

Lecture 28: Similar matrices and Jordan form

We've nearly covered the entire heart of linear algebra – once we've finished singular value decompositions we'll have seen all the most central topics.

$A^T A$ is positive definite

A matrix is *positive definite* if $\mathbf{x}^T A \mathbf{x} > 0$ for all $\mathbf{x} \neq \mathbf{0}$. This is a very important class of matrices; positive definite matrices appear in the form of $A^T A$ when computing least squares solutions. In many situations, a rectangular matrix is multiplied by its transpose to get a square matrix.

Given a symmetric positive definite matrix A , is its inverse also symmetric and positive definite? Yes, because if the (positive) eigenvalues of A are $\lambda_1, \lambda_2, \dots, \lambda_d$ then the eigenvalues $1/\lambda_1, 1/\lambda_2, \dots, 1/\lambda_d$ of A^{-1} are also positive.

If A and B are positive definite, is $A + B$ positive definite? We don't know much about the eigenvalues of $A + B$, but we can use the property $\mathbf{x}^T A \mathbf{x} > 0$ and $\mathbf{x}^T B \mathbf{x} > 0$ to show that $\mathbf{x}^T (A + B) \mathbf{x} > 0$ for $\mathbf{x} \neq \mathbf{0}$ and so $A + B$ is also positive definite.

Now suppose A is a rectangular (m by n) matrix. A is almost certainly not symmetric, but $A^T A$ is square and symmetric. Is $A^T A$ positive definite? We'd rather not try to find the eigenvalues or the pivots of this matrix, so we ask when $\mathbf{x}^T A^T A \mathbf{x}$ is positive.

Simplifying $\mathbf{x}^T A^T A \mathbf{x}$ is just a matter of moving parentheses:

$$\mathbf{x}^T (A^T A) \mathbf{x} = (A \mathbf{x})^T (A \mathbf{x}) = |A \mathbf{x}|^2 \geq 0.$$

The only remaining question is whether $A \mathbf{x} = \mathbf{0}$. If A has rank n (independent columns), then $\mathbf{x}^T (A^T A) \mathbf{x} = \mathbf{0}$ only if $\mathbf{x} = \mathbf{0}$ and A is positive definite.

Another nice feature of positive definite matrices is that you never have to do row exchanges when row reducing – there are never 0's or unsuitably small numbers in their pivot positions.

Similar matrices A and $B = M^{-1} A M$

Two square matrices A and B are *similar* if $B = M^{-1} A M$ for some matrix M . This allows us to put matrices into families in which all the matrices in a family are similar to each other. Then each family can be represented by a diagonal (or nearly diagonal) matrix.

Distinct eigenvalues

If A has a full set of eigenvectors we can create its eigenvector matrix S and write $S^{-1} A S = \Lambda$. So A is similar to Λ (choosing M to be this S).

If $A = \begin{bmatrix} 2 & 1 \\ 1 & 2 \end{bmatrix}$ then $\Lambda = \begin{bmatrix} 3 & 0 \\ 0 & 1 \end{bmatrix}$ and so A is similar to $\begin{bmatrix} 3 & 0 \\ 0 & 1 \end{bmatrix}$. But A is also similar to:

$$\begin{aligned} M^{-1} \begin{bmatrix} 1 & -4 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} A \\ \begin{bmatrix} 3 & 0 \\ 0 & 1 \end{bmatrix} \end{bmatrix} M \begin{bmatrix} 1 & 4 \\ 0 & 1 \end{bmatrix} &= \begin{bmatrix} 1 & -4 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} B \\ \begin{bmatrix} 2 & 9 \\ 1 & 6 \end{bmatrix} \end{bmatrix} \\ &= \begin{bmatrix} -2 & -15 \\ 1 & 6 \end{bmatrix}. \end{aligned}$$

In addition, B is similar to Λ . All these similar matrices have the same eigenvalues, 3 and 1; we can check this by computing the trace and determinant of A and B .

Similar matrices have the same eigenvalues!

In fact, the matrices similar to A are all the 2 by 2 matrices with eigenvalues 3 and 1. Some other members of this family are $\begin{bmatrix} 3 & 7 \\ 0 & 1 \end{bmatrix}$ and $\begin{bmatrix} 1 & 7 \\ 0 & 3 \end{bmatrix}$. To prove that similar matrices have the same eigenvalues, suppose $A\mathbf{x} = \lambda\mathbf{x}$. We modify this equation to include $B = M^{-1}AM$:

$$\begin{aligned} AMM^{-1}\mathbf{x} &= \lambda\mathbf{x} \\ M^{-1}AMM^{-1}\mathbf{x} &= \lambda M^{-1}\mathbf{x} \\ BM^{-1}\mathbf{x} &= \lambda M^{-1}\mathbf{x}. \end{aligned}$$

The matrix B has the same λ as an eigenvalue. $M^{-1}\mathbf{x}$ is the eigenvector.

If two matrices are similar, they have the same eigenvalues and the same number of independent eigenvectors (but probably not the same eigenvectors).

When we diagonalize A , we're finding a diagonal matrix Λ that is similar to A . If two matrices have the same n distinct eigenvalues, they'll be similar to the same diagonal matrix.

Repeated eigenvalues

If two eigenvalues of A are the same, it may not be possible to diagonalize A . Suppose $\lambda_1 = \lambda_2 = 4$. One family of matrices with eigenvalues 4 and 4 contains only the matrix $\begin{bmatrix} 4 & 0 \\ 0 & 4 \end{bmatrix}$. The matrix $\begin{bmatrix} 4 & 1 \\ 0 & 4 \end{bmatrix}$ is not in this family.

There are two families of similar matrices with eigenvalues 4 and 4. The larger family includes $\begin{bmatrix} 4 & 1 \\ 0 & 4 \end{bmatrix}$. Each of the members of this family has only one eigenvector.

The matrix $\begin{bmatrix} 4 & 0 \\ 0 & 4 \end{bmatrix}$ is the only member of the other family, because:

$$M^{-1} \begin{bmatrix} 4 & 0 \\ 0 & 4 \end{bmatrix} M = 4M^{-1}M = \begin{bmatrix} 4 & 0 \\ 0 & 4 \end{bmatrix}$$

for any invertible matrix M .

Jordan form

Camille Jordan found a way to choose a “most diagonal” representative from each family of similar matrices; this representative is said to be in *Jordan normal form*. For example, both $\begin{bmatrix} 4 & 1 \\ 0 & 4 \end{bmatrix}$ and $\begin{bmatrix} 4 & 0 \\ 0 & 4 \end{bmatrix}$ are in Jordan form. This form used to be the climax of linear algebra, but not any more. Numerical applications rarely need it.

We can find more members of the family represented by $\begin{bmatrix} 4 & 1 \\ 0 & 4 \end{bmatrix}$ by choosing diagonal entries to get a trace of 4, then choosing off-diagonal entries to get a determinant of 16:

$$\begin{bmatrix} 4 & 1 \\ 0 & 4 \end{bmatrix}, \begin{bmatrix} 5 & 1 \\ -1 & 3 \end{bmatrix}, \begin{bmatrix} 4 & 0 \\ 17 & 4 \end{bmatrix}, \begin{bmatrix} a & b \\ (8a - a^2 - 16)/b & 8 - a \end{bmatrix}.$$

(None of these are diagonalizable, because if they were they would be similar to $\begin{bmatrix} 4 & 0 \\ 0 & 4 \end{bmatrix}$. That matrix is only similar to itself.) What about this one?

$$A = \begin{bmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}$$

Its eigenvalues are four zeros. Its rank is 2 so the dimension of its nullspace is $4 - 2 = 2$. It will have two independent eigenvectors and two “missing” eigenvectors. When we look instead at

$$\begin{bmatrix} 0 & 1 & 7 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix},$$

its rank and the dimension of its nullspace are still 2, but it’s not as nice as A . B is similar to A , which is the Jordan normal form representative of this family. A has a 1 above the diagonal for every missing eigenvector and the rest of its entries are 0.

Now consider:

$$C = \begin{bmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 \end{bmatrix}.$$

Again it has rank 2 and its nullspace has dimension 2. Its four eigenvalues are 0. Surprisingly, it is not similar to A . We can see this by breaking the matrices

into their *Jordan blocks*:

$$A = \left[\begin{array}{ccc|c} 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 \\ \hline 0 & 0 & 0 & 0 \end{array} \right], \quad C = \left[\begin{array}{ccc|c} 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ \hline 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 \end{array} \right].$$

A Jordan block J_i has a repeated eigenvalue λ_i on the diagonal, zeros below the diagonal and in the upper right hand corner, and ones above the diagonal:

$$J_i = \begin{bmatrix} \lambda_i & 1 & 0 & \cdots & 0 \\ 0 & \lambda_i & 1 & & 0 \\ \vdots & & \ddots & & \vdots \\ 0 & 0 & & \lambda_i & 1 \\ 0 & 0 & \cdots & 0 & \lambda_i \end{bmatrix}.$$

Two matrices may have the same eigenvalues and the same number of eigenvectors, but if their Jordan blocks are different sizes those matrices can not be similar.

Jordan's theorem says that every square matrix A is similar to a Jordan matrix J , with Jordan blocks on the diagonal:

$$J = \begin{bmatrix} J_1 & 0 & \cdots & 0 \\ 0 & J_2 & \cdots & 0 \\ \vdots & & \ddots & \vdots \\ 0 & 0 & \cdots & J_d \end{bmatrix}.$$

In a Jordan matrix, the eigenvalues are on the diagonal and there may be ones above the diagonal; the rest of the entries are zero. The number of blocks is the number of eigenvectors – there is one eigenvector per block.

To summarize:

- If A has n distinct eigenvalues, it is diagonalizable and its Jordan matrix is the diagonal matrix $J = \Lambda$.
- If A has repeated eigenvalues and “missing” eigenvectors, then its Jordan matrix will have $n - d$ ones above the diagonal.

We have not learned to compute the Jordan matrix of a matrix which is missing eigenvectors, but we do know how to diagonalize a matrix which has n distinct eigenvalues.