

Eigenvectors and Diagonalizing Matrices

E. L. Lady

Let A be an $n \times n$ matrix and suppose there exists a basis $\mathbf{v}_1, \dots, \mathbf{v}_n$ for \mathbb{R}^n such that for each i , $A\mathbf{v}_i = \lambda_i\mathbf{v}_i$ for some scalar λ . (I. e. \mathbf{v}_i is an **eigenvector** for A corresponding to the **eigenvalue** λ_i .)

Let P be the $n \times n$ matrix whose columns are the basis vectors $\mathbf{v}_1, \dots, \mathbf{v}_n$, i. e. $P = [\mathbf{v}_1 \ \mathbf{v}_2 \ \dots \ \mathbf{v}_n]$. The fact that the columns of P are a basis for \mathbb{R}^n implies that the column space of P is all of \mathbb{R}^n , so $\text{rank } P = n$ and therefore P is non-singular.

Now notice that $P^{-1}\mathbf{v}_i$ is the i^{th} column of $P^{-1}P$, i. e. $P^{-1}\mathbf{v}_i = \mathbf{e}_i$, where \mathbf{e}_i is the i^{th} column of the identity matrix I_n .

Theorem. $P^{-1}AP$ is the diagonal matrix with $\lambda_1, \dots, \lambda_n$ on the diagonal.

PROOF:

$$\begin{aligned} P^{-1}AP &= P^{-1}A [\mathbf{v}_1 \ \dots \ \mathbf{v}_n] \\ &= P^{-1} [A\mathbf{v}_1 \ \dots \ A\mathbf{v}_n] \\ &= P^{-1} [\lambda_1\mathbf{v}_1 \ \dots \ \lambda_n\mathbf{v}_n] \\ &= [\lambda_1P^{-1}\mathbf{v}_1 \ \dots \ \lambda_nP^{-1}\mathbf{v}_n] \\ &= [\lambda_1\mathbf{e}_1 \ \dots \ \lambda_n\mathbf{e}_n] \\ &= \begin{bmatrix} \lambda_1 & 0 & \dots & 0 \\ 0 & \lambda_2 & \dots & 0 \\ \vdots & \vdots & & \vdots \\ 0 & 0 & \dots & \lambda_n \end{bmatrix}. \end{aligned}$$

Example. $A = \begin{bmatrix} 19 & 20 & -16 \\ 20 & 13 & 4 \\ -16 & 4 & 31 \end{bmatrix}$, and the \mathbf{v}_i are $\begin{bmatrix} 1 \\ 2 \\ 2 \end{bmatrix}$, $\begin{bmatrix} -2 \\ -1 \\ 2 \end{bmatrix}$, $\begin{bmatrix} 2 \\ -2 \\ 1 \end{bmatrix}$. Then a

calculation shows that $A\mathbf{v}_i = \lambda_i\mathbf{v}_i$, where $\lambda_1 = 27$, $\lambda_2 = 45$, and $\lambda_3 = -9$. Now let

$P = \begin{bmatrix} 1 & -2 & 2 \\ 2 & -1 & -2 \\ 2 & 2 & 1 \end{bmatrix}$. Then $P^{-1} \begin{bmatrix} 1 & -2 & 2 \\ 2 & -1 & -2 \\ 2 & 2 & 1 \end{bmatrix} = P^{-1}P = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$. Thus

$$P^{-1}AP = P^{-1}A \begin{bmatrix} 1 & -2 & 2 \\ 2 & -1 & -2 \\ 2 & 2 & 1 \end{bmatrix} = P^{-1} \begin{bmatrix} \lambda_1 & -2\lambda_2 & 2\lambda_3 \\ 2\lambda_1 & -\lambda_2 & -2\lambda_3 \\ 2\lambda_1 & 2\lambda_2 & \lambda_3 \end{bmatrix} = \begin{bmatrix} \lambda_1 & 0 & 0 \\ 0 & \lambda_2 & 0 \\ 0 & 0 & \lambda_3 \end{bmatrix}.$$

(For the record, $P^{-1} = \frac{1}{9} \begin{bmatrix} 1 & 2 & 2 \\ -2 & -1 & 2 \\ 2 & -2 & 1 \end{bmatrix}$. If you need to use this to follow the

calculation, then you've missed the point.)

Another interesting thing about the eigenvectors given above is that they are mutually orthogonal (perpendicular) to each other, as you can easily verify by computing the dot products. This is not something that is universally true for eigenvectors, but it is also not an accident in this case.

The matrix

$$B = \begin{bmatrix} 23 & 6 & 25 \\ 5 & 4 & 5 \\ -20 & -6 & -22 \end{bmatrix}$$

has three distinct eigenvalues, namely 3, 4, and -2 . A set of eigenvectors for these is

$$\begin{bmatrix} 1 \\ 5 \\ -2 \end{bmatrix}, \quad \begin{bmatrix} 1 \\ 1 \\ -1 \end{bmatrix}, \quad \begin{bmatrix} -1 \\ 0 \\ 1 \end{bmatrix}.$$

Taking dot products shows that these three vectors are **not** orthogonal to each other. However, they are linearly independent.

The rule is as follows:

Theorem. (1) If $\mathbf{v}_1, \dots, \mathbf{v}_r$ are eigenvectors for a matrix A and if the corresponding eigenvalues are all different, then $\mathbf{v}_1, \dots, \mathbf{v}_r$ must be linearly independent.

(2) If the $n \times n$ matrix A is **symmetric** then eigenvectors corresponding to **different** eigenvalues must be orthogonal to each other. Furthermore, in this case there will exist n linearly independent eigenvectors for A , so that A will be diagonalizable.

(3) In the case of a symmetric matrix, the n different eigenvectors will not necessarily all correspond to different eigenvalues, so they may not automatically be orthogonal to each other. However (if the entries in A are all real numbers, as is always the case in this course), it's always possible to find some set of n eigenvectors which are mutually orthogonal.

The reason why eigenvectors corresponding to distinct eigenvalues of a **symmetric** matrix must be orthogonal is actually quite simple. In fact, it is a special case of the following fact:

Proposition. Let A be any $n \times n$ matrix. If \mathbf{v} is an eigenvector for A^T and if \mathbf{w} is an eigenvector for A , and if the corresponding eigenvalues are different, then \mathbf{v} and \mathbf{w} must be orthogonal.

Of course in the case of a symmetric matrix, $A^T = A$, so this says that eigenvectors for A corresponding to different eigenvalues must be orthogonal.

This proposition is the result of a Lemma which is an easy exercise in summation notation. (There's also a very fast slick proof.)

Lemma. If A is an $n \times n$ matrix and \mathbf{v} and \mathbf{w} are vectors in \mathbb{R}^n , then the dot product of \mathbf{v} and $A\mathbf{w}$ is the same as the dot product of $A^T\mathbf{v}$ and \mathbf{w} .

Now, taking this Lemma for granted, suppose that \mathbf{v} is an eigenvector for A^T and \mathbf{w} is an eigenvector for A , and suppose that the eigenvalues are different. To be specific, let's suppose that the two eigenvalues are 2 and 5, respectively. Then from the Lemma we get

$$2\mathbf{v} \cdot \mathbf{w} = (A^T \mathbf{v}) \cdot \mathbf{w} = \mathbf{v} \cdot (A\mathbf{w}) = \mathbf{v} \cdot 5\mathbf{w} = 5\mathbf{v} \cdot \mathbf{w}.$$

But since 2 and 5 are scalars (and, for that matter, so is $\mathbf{v} \cdot \mathbf{w}$), the only way that $2\mathbf{v} \cdot \mathbf{w} = 5\mathbf{v} \cdot \mathbf{w}$ can be possible is for $\mathbf{v} \cdot \mathbf{w}$ to be 0.

Returning to the original example (at the top of p. 2), the fact that the eigenvectors of A are orthogonal to each other means that the columns of the matrix P are orthogonal to each other. And it's very easy to see that a consequence of this is that the product $P^T P$ is a diagonal matrix. In fact,

$$P^T P = \begin{bmatrix} 1 & 2 & 2 \\ -2 & -1 & 2 \\ 2 & -2 & 1 \end{bmatrix} \begin{bmatrix} 1 & -2 & 2 \\ 2 & -1 & -2 \\ 2 & 2 & 1 \end{bmatrix} = \begin{bmatrix} 9 & 0 & 0 \\ 0 & 9 & 0 \\ 0 & 0 & 9 \end{bmatrix}.$$

(The entries on the diagonal are the squares of the magnitudes of the columns — in this case, all 9.) This is more than just a curiosity, for one can also write $\frac{1}{9}P^T P = I_3$, showing that $\frac{1}{9}P^T = P^{-1}$. In general,

If the columns of a matrix are orthogonal to each other and non-zero, then the inverse can be obtained from the transpose by dividing each row of the transpose by the square of its magnitude.

Since it's usually a lot easier to take the transpose of a matrix than to compute the inverse, this is certainly a useful fact. But one can do even better.

Remember that when it comes to eigenvectors, the magnitude is not really that relevant. What counts is the direction of the vector. So it's permissible to "unitize" a set of eigenvectors and get new eigenvectors all having magnitude 1.

For instance, in the original example above, all the eigenvectors originally given have magnitude 3 (as one can easily check). One can get a new set of eigenvectors

$$\mathbf{v}'_1 = \begin{bmatrix} 1/3 \\ 2/3 \\ 2/3 \end{bmatrix}, \quad \mathbf{v}'_2 = \begin{bmatrix} -2/3 \\ -1/3 \\ 2/3 \end{bmatrix}, \quad \mathbf{v}'_3 = \begin{bmatrix} 2/3 \\ -2/3 \\ 1/3 \end{bmatrix}$$

all with magnitude 1. If one makes these vectors into the columns of a matrix, one gets

$$P_1 = \begin{bmatrix} 1/3 & -2/3 & 2/3 \\ 2/3 & -1/3 & -2/3 \\ 2/3 & 2/3 & 1/3 \end{bmatrix}$$

with the property that $P_1^T P_1 = I$, which shows that $P_1^T = P_1^{-1}$.

Definition. A $n \times n$ matrix whose inverse is the same as its transpose is called an **orthogonal matrix**.

In order to be orthogonal, it is necessary that the columns of a matrix be orthogonal to each other. But it is also necessary that all the columns have magnitude 1.

Exercise (Easy!). If P is an orthogonal matrix, then the rows of P are also orthogonal to each other and all have magnitude 1.

Putting all these pieces together (including some parts that were not actually proved), we get the following:

Theorem. If A is a **symmetric** $n \times n$ matrix whose entries are all real numbers, then there exists an orthogonal matrix P such that $P^T A P$ is a diagonal matrix.