The Cayley-Hamilton theorem

Theorem 1. *Let A be a $n \times n$ matrix, and let $p(\lambda) = \det(\lambda I - A)$ be the characteristic polynomial of A. Then $p(A) = 0$.*

Proof. *Step 1:* Assume first that A is diagonalizable. In this case, we can find an invertible matrix S and a diagonal matrix

$$D = \begin{bmatrix} \lambda_1 & 0 & \dots & 0 \\ 0 & \ddots & & \vdots \\ \vdots & & \ddots & 0 \\ 0 & \dots & 0 & \lambda_n \end{bmatrix}$$

such that $A = SDS^{-1}$. The k-th power of D is given by

$$D^k = \begin{bmatrix} \lambda_1^k & 0 & \dots & 0 \\ 0 & \ddots & & \vdots \\ \vdots & & \ddots & 0 \\ 0 & \dots & 0 & \lambda_n^k \end{bmatrix}.$$

This implies

$$p(D) = \begin{bmatrix} p(\lambda_1) & 0 & \dots & 0 \\ 0 & \ddots & & \vdots \\ \vdots & & \ddots & 0 \\ 0 & \dots & 0 & p(\lambda_n) \end{bmatrix}.$$

For each $j = 1, \dots, n$, the number λ_j is an eigenvalue of A. This implies $p(\lambda_j) = 0$ for $j = 1, \dots, n$. Thus, we conclude that $p(D) = 0$.

On the other hand, the identity $A = SDS^{-1}$ implies $A^k = SD^kS^{-1}$ for all k. Therefore, we have $p(A) = S\,p(D)\,S^{-1}$. Since $p(D) = 0$, we conclude that $p(A) = 0$. This completes the proof of the Cayley-Hamilton theorem in this special case.

Step 2: To prove the Cayley-Hamilton theorem in general, we use the fact that any matrix $A \in \mathbb{C}^{n \times n}$ can be approximated by diagonalizable matrices. More precisely, given any matrix $A \in \mathbb{C}^{n \times n}$, we can find a sequence of matrices $\{A_k : k \in \mathbb{N}\}$ such that $A_k \to A$ as $k \to \infty$ and each matrix A_k has n distinct eigenvalues. Hence, the matrix A_k is diagonalizable for each

$k \in \mathbb{N}$. Therefore, it follows from our results in Step 1 that $p_k(A_k) = 0$, where $p_k(\lambda) = \det(\lambda I - A_k)$ denotes the characteristic polynomial of A_k.

Note that each entry of the matrix $p(A)$ can be written as a polynomial in the entries of A. Since $\lim_{k\to\infty} A_k = A$, we conclude that $\lim_{k\to\infty} p_k(A_k) = p(A)$. Since $p_k(A_k) = 0$ for every $k \in \mathbb{N}$, we must have $p(A) = 0$.

Decomposition into generalized eigenspaces

We'll need the following tool from algebra:

Theorem 2. *Suppose that $f(\lambda)$ and $g(\lambda)$ are two polynomials that are relatively prime. (This means that any polynomial that divides both $f(\lambda)$ and $g(\lambda)$ must be constant, i.e. of degree 0.) Then we can find polynomials $p(\lambda)$ and $q(\lambda)$ such that $p(\lambda)\, f(\lambda) + q(\lambda)\, g(\lambda) = 1$.*

This is standard result in algebra. The polynomials $p(\lambda)$ and $q(\lambda)$ can be found using the Euclidean algorithm. A proof can be found in most algebra textbooks.

This result is the key ingredient in the proof of the following theorem:

Theorem 3. *Let A be an $n \times n$ matrix, and let $f(\lambda)$ and $g(\lambda)$ be two polynomials that are relatively prime. Moreovr, let x be a vector satisfying $f(A)\, g(A)\, x = 0$. Then there exists a unique pair of vectors y, z such that $f(A)\, y = 0$, $g(A)\, z = 0$, and $y + z = x$. In other words,* $\ker(f(A)\, g(A)) = \ker f(A) \oplus \ker g(A)$.

Proof. Since the polynomials $f(\lambda)$ and $g(\lambda)$ are relatively prime, we can find polynomials $p(\lambda)$ and $q(\lambda)$ such that

$$p(\lambda)\, f(\lambda) + q(\lambda)\, g(\lambda) = 1.$$

This implies

$$p(A)\, f(A) + q(A)\, g(A) = I.$$

In order to prove the existence part, we define vectors y, z by $y = q(A)\, g(A)\, x$ and $z = p(A)\, f(A)\, x$. Then

$$f(A)\, y = f(A)\, q(A)\, g(A)\, x = q(A)\, f(A)\, g(A)\, x = 0,$$

$$g(A)\, z = g(A)\, p(A)\, f(A)\, x = p(A)\, f(A)\, g(A)\, x = 0,$$

and

$$y + z = (p(A)\, f(A) + q(A)\, g(A))\, x = x.$$

Therefore, the vectors y, z have all the required properties.

In order to prove the uniqueness part, it suffices to show that $\ker f(A) \cap \ker g(A) = \{0\}$. Assume that x lies in the intersection of $\ker f(A)$ and $\ker g(A)$, so that $f(A)\, x = 0$ and $g(A)\, x = 0$. This implies $p(A)\, f(A)\, x = 0$ and $q(A)\, g(A)\, x = 0$. Adding both equations, we obtain $x = (p(A)\, f(A) + q(A)\, g(A))\, x = 0$. This shows that show that $\ker f(A) \cap \ker g(A) = \{0\}$, as claimed.

Let A be a $n \times n$ matrix, and denote by $p(\lambda) = \det(\lambda I - A)$ the characteristic polynomial of A. By virtue of the fundamental theorem of algebra, we may write the polynomial $p(\lambda)$ in the form

$$p(\lambda) = (\lambda - \lambda_1)^{\alpha_1} \cdots (\lambda - \lambda_m)^{\alpha_m},$$

where $\lambda_1, \ldots, \lambda_m$ are the *distinct* eigenvalues of A and $\alpha_1, \ldots, \alpha_m$ denote their respective algebraic multiplicities. (Note that we do not require A to have n distinct eigenvalues! Some of the numbers $\alpha_1, \ldots, \alpha_m$ may be greater than 1.)

For abbreviation, write $p(\lambda) = g_1(\lambda) \cdots g_m(\lambda)$, where $g_j(\lambda) = (\lambda - \lambda_j)^{\alpha_j}$ for $j = 1, \ldots, m$. Repeated application of the previous theorem yields the direct sum decomposition

$$\ker p(A) = \ker g_1(A) \oplus \ldots \oplus \ker g_m(A),$$

i.e.

$$\ker p(A) = \ker(A - \lambda_1 I)^{\alpha_1} \oplus \ldots \oplus (A - \lambda_m I)^{\alpha_m}.$$

The spaces $\ker(A - \lambda_1 I)^{\alpha_1}, \ldots, (A - \lambda_m I)^{\alpha_m}$ are called the *generalized eigenspaces* of A.

At this point, we can use the Cayley-Hamilton theorem to our advantage: according to that theorem, we have $p(A) = 0$, hence $\ker p(A) = \mathbb{C}^n$. As a result, we obtain the following decomposition of $\mathbb{C}^n$ into generalized eigenspaces:

$$\mathbb{C}^n = \ker(A - \lambda_1 I)^{\alpha_1} \oplus \ldots \oplus (A - \lambda_m I)^{\alpha_m}.$$

Theorem 4. *Let $A \in \mathbb{C}^{n\times n}$ be given. Then we can find matrices $L, N \in \mathbb{C}^n$ with the following properties:*
(i) $L + N = A$
(ii) $LN = NL$
(iii) L is diagonalizable
(iv) N is nilpotent, i.e. $N^n = 0$.
Moreover, the matrices L and N are unique (i.e. there exists only one pair of matrices with that property).

Proof. Existence: Consider the decomposition of $\mathbb{C}^n$ into generalized eigenspaces:

$$\mathbb{C}^n = \ker(A - \lambda_1 I)^{\alpha_1} \oplus \ldots \oplus (A - \lambda_m I)^{\alpha_m}.$$

Consider the linear transformation from $\mathbb{C}^n$ into itself that sends a vector $x \in \ker(A-\lambda_j I)^{\alpha_j}$ to $\lambda_j x$ ($j = 1, \ldots, m$). Let L be the $n\times n$ matrix associated with this linear transformation. This implies $Lx = \lambda_j x$ for all $x \in \ker(A - \lambda_j I)^{\alpha_j}$. Clearly, $\ker(L - \lambda_j I) = \ker(A - \lambda_j I)^{\alpha_j}$ for $j = 1, \ldots, m$. Therefore, there exists a basis of $\mathbb{C}^n$ that consists of eigenvectors of L. Consequently, L is diagonalizable.

We claim that A and L commute, i.e. $LA = AL$. It suffices to show that $LAx = ALx$ for all vectors $x \in \ker(A - \lambda_j I)^{\alpha_j}$ and all $j = 1, \ldots, m$. Indeed, if x belongs to the generalized eigenspace $\ker(A - \lambda_j I)^{\alpha_j}$, then Ax lies in the same generalized eigenspace. Therefore, we have $Lx = \lambda_j x$ and $LAx = \lambda_j Ax$. Putting these facts together, we obtain $LAx = \lambda_j Ax = ALx$, as claimed. Therefore, we have $LA = AL$.

We now put $N = A - L$. Clearly, $L + N = A$ and $LN = LA - L^2 = AL - L^2 = NL$. Hence, it remains to show that $N^n = 0$. As above, it is enough to show that $N^n x = 0$ for all vectors $x \in \ker(A - \lambda_j I)^{\alpha_j}$ and all $j = 1, \ldots, m$. By definition of L and N, we have $Nx = Ax - Lx = (A - \lambda_j I)x$ for all $x \in \ker(A - \lambda_j I)^{\alpha_j}$. From this it is easy to see that $N^n x = (A - \lambda_j I)^n x$. However, $(A - \lambda_j I)^n x = 0$ since $x \in \ker(A - \lambda_j I)^{\alpha_j}$ and $\alpha_j \leq n$. Thus, we conclude that $N^n x = 0$ for all $x \in \ker(A - \lambda_j I)^{\alpha_j}$. This completes the proof of the existence part.

Uniqueness: We next turn to the proof of the uniqueness statement. Suppose that $L, N \in \mathbb{C}^{n\times n}$ satsify (i) – (iv). We claim that $Lx = \lambda_j x$ for all vectors $x \in \ker(A - \lambda_j I)^{\alpha_j}$ and all $j = 1, \ldots, m$. To this end, we use the formula $L-\lambda_j I = (A-\lambda_j I) - N$. Since N commutes with $A-\lambda_j I$, it follows that

$$(L-\lambda_j I)^{2n} = \sum_{l=0}^{2n} \binom{2n}{l} (-N)^l \, (A-\lambda_j I)^{2n-l}.$$

Using the identity $N^n = 0$, we obtain

$$(L-\lambda_j I)^{2n} = \sum_{l=0}^{n-1} \binom{2n}{l} (-N)^l \, (A-\lambda_j I)^{2n-l}.$$

Suppose that $x \in \ker(A-\lambda_j I)^{\alpha_j}$. Since $\alpha_j \leq n$, we have $(A-\lambda_j I)^{2n-l}x = 0$ for all $l = 0, \ldots, n-1$. This implies $(L-\lambda_j I)^{2n}x = 0$. Since L is diagonalizable, we it follows that $(L-\lambda_j I)x = 0$. Thus, we conclude that $Lx = \lambda_j x$ for all vectors $x \in \ker(A-\lambda_j I)^{\alpha_j}$ and all $j = 1, \ldots, m$.

Since

$$\mathbb{C}^n = \ker(A-\lambda_1 I)^{\alpha_1} \oplus \ldots \oplus (A-\lambda_m I)^{\alpha_m},$$

there is exactly one matrix L such that $Lx = \lambda_j x$ for $x \in \ker(A-\lambda_j I)^{\alpha_j}$ and $j = 1, \ldots, m$. This completes the proof of the uniqueness statement.