

We can also project onto a plane that does not go through the origin
(i.e. a plane that is not a subspace of $\mathbb{R}^3$).

Such a plane has formula $\vec{n}^T (\vec{v} - \vec{v}_0) = 0$

- Translate $\vec{v}_0$ to the origin so the plane goes through 0
- Project
- Translate the origin back to $\vec{v}_0$.

$$\begin{bmatrix} \mathbf{I} & \begin{smallmatrix} 0 \\ 0 \\ 0 \end{smallmatrix} \\ -\vec{v}_0 & 1 \end{bmatrix} \begin{bmatrix} \mathbf{I} - \vec{n}\vec{n}^T & \begin{smallmatrix} 0 \\ 0 \\ 0 \end{smallmatrix} \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} \mathbf{I} & \begin{smallmatrix} 0 \\ 0 \\ 0 \end{smallmatrix} \\ \vec{v}_0 & 1 \end{bmatrix} = \text{projection onto } \vec{n}^T (\vec{v} - \vec{v}_0) = 0$$

Eg. Projecting onto $x + y + 2z = 1$.

$$\begin{bmatrix} 1 \\ 1 \\ 2 \end{bmatrix}^T \begin{bmatrix} x \\ y \\ z \end{bmatrix} = 1$$

Choose $\vec{v}_0 = \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}$, get

$$\begin{bmatrix} 1 \\ 1 \\ 2 \end{bmatrix}^T \left(\begin{bmatrix} x \\ y \\ z \end{bmatrix} - \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} \right) = 0$$

$$T-PT_+ = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ -1 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 0 & -1 & -2 & 0 \\ -1 & 0 & -2 & 0 \\ -2 & -2 & -3 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 1 \end{bmatrix}$$

$$\vec{n}\vec{n}^T = \begin{bmatrix} 1 \\ 1 \\ 2 \end{bmatrix} \begin{bmatrix} 1 & 1 & 2 \end{bmatrix} = \begin{bmatrix} 1 & 1 & 2 \\ 1 & 1 & 2 \\ 2 & 2 & 4 \end{bmatrix}$$

Class 22: 11/27

Today: SVD

(Singular value Decomposition)

Question: To what extent can we diagonalize a non-square matrix?

When $A = S\Lambda S^{-1}$, we had $AS = SA$

$$\text{so } \begin{aligned} A\vec{x}_1 &= \lambda_1 \vec{x}_1 \\ A\vec{x}_2 &= \lambda_2 \vec{x}_2 \text{ etc.} \end{aligned}$$

But that required A to be square.

For SVD, equations are

$$\begin{aligned} A\vec{u}_1 &= \sigma_1 \vec{u}_1 \\ A\vec{u}_2 &= \sigma_2 \vec{u}_2 \end{aligned}$$

$$A\vec{u}_r = \sigma_r \vec{u}_r$$

(u 's and v 's are called singular vectors)

and σ 's are singular values)

SVD: $A = U \Sigma V^T$ where U, V are orthogonal, Σ diagonal.

(So $AV = U \Sigma$) We make Σ have $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_r > 0$

Note: this says $A = \vec{u}_1 \sigma_1 \vec{v}_1^T + \dots + \vec{u}_r \sigma_r \vec{v}_r^T$

Two issues:

1. How do we find U, Σ, V ?
2. Why do we care?

1. To find U, Σ, V , use that if $A = U \Sigma V^T$, then
 $ATA = V \Sigma U^T U \Sigma V^T = V \Sigma^2 V^T$
and $AA^T = U \Sigma V^T V \Sigma U^T = U \Sigma^2 U^T$

So diagonalize ATA to find V and Σ^2

Then we can diagonalize AA^T to find U

or use $AV = U \Sigma$ to find U .

(Note: ATA is symmetric, so V is orthogonal)

Example:

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

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

$$\begin{aligned} ATA's \text{ eigenvalues come from } \det(ATA - \lambda I) &= \det \begin{bmatrix} 1-\lambda & 0 & 0 \\ 0 & 1-\lambda & 1 \\ 0 & 1 & 1-\lambda \end{bmatrix} \\ &= (1-\lambda) \begin{vmatrix} 1-\lambda & 1 \\ 1 & 1-\lambda \end{vmatrix} = (1-\lambda)[(1-\lambda)^2 - 1] \\ &= (1-\lambda)(-2\lambda + \lambda^2) = (1-\lambda)\lambda(\lambda-2) \end{aligned}$$

$$\text{So } ATA = V \begin{bmatrix} 2 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{bmatrix} V^T \text{ and } \Sigma = \begin{bmatrix} \sqrt{2} & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{bmatrix}$$

$$\text{Find } V: \lambda_1 = 2: \begin{bmatrix} -1 & 0 & 0 \\ 0 & -1 & 1 \\ 0 & 1 & -1 \end{bmatrix} \vec{x} = 0 \rightsquigarrow \vec{x} = \frac{1}{\sqrt{2}} \begin{bmatrix} 0 \\ 1 \\ 1 \end{bmatrix}$$

$$\lambda_2 = 1: \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{bmatrix} \vec{x} = 0 \rightsquigarrow \vec{x} = \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}$$

$$\lambda_3 = 0: \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 1 \\ 0 & 1 & 1 \end{bmatrix} \vec{x} = 0 \rightsquigarrow \vec{x} = \frac{1}{\sqrt{2}} \begin{bmatrix} 0 \\ 1 \\ -1 \end{bmatrix}$$

$$\text{So } V = \begin{bmatrix} 0 & 1 & 0 \\ \frac{1}{\sqrt{2}} & 0 & \frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} & 0 & -\frac{1}{\sqrt{2}} \end{bmatrix}$$

$$\text{Now } AV = U\Sigma, \text{ so } \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 1 \end{bmatrix} \begin{bmatrix} 0 & 1 & 0 \\ \frac{1}{\sqrt{2}} & 0 & \frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} & 0 & -\frac{1}{\sqrt{2}} \end{bmatrix} = U \begin{bmatrix} \sqrt{2} & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix}$$

$$\begin{bmatrix} 0 & 1 & 0 \\ \frac{1}{\sqrt{2}} & 0 & 0 \end{bmatrix} = [\vec{u}_1 \quad \vec{u}_2] \begin{bmatrix} \sqrt{2} & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix} \\ = \begin{bmatrix} \sqrt{2} \vec{u}_1 & \vec{u}_2 & 0 \end{bmatrix}$$

$$\text{So } \vec{u}_1 = \begin{bmatrix} 0 \\ 1 \\ 1 \end{bmatrix}, \vec{u}_2 = \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}$$

$$A = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} \begin{bmatrix} \sqrt{2} & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix} \begin{bmatrix} 0 & 1 & 0 \\ \frac{1}{\sqrt{2}} & 0 & \frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} & 0 & -\frac{1}{\sqrt{2}} \end{bmatrix}$$

2. Why do we care?

Two reasons.

- 1) Easily read off bases for 4 fundamental subspaces
- 2) Data compression

- 1) Fact: First r columns of V are a basis for the row space of A
 Last $n-r$ columns of V are a basis for the nullspace of A
 First r columns of U are a basis for the column space of A
 Last $n-r$ columns of U are a basis for the left nullspace of A

Why? Use $AV = U\Sigma$.

If A is $m \times n$, this says that A takes an orthonormal basis of $\mathbb{R}^n$ to an orthogonal basis of $\mathbb{R}^m$.

$$A \begin{bmatrix} \vec{v}_1 & \vec{v}_2 & \dots & \vec{v}_n \end{bmatrix} = \begin{bmatrix} \vec{u}_1 & \dots & \vec{u}_m \end{bmatrix} \begin{bmatrix} \sigma_1 & & & \\ & \sigma_2 & & \\ & & \ddots & \\ & & & \sigma_r & & \\ & & & & & \text{---} \\ & & & & & & \text{---} \\ & & & & & & & \text{---} \end{bmatrix}$$

n-r columns of 0's

m-r rows of 0's

So $A\vec{v}_1, A\vec{v}_2, \dots, A\vec{v}_r$ (which are $\vec{u}_1\sigma_1, \vec{u}_2\sigma_2, \dots, \vec{u}_r\sigma_r$) form a basis for $C(A)$. $A\vec{v}_{r+1}, \dots, A\vec{v}_n = 0$, so they don't contribute.

* First r columns of U form an orthonormal basis for $C(A)$.
 But the left nullspace of A is $C(A)^\perp$, so the last $m-r$ columns of U form an orthonormal basis for the left nullspace.

(Repeat argument with A^T to get results for V).

2) Data compression

An image is made of pixels (say 2^9 by 2^{10} pixels). Each pixel has a grayscale value. You can keep track of this with a $2^9 \times 2^{10}$ matrix. This has 2^{19} entries. No big deal!

Until you start making movies with 30 images per second, or doing a CT scan which has tons of cross section images. Now there are lots of $2^9 \times 2^{10}$ matrices to keep track of.

Options:

- 1) Bigger pixels. Double pixel height + width and reduce the number of pixels by a factor of 4.
Downside: $\frac{1}{4}$ of 2^{19} (times many images per second) is still huge, and picture quality decreases.

- 2) Use SVD: If A is your pixel matrix, approximate A using the largest singular values
$$A \approx \sigma_1 \vec{u}_1 \vec{v}_1^T \quad (\vec{u}_1 \text{ has } m \text{ entries, } \vec{v}_1 \text{ has } n \text{ entries})$$

$$2^9 \quad \quad \quad 2^{10}$$

So this means we only need to keep track of $2^9 + 2^{10}$ numbers!

In reality, use the first few singular values (until the approx. is pretty good).

Say $A \approx \sigma_1 \vec{u}_1 \vec{v}_1^T + \sigma_2 \vec{u}_2 \vec{v}_2^T + \sigma_3 \vec{u}_3 \vec{v}_3^T$

Keep track of $3(2^9 + 2^{10})$ numbers.
Way less than 2^{19} !