Queen Mary University of London

MTH6140 Linear Algebra II

Notes 5 **18th November 2010**

5 Linear maps on a vector space

In this chapter we consider a linear map T from a vector space V to itself. If $\dim(V) = n$ then, as in the last chapter, we can represent T by an $n \times n$ matrix relative to any basis for V. However, this time we have less freedom: instead of having two bases to choose, there is only one. This makes the theory much more interesting!

5.1 Projections and direct sums

We begin by looking at a particular type of linear map whose importance will be clear later on.

Definition 5.1 The linear map $P : V \to V$ is a *projection* if $P^2 = P$ (where, as usual, P^2 is defined by $P^2(v) = P(P(v))$).

Proposition 5.1 If $P : V \to V$ is a projection, then $V = \mathrm{Im}(P) \oplus \mathrm{Ker}(P)$.

Proof We have two things to do:

$\mathrm{Im}(P) + \mathrm{Ker}(P) = V$: Take any vector $v \in V$, and let $w = P(v) \in \mathrm{Im}(P)$. We claim that $v - w \in \mathrm{Ker}(P)$. This holds because

$$P(v-w) = P(v) - P(w) = P(v) - P(P(v)) = P(v) - P^2(v) = 0,$$

since $P^2 = P$. Now $v = w + (v - w)$ is the sum of a vector in $\mathrm{Im}(P)$ and one in $\mathrm{Ker}(P)$.

$\mathrm{Im}(P) \cap \mathrm{Ker}(P) = \{0\}$: Take $v \in \mathrm{Im}(P) \cap \mathrm{Ker}(P)$. Then $v = P(w)$ for some vector w; and

$$0 = P(v) = P(P(w)) = P^2(w) = P(w) = v,$$

as required (the first equality holding because $v \in \mathrm{Ker}(P)$).

It goes the other way too: if $V = U \oplus W$, then there is a projection $P : V \to V$ with $\mathrm{Im}(P) = U$ and $\mathrm{Ker}(P) = W$. For every vector $v \in V$ can be uniquely written as $v = u + w$, where $u \in U$ and $w \in W$; we define P by the rule that $P(v) = u$. This is a well-defined map, and is easily checked to be linear. Now the assertions are clear.

The diagram in Figure 1 shows geometrically what a projection is. It moves any vector v in a direction parallel to $\mathrm{Ker}(P)$ to a vector lying in $\mathrm{Im}(P)$.

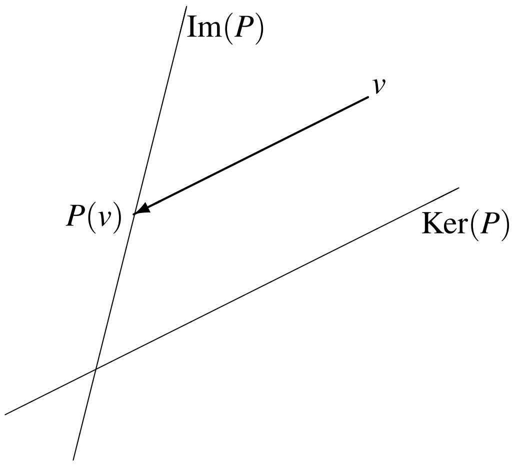

Figure 1: A projection

We can extend this to direct sums with more than two terms. First, notice that if P is a projection and $P' = I - P$ (where I is the identity map), then P' is also a projection, since

$$(P')^2 = (I-P)^2 = I - 2P + P^2 = I - 2P + P = I - P = P';$$

and $P + P' = I$; also $PP' = P(I - P) = P - P^2 = O$. Finally, we see that $\mathrm{Ker}(P) = \mathrm{Im}(P')$; so $V = \mathrm{Im}(P) \oplus \mathrm{Im}(P')$. In this form the result extends:

Proposition 5.2 Suppose that $P_1, P_2, \ldots, P_r$ are projections on V satisfying

(a) $P_1 + P_2 + \cdots + P_r = I$, where I is the identity transformation;

(b) $P_i P_j = O$ for $i \neq j$.

Then $V = U_1 \oplus U_2 \oplus \cdots \oplus U_r$, where $U_i = \mathrm{Im}(P_i)$.

Proof We have to show that any vector v can be *uniquely* written in the form $v = u_1 + u_2 + \cdots + u_r$, where $u_i \in U_i$ for $i = 1, \ldots, r$. We have

$$v = I(v) = P_1(v) + P_2(v) + \cdots + P_r(v) = u_1 + u_2 + \cdots + u_r,$$

where $u_i = P_i(v) \in \mathrm{Im}(P_i)$ for $i = 1, \ldots, r$. So any vector can be written in this form. Now suppose (in order to obtain a contradiction) that we have any expression

$$v = u'_1 + u'_2 + \cdots + u'_r,$$

with $u'_i \in U_i$ for $i = 1, \ldots, r$. Since $u'_i \in U_i = \mathrm{Im}(P_i)$, we have $u'_i = P_i(v_i)$ for some v_i; then

$$P_i(u'_i) = P_i^2(v_i) = P_i(v_i) = u'_i.$$

On the other hand, for $j \neq i$, we have

$$P_i(u'_j) = P_i P_j(v_j) = 0,$$

since $P_i P_j = O$. So applying P_i to the expression for v, we obtain

$$P_i(v) = P_i(u'_1) + P_i(u'_2) + \cdots + P_i(u'_r) = P_i(u'_i) = u'_i,$$

since all terms in the sum except the ith are zero. So the only possible expression is given by $u_i = P_i(v)$, and the proof is complete.

Conversely, if $V = U_1 \oplus U_2 \oplus \cdots \oplus U_r$, then we can find projections $P_i, P_2, \ldots, P_r$ satisfying the conditions of the above Proposition. For any vector $v \in V$ has a unique expression as

$$v = u_1 + u_2 + \cdots + u_r$$

with $u_i \in U_i$ for $i = 1, \ldots, r$; then we define $P_i(v) = u_i$.

The point of this is that projections give us another way to recognise and describe direct sums.

5.2 Linear maps and matrices

Let $T : V \to V$ be a linear map. If we choose a basis $v_1, \ldots, v_n$ for V, then V can be written in coordinates as K^n, and T is represented by a matrix A, say, where

$$T(v_i) = \sum_{j=1}^{n} A_{ji} v_j$$

for each $i \in \{1, \ldots, n\}$. Then just as in the last section, the action of T on V is represented by the action of A on K^n: $T(v)$ is represented by the product Av. Also, as in the last chapter, sums and products (and hence arbitrary polynomials) of linear maps are represented by sums and products of the representing matrices: that is, for any polynomial $f(x)$, the map $f(T)$ is represented by the matrix $f(A)$.

What happens if we change the basis? This also follows from the formula we worked out in the last chapter. However, there is only *one* basis to change.

Proposition 5.3 Let T be a linear map on V which is represented by the matrix A relative to a basis B, and by the matrix A' relative to a basis B'. Let $P = P_{B,B'}$ be the transition matrix between the two bases. Then

$$A' = P^{-1}AP.$$

Proof This is just Proposition 3.6, since P and Q are the same here.

Definition 5.2 Two $n \times n$ matrices A and A' are said to be *similar* if $A' = P^{-1}AP$ for some invertible matrix P.

Thus similarity is an equivalence relation, and two matrices are similar if and only if they represent the same linear map with respect to different bases.

The crucial difference from Chapter 3 is the following. In the present Chapter 5 we have less freedom: we only allow ourselves *one* basis. Recall that in Chapter 3 we had *two* bases: one for the vector space V, and one for the vector space W (we allowed the possibility that $W = V$, but in that case we still allowed ourselves to use *two* bases for V).
Consequently the theory of *similarity* of square matrices is very different from the theory of *equivalence*. Roughly speaking it is 'easy' for two square matrices to be equivalent (since for this they only have to have the same *rank*), but 'difficult' for them to be similar.
In particular, if two matrices are similar then they are certainly equivalent.

There is no simple canonical form for similarity like the one for equivalence that we met earlier. For the rest of this section we look at a special class of matrices or linear maps, the "diagonalisable" ones, where we do have a nice simple representative of the similarity class.

5.3 Eigenvalues and eigenvectors

Definition 5.3 Let T be a linear map on V. A vector $v \in V$ is said to be an *eigenvector* of T, with *eigenvalue* $\lambda \in K$, if $v \neq 0$ and $T(v) = \lambda v$. The set $\{v : T(v) = \lambda v\}$ consisting of the zero vector and the eigenvectors with eigenvalue λ is called the λ-*eigenspace* of T. A subspace U of V is called an *eigenspace* for T if it is the λ-eigenspace of T for some $\lambda \in K$.

Note that we require that $v \neq 0$; otherwise the zero vector would be an eigenvector for any value of λ. With this requirement, each eigenvector has a unique eigenvalue: for if $T(v) = \lambda v = \mu v$, then $(\lambda - \mu)v = 0$, and so (since $v \neq 0$) we have $\lambda = \mu$.

The name *eigenvalue* is a mixture of German and English; it means "characteristic value" or "proper value" (here "proper" is used in the sense of "property"). Another term used in older books is "latent root". Here "latent" means "hidden": the idea is that the eigenvalue is somehow hidden in a matrix representing T, and we have to extract it by some procedure. We'll see how to do this soon.

Example 5.1 Let

$$A = \begin{bmatrix} -6 & 6 \\ -12 & 11 \end{bmatrix}.$$

The vector $v = \begin{bmatrix} 3 \\ 4 \end{bmatrix}$ satisfies

$$\begin{bmatrix} -6 & 6 \\ -12 & 11 \end{bmatrix} \begin{bmatrix} 3 \\ 4 \end{bmatrix} = 2 \begin{bmatrix} 3 \\ 4 \end{bmatrix},$$

so is an eigenvector with eigenvalue 2. Similarly, the vector $w = \begin{bmatrix} 2 \\ 3 \end{bmatrix}$ is an eigenvector with eigenvalue 3.

If we knew that, for example, 2 is an eigenvalue of A, then we could find a corresponding eigenvector $\begin{bmatrix} x \\ y \end{bmatrix}$ by solving the linear equations

$$\begin{bmatrix} -6 & 6 \\ -12 & 11 \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} = 2 \begin{bmatrix} x \\ y \end{bmatrix}.$$

In the next-but-one section, we will see how to find the eigenvalues, and the fact that there cannot be more than n of them for an $n \times n$ matrix.

5.4 Diagonalisability

A square matrix A is said to be *diagonal* if $A_{ij} = 0$ whenever $i \neq j$ (i.e. its only non-zero entries appear on the diagonal). Some linear maps have a particularly simple representation by matrices:

Definition 5.4 The linear map T on V is *diagonalisable* if there is a basis of V relative to which the matrix representing T is a diagonal matrix.

We have the analogous definition for matrices:

Definition 5.5 A square matrix is called *diagonalisable* if it is similar to a diagonal matrix (i.e. there is an invertible matrix P such that $P^{-1}AP$ is a diagonal matrix).

Proposition 5.4 The linear map T on V is diagonalisable if and only if there is a basis of V consisting of eigenvectors of T.

Proof Suppose that $v_1,\dots,v_n$ is such a basis showing that T is diagonalisable. Then $T(v_i) = A_{ii}v_i$ for $i = 1,\dots,n$, where A_{ii} is the ith diagonal entry of the diagonal matrix A. Thus, the basis vectors are eigenvectors. Conversely, if we have a basis $v_1,\dots,v_n$ of eigenvectors, with corresponding eigenvalues $\lambda_1,\dots,\lambda_n$, then the matrix representing T (with respect to that basis) is diagonal (i.e. it is the diagonal matrix with entries $\lambda_1,\dots,\lambda_n$ along the diagonal).

Example 5.2 The matrix $\begin{bmatrix} 1 & 2 \\ 0 & 1 \end{bmatrix}$ is not diagonalisable. It is easy to see that its only eigenvalue is 1 (see the Remark after Theorem 5.7), and the only eigenvectors are scalar multiples of $\begin{bmatrix} 1 \\ 0 \end{bmatrix}$. So we cannot find a basis of eigenvectors.

Theorem 5.5 Let $T : V \to V$ be a linear map. Then the following are equivalent:

(a) T is diagonalisable;

(b) V is the direct sum of the eigenspaces of T;

(c) $T = \lambda_1 P_1 + \cdots + \lambda_r P_r$, where $\lambda_1,\dots,\lambda_r$ are the distinct eigenvalues of T, and $P_1,\dots,P_r$ are projections satisfying $P_1 + \cdots + P_r = I$ and $P_iP_j = 0$ for $i \neq j$.

Proof Let $\lambda_1,\dots,\lambda_r$ be the distinct eigenvalues of T, and let $v_{i1},\dots,v_{im_i}$ be a basis for the λ_i-eigenspace of T. Then T is diagonalisable if and only if the union of these bases is a basis for V. So (a) and (b) are equivalent.

Now suppose that (b) holds. The converse to Proposition 5.2 shows that there are projections $P_1,\dots,P_r$ satisfying the conditions of (c) where $\mathrm{Im}(P_i)$ is the λ_i-eigenspace. Now in this case it is easily checked that T and $\sum \lambda_i P_i$ agree on every vector in V, so they are equal. So (b) implies (c).

Finally, if $T = \sum \lambda_i P_i$, where the P_i satisfy the conditions of (c), then V is the direct sum of the spaces $\mathrm{Im}(P_i)$, and $\mathrm{Im}(P_i)$ is the λ_i-eigenspace. So (c) implies (b), and we are done.

Example 5.3 The matrix $A = \begin{bmatrix} -6 & 6 \\ -12 & 11 \end{bmatrix}$ from Example 5.1 is diagonalisable, since the eigenvectors $\begin{bmatrix} 3 \\ 4 \end{bmatrix}$ and $\begin{bmatrix} 2 \\ 3 \end{bmatrix}$ are linearly independent, and so form a basis for $\mathbb{R}^2$. Indeed, we see that

$$\begin{bmatrix} -6 & 6 \\ -12 & 11 \end{bmatrix} \begin{bmatrix} 3 & 2 \\ 4 & 3 \end{bmatrix} = \begin{bmatrix} 3 & 2 \\ 4 & 3 \end{bmatrix} \begin{bmatrix} 2 & 0 \\ 0 & 3 \end{bmatrix},$$

so that $P^{-1}AP$ is diagonal, where P is the matrix whose columns are the eigenvectors of A.

Furthermore, one can find two projection matrices whose column spaces are the eigenspaces, namely

$$P_1 = \begin{bmatrix} 9 & -6 \\ 12 & -8 \end{bmatrix}, \qquad P_2 = \begin{bmatrix} -8 & 6 \\ -12 & 9 \end{bmatrix}.$$

Check directly that $P_1^2 = P_1$, $P_2^2 = P_2$, $P_1P_2 = P_2P_1 = 0$, $P_1 + P_2 = I$, and $2P_1 + 3P_2 = A$.

5.5 Characteristic and minimal polynomials

We defined the determinant of a square matrix A. Now we want to define the determinant of a linear map T. The obvious way to do this is to take the determinant of any matrix representing T. For this to be a good definition, we need to show that it doesn't matter which matrix we take; in other words, that $\det(A') = \det(A)$ if A and A' are similar. But, if $A' = P^{-1}AP$, then

$$\det(P^{-1}AP) = \det(P^{-1})\det(A)\det(P) = \det(A),$$

since $\det(P^{-1})\det(P) = (\det(P))^{-1}\det(P) = 1$. So our plan will succeed:

Definition 5.6 (a) The *determinant* $\det(T)$ of a linear map $T : V \to V$ is the determinant of any matrix representing T.

(b) The *characteristic polynomial* $c_T(x)$ of a linear map $T : V \to V$ is the characteristic polynomial of any matrix representing T.

(c) The *minimal polynomial* $m_T(x)$ of a linear map $T : V \to V$ is the monic polynomial of smallest degree which is satisfied by T.

(d) Similarly, the *minimal polynomial* $m_A(x)$ of a square matrix A is the monic polynomial of smallest degree which is satisfied by A.

The second part of the definition is OK, by the same reasoning as the first (since $c_A(x)$ is just a determinant). But the third part also creates a bit of a problem: how do we know that T satisfies any polynomial? Well, the Cayley–Hamilton Theorem tells us that $c_A(A) = O$ for any matrix A representing T. Now $c_A(A)$ represents $c_T(T)$, and $c_A = c_T$ by definition; so $c_T(T) = O$. Indeed, the Cayley–Hamilton Theorem can be stated in the following form:

Proposition 5.6 For any linear map T on V, $m_T(x)$ divides $c_T(x)$ (as polynomials) [i.e. $c_T(x)/m_T(x)$ is itself a polynomial].
The analogous fact holds for matrices: if A is a square matrix then $m_A(x)$ divides $c_A(x)$.

Proof Suppose not; then we can divide $c_T(x)$ by $m_T(x)$, getting a quotient $q(x)$ and non-zero remainder $r(x)$; that is,

$$c_T(x) = m_T(x)q(x) + r(x).$$

Substituting T for x, using the fact that $c_T(T) = m_T(T) = O$, we find that $r(T) = 0$. But r is the *remainder*, so its degree is less than the degree of m_T; this contradicts the definition of m_T as the polynomial of least degree satisfied by T.

Theorem 5.7 Let T be a linear map on V. Then the following conditions are equivalent for an element λ of K:

(a) λ is an eigenvalue of T;

(b) λ is a root of the characteristic polynomial of T;

(c) λ is a root of the minimal polynomial of T.

In other words,

$$m_T(\lambda) = 0 \iff c_T(\lambda) = 0 \iff T(v) = \lambda v,\ v \neq 0,$$

and of course the analogous fact holds for square matrices A:

$$m_A(\lambda) = 0 \iff c_A(\lambda) = 0 \iff Av = \lambda v,\ v \neq 0,$$

Remark: Hence to find the eigenvalues of T: take a matrix A representing T; write down its characteristic polynomial $c_A(x) = \det(xI - A)$; and find the roots of this polynomial. In our earlier example,

$$\begin{vmatrix} x+6 & -6 \\ 12 & x-11 \end{vmatrix} = (x+6)(x-11)+72 = x^2 - 5x + 6 = (x-2)(x-3),$$

so the eigenvalues are 2 and 3, as we found.

Proof (b) implies (a): Suppose that $c_T(\lambda) = 0$, that is, $\det(\lambda I - T) = 0$. Then $\lambda I - T$ is not invertible, so its kernel is non-zero. Pick a non-zero vector v in $\mathrm{Ker}(\lambda I - T)$. Then $(\lambda I - T)v = 0$, so that $T(v) = \lambda v$; that is, λ is an eigenvalue of T.

(c) implies (b): Suppose that λ is a root of $m_T(x)$. Then $(x-\lambda)$ divides $m_T(x)$. But $m_T(x)$ divides $c_T(x)$, by the Cayley–Hamilton Theorem (i.e. Proposition 5.6): so $(x-\lambda)$ divides $c_T(x)$, whence λ is a root of $c_T(x)$.

(a) implies (c): Let λ be an eigenvalue of A with eigenvector v. We have $T(v) = \lambda v$. By induction, $T^k(v) = \lambda^k v$ for any k, and so $f(T)(v) = f(\lambda)v$ for any polynomial f. Choosing $f = m_T$, we have $m_T(T) = 0$ by definition, so $m_T(\lambda)v = 0$; since $v \neq 0$, we have $m_T(\lambda) = 0$, as required.

Using this result, we can give a necessary and sufficient condition for T to be diagonalisable. First, a lemma.

Lemma 5.8 Let $v_1,\ldots,v_r$ be eigenvectors of T with distinct eigenvalues $\lambda_1,\ldots,\lambda_r$. Then $v_1,\ldots,v_r$ are linearly independent.

Proof Suppose (in order to obtain a contradiction) that $v_1,\ldots,v_r$ are linearly dependent, so that there exists a linear relation

$$c_1v_1 + \cdots + c_rv_r = 0,$$

with coefficients c_i not all zero. Some of these coefficients may be zero; choose a relation with the smallest possible number of non-zero coefficients. Without loss of generality, suppose that $c_1 \neq 0$ (if $c_1 = 0$ just re-label the coefficients). Now acting on the given relation with T, using the fact that $T(v_i) = \lambda_i v_i$, we get

$$c_1\lambda_1v_1 + \cdots + c_r\lambda_rv_r = T(0) = 0.$$

Subtracting λ_1 times the first equation from the second, we get

$$c_2(\lambda_2-\lambda_1)v_2 + \cdots + c_r(\lambda_r-\lambda_1)v_r = 0.$$

Now this equation has *fewer* non-zero coefficients than the one we started with (since $\lambda_i - \lambda_1 \neq 0$ for all $2 \leq i \leq r$, by hypothesis), which was assumed to have the smallest possible number. So the coefficients in this equation must all be zero. That is, $c_i(\lambda_i - \lambda_1) = 0$, so $c_i = 0$ (since $\lambda_i \neq \lambda_1$), for $i = 2,\ldots,r$. This doesn't leave much of the original equation, only $c_1v_1 = 0$, from which we conclude that $c_1 = 0$, contrary to our assumption. So the vectors must have been linearly independent.

Theorem 5.9 The linear map T on V is diagonalisable if and only if its minimal polynomial is the product of distinct linear factors, that is, its roots all have multiplicity 1. (The analogous fact is true for matrices: a square matrix A is diagonalisable if and only if its minimal polynomial is the product of distinct linear factors, that is, its roots all have multiplicity 1).

Proof Suppose first that T is diagonalisable, with eigenvalues $\lambda_1, \dots, \lambda_r$. Then (see Proposition 5.4) there is a basis such that T is represented by a diagonal matrix D whose diagonal entries are the eigenvalues. Now for any polynomial f, $f(T)$ is represented by $f(D)$, a diagonal matrix whose diagonal entries are $f(\lambda_i)$ for $i = 1, \dots, r$. Choose

$$f(x) = (x - \lambda_1) \cdots (x - \lambda_r).$$

Then all the diagonal entries of $f(D)$ are zero; so $f(D) = 0$. We claim that f is the minimal polynomial of T; clearly it has no repeated roots, so we will be done. We know that each λ_i is a root of $m_T(x)$, so that $f(x)$ divides $m_T(x)$; and we also know that $f(T) = 0$, so that the degree of f cannot be smaller than that of m_T. So the claim follows.

Conversely, we have to show that if m_T is a product of distinct linear factors then T is diagonalisable. This is a little argument with polynomials. Let $f(x) = \prod(x - \lambda_i)$ be the minimal polynomial of T, with the roots λ_i all distinct. Let $h_i(x) = f(x)/(x - \lambda_i)$. Then the polynomials $h_1, \dots, h_r$ have no common factor except 1; for the only possible factors are $(x - \lambda_i)$, but this fails to divide h_i. Now the Euclidean algorithm shows that we can write the h.c.f. as a linear combination:

$$1 = \sum_{i=1}^{r} h_i(x) k_i(x).$$

Let $U_i = \mathrm{Im}(h_i(T))$. The vectors in U_i are eigenvectors of T with eigenvalue λ_i; for if $u \in U_i$, say $u = h_i(T)v$, then

$$(T - \lambda_i I)u_i = (T - \lambda_i I)h_i(T)(v) = f(T)v = 0,$$

so that $T(v) = \lambda_i(v)$. Moreover every vector can be written as a sum of vectors from the subspaces U_i. For, given $v \in V$, we have

$$v = Iv = \sum_{i=1}^{r} h_i(T)(k_i(T)v),$$

with $h_i(T)(k_i(T)v) \in \mathrm{Im}(h_i(T)$. The fact that the expression is unique follows from the lemma, since the eigenvectors are linearly independent.

So how, in practice, do we "diagonalise" a matrix A, that is, find an invertible matrix P such that $P^{-1}AP = D$ is diagonal? We saw an example of this earlier (see Example 5.3). The matrix equation can be rewritten as $AP = PD$, from which we see that the columns of P are the eigenvectors of A. So the procedure is: Find the eigenvalues of A, and find a basis of eigenvectors; then let P be the matrix which has the eigenvectors as columns, and D the diagonal matrix whose diagonal entries are the eigenvalues. Then $P^{-1}AP = D$. It is worth remembering that the matrix P is *not unique*: we have the freedom to interchange its columns, and also multiply any column by a non-zero scalar (because the resulting column is still an eigenvector for the same eigenvalue).

How do we find the minimal polynomial of a matrix? We know that it divides the characteristic polynomial, and that every root of the characteristic polynomial is a root of the minimal polynomial; after that it's trial and error.

For example, if the characteristic polynomial is $(x-1)^2(x-2)^3$, then the minimal polynomial must be one of $(x-1)(x-2)$ (this would correspond to the matrix being diagonalisable), $(x-1)^2(x-2)$, $(x-1)(x-2)^2$, $(x-1)^2(x-2)^2$, $(x-1)(x-2)^3$ or $(x-1)^2(x-2)^3$. If we try them in this order, the first one to be satisfied by the matrix is the minimal polynomial.

For example, the characteristic polynomial of $A = \begin{bmatrix} 1 & 2 \\ 0 & 1 \end{bmatrix}$ is $(x-1)^2$; its minimal polynomial is not $(x-1)$ (since $A \neq I$); so it is $(x-1)^2$. The fact that it has repeated roots confirms our earlier finding (see Example 5.2) that this matrix is *not* diagonalisable.