Lecture 4: Review of Linear Algebra

Mark Hasegawa-Johnson

ECE 417: Multimedia Signal Processing, Fall 2021

Outline

A linear transform $\vec{y} = A\vec{x}$ maps vector space $\vec{x}$ onto vector space $\vec{y}$. For example: the matrix $A = \begin{bmatrix} 1 & 1 \\ 0 & 2 \end{bmatrix}$ maps the vectors $\vec{x}_0, \vec{x}_1, \vec{x}_2, \vec{x}_3 =$

$$\begin{bmatrix} 1 \\ 0 \end{bmatrix}, \begin{bmatrix} \frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \end{bmatrix}, \begin{bmatrix} -\frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} \end{bmatrix}$$

to the vectors $\vec{y}_0, \vec{y}_1, \vec{y}_2, \vec{y}_3 =$

$$\begin{bmatrix} 1 \\ 0 \end{bmatrix}, \begin{bmatrix} \sqrt{2} \\ \sqrt{2} \end{bmatrix}, \begin{bmatrix} 1 \\ 2 \end{bmatrix}, \begin{bmatrix} 0 \\ \sqrt{2} \end{bmatrix}$$

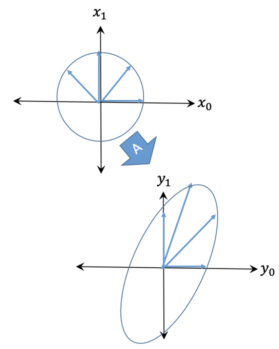

A linear transform $\vec{y} = A\vec{x}$ maps vector space $\vec{x}$ onto vector space $\vec{y}$. The absolute value of the determinant of A tells you how much the area of a unit circle is changed under the transformation.

For example, if $A = \begin{bmatrix} 1 & 1 \\ 0 & 2 \end{bmatrix}$, then the unit circle in $\vec{x}$ (which has an area of π) is mapped to an ellipse with an area that is $\text{abs}(|A|) = 2$ times larger, i.e., i.e., $\pi\text{abs}(|A|) = 2\pi$.

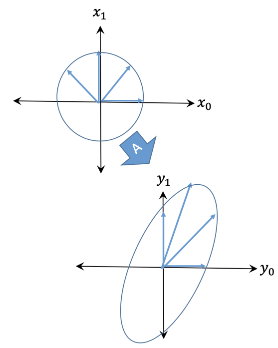

For a D-dimensional square matrix, there may be up to D different directions $\vec{x} = \vec{v}_d$ such that, for some scalar λ_d, $A\vec{v}_d = \lambda_d \vec{v}_d$. For example, if $A = \begin{bmatrix} 1 & 1 \\ 0 & 2 \end{bmatrix}$, then the eigenvectors are

$$\vec{v}_0 = \begin{bmatrix} 1 \\ 0 \end{bmatrix}, \quad \vec{v}_1 = \begin{bmatrix} \frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} \end{bmatrix},$$

and the eigenvalues are $\lambda_0 = 1$, $\lambda_1 = 2$. Those vectors are red and extra-thick, in the figure to the left. Notice that one of the vectors gets scaled by $\lambda_0 = 1$, but the other gets scaled by $\lambda_1 = 2$.

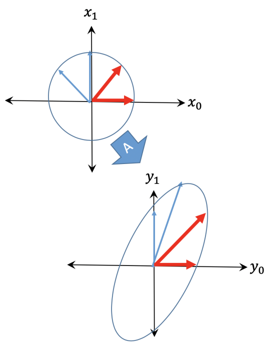

An eigenvector is a direction, not just a vector. That means that if you multiply an eigenvector by any scalar, you get the same eigenvector: if $A\vec{v}_d = \lambda_d \vec{v}_d$, then it's also true that $cA\vec{v}_d = c\lambda_d \vec{v}_d$ for any scalar c. For example: the following are the same eigenvector as $\vec{v}_1$

$$\sqrt{2}\vec{v}_1 = \begin{bmatrix} 1 \\ 1 \end{bmatrix}, \quad -\vec{v}_1 = \begin{bmatrix} -\frac{1}{\sqrt{2}} \\ -\frac{1}{\sqrt{2}} \end{bmatrix}$$

Since scale and sign don't matter, by convention, we normalize so that an eigenvector is always unit-length ($\|\vec{v}_d\| = 1$) and the first nonzero element is non-negative ($v_{d0} > 0$).

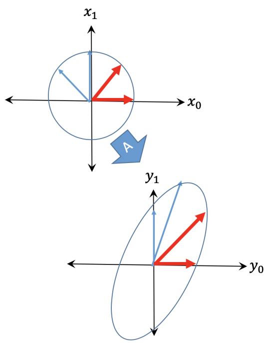

Eigenvalues: Before you find the eigenvectors, you should first find the eigenvalues. You can do that using this fact:

$$A\vec{v}_d = \lambda_d \vec{v}_d$$
$$A\vec{v}_d = \lambda_d I \vec{v}_d$$
$$A\vec{v}_d - \lambda_d I \vec{v}_d = \vec{0}$$
$$(A - \lambda_d I)\vec{v}_d = \vec{0}$$

That means that when you use the linear transform $(A - \lambda_d I)$ to transform the unit circle, the result has an area of $|A - \lambda I| = 0$.

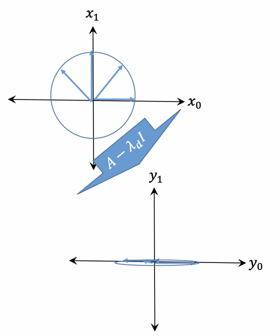

Example:

$$|A - \lambda I| = \begin{vmatrix} 1-\lambda & 1 \\ 0 & 2-\lambda \end{vmatrix}$$
$$= 2 - 3\lambda + \lambda^2$$

which has roots at $\lambda_0 = 1$, $\lambda_1 = 2$

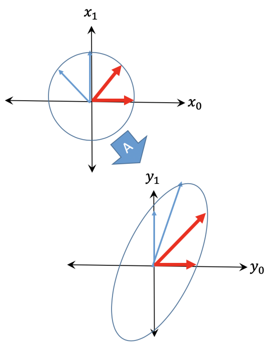

There are always D eigenvalues

- The determinant $|A - \lambda I|$ is a D^{th}-order polynomial in λ.
- By the fundamental theorem of algebra, the equation

$$|A - \lambda I| = 0$$

  has exactly D roots (counting repeated roots and complex roots).
- Therefore, **any square matrix has exactly D eigenvalues** (counting repeated eigenvalues, and complex eigenvalues.

There are not always D eigenvectors

Not every square matrix has D eigenvectors. Some of the most common exceptions are:

- **Repeated eigenvalues:** if two of the roots of the polynomial are the same ($\lambda_j = \lambda_i$), then that means there is a two-dimensional subspace, $\vec{v}$, such that $A\vec{v} = \lambda_i\vec{v}$. You can arbitrarily choose any two orthogonal vectors from this subspace to be the eigenvectors.
- **Complex eigenvalues** correspond to complex eigenvalues. For example, the matrix

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

has the eigenvalues $\lambda = \pm j$, and the corresponding eigenvectors

$$\vec{v}_1 = \frac{1}{\sqrt{2}}\begin{bmatrix} 1 \\ j \end{bmatrix}, \quad \vec{v}_2 = \frac{1}{\sqrt{2}}\begin{bmatrix} 1 \\ -j \end{bmatrix}$$

Complex eigenvalues & vectors require a little bit of extra

Outline

Review: Eigenvalues and eigenvectors

The eigenvectors of a $D \times D$ square matrix, A, are the vectors $\vec{v}$ such that

$$A\vec{v} = \lambda\vec{v} \tag{1}$$

The scalar, λ, is called the eigenvalue. It's only possible for Eq. (1) to have a solution if

$$|A - \lambda I| = 0 \tag{2}$$

Left and right eigenvectors

We've been working with right eigenvectors and right eigenvalues:

$$A\vec{v}_d = \lambda_d \vec{v}_d$$

There may also be left eigenvectors, which are row vectors $\vec{u}_d$ and corresponding left eigenvalues κ_d:

$$\vec{u}_d^T A = \kappa_d \vec{u}_d^T$$

Eigenvectors on both sides of the matrix

You can do an interesting thing if you multiply the matrix by its eigenvectors both before and after:

$$\vec{u}_i^T(A\vec{v}_j) = \vec{u}_i^T(\lambda_j\vec{v}_j) = \lambda_j\vec{u}_i^T\vec{v}_j$$

...but...

$$(\vec{u}_i^T A)\vec{v}_j = (\kappa_i\vec{u}_i^T)\vec{v}_j = \kappa_i\vec{u}_i^T\vec{v}_j$$

There are only two ways that both of these things can be true. Either

$$\kappa_i = \lambda_j \quad \text{or} \quad \vec{u}_i^T\vec{v}_j = 0$$

Left and right eigenvectors must be paired!!

There are only two ways that both of these things can be true. Either

$$\kappa_i = \lambda_j \quad \text{or} \quad \vec{u}_i^T \vec{v}_j = 0$$

Remember that eigenvalues solve $|A - \lambda_d I| = 0$. In almost all cases, the solutions are all distinct (A has distinct eigenvalues), i.e., $\lambda_i \neq \lambda_j$ for $i \neq j$. That means there is **at most one** λ_i that can equal each κ_i:

$$\begin{cases} i \neq j & \vec{u}_i^T \vec{v}_j = 0 \\ i = j & \kappa_i = \lambda_i \end{cases}$$

Outline

Symmetric matrices: left=right

If A is symmetric ($A = A^T$), then the left and right eigenvectors and eigenvalues are the same, because

$$\lambda_i \vec{u}_i^T = \vec{u}_i^T A = (A^T \vec{u}_i)^T = (A\vec{u}_i)^T$$

. . . and that last term is equal to $\lambda_i \vec{u}_i^T$ if and only if $\vec{u}_i = \vec{v}_i$.

Symmetric matrices: eigenvectors are orthonormal

Let's combine the following facts:

- $\vec{u}_i^T \vec{v}_j = 0$ for $i \neq j$ — any square matrix with distinct eigenvalues
- $\vec{u}_i = \vec{v}_i$ — symmetric matrix
- $\vec{v}_i^T \vec{v}_i = 1$ — standard normalization of eigenvectors for any matrix (this is what $\|\vec{v}_i\| = 1$ means).

Putting it all together, we get that

$$\vec{v}_i^T \vec{v}_j = \begin{cases} 1 & i = j \\ 0 & i \neq j \end{cases}$$

The eigenvector matrix

So if A is symmetric with distinct eigenvalues, then its eigenvectors are orthonormal:

$$\vec{v}_i^T \vec{v}_j = \begin{cases} 1 & i = j \\ 0 & i \neq j \end{cases}$$

We can write this as

$$V^T V = I$$

where

$$V = [\vec{v}_0, \ldots, \vec{v}_{D-1}]$$

The eigenvector matrix is orthonormal

$$V^T V = I$$

...and it also turns out that

$$VV^T = I$$

Proof: $VV^T = VIV^T = V(V^T V)V^T = (VV^T)^2$, but the only matrix that satisfies $VV^T = (VV^T)^2$ is $VV^T = I$.

Eigenvectors orthogonalize a symmetric matrix

So now, suppose A is symmetric:

$$\vec{v}_i^T A \vec{v}_j = \vec{v}_i^T (\lambda_j \vec{v}_j) = \lambda_j \vec{v}_i^T \vec{v}_j = \begin{cases} \lambda_j, & i = j \\ 0, & i \neq j \end{cases}$$

In other words, if a symmetric matrix has D eigenvectors with distinct eigenvalues, then its eigenvectors orthogonalize A:

$$V^T A V = \Lambda$$

$$\Lambda = \begin{bmatrix} \lambda_0 & 0 & 0 \\ 0 & \dots & 0 \\ 0 & 0 & \lambda_{D-1} \end{bmatrix}$$

A symmetric matrix is the weighted sum of its eigenvectors:

One more thing. Notice that

$$A = VV^TAVV^T = V\Lambda V^T$$

The last term is

$$[\vec{v}_0, \ldots, \vec{v}_{D-1}] \begin{bmatrix} \lambda_0 & 0 & 0 \\ 0 & \ldots & 0 \\ 0 & 0 & \lambda_{D-1} \end{bmatrix} \begin{bmatrix} \vec{v}_0^T \\ \vdots \\ \vec{v}_{D-1}^T \end{bmatrix} = \sum_{d=0}^{D-1} \lambda_d \vec{v}_d \vec{v}_d^T$$

Summary: properties of symmetric matrices

If A is symmetric with D eigenvectors, and D distinct eigenvalues, then

$$A = V\Lambda V^T$$

$$\Lambda = V^T AV$$

$$VV^T = V^T V = I$$

Outline

In-Lecture Written Example Problem

Pick an arbitrary 2×2 symmetric matrix. Find its eigenvalues and eigenvectors. Show that $\Lambda = V^T A V$ and $A = V \Lambda V^T$.

In-Lecture Jupyter Example Problem

Create a jupyter notebook. Pick an arbitrary 2×2 matrix. Plot a unit circle in the $\vec{x}$ space, and show what happens to those vectors after transformation to the $\vec{y}$ space. Calculate the determinant of the matrix, and its eigenvalues and eigenvectors. Show that $A\vec{v} = \lambda\vec{v}$.

Outline

Summary

- A linear transform, A, maps vectors in space $\vec{x}$ to vectors in space $\vec{y}$.
- The determinant, $|A|$, tells you how the volume of the unit sphere is scaled by the linear transform.
- Every $D \times D$ linear transform has D eigenvalues, which are the roots of the equation $|A - \lambda I| = 0$.
- Left and right eigenvectors of a matrix are either orthogonal ($\vec{u}_i^T \vec{v}_j = 0$) or share the same eigenvalue ($\kappa_i = \lambda_j$).
- For a symmetric matrix, the left and right eigenvectors are the same. If the eigenvalues are distinct and real, then:

$$A = V \Lambda V^T, \quad \Lambda = V^T A V, \quad V V^T = V^T V = I$$