

As you get settled...

we'll start @ 11:47
lecture ends: 1:25

- Get out your notes
- Get out a place to do today's ICA (4)

Now playing:
"3 nights", Dominic Fike
"Starstruck", Years & Years

- Topics we'll start with: dot products, writing solutions as vectors

Can you see this text? ¹



Vectors and Machine Learning

Warm-up: which of the following vectors is "most similar" to the vector:

$$\vec{a} = \begin{bmatrix} 1 \\ 3 \end{bmatrix}$$

Choices:

$$\vec{b} = \begin{bmatrix} 2 \\ 4 \end{bmatrix} \quad \vec{c} = \begin{bmatrix} 0 \\ 3 \end{bmatrix} \quad \vec{d} = \begin{bmatrix} -1 \\ -3 \end{bmatrix}$$

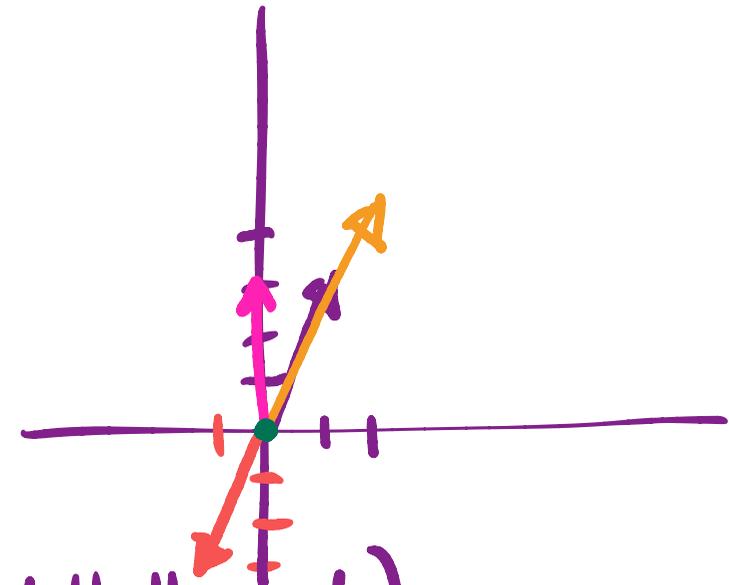
$$\vec{e} = \begin{bmatrix} 0 \\ 0 \end{bmatrix}$$

$\vec{a} \cdot \vec{b} = 1 \cdot 2 + 3 \cdot 4 = 14$ *↳ dot product*

$\vec{a} \cdot \vec{c} = 9$

$\vec{a} \cdot \vec{d} = -10$

argument for working w/ unit vectors (vec. w/ ||v|| = 1)



Visualizing our solutions space (many solutions)

ICA Question 1: solve the following system of equations (make it into RREF form), then express the solution as vectors and describe the solution space in terms of shape (doesn't exist, a point, a line, a 2D plane, etc).

$$\begin{bmatrix} a \\ b \\ c \\ d \\ e \end{bmatrix} = \begin{bmatrix} 3 \\ 2 \\ 0 \\ 0 \\ 0 \end{bmatrix} + \begin{bmatrix} \\ \\ \\ \\ \end{bmatrix} c$$

$r_0' = r_0 - r_2 \rightarrow$

$$\left[\begin{array}{ccccc|c} 1 & 0 & -1 & -2 & -3 & 3 \\ 0 & 1 & 2 & 3 & 4 & 2 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{array} \right]$$

$a - c - 2d - 3e = 3$
 $b + 2c + 3d + 4e = 2$

$$\begin{aligned} c &= c \\ d &= d \\ e &= e \end{aligned}$$

$$\begin{aligned}
 a - c - 2d - 3e &= 3 \rightarrow a = 3 + c + 2d + 3e \\
 b + 2c + 3d + 4e &= 2 \rightarrow b = 2 - 2c - 3d - 4e \\
 c &= c & c &= 0 + c \\
 d &= d & d &= 0 + 0c + d \\
 e &= e & e &= 0 + 0c + 0d + e
 \end{aligned}$$

$$\begin{bmatrix} a \\ b \\ c \\ d \\ e \end{bmatrix} = \begin{bmatrix} 3 \\ 2 \\ 0 \\ 0 \\ 0 \end{bmatrix} + \begin{bmatrix} 1 \\ -2 \\ 1 \\ 0 \\ 0 \end{bmatrix} c + \begin{bmatrix} 2 \\ -3 \\ 0 \\ 1 \\ 0 \end{bmatrix} d + \begin{bmatrix} 3 \\ -4 \\ 0 \\ 0 \\ 1 \end{bmatrix} e$$

Machine Learning

- What is *machine learning* anyways?

↳ data!

↳ algorithms → strategy

↳ "learn" → improving, identifying patterns

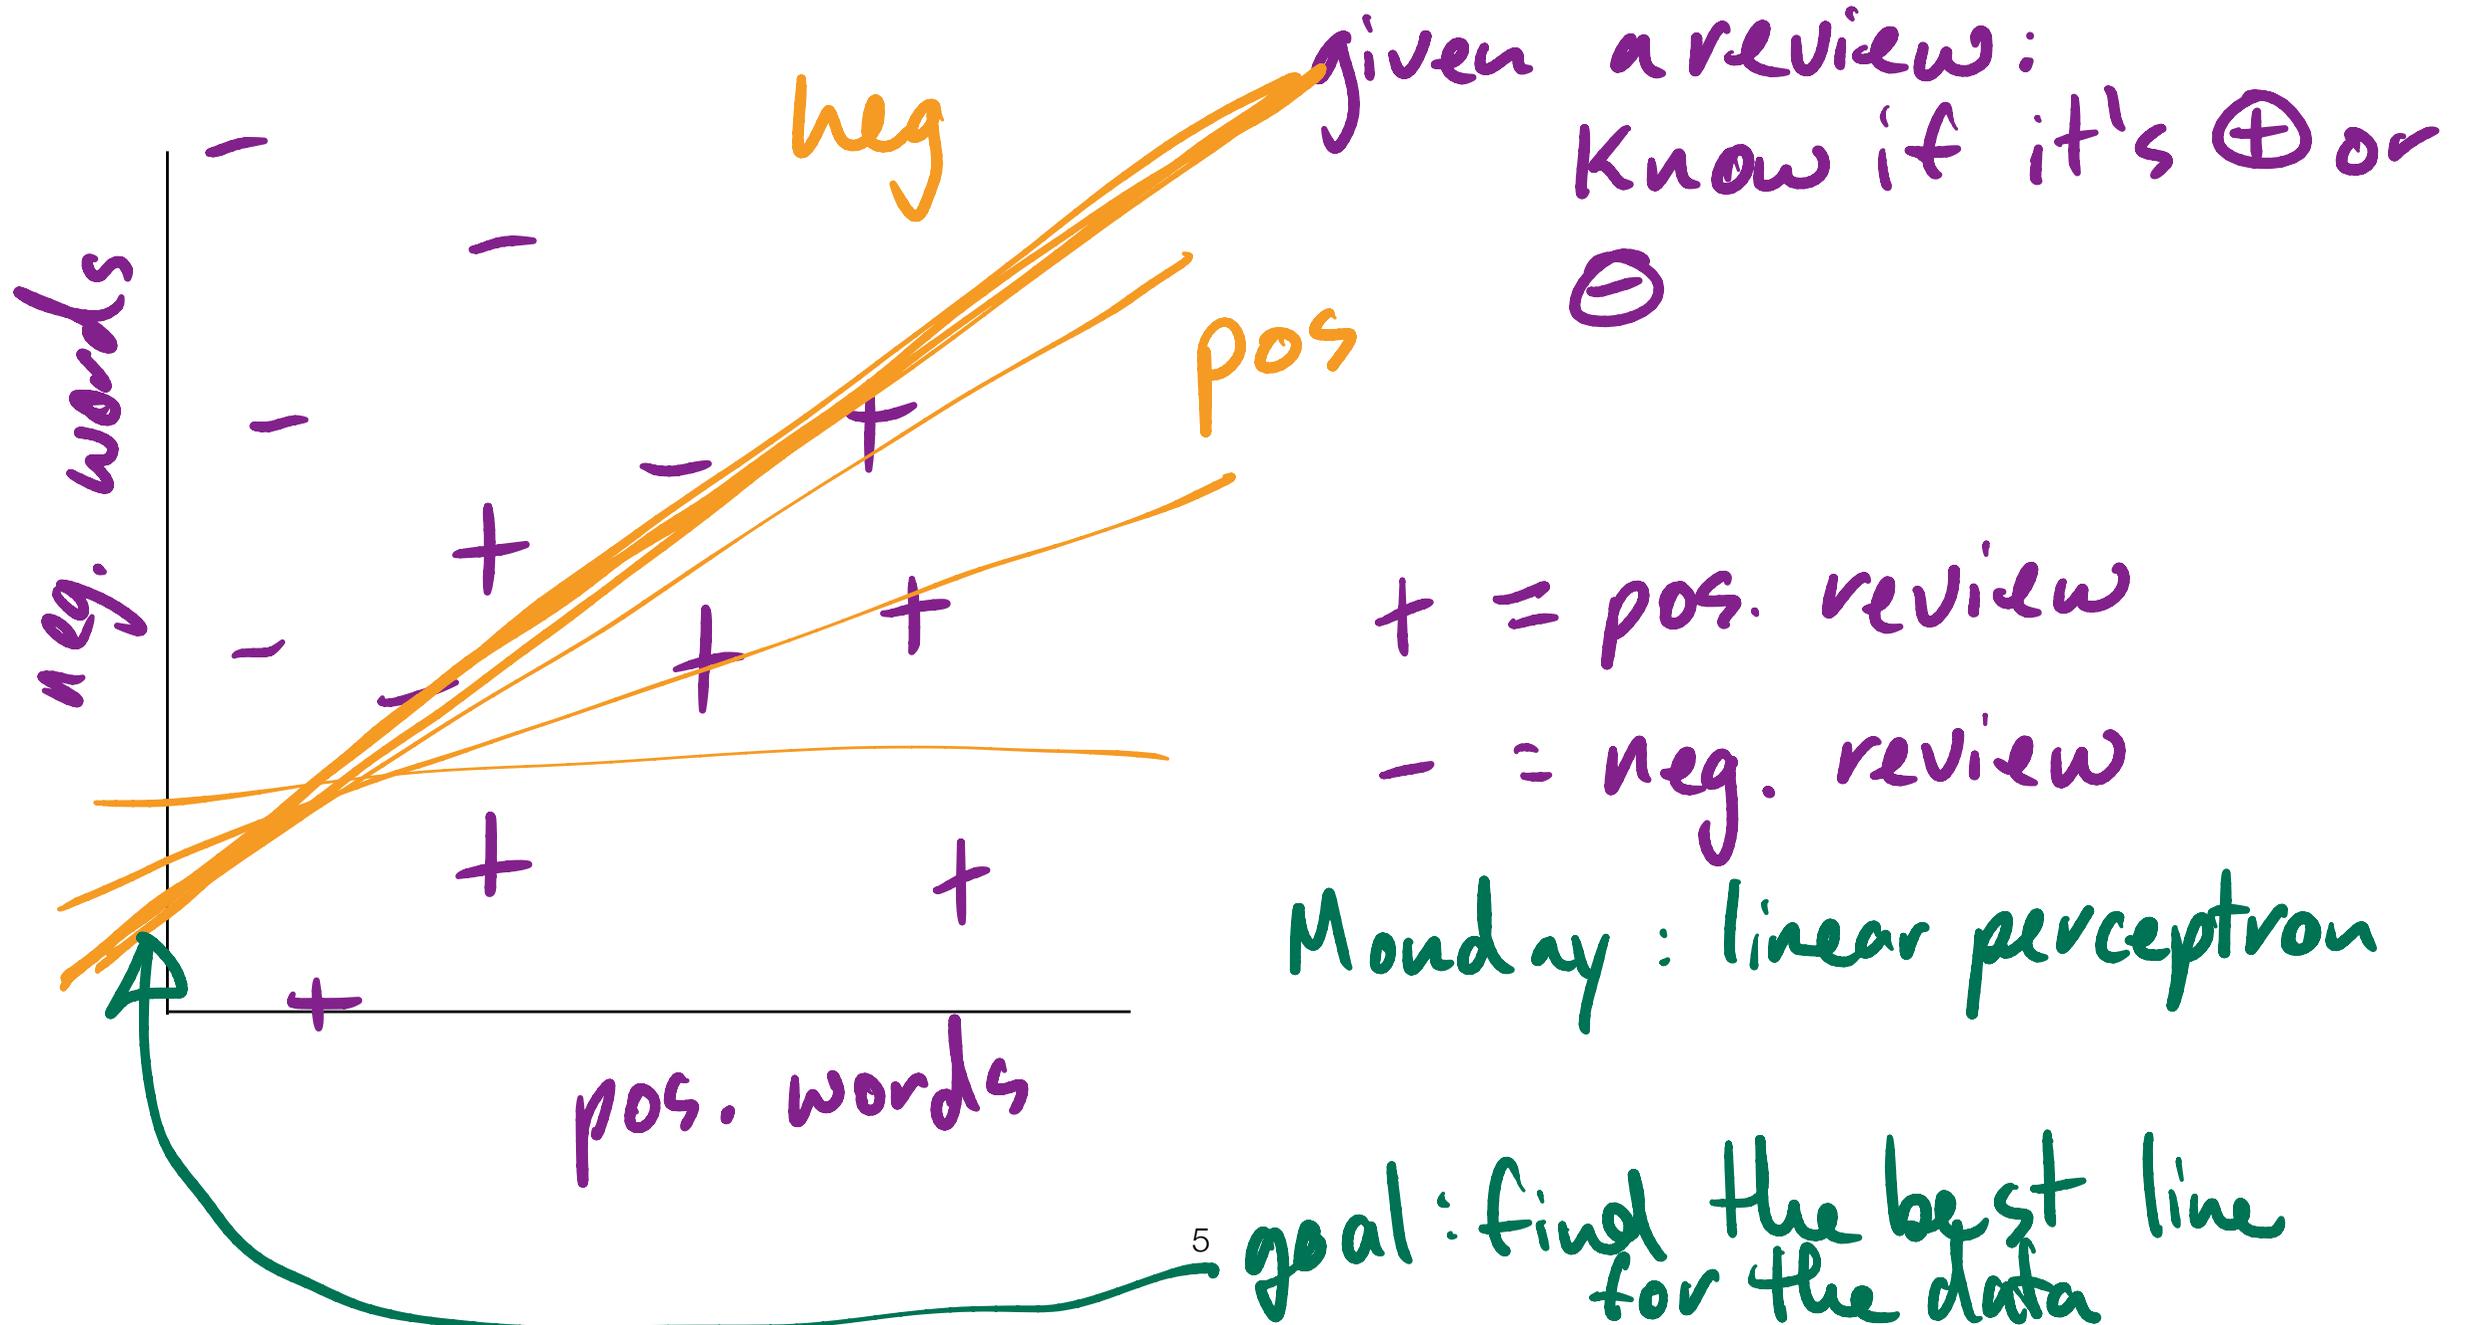
↳ linear regression → ex. algorithm

↳ making predictions → made by model

↳ AI

↳ NLP

Visualizing Linear Machine Learning Algorithms



Machine Learning - from data to features

- Some of the data that we model in the real world is **numeric**, some is not.

numeric

temperature

stocks

height/age

vaccination rate
housing \$\$\$

not numeric

food descriptions

word-based

book titles

music

images

- In all cases, we need to **featurize** our datasets:

pos words neg words

label	feature1	feature2	feature3	...	featureN
pos	10	2	0		
neg	2	4	0		
			1		
			1		

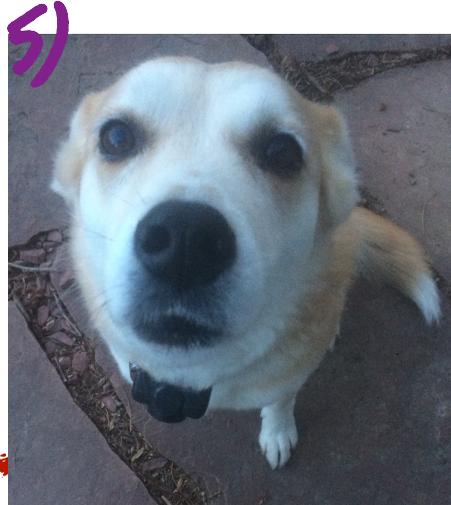
contains an "!"

↳ category we're trying to predict

Machine Learning - from data to features

ICA Question 2: suppose that you are given the following data. **Featurize** the data by choosing two features and filling in the table. Make sure that your features **are numeric!**

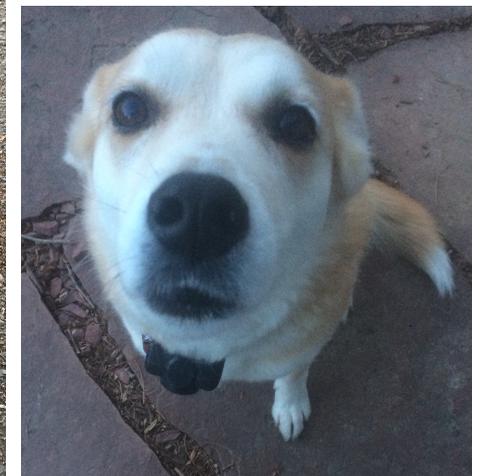
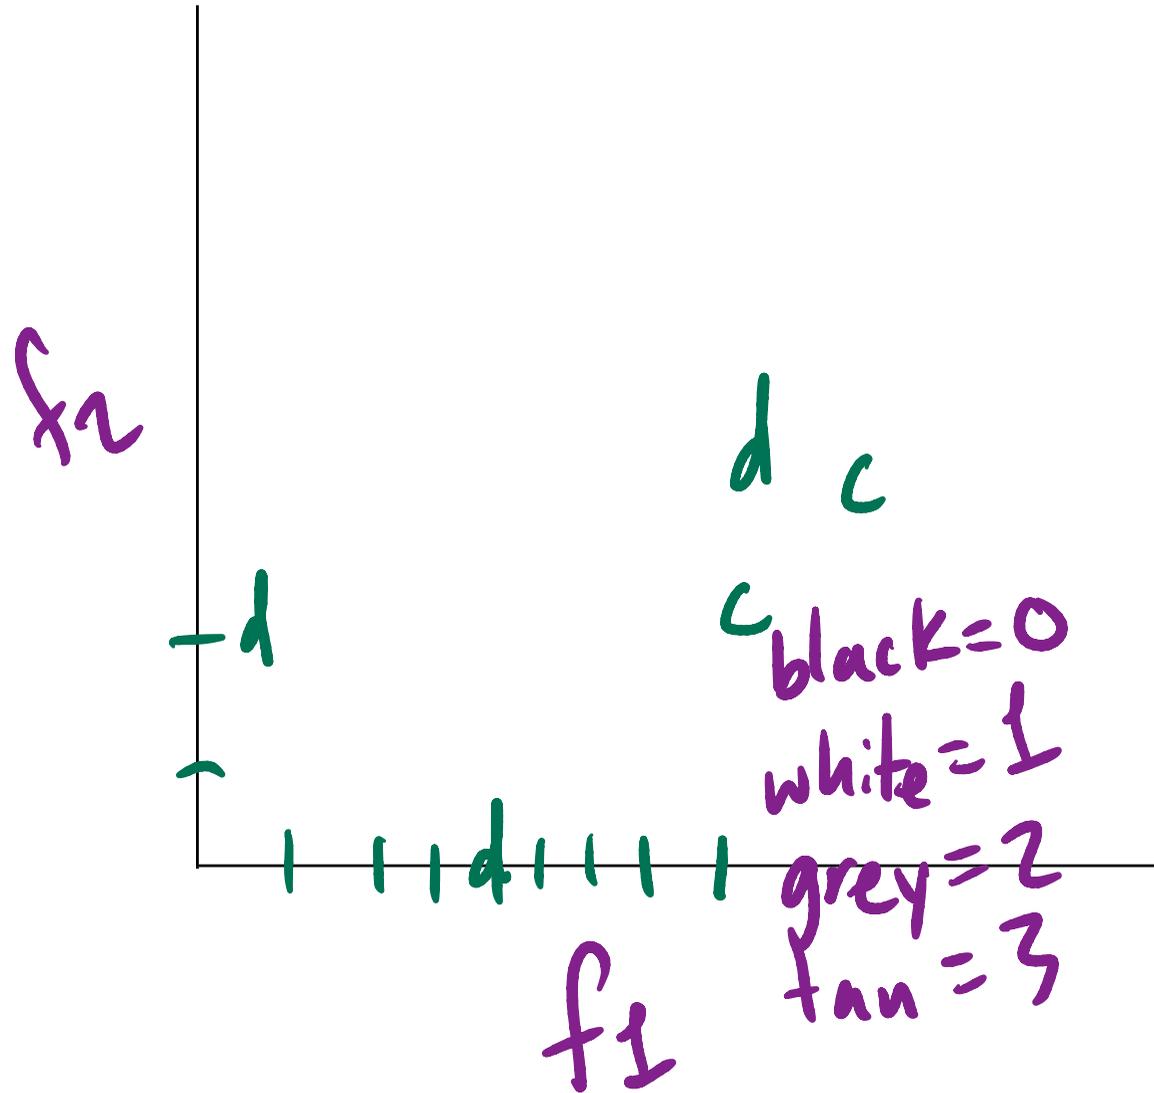
Data:



label	feature1	feature2
cat		
cat		
dog		
dog		
dog		

Visualizing Linear Machine Learning Algorithms color

label	feature1 age	feature2
cat	8	2
cat	9	3
dog	4	0
dog	0.5	2
dog	8	3



Machine Learning & Ethics

break until : 12:59

- What are different contexts that we can think of where ML is used or likely used in real life?

↳ self driving cars

↳ spam detection

↳ facial recognition

↳ recitivism

↳ music recs

↳ speech analysis

↳ Q A systems

↳ financial analysis

Machine Learning & Ethics

- What is a **stereotype**?

↳ a generalization of a group based on one example / identity / cultural assumptions

Machine Learning & Ethics

- What is a **stereotype**?
 - Generalization
 - Culturally shared or widespread
 - Negative, neutral, or positive
 - Concerns: types of people, groups, ways of being, cultures, cultural products, intellectual products

Machine Learning & Ethics

- How can stereotypes be harmful?

↳ denial of services b/c of stereotype

↳ alienate people

↳ lead to discrimination in hiring

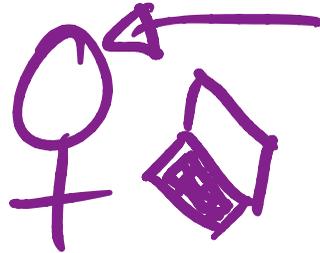
↳ make accessing healthcare more difficult

Representational vs. Allocative Bias

- A **representational** bias is when a system detracts from the representation of certain groups and their identities

google translate: "She is a doctor" → "He is a doctor"

image captioning:



id'd as a man b/c of
computer even if she's
a woman

- A **allocation/allocative** bias is when a system unfairly allocates resources to certain groups over others

Representational bias

*remember to turn in your
ICA on grade scope!*

ICA Question 3: you are working for Northeastern University. Suppose that you are given an ML model that helps you review applications.

- 1) Choose 3 numeric features that you might use to featurize an applicant.
- 2) Identify 3 places where an ML model reviewing applicants (regardless of features) might produce bias. Label each instance of bias as Representational (R) or Allocative (A).
- 3) If you finish, do some research online—do you find evidence that college applications are being automatically reviewed with ML models? What about job applications?

Machine Learning & Ethics

- Since machine learning algorithms "learn" from data and data is not produced in a vacuum...
- Where does the data come from?
- How was it labeled?
- What is the behavior of our model in the real world?

↳ is it biased?

Admin Stuff

- When you post on piazza, please **include HW #, Q #, and topic** in the title (when applicable), and **use follow-ups** to expand on a question that another student has asked!
- (this will help keep piazza most useful to you all!)
- Please use my Calendly office hours for higher-level "let's talk about this math/etc topic questions" and khoury office hours for HW questions (I have OH in both places).

Admin Stuff

- Where are you on HW 1?
 - A. I haven't looked at it
 - B. I've glanced at the problems
 - C. I've gotten started but I'm not very far
 - D. I'm probably half way through
 - E. I'm finished/almost finished

Schedule

Turn in ICA 4 on Gradescope
We are remote until Feb 5th.

wow, so many
office hours now!

khouryofficehours.com

Mon	Tue	Wed	Thu	Fri	Sat	Sun
January 24th Lecture 3 - Matrices & vector geometry	Felix OH Calendly		Lecture 4 - ML, linear perceptron Felix OH Khoury Office Hours			
January 31st Lecture 5 - Linear Perceptron	Felix OH Calendly	HW 1 due @ 11:59pm	Lecture 4 - matrix multiplication, transforms Felix OH Khoury Office Hours			