Lecture 1: Review of Linear Algebra

Mark Hasegawa-Johnson

ECE 417: Multimedia Signal Processing, Fall 2023

Outline

Welcome to ECE 417, Multimedia Signal Processing!

- This course is about video and audio signals.
- At this point, let's talk about the web page: `https://courses.grainger.illinois.edu/ece417/fa2023/`

CampusWire and GradeScope

- If you're not yet added to the CampusWire or GradeScope pages, please add yourself.
- The CampusWire link is `https://campuswire.com/p/G4B80E16A`, with code 8237.
- The GradeScope link is `https://www.gradescope.com/courses/560497`, with code K3EX68.

Outline

Reading: `https://math.mit.edu/~gs/linearalgebra/ila6/ila6_6_1.pdf`

A linear transform $y = Ax$ maps vector space x onto vector space y. For example: the matrix $A = \begin{bmatrix} 1 & 1 \\ 0 & 2 \end{bmatrix}$ maps the vectors $x_0, x_1, x_2, x_3 =$

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

to the vectors $y_0, y_1, y_2, y_3 =$

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

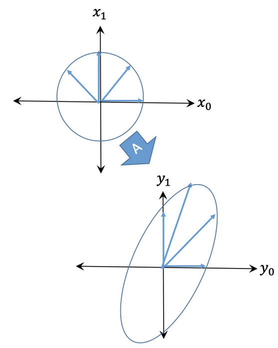

A linear transform $y = Ax$ maps vector space x onto vector space y. The absolute value of the determinant of A tells you how much the area of a unit circle is changed under the transformation.

For example, if $A = \begin{bmatrix} 1 & 1 \\ 0 & 2 \end{bmatrix}$, then the unit circle in 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., $\pi\text{abs}(|A|) = 2\pi$.

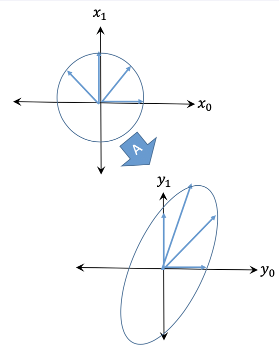

For a d-dimensional square matrix, there may be up to d different directions $x = v_i$ such that, for some scalar λ_i, $Av_i = \lambda_i v_i$. For example, if $A = \begin{bmatrix} 1 & 1 \\ 0 & 2 \end{bmatrix}$, then the eigenvectors are

$$v_0 = \begin{bmatrix} 1 \\ 0 \end{bmatrix}, \quad v_1 = \begin{bmatrix} \frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} \end{bmatrix},$$

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 $\lambda_1 = 2$.

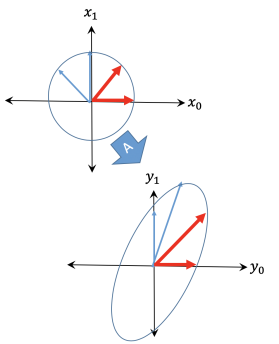

An eigenvector is a direction, not just a vector. That means that if you multiply an eigenvector by any scalar, you get the same eigenvector: if $A v_i = \lambda_i v_i$, then it's also true that $cAv_i = c\lambda_i v_i$ for any scalar c. For example: the following are the same eigenvector as v_1

$$\sqrt{2}v_1 = \begin{bmatrix} 1 \\ 1 \end{bmatrix}, \quad -v_1 = \begin{bmatrix} -\frac{1}{\sqrt{2}} \\ -\frac{1}{\sqrt{2}} \end{bmatrix}$$

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

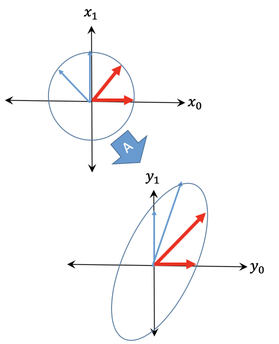

Eigenvalues: Before you find the eigenvectors, you should first find the eigenvalues. You can do that using this fact:

$$
\begin{aligned}
Av_i &= \lambda_i v_i \\
Av_i &= \lambda_i I v_i \\
Av_i - \lambda_i I v_i &= 0 \\
(A - \lambda_i I) v_i &= 0
\end{aligned}
$$

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

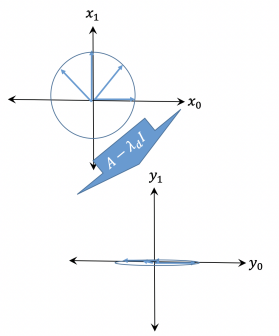

Example:

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

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

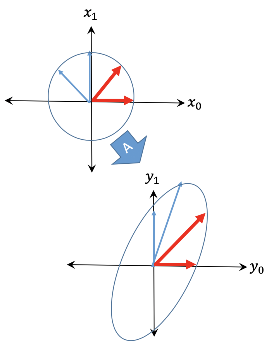

There are always d eigenvalues

- 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 unique real eigenvectors

Not every square matrix has d uniquely-defined, real-valued eigenvectors. Some of the most common exceptions are **repeated eigenvalues** and **complex eigenvalues**.

- **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, v, such that $Av = \lambda_i v$. SOLUTION: You can arbitrarily choose any two orthogonal vectors from this subspace to be the eigenvectors. These are not uniquely defined, but you can choose a set which is convenient.

There are not always d unique real eigenvectors

- **Complex eigenvalues:** A real-valued matrix can have complex eigenvalues only if the corresponding eigenvectors are also complex. Usually this means that there is some sort of periodic sinusoidal transformation of any real-valued vector. For example, consider this matrix:

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

Any real-valued vector $x = [x_1, x_2]^T$ has its elements swapped, i.e., $Ax = [x_2, -x_1]^T$. However, this matrix has complex eigenvalues $\lambda = \pm j$, and corresponding complex eigenvectors such that $Av_i = \lambda_i v_i$:

$$v_1 = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ j \end{bmatrix}, \quad v_2 = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ -j \end{bmatrix}$$

Outline

Left and right eigenvectors

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

$$Av_i = \lambda_i v_i$$

There may also be left eigenvectors, which are row vectors u_i and corresponding left eigenvalues κ_i:

$$u_i^T A = \kappa_i u_i^T$$

It turns out that (1) the eigenvalues are the same, $\kappa_i = \lambda_i$, (2) the eigenvectors might not be the same, but (3) unpaired eigenvectors are orthogonal.

Proof: Right & Left Eigenvalues are the same

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

$$u_i^T(Av_j) = u_i^T(\lambda_j v_j) = \lambda_j u_i^T v_j$$

...but...

$$(u_i^T A)v_j = (\kappa_i u_i^T)v_j = \kappa_i u_i^T v_j$$

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

$$\kappa_i = \lambda_j \quad \text{or} \quad u_i^T v_j = 0$$

Summary: Left and right eigenvalues must be paired!!

Summary: for an arbitrary square matrix A,

- Left and right eigenvalues are the same, $\lambda_i = \kappa_i \forall i$.
- Eigenvectors might NOT be the same
- Left and right eigenvectors of unpaired eigenvalues are orthogonal, $\lambda_i \neq \lambda_j \Rightarrow u_i^T v_j = 0$.

Outline

Symmetric matrices: left=right

Suppose that $A \in \Re^{m\times n}$ is any arbitrary matrix, not even square ($m \neq n$). The product A^TA is both square and symmetric. For example:

$$A^TA = \left[\begin{array}{ccc} a_{1,1} & a_{1,2} & a_{1,3} \\ a_{2,1} & a_{2,2} & a_{2,3} \end{array}\right] \left[\begin{array}{cc} a_{1,1} & a_{2,1} \\ a_{1,2} & a_{2,2} \\ a_{1,3} & a_{2,3} \end{array}\right]$$

$$= \left[\begin{array}{cc} \sum_j a_{1,j}^2 & \sum_j a_{1,j}a_{2,j} \\ \sum_j a_{1,j}a_{2,j} & \sum_j a_{2,j}^2 \end{array}\right] = \left[\begin{array}{cc} \|a_1\|^2 & \sum_j a_1^Ta_2 \\ a_1^Ta_2 & \|a_2\|^2 \end{array}\right]$$

where the last row uses a_i to mean the i^{th} column of A. The matrix of A^TA is thus the matrix of inner-products of the columns of A; this is called the **gram matrix**, so we'll use the notation $G = A^TA$.

Positive semi-definite matrices

A gram matrix is also positive semi-definite (notation: $G \succeq 0$), meaning that

- Its determinant is non-negative, $|G| \geq 0$, and
- all of its eigenvalues are non-negative, $\lambda_i \geq 0$.

Intuitive explanation (not quite a proof): The elements on the main diagonal of S are larger than the other elements in the sense that

$$a_i^T a_j = \|a_i\| \cdot \|a_j\| \cos\left(\angle(a_i, a_j)\right) \leq \|a_i\| \cdot \|a_j\|$$

Symmetric matrices: left=right

Suppose $G = A^T A$ is any symmetric square matrix: then its left and right eigenvectors and eigenvalues are the same.

- The right eigenvectors are $\lambda_i v_i = G v_i$
- The left eigenvectors are $\lambda_i u_i^T = u_i^T G$
- ...but transposing $G v_i$ gives:

$$(G v_i)^T = v_i^T G^T = v_i^T G$$

...so it must be the case that $v_i = u_i$.

Positive semidefinite (PSD) matrices: real generalized eigenvectors

Suppose $G = A^T A \succeq 0$. Then every eigenvalue has an associated generalized eigenvector:

- If λ_i is unique, then there is an associated real eigenvector, $\lambda_i v_i = G v_i$.
- If $\lambda_i = \lambda_{i+1} = \cdots \lambda_{i+k=1}$, then there is a k-dimensional subspace whose vectors v all satisfy $\lambda_i v = Gv$. We can choose an arbitrary orthonormal basis of that subspace, and call those the "generalized eigenvectors" $v_i, \cdots, v_{i+k-1}$ of $\lambda_i, \cdots, \lambda_{i+k-1}$.
  - Most common example: if $A \in \Re^{mxn}$, $n > m$, then at least $n - m$ of the eigenvalues of G are zero.

Symmetric matrices: eigenvectors are orthonormal

Let's combine the following facts:

- $u_i^T v_j = 0$ for $i \neq j$ — any square matrix with distinct eigenvalues
- $u_i = v_i$ — symmetric matrix
- $v_i^T v_i = 1$ — standard normalization of eigenvectors for any matrix (this is what $\|v_i\| = 1$ means).

Putting it all together, we get that

$$v_i^T v_j = \begin{cases} 1 & i = j \\ 0 & i \neq j \end{cases}$$

The eigenvector matrix

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

$$v_i^T v_j = \begin{cases} 1 & i = j \\ 0 & i \neq j \end{cases}$$

We can write this as

$$V^T V = I$$

where

$$V = [v_1, \ldots, v_d]$$

The eigenvector matrix is orthonormal

$$V^T V = I$$

. . . and it also turns out that

$$VV^T = I$$

Eigenvectors orthogonalize a symmetric matrix

$$v_i^T G v_j = v_i^T(\lambda_j v_j) = \lambda_j v_i^T v_j = \begin{cases} \lambda_j, & i = j \\ 0, & i \neq j \end{cases}$$

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

$$V^T G V = \Lambda$$

$$\Lambda = \begin{bmatrix} \lambda_1 & 0 & 0 \\ 0 & \dots & 0 \\ 0 & 0 & \lambda_d \end{bmatrix}$$

Summary: symmetric positive semi-definite matrices

If G is symmetric and positive semi-definite, then

$$\Lambda = V^T G V$$

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

Putting those two together, we also get this statement, which says that you can reconstruct G from the scaled outer products of its eigenvectors:

$$V \Lambda V^T = VV^T G V V^T = G$$

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 x space, and show what happens to those vectors after transformation to the y space. Calculate the determinant of the matrix, and its eigenvalues and eigenvectors. Show that $Av = \lambda v$.

Outline

Summary

- A linear transform, A, maps vectors in space x to vectors in space 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 ($u_i^T v_j = 0$) or share the same eigenvalue ($\kappa_i = \lambda_j$).
- For a symmetric positive semidefinite matrix $G = A^T A$, the left and right eigenvectors are the same. If the eigenvalues are distinct and real, then:

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