Lecture 18: Algebraic and geometric multuplicity

March 26, 2025

Lecture 18: Algebraic and Geometric Multiplicity

Eigenvectors and Diagonalizability

Definition An eigenvector of a square matrix A is a non-zero column vector v for which $Av = \lambda v$ for some scalar λ, called the eigenvalue of A to which v corresponds.

The eigenvalues of A are the roots of its characteristic polynomial $\det(xI_n - A)$.

The eigenspace corresponding to a particular eigenvalue λ is the set of all vectors v satsfying $Av = \lambda v$. It is a subpsace of the relevant $\mathbb{R}^n$, of dimension at least 1.

The matrix $A \in M_n(\mathbb{R})$ is diagonalizable if and only if $\mathbb{R}^n$ has a basis consisting of eigenvectors of A. In this case $P^{-1}AP$ is diagonal, where P is a matrix whose n columns are linearly independent eigenvectors of A. The diagonal entries of $P^{-1}AP$ are the corresponding eigenvalues.

Non-diagonalizabilty (two examples)

For $A \in M_n(\mathbb{R})$, it does not always happen that $\mathbb{R}^n$ has a basis consisting of eigenvectors of A.

Examples

1. The matrix $A = \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix}$ is diagonalizable in $M_2(\mathbb{C})$ but not in $M_2(\mathbb{R})$.
   This matrix represents an anti-clockwise rotation through 90° about the origin. It does not fix any line in $\mathbb{R}^2$. Its characteristic polynomial is $x^2 + 1$.
2. The matrix $B = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix}$ is not diagonalizable even over $\mathbb{C}$.
   This matrix represents a horizontal shear. Its characteristic polynomial is $(x-1)^2$ but its 1-eigenspace consists only of the X-axis. It does not have two linearly independent eigenvectors.

A shear in $\mathbb{R}^2$

Example (from Lecture 16) $B = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix}$. The linear transformation T described by B sends $(x, y) \in \mathbb{R}^2$ to $(x + y, y)$. This is a horizontal shear: it shifts every point horizontally by its y-coordinate.

For every point $v \in \mathbb{R}^2$, $T(v)$ is on the same horizontal line as v. It follows that $T(v)$ is a scalar multiple of v only if v lies on the X-axis. In this case $T(v) = v$.

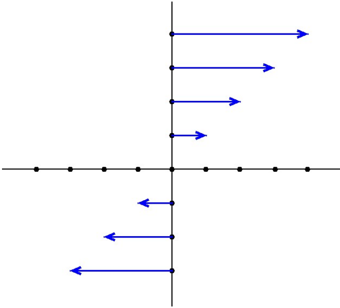

The characteristic polynomial of B (and T) is $(\lambda - 1)^2$.
The only eigenvalue is 1, and it has algebraic multiplicity 2, meaning it appears twice as a root of the characteristic polynomial.

But its geometric multiplicity is only 1, meaning its corresponding eigenspace is 1-dimensional, just the line $y = 0$.

Repeated or distinct eigenvalues

The "shear" example shows that $\mathbb{R}^2$ does not have a basis consisting of eignevectors of $B = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix}$, so B is not similar to a diagonal matrix.

Also B has 1 as a repeated eigenvalue (double root of its characteristic polynomial).

We will show that a matrix is diagonalizable [1] if its characteristic polynomial has distinct roots.

Example $A = \begin{bmatrix} -4 & 7 \\ -2 & 5 \end{bmatrix}$. $\det(\lambda I - A) = \lambda^2 - \lambda - 6 = (\lambda + 2)(\lambda - 3)$: distinct roots. Distinct eigenvalues $-2, 3$.

Respective corresponding eigenvectors: $\begin{bmatrix} 7 \\ 2 \end{bmatrix}, \begin{bmatrix} 1 \\ 1 \end{bmatrix}$.

Note these are linearly independent, so form a basis of $\mathbb{R}^2$.

Conclusion $P^{-1}AP = \begin{bmatrix} -2 & 0 \\ 0 & 3 \end{bmatrix}$, where $P = \begin{bmatrix} 7 & 1 \\ 2 & 1 \end{bmatrix}$.

[1] Small print: possibly considered as a matrix in $M_n(\mathbb{C})$ if its eigenvalues are not real

Eigenvectors for distinct eigenvalues are independent

Theorem Let $A \in M_n(\mathbb{R})$ and let $v_1, \dots, v_k$ be eigenvectors of A in $\mathbb{R}^n$, corresponding to *distinct* eigenvalues $\lambda_1, \dots, \lambda_k$ of A. Then $\{v_1, \dots, v_k\}$ is a linearly independent subset of $\mathbb{R}^n$.

Proof (for $k = 3$.) First note that no two of v_1, v_2, v_3 are scalar multiples of each other, since they correspond to different eigenvalues.

Now suppose $a_1 v_1 + a_2 v_2 + a_3 v_3 = 0$, for scalars a_1, a_2, a_3 in $\mathbb{R}$. We need to show $a_1 = a_2 = a_3 = 0$.

Multiplying $a_1 v_1 + a_2 v_2 + a_3 v_3$ on the left by A, we have

$$a_1 A v_1 + a_1 A v_2 + a_3 A v_3 = 0 \Longrightarrow a_1 \lambda_1 v_1 + a_2 \lambda_2 v_2 + a_3 \lambda_3 v_3 = 0.$$

Multiply $a_1 v_1 + a_2 v_2 + a_3 v_3$ by λ_1: $a_1 \lambda_1 v_1 + a_2 \lambda_1 v_2 + a_3 \lambda_1 v_3 = 0$.

Subtract to get $\boxed{a_2 \underbrace{(\lambda_1 - \lambda_2)}_{\neq 0} v_2 + a_3 \underbrace{(\lambda_1 - \lambda_3)}_{\neq 0} v_3 = 0.}$

Since v_2 and v_3 are linearly independent and $\lambda_1 - \lambda_2 \neq 0$, and $\lambda_1 = \lambda_3 \neq 0$, it follows that $a_2 = a_3 = 0$, and hence that $a_1 = 0$ also.

At most n distinct eigenvalues

The following consequence of the theorem shows that a matrix cannot have too many distinct eigenvalues. We already knew this, since the eigenvalues are roots of a polynomial of degree n, but here we deduce it without having to appeal to any theory about polynomial equations.

Corollary Let $A \in M_n(\mathbb{R})$. Then A has at most n distinct eigenvalues in $\mathbb{R}$.

Proof If A has k distinct eigenvalues, with corresponding eigenvectors $v_1, \dots, v_k$ in $\mathbb{R}^n$, then k cannot exceed the dimension of $\mathbb{R}^n$, since $\{v_1, \dots, v_k\}$ is a linearly independent set in $\mathbb{R}^n$. Hence $k \leq n$.

Another Corollary If $A \in M_n(\mathbb{R})$ has n distinct eigenvalues, then A is diagonalizable.

Proof A set consisting of one eigenvector for each of the n eigenvalues is linearly independent and hence is a basis.

Notes about determinants and characteristic polynomials

1. The characteristic polynomial of the square matrix $A \in M_n(\mathbb{R})$ is the determinant of $\lambda I_n - A$.
2. If t is a root of this polynomial, the t-eigenspace of A is the nullspace of the matrix $tI_n - A$.
3. The determinant of a diagonal or upper triangular matrix is the product of the entries on its main diagonal.
4. A square matrix is block diagonal if its non-zero entries are all contained in square blocks along its diagonal. The determinant of a block diagonal matrix is the product of the determinants of its diagonal blocks.
5. Similar matrices have the same characteristic polynomial and the same eigenvalues and eigenspace dimensions, since they represent the same linear transformation.

Similar Matrices have the same Characteristic Polynomial

Suppose that A and B are similar square matrices, so that

$$B = P^{-1}AP$$

for some invertible matrix P. Then

$$\begin{aligned}
\det(xI - B) &= \det(xI - P^{-1}AP) \\
&= \det(P^{-1}(xI)P - P^{-1}AP) \\
&= \det\left(P^{-1}(xI - A)P\right) \\
&= \det P^{-1} \det(xI - A) \det P \\
&= \det(xI - A).
\end{aligned}$$

If two matrices have the same characteristic polynomial, they are not necessarily similar. For example $\begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$ and $\begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix}$ have the same characteristic polynomial but are not similar.

Multiplicity of Eigenvalues

Let λ be an eigenvalue of a matrix $A \in M_n(\mathbb{R})$. The algebraic multiplicity of λ is the number of times that λ occurs as a root of the characteristic polynomial. The geometric multiplicity is the dimension of the t-eigenspace of A.

Example The matrix $A = \begin{bmatrix} 3 & 0 & 0 & 0 \\ 0 & 3 & 0 & 0 \\ 0 & 0 & 4 & 1 \\ 0 & 0 & 0 & 4 \end{bmatrix}$ has two distinct eigenvalues, 3 and 4. Both have algebraic multiplicity 2; the characteristic polynomial is $(\lambda - 3)^3(\lambda - 4)^2$.

The 3-eigenspace has dimension 2, its elements are $\begin{bmatrix} a \\ b \\ 0 \\ 0 \end{bmatrix}$, for $a, b \in \mathbb{R}$.

The 4-eigenspace only has dimension 1, its elements are $\begin{bmatrix} 0 \\ 0 \\ c \\ 0 \end{bmatrix}$, for $c \in \mathbb{R}$.

This A is not diagonalizable since it does not have four independent eigenvectors.

Geometric Multiplicity $\leq$ Algebraic Multiplicity

Theorem The geometric multplicity of an eigenvalue is at most equal to its algebraic multiplicity.

Proof: Suppose that t has geometric multiplicity k as an eigenvalue of the square matrix $A \in M_n(\mathbb{R})$, and let $\{v_1, \dots, v_k\}$ be a basis for the t-eigenspace of A. Extend this to a basis $\mathcal{B}$ of $\mathbb{R}^n$, and let P be the matrix whose columns are the elements of $\mathcal{B}$. Then the first k columns of $P^{-1}AP$ have t in the diagonal position and zeros elsewhere. It follows that $\lambda - t$ occurs at least k times as a factor of $\det(\lambda I_n - P^{-1}AP)$, so the algebraic multiplicity of t is at least k.

Corollary A matrix is diagonalizable if and only if the geometric multiplicity of each of its eigenvalues is equal to the algebraic multiplicity.