j} + r_j$, with $\deg r_j < \deg \mu_{T,v_j}$. Then

$$0 = \mu_T(T)v_j = q_j(T)(\mu_{T,v_j}(T)v_j) + r_j(T)v_j = r_j(T)v_j,$$

and the minimality of the degree of μ_{T,v_j} forces $r_j \equiv 0$. Since every μ_{T,v_j} divides μ_T , their least common multiple p also divides μ_T , and hence $p = \mu_T$. \square

Thus one method to compute the minimal polynomial is to compute the polynomials $\mu_{T,v_1}, \dots, \mu_{T,v_n}$ and then their least common multiple. To avoid unnecessary calculations, it could be useful to remember that $\deg \mu_T \leq n$.

Example 1. Let us compute the minimal polynomials of the matrices A and B of the introduction. Let $\mathcal{B} = \{e_1, e_2, e_3\}$ be the canonical basis of \mathbb{R}^3 . We have

$$\begin{aligned} Ae_1 = Be_1 &= \begin{vmatrix} -1 \\ -1 \\ 0 \end{vmatrix}, & A^2e_1 = B^2e_1 &= \begin{vmatrix} 2 \\ 1 \\ 1 \end{vmatrix}, \\ A^3e_1 &= \begin{vmatrix} -1 \\ -1 \\ 0 \end{vmatrix} = Ae_1, & B^3e_1 &= \begin{vmatrix} -1 \\ -1 \\ -2 \end{vmatrix} = -2B^2e_1 - Be_1 + 2e_1; \end{aligned}$$

therefore

$$\mu_{A,e_1}(t) = t^3 - t, \quad \mu_{B,e_1}(t) = t^3 + 2t^2 + t - 2.$$

Since $\deg \mu_{A,e_1} = 3$ and the minimal polynomial of A should be a monic multiple of μ_{A,e_1} of degree at most three, we can conclude that $\mu_A = \mu_{A,e_1}$ without computing μ_{A,e_2} and μ_{A,e_3} (and it is easy to check that $\mu_{A,e_2}(t) = t^2 - t$ and $\mu_{A,e_3}(t) = t^3 - t$). For the same reason we have $\mu_B = \mu_{B,e_1}$.

Let $\lambda_1, \dots, \lambda_k \in \mathbb{K}$ be the distinct eigenvalues of T . If T is diagonalizable, then Proposition 1.2 immediately yields $\mu_T(t) = (t - \lambda_1) \cdots (t - \lambda_k)$. This is the standard characterization of diagonalizable linear operators:

Theorem 1.3: *Let $T: V \rightarrow V$ be a linear operator on a finite-dimensional vector space V over the field \mathbb{K} . Then T is diagonalizable if and only if μ_T is of the form*

$$\mu_T(t) = (t - \lambda_1) \cdots (t - \lambda_k), \tag{1.1}$$

where $\lambda_1, \dots, \lambda_k$ are distinct elements of \mathbb{K} .

Therefore to decide whether a given linear operator on a finite-dimensional vector space is diagonalizable it suffices to check whether its minimal polynomial is of the form (1.1).

2. The procedure

Our aim now is to find an effective procedure to decide whether a given polynomial $p \in \mathbb{K}[t]$ can be written in the form (1.1). To do so, we need to know when all the roots of p have multiplicity one, and when they all belong to the field \mathbb{K} . The first question has a standard answer:

Proposition 2.1: *Let $p \in \mathbb{K}[t]$ be a non-constant polynomial, and let $p' \in \mathbb{K}[t]$ denote its derivative. Then the following assertions are equivalent:*

- (i) p admits a root in \mathbb{K} of multiplicity greater than 1;
- (ii) p and p' have a common root in \mathbb{K} ;
- (iii) the greatest common divisor g.c.d. (p, p') of p and p' has a root in \mathbb{K} .

Recalling Theorem 1.3 we get the following

Corollary 2.2: *Let $T: V \rightarrow V$ be a linear operator on a finite-dimensional vector space V over the field \mathbb{K} . Then:*

- (i) *If \mathbb{K} is algebraically closed (e.g., $\mathbb{K} = \mathbb{C}$), then T is diagonalizable if and only if $\text{g.c.d.}(\mu_T, \mu'_T) = 1$;*
- (ii) *If \mathbb{K} is not algebraically closed (e.g., $\mathbb{K} = \mathbb{R}$), then T is diagonalizable if and only if all the roots of μ_T are in \mathbb{K} and $\text{g.c.d.}(\mu_T, \mu'_T) = 1$.*

To decide whether a complex linear operator T is diagonalizable it then suffices to compute the greatest common divisor of μ_T and μ'_T . On the other hand, if $\mathbb{K} = \mathbb{R}$ this is not enough; to complete the picture we need Sturm's theorem—and to state it we need a few more definitions.

Let $\mathbf{c} = (c_0, \dots, c_s) \in \mathbb{R}^{s+1}$ be a finite sequence of real numbers. If $c_0 \cdots c_s \neq 0$, the *number of variations in sign* of \mathbf{c} is the number of indices $1 \leq j \leq s$ such that $c_{j-1}c_j < 0$ (that is, such that c_{j-1} and c_j have opposite sign). If some element of \mathbf{c} is zero, then the *number of variations in sign* of \mathbf{c} is the number of variations in sign of the sequence of non-zero elements of \mathbf{c} . We denote the number of variations in sign of \mathbf{c} by $V_{\mathbf{c}}$.

Now let $p \in \mathbb{R}[t]$ be a non-constant polynomial. The *standard sequence* associated with p is the sequence $p_0, \dots, p_s \in \mathbb{R}[t]$ defined by

$$\begin{aligned} p_0 &= p, & p_1 &= p', \\ p_0 &= q_1 p_1 - p_2, & & \text{with } \deg p_2 < \deg p_1, \\ & \vdots & & \vdots \\ p_{j-1} &= q_j p_j - p_{j+1}, & & \text{with } \deg p_{j+1} < \deg p_j, \\ & \vdots & & \vdots \\ p_{s-1} &= q_s p_s, & & \text{(that is, } p_{s+1} \equiv 0). \end{aligned}$$

In other words, the standard sequence is obtained changing the sign in the remainder term of the Euclidean algorithm for the computation of $\text{g.c.d.}(p, p')$. In particular, $\text{g.c.d.}(p, p') = 1$ if and only if p_s is constant.

Sturm's theorem then says:

Theorem 2.3: *Let $p \in \mathbb{R}[t]$ be a polynomial such that $\text{g.c.d.}(p, p') = 1$, and take $a < b$ such that $p(a)p(b) \neq 0$. Let $p_0, \dots, p_s \in \mathbb{R}[t]$ be the standard sequence associated with p . Then the number of roots of p in $[a, b]$ is equal to $V_{\mathbf{a}} - V_{\mathbf{b}}$, where $\mathbf{a} = (p_0(a), \dots, p_s(a))$ and $\mathbf{b} = (p_0(b), \dots, p_s(b))$.*

For a proof see [1, pp. 295–299].

Now, for any polynomial $p(t) = a_d t^d + \cdots + a_0 \in \mathbb{R}[t]$ there exists $M > 0$ such that $p(t)$ has the same sign as a_d , the *leading coefficient* of p , if $t \geq M$ and the same sign as $(-1)^d a_d$ if $t \leq -M$. In particular, all the roots of p are contained in $[-M, M]$, and Sturm's theorem implies the following:

Corollary 2.4: *Let $p \in \mathbb{R}[t]$ be a non-constant polynomial such that $\text{g.c.d.}(p, p') = 1$. Let $p_0, \dots, p_s \in \mathbb{R}[t]$ be the standard sequence associated with p , and let d_j be the degree and $c_j \in \mathbb{R}$ the leading coefficient of p_j for $j = 0, \dots, s$. Then the number of real*

roots of p is given by $V_- - V_+$, where V_- is the number of variations in sign of the sequence $((-1)^{d_0}c_0, \dots, (-1)^{d_s}c_s)$, and V_+ is the number of variations in sign of the sequence (c_0, \dots, c_s) .

Proof: It suffices to choose $M > 0$ large enough so that $p_j(t)$ has the same sign as c_j when $t \geq M$ and the same sign as $(-1)^{d_j}c_j$ when $t \leq -M$, for each $j = 0, \dots, s$, and then apply Sturm's theorem with $a = -M$ and $b = M$. \square

We finally have all the ingredients necessary to state the desired procedure. Let $T: V \rightarrow V$ be a linear operator on a finite-dimensional vector space V over $\mathbb{K} = \mathbb{R}$ or \mathbb{C} . Then:

- (1) Compute the minimal polynomial μ_T .
- (2) Compute the standard sequence p_0, \dots, p_s associated with μ_T . If p_s is not constant, then T is not diagonalizable. If p_s is constant and $\mathbb{K} = \mathbb{C}$, then T is diagonalizable. If p_s is constant and $\mathbb{K} = \mathbb{R}$, go to Step (3).
- (3) Compute V_- and V_+ for μ_T . Then T is diagonalizable if and only if $V_- - V_+ = \deg \mu_T$.

Thus we are always able to decide whether a given linear operator on a finite-dimensional real or complex vector space is diagonalizable or not. One feature that I find particularly interesting in this procedure is that the solution of a typical linear algebra problem is reduced to an apparently totally unrelated manipulation of polynomials, showing in a simple case how different parts of mathematics can be connected in unexpected ways.

We end this note with some examples of application of our procedure.

Example 2. First of all, we solve the original homework. We have already computed the minimal polynomial $\mu_B(t) = t^3 + 2t^2 + t - 2$. The standard sequence associated with μ_B is

$$p_0(t) = t^3 + 2t^2 + t - 2, \quad p_1(t) = 3t^2 + 4t + 1, \quad p_2(t) = \frac{2}{9}t + \frac{20}{9}, \quad p_3(t) = -261.$$

Since p_3 is constant, B is diagonalizable over \mathbb{C} . To compute $V_- - V_+$ we count the number of variations in sign of the sequences $(-1, 3, -\frac{2}{9}, -261)$ and $(1, 3, \frac{2}{9}, -261)$. We obtain

$$V_- - V_+ = 2 - 1 = 1 < 3 = \deg \mu_B,$$

and so B is not diagonalizable over \mathbb{R} . On the other hand, the standard sequence associated with μ_A is

$$p_0(t) = t^3 - t, \quad p_1(t) = 3t^2 - 1, \quad p_2(t) = \frac{2}{3}t, \quad p_3(t) = 1.$$

The number of variations in sign of $(-1, 3, -\frac{2}{3}, 1)$ is $V_- = 3$, and of $(1, 3, \frac{2}{3}, 1)$ is $V_+ = 0$; therefore $V_- - V_+ = 3 - 0 = 3$, and thus A is diagonalizable over \mathbb{R} (as it should be).

Both these matrices were diagonalizable over \mathbb{C} ; since their minimal polynomials have degree 3, necessarily their (complex) eigenvalues are all distinct. In the next example this is not the case:

Example 3. Let

$$C = \begin{vmatrix} 0 & -2 & 2 & 6 \\ 2 & 4 & -1 & -5 \\ -3 & -4 & 3 & 7 \\ 1 & 2 & -1 & -3 \end{vmatrix}.$$

To compute the minimal polynomial of C we start, as in Example 1, by applying the iterates of C to e_1 . We get

$$Ce_1 = \begin{vmatrix} 0 \\ 2 \\ -3 \\ 1 \end{vmatrix}, \quad C^2e_1 = \begin{vmatrix} -4 \\ 6 \\ -10 \\ 4 \end{vmatrix}, \quad C^3e_1 = \begin{vmatrix} -8 \\ 6 \\ -14 \\ 6 \end{vmatrix}, \quad C^4e_1 = \begin{vmatrix} -4 \\ -8 \\ 0 \\ 0 \end{vmatrix}.$$

It is easy to check that $\{e_1, Ce_1, C^2e_1, C^3e_1\}$ are linearly independent and that

$$C^4e_1 - 4C^3e_1 + 8C^2e_1 - 8Ce_1 + 4e_1 = 0;$$

therefore $\deg \mu_{C, e_1} = 4$ and, as in Example 1, we can conclude that

$$\mu_C(t) = \mu_{C, e_1}(t) = t^4 - 4t^3 + 8t^2 - 8t + 4.$$

The standard sequence associated with μ_C starts with

$$p_0(t) = t^4 - 4t^3 + 8t^2 - 8t + 4, \quad p_1(t) = 4t^3 - 12t^2 + 16t - 8, \quad p_2(t) = -t^2 + 2t - 2.$$

Since p_2 divides p_1 , it is the last polynomial in the sequence; since it is not constant, we conclude that C is not diagonalizable even over \mathbb{C} (and in particular it cannot have four distinct eigenvalues).

Our final example involves a minimal polynomial of degree strictly less than the dimension of the space:

Example 4. Let

$$D = \begin{vmatrix} 2 & -2 & 2 & 8 \\ -2 & 4 & -1 & -9 \\ 1 & -4 & 3 & 11 \\ -1 & 2 & -1 & -5 \end{vmatrix}.$$

To compute the minimal polynomial of D we start again by applying the iterates of D to e_1 . We get

$$De_1 = \begin{vmatrix} 2 \\ -2 \\ 1 \\ -1 \end{vmatrix}, \quad D^2e_1 = \begin{vmatrix} 2 \\ -4 \\ 2 \\ -2 \end{vmatrix} = 2De_1 - 2e_1;$$

therefore $\mu_{D,e_1}(t) = t^2 - 2t + 2$, and we cannot conclude right now that $\mu_{D,e_1} = \mu_D$. Proceeding with the computations we get

$$De_2 = \begin{pmatrix} -2 \\ 4 \\ -4 \\ 2 \end{pmatrix}, \quad D^2e_2 = \begin{pmatrix} -4 \\ 6 \\ -8 \\ 4 \end{pmatrix} = 2De_2 - 2e_2; \quad De_3 = \begin{pmatrix} 2 \\ -1 \\ 3 \\ -1 \end{pmatrix}, \quad D^2e_3 = \begin{pmatrix} 4 \\ -2 \\ 4 \\ -2 \end{pmatrix} = 2De_3 - 2e_3;$$

$$De_4 = \begin{pmatrix} 8 \\ -9 \\ 11 \\ -5 \end{pmatrix}, \quad D^2e_4 = \begin{pmatrix} 16 \\ -18 \\ 22 \\ -12 \end{pmatrix} = 2De_4 - 2e_4;$$

hence we have $\mu_{D,e_2} = \mu_{D,e_3} = \mu_{D,e_4} = \mu_{D,e_1}$ and $\mu_D(t) = t^2 - 2t + 2$. In particular, D has (at most) two distinct complex eigenvalues.

The standard sequence associated with μ_D is

$$p_0(t) = t^2 - 2t + 2, \quad p_1(t) = 2t - 2, \quad p_2(t) = -1.$$

Since p_2 is constant, D is diagonalizable over \mathbb{C} . The number of variations in sign of the sequences $(1, -2, -1)$ and $(1, 2, -1)$ is

$$V_- - V_+ = 1 - 1 = 0 < 2 = \deg \mu_D,$$

and so D is not diagonalizable over \mathbb{R} .

References

- [1] N. Jacobson, *Basic Algebra, I*. Freeman, San Francisco, 1974.

Marco Abate
Dipartimento di Matematica
Università di Ancona
Via Brezze Bianche
60131 Ancona
Italy
abate@anva:1.unian.it

When Is A Linear Operator Diagonalizable?

Marco Abate

Brief Descriptive Summary

Have you ever wondered whether that mystifying 10×10 matrix were diagonalizable? Computing the characteristic polynomial is useless; there are no self-evident eigenvalues in view. And you don't know how to write a program to make the computer do the work. And you are losing your sleep about it (well, almost). Grieve no more! We are proud to present an explicit pen-and-paper procedure to let you decide whether any given square matrix is diagonalizable, both over the complex and over the real numbers! Read and try yourself; your sleep won't be troubled anymore.

[If this is too much of a joke, the following summary can be used instead.]

To decide whether a given square matrix is diagonalizable the usual techniques depend on precisely computing the eigenvalues, which often is a hopeless task. This paper describes instead an explicit pen-and-paper procedure to check diagonalizability without using any information about the eigenvalues; the most difficult operation needed is just the Euclidean algorithm for computing the greatest common divisor of two polynomials.

E-mail: abate@anvax1.unian.it

When Is A Linear Operator Diagonalizable?

Marco Abate

Biographical Sketch

MARCO ABATE was born in Milan in 1962, got his Ph.D. from Scuola Normale Superiore, Pisa in 1988, and now is full professor of Geometry at the University of Ancona (placing him among the ten youngest italian full professors). His interests include geometrical function theory, holomorphic dynamical systems, complex differential geometry and traveling around the world. In his spare time, he writes comics, sometimes even about mathematics.