POTENTIAL DIAGONALIZABILITY

KEITH CONRAD

When we work over a field that is not algebraically closed, we should distinguish two reasons a matrix doesn't diagonalize: 1) it diagonalizes over a larger field, but just not over the field in which we are working, and 2) it doesn't diagonalize even if we make the scalar field larger.

Example 1. The matrix $\left(\begin{smallmatrix} 0 & -1 \\ 1 & 0 \end{smallmatrix}\right)$ in $\mathrm{M}_2(\mathbf{R})$ is not diagonalizable, but it becomes diagonalizable in $\mathrm{M}_2(\mathbf{C})$ since its characteristic polynomial splits with distinct roots in $\mathbf{C}[T]$.

Example 2. The matrix $\left(\begin{smallmatrix} 1 & 1 \\ 0 & 1 \end{smallmatrix}\right)$ is not diagonalizable over all fields. Indeed, its only eigenvalue is 1 and its only eigenvectors are scalar multiples of $\binom{1}{0}$, so there is never a basis of eigenvectors for this matrix.

The second example is a more intrinsic kind of non-diagonalizability,[1] while the first example could be chalked up to the failure to work over the "right" field. If passing to a larger field makes a matrix diagonalizable, we get many of the benefits of diagonalizability, such as computing powers, if we are willing to work over the larger field.

To characterize matrices in $\mathrm{M}_n(F)$ that diagonalize after we enlarge the scalar field, we first show that certain concepts related to $\mathrm{M}_n(F)$ are insensitive to replacing F with a larger field: the minimal polynomial of a matrix and whether two matrices are conjugate to each other.

Lemma 3. *Let E/F be a field extension and $v_1, \ldots, v_m$ be vectors in F^n. Then, viewing $v_1, \ldots, v_m$ inside E^n, the v_i's are linearly independent over F if and only if they are linearly independent over E. If $v_1, \ldots, v_r$ are linearly dependent over E and $c_1v_1 + \cdots + c_mv_m = 0$ is an E-linear relation with a particular v_i having a nonzero coefficient, there is also an F-linear relation where v_i has a nonzero coefficient.*

Proof. Assuming $v_1, \ldots, v_m$ have no nontrivial E-linear relations, they certainly have no nontrivial F-linear relations.

Now suppose there is an E-linear relation

$$c_1v_1 + \cdots + c_mv_m = 0$$

where $c_i \in E$. Thinking about E as an F-vector space, the coefficients $c_1, \ldots, c_m$ have a finite-dimensional F-span inside of E. Let $\alpha_1, \ldots, \alpha_d \in E$ be an F-basis of this, and write $c_i = a_{i1}\alpha_1 + \cdots + a_{id}\alpha_d$ where $a_{ij} \in F$. Then in E^n we have

$$0 = \sum_{i=1}^{m} c_iv_i = \sum_{i=1}^{m} \left(\sum_{j=1}^{d} a_{ij}\alpha_j \right) v_i = \sum_{j=1}^{d} \alpha_j \left(\sum_{i=1}^{m} a_{ij}v_i \right).$$

If we look carefully at what this says in each of the n coordinates, it tells us that the F-linear combination of the α_j's using coefficients from the sums $\sum_{i=1}^{m} a_{ij}v_i$ in a common coordinate is 0, so by linear independence of the α_j's over F the sums $\sum_{i=1}^{m} a_{ij}v_i$ have all coordinates equal to 0. That means all the sums $\sum_{i=1}^{m} a_{ij}v_i = 0$ are 0 in F^n.

[1] Matrices that don't diagonalize over a larger field can be brought into a nearly diagonal form. This is the Jordan canonical form of the matrix, and is not discussed here.

Now if the v_i's are linearly independent over F then from $\sum_{i=1}^m a_{ij}v_i = 0$ we see that every a_{ij} must be 0, so every c_i is 0 and we have proved the v_i's are linearly independent over E. On the other hand, if the v_i's are linearly dependent over F and we started with a nontrivial E-linear relation where a particular coefficient c_i is nonzero, then some a_{ij} is nonzero so for that j the vanishing of $\sum_{i=1}^m a_{ij}v_i$ gives us a nontrivial F-linear relation where the coefficient of v_i is nonzero. □

Theorem 4. *Let E/F be a field extension.*

1. *For each $A \in \mathrm{M}_n(F)$, its minimal polynomial in $F[T]$ is its minimal polynomial in $E[T]$.*
2. *Two matrices in $\mathrm{M}_n(F)$ are conjugate in $\mathrm{M}_n(F)$ if and only if they are conjugate in $\mathrm{M}_n(E)$.*

Proof. **(1)** Let $m(T)$ be the minimal polynomial of A in $F[T]$. Since $m(T)$ is in $E[T]$ and kills A, $m(T)$ is divisible by the minimal polynomial of A in $E[T]$. Next we will show if $f(T) \in E[T]$ is nonzero and $f(A) = O$ then there is a polynomial in $F[T]$ of the same degree that kills A, so $\deg f \geq \deg m$. Therefore $m(T)$ is the minimal polynomial of A in $E[T]$.

Suppose

$$c_r A^r + c_{r-1}A^{r-1} + \cdots + c_1 A + c_0 I_n = O,$$

where $c_j \in E$ and $c_r \neq 0$. This gives us a nontrivial E-linear relation on $I_n, A, A^2, \ldots, A^r$. By Lemma 3 applied to $\mathrm{M}_n(F)$ as an F-vector space (viewed as F^{n^2}), $I_n, A, A^2, \ldots, A^r$ must have a nontrivial F-linear relation, and since $c_r \neq 0$ there is such a relation over F where the coefficient of A^r is again nonzero. This linear relation over F gives us a polynomial in $F[T]$ of degree r that kills A, and settles (1).

(2) Let $A, B \in \mathrm{M}_n(F)$ satisfy $A = PBP^{-1}$ for some $P \in \mathrm{GL}_n(E)$. We want to show $A = QBQ^{-1}$ for some $Q \in \mathrm{GL}_n(F)$. Rewrite $A = PBP^{-1}$ as $AP = PB$. Inside E, the F-span of the matrix entries of P has a finite basis, say $\alpha_1, \ldots, \alpha_d \in E$. Write $P = \alpha_1 C_1 + \cdots + \alpha_d C_d$, where $C_j \in \mathrm{M}_n(F)$. (Note each C_j is not O, since if some $C_j = O$ then the matrix entries of P are spanned over F without needing α_j.) Then

$$AP = \alpha_1 AC_1 + \cdots + \alpha_d AC_d, \quad PB = \alpha_1 C_1 B + \cdots + \alpha_d C_d B.$$

Since $AP = PB$ and $\alpha_1, \ldots, \alpha_d$ are linearly independent over F,

$$AC_j = C_j B \quad \text{for all } j.$$

Then for *all* $x_1, \ldots, x_d \in F$,

$$A(x_1 C_1 + \cdots + x_d C_d) = (x_1 C_1 + \cdots + x_d C_d)B. \tag{1}$$

Define

$$f(X_1, \ldots, X_d) = \det(X_1 C_1 + \cdots + X_d C_d) \in F[X_1, \ldots, X_d].$$

Since $f(\alpha_1, \ldots, \alpha_d) = \det(P) \neq 0$, f is not the zero polynomial. If we can find $a_1, \ldots, a_d \in F$ (not just in E!) such that $f(a_1, \ldots, a_d) \neq 0$, then the matrix $Q := \sum a_j C_j \in \mathrm{M}_n(F)$ has $\det(Q) \neq 0$ and (1) becomes $AQ = QB$, so $A = QBQ^{-1}$, which shows A and B are conjugate in $\mathrm{M}_n(F)$. To make this argument complete we need to find $a_1, \ldots, a_d \in F$ such that $f(a_1, \ldots, a_d) \neq 0$. If F is infinite then a general theorem on multivariable polynomials says *every* nonzero element of $F[X_1, \ldots, X_d]$ takes a nonzero value at some n-tuple from F, so we're done.

What if F is finite? Part (2) is still true, but it is a subtle issue because some polynomials over a finite field can be nonzero as abstract polynomials (that is, have a nonzero coefficient somewhere) while being zero as a polynomial function on he finite field (having value 0 at all substitutions from the field). For example, if F is a field of order q then $X^q - X$ is not 0 in $F[X]$ but its value at each element of F is 0. Therefore it is not immediately clear if

the proof we gave above is still valid when F is finite. To prove part (2) for both finite and infinite fields in a uniform manner, we will use the rational canonical form. Matrices A and B in $\mathrm{M}_n(F)$ have unique rational canonical forms R and S in $\mathrm{M}_n(F)$. The matrices R and S are in rational canonical form when viewed in $\mathrm{M}_n(E)$, so they are the rational canonical forms of A and B in $\mathrm{M}_n(E)$ too. If A and B are conjugate in $\mathrm{M}_n(E)$ then $R = S$, and viewing that equation in $\mathrm{M}_n(F)$ tells us A and B are conjugate in $\mathrm{M}_n(F)$. □

Corollary 5. *If a matrix in $\mathrm{M}_n(F)$ becomes diagonal over some field extension of F then it does so over the field generated by F and the eigenvalues of the matrix. In particular, the matrix diagonalizes over a finite extension of F.*

Proof. Let $A \in \mathrm{M}_n(F)$ and assume A diagonalizes in $\mathrm{M}_n(L)$ for some field extension L/F. Since A is diagonalizable in $\mathrm{M}_n(L)$, the minimal polynomial of A in $L[T]$ splits in $L[T]$ with distinct roots. These roots are the eigenvalues of A. We want to show that A diagonalizes over the field generated by F and the eigenvalues of A. Let this field be K, so $K \subset L$ and K/F is finite since the eigenvalues are algebraic over F.

By hypothesis, there is some $P \in \mathrm{GL}_n(L)$ such that $D := PAP^{-1}$ is a diagonal matrix. The diagonal entries of D are the eigenvalues of A, so $D \in \mathrm{M}_n(K)$. Since D and A both lie in $\mathrm{M}_n(K)$, their conjugacy in $\mathrm{M}_n(L)$ implies conjugacy in $\mathrm{M}_n(K)$ by Theorem 4(2). Therefore A is diagonalizable in $\mathrm{M}_n(K)$. □

Recall that a polynomial is called *separable* when it has no repeated roots (in a splitting field). Separability of a polynomial $f(T)$ can be checked without looking for the roots in a splitting field: $f(T)$ is separable if and only if $(f(T), f'(T)) = 1$, and this can be checked using Euclid's algorithm in $F[T]$.

Theorem 6. *A matrix in $\mathrm{M}_n(F)$ is diagonalizable over some extension field of F if and only its minimal polynomial in $F[T]$ is separable.*

Proof. Let A be a matrix in $\mathrm{M}_n(F)$ with minimal polynomial $m_A(T)$ in $F[T]$. Suppose A diagonalizes over some extension L/F. Since $m_A(T)$ is also the minimal polynomial of A in $L[T]$ (Theorem 4(1)), diagonalizability of A in $\mathrm{M}_n(L)$ implies $m_A(T)$ has no repeated roots, so $m_A(T)$ is separable.

Conversely, suppose $m_A(T)$ is separable. Let E be the splitting field of $m_A(T)$ over F. Then $m_A(T)$ splits in $E[T]$ with distinct roots, so A is diagonalizable in $\mathrm{M}_n(E)$. □

Definition 7. A linear operator on an n-dimensional F-vector space is called *potentially diagonalizable* when a matrix representation for it in $\mathrm{M}_n(F)$ is diagonalizable over some extension field of F.

Since all matrix representations of a linear operator have the same minimal polynomial (this is the minimal polynomial of the abstract linear operator) and Theorem 6 tells us that potential diagonalizability is determined by the minimal polynomial, if one matrix representation is diagonalizable over an extension field then all matrix representations are.

Example 8. The matrix $\left(\begin{smallmatrix} 0 & -1 \\ 1 & 0 \end{smallmatrix}\right)$ in $\mathrm{M}_2(\mathbf{R})$ is potentially diagonalizable, since it is diagonalizable in $\mathrm{M}_2(\mathbf{C})$. A linear operator on a 2-dimensional real vector space with matrix representation $\left(\begin{smallmatrix} 0 & -1 \\ 1 & 0 \end{smallmatrix}\right)$ is potentially diagonalizable.

The terminology in Definition 7 is not standard in this context, but the adjective potential is used in other contexts to refer to a property that is achieved only after an extension of the field, so its use here seems unobjectionable. The terminology used for this concept in Bourbaki [1] is "absolutely semisimple." In [2], Godement calls this property "semisimplicity," but the meaning of semisimplicity as used today means something slightly different, which we don't discuss here.

In terms of minimal polynomials,

$$\begin{aligned} A \text{ is upper-triangularizable} &\iff m_A(T) \text{ splits in } F[T], \\ A \text{ is diagonalizable} &\iff m_A(T) \text{ splits in } F[T] \text{ with distinct roots}, \\ A \text{ is potentially diagonalizable} &\iff m_A(T) \text{ is separable}. \end{aligned}$$

Example 9. Consider the three real matrices $\left(\begin{smallmatrix} 1 & 0 \\ 0 & 1 \end{smallmatrix}\right)$, $\left(\begin{smallmatrix} 0 & -1 \\ 1 & 0 \end{smallmatrix}\right)$, and $\left(\begin{smallmatrix} 1 & 1 \\ 0 & 1 \end{smallmatrix}\right)$ as operators on $\mathbf{R}^2$. The first is diagonalizable on $\mathbf{R}^2$, the second is not diagonalizable (on $\mathbf{R}^2$!) but is potentially diagonalizable, and the third is not even potentially diagonalizable. This is consistent with their minimal polynomials, which are $T-1$, T^2+1, and $(T-1)^2$.

When two linear operators $V \to V$ are diagonalizable and commute, every polynomial expression in the two operators is also diagonalizable. This result extends from diagonalizable to potentially diagonalizable operators. To prove this we use two lemmas.

Lemma 10. *For two matrices A and A' in $\mathrm{M}_n(F)$ and a finite extension field E/F, $A' \in E[A]$ if and only if $A' \in F[A]$, where $F[A]$ and $E[A]$ are the rings generated over F and E by A.*

Proof. Since $F[A] \subset E[A]$, if $A' \in F[A]$ then $A' \in E[A]$. Now assume $A' \in E[A]$. Write

$$A' = c_r A^r + c_{r-1}A^{r-1} + \cdots + c_0 I_n,$$

where $c_j \in E$. Rewrite this as

$$c_r A^r + c_{r-1}A^{r-1} + \cdots + c_0 I_n - A' = O.$$

This provides an E-linear dependence relation on $A', I_n, A, A^2, \ldots, A^r$ where the coefficient of A' is not 0. Therefore by Lemma 3 (applied to the vector space $\mathrm{M}_n(F)$ viewed as F^{n^2}), there is an F-linear dependence relation on $A', I_n, A, A^2, \ldots, A^r$ where the coefficient of A' is not 0, so $A' \in F[A]$. □

Lemma 11 (Lagrange)**.** *If $a_1, \ldots, a_n$ are distinct in F and $b_1, \ldots, b_n \in F$, there is a unique polynomial $f(T)$ in $F[T]$ of degree less than n such that $f(a_i) = b_i$.*

Proof. For uniqueness, if $f(T)$ and $g(T)$ both fit the conditions of the lemma then their difference $f(T) - g(T)$ has degree less than n and vanishes at each a_i. A nonzero polynomial doesn't have more roots than its degree, so $f(T) - g(T) = 0$, hence $f(T) = g(T)$. That settles uniqueness.

As for existence, it suffices to write down a polynomial of degree less than n that is 1 at a_i and 0 at a_j for each $j \neq i$. Then a linear combination of these polynomials with coefficients b_i will equal b_i at each a_i. The polynomial

$$\prod_{\substack{j=1 \\ j \neq i}}^{n} \frac{T - a_j}{a_i - a_j}$$

has the desired property: at a_i it is 1 and at every other a_j it is 0. Its degree is $n-1$. □

Lemma 11 is called Lagrange interpolation.

Theorem 12. *Let V be an F-vector space and let A and B be F-linear operators $V \to V$ that commute and are potentially diagonalizable.*

1. *Every element of $F[A, B]$ is potentially diagonalizable. In particular, $A+B$ and AB are potentially diagonalizable.*
2. *If F is infinite then for all but finitely many $c \in F$, $F[A, B] = F[A + cB]$.*

Proof. Pick a basis for V to identify A and B with (commuting) matrices in $\mathrm{M}_n(F)$, where $n = \dim_F(V)$. This passage from operators to matrices doesn't change minimal polynomials. Now $F[A, B] \subset \mathrm{M}_n(F)$. Let E/F be a field extension in which $m_A(T)$ and $m_B(T)$ both split. They each have no repeated roots, so in $\mathrm{M}_n(E)$ the matrices A and B are simultaneously diagonalizable.

(1) We want to show every matrix in $F[A, B]$ is potentially diagonalizable. Since A and B diagonalize over E and commute, they are simultaneously diagonalizable, so every matrix in $E[A, B]$ is diagonalizable. Therefore every matrix in $F[A, B] \subset E[A, B]$ is diagonalizable in $\mathrm{M}_n(E)$ and hence is potentially diagonalizable when regarded as a matrix in $\mathrm{M}_n(F)$.

(2) We want to show for all but finitely many $c \in F$ that A and B are in $F[A + cB]$. It suffices to show A and B are in $E[A + cB]$ by Lemma 10. The advantage to working over E is that A and B diagonalize in $\mathrm{M}_n(E)$, and simultaneously at that since A and B commute. So some $P \in \mathrm{GL}_n(E)$ simultaneously conjugates A and B into diagonal matrices:

$$PAP^{-1} = \begin{pmatrix} \lambda_1 & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & \lambda_n \end{pmatrix}, \quad PBP^{-1} = \begin{pmatrix} \mu_1 & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & \mu_n \end{pmatrix}.$$

The diagonal entries are the eigenvalues of A and B. There could be repetitions among these eigenvalues (consider the case when A is a scaling operator).

For all $c \in F$,

$$P(A + cB)P^{-1} = \begin{pmatrix} \lambda_1 + c\mu_1 & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & \lambda_n + c\mu_n \end{pmatrix}.$$

When could these diagonal entries coincide? If $\lambda_i + c\mu_i = \lambda_j + c\mu_j$ and $\mu_i \neq \mu_j$ then $c = (\lambda_i - \lambda_j)/(\mu_j - \mu_i)$. This is only a finite number of possibilities for c (as i and j vary), so as long as we avoid these finitely many values for c, which is possible since F is infinite, we have

$$\lambda_i + c\mu_i = \lambda_j + c\mu_j \Longrightarrow \mu_i = \mu_j \Longrightarrow \lambda_i = \lambda_j.$$

By Lagrange interpolation, there is a polynomial $h(T) \in E[T]$ such that $h(\lambda_i + c\mu_i) = \lambda_i$ for all i. At first you might think there is a well-definedness issue here, because the $\lambda_i + c\mu_i$'s may not all be distinct. (There is no polynomial satisfying $h(a) = 0$ and $h(b) = 1$ if $a = b$.) But because equal $\lambda_i + c\mu_i$'s implies equal λ_i's, the interpolation we set up makes sense. Now

$$\begin{aligned} h(P(A + cB)P^{-1}) &= \begin{pmatrix} h(\lambda_1 + c\mu_1) & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & h(\lambda_n + c\mu_n) \end{pmatrix} \\ &= \begin{pmatrix} \lambda_1 & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & \lambda_n \end{pmatrix} \\ &= PAP^{-1}. \end{aligned}$$

Since $h(P(A + cB)P^{-1}) = Ph(A + cB)P^{-1}$, we get $h(A + cB) = A$, so $A \in E[A + cB]$. In a similar way we get $B \in E[A + cB]$. □

We will now see a lovely application (taken from [3]) of Theorem 12 to field extensions.

Theorem 13. *Let L/K be a finite extension of fields and α and β be in L.*

(1) *If α and β have separable minimal polynomials in $K[T]$ then every element of the field $K(\alpha, \beta)$ has a separable minimal polynomial in $K[T]$.*
(2) *If every element of L has a separable minimal polynomial in $K[T]$ then $L = K(\gamma)$ for some $\gamma \in L$.*

Proof. The link between this theorem and diagonalizability is based on the interpretation of each element of L as a K-linear map $L \to L$ using multiplication by the element: for $\alpha \in L$, let[2] $m_\alpha \colon L \to L$ by $m_\alpha(x) := \alpha x$. Since m_α is K-linear on L, the correspondence $\alpha \mapsto m_\alpha$ is a function $L \to \mathrm{Hom}_K(L, L)$. Since $(\alpha + \beta)x = \alpha x + \beta x$ and $\alpha(\beta x) = (\alpha\beta)x$, we have

$$m_{\alpha+\beta} = m_\alpha + m_\beta, \quad m_\alpha \circ m_\beta = m_{\alpha\beta}.$$

Also $m_{c\alpha} = cm_\alpha$ for $c \in K$, and $m_1 = \mathrm{id}_L$. Thus the function $L \to \mathrm{Hom}_K(L, L)$ where $\alpha \mapsto m_\alpha$ is K-linear and a ring homomorphism. It is injective since we can recover α from m_α by acting on the distinguished element 1 in L: $m_\alpha(1) = \alpha \cdot 1 = \alpha$. In particular, $m_\alpha = O$ if and only if $\alpha = 0$.

Since $\alpha \mapsto m_\alpha$ is a ring homomorphism fixing K, for all $f(T) \in K[T]$ we have $f(m_\alpha) = m_{f(\alpha)}$. Thus $f(m_\alpha) = O$ if and only if $f(\alpha) = 0$, so the minimal polynomials of α and m_α in $K[T]$ are the same for all $\alpha \in L$. This will let us exploit the link between separable polynomials and potential diagonalizability.

(1) Since L/K is a finite extension, α and β are algebraic over K, so the field $K(\alpha, \beta)$ equals $K[\alpha, \beta]$. The embedding of L into $\mathrm{Hom}_K(L, L)$ identifies the field $K[\alpha, \beta]$ with the ring $K[m_\alpha, m_\beta]$ and preserves minimal polynomials, so it suffices to show each operator in $K[m_\alpha, m_\beta]$ has a separable minimal polynomial.

We can apply Theorem 12(1) to the vector space L over the field K and the operators $A = m_\alpha$, and $B = m_\beta$ on L. These operators commute since α and β commute. The minimal polynomials of A and B in $K[T]$ equal those of α and β, so the polynomials are separable by hypothesis. Therefore A and B are potentially diagonalizable, so every operator in $K[A, B]$ is potentially diagonalizable by Theorem 12(1), so the minimal polynomial of every operator in $K[A, B]$ is separable. In particular, for every $\gamma \in K[\alpha, \beta]$ the minimal polynomial of $m_\gamma \in K[A, B]$ is separable in $K[T]$. This is the minimial polynomial of γ in $K[T]$, so γ is separable over K.

(2) Since $[L : K]$ is finite, it suffices to show that for α and β in L, $K(\alpha, \beta) = K(\gamma)$ for some γ. First suppose K is infinite. As in (1), identify $K(\alpha, \beta) = K[\alpha, \beta]$ with $K[m_\alpha, m_\beta]$ in $\mathrm{Hom}_K(L, L)$. By Theorem 12(2), $K[m_\alpha, m_\beta] = K[m_\alpha + cm_\beta]$ for some $c \in K$ (in fact, all but finitely many $c \in K$ will work). As $m_\alpha + cm_\beta = m_{\alpha+c\beta}$, we get $K[m_\alpha, m_\beta] = K[m_{\alpha+c\beta}]$, so $K[\alpha, \beta] = K[\alpha + c\beta]$.

Now suppose K is finite. Then L is finite, so from the theory of finite fields $L^\times$ is cyclic: $L^\times = \langle\gamma\rangle$. Therefore $L = K(\gamma)$. □

References

[1] N. Bourbaki, *Algebra II*, Springer-Verlag, Berlin, 2003.
[2] R. Godement, *Algebra*, Houghton-Mifflin, Boston, 1968.
[3] F. Richman, "Separable Extensions and Diagonalizability," Amer. Math. Monthly **97** (1990), 395–398.

[2]The m-notation in this proof does not mean minimial polynomial!