As you walk in

Welcome!

1) Help draw some aliens for our dataset today!
   - Grab a sharpie and some sticky notes
   - Just stick figures, nothing quite this fancy →

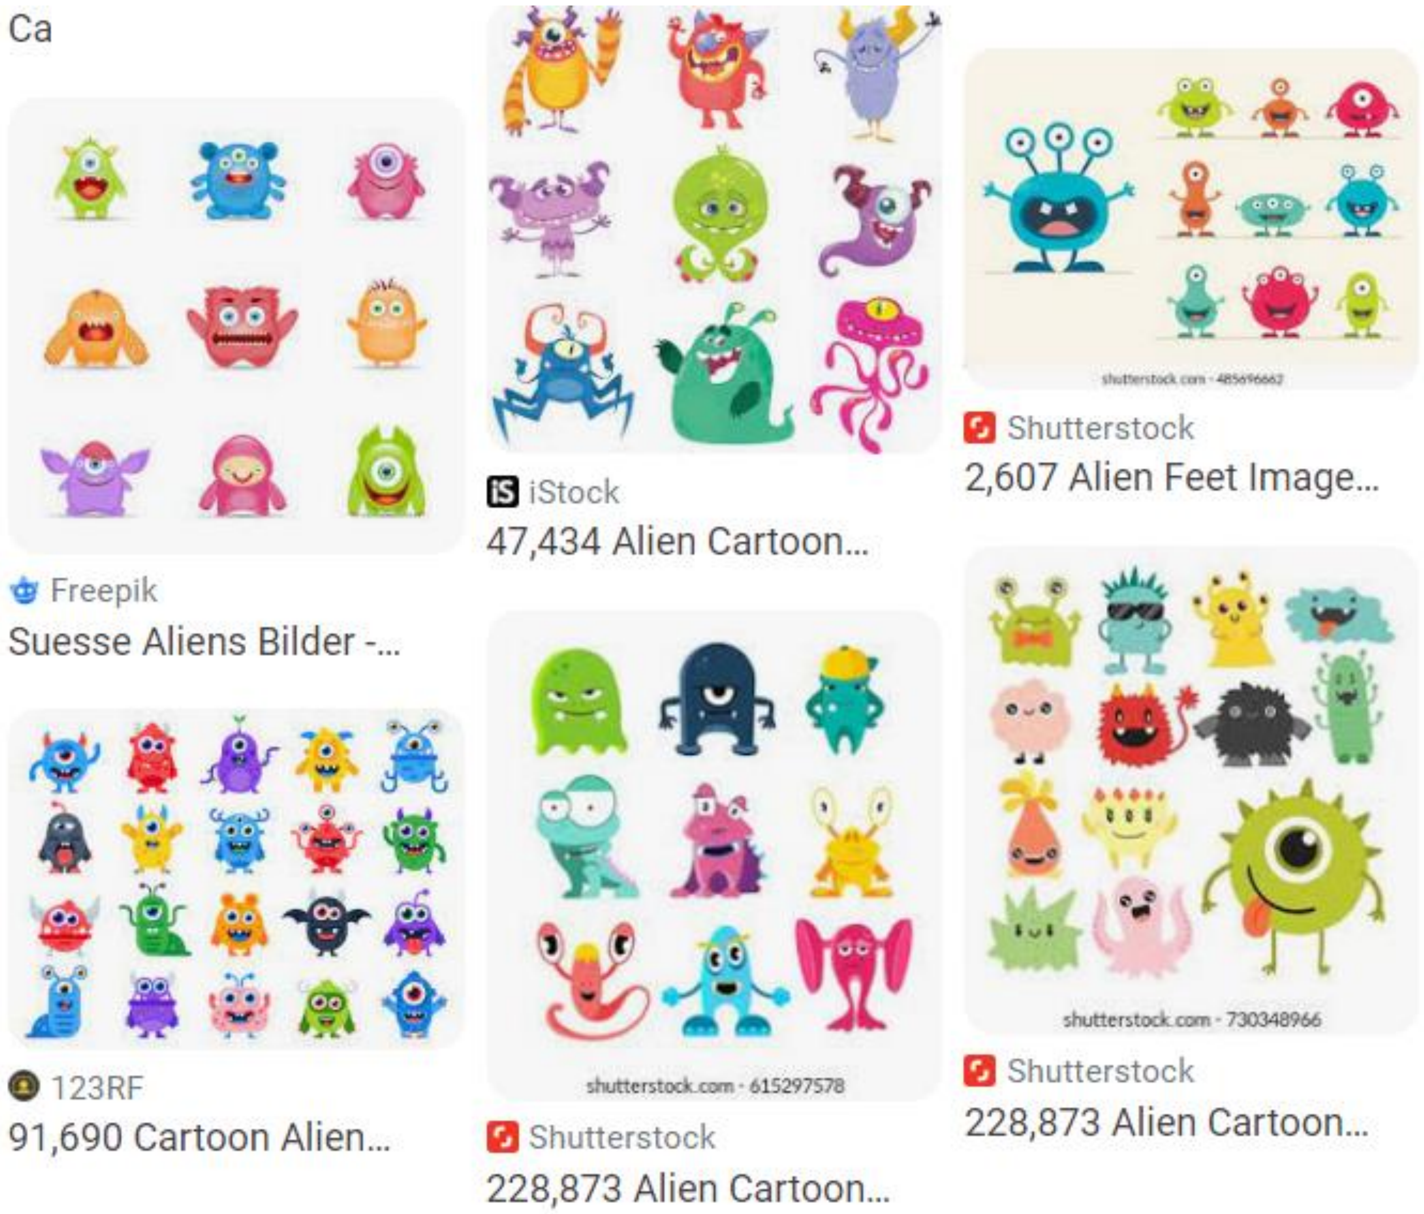

Exercise: Human-defined Feature Space

Try to organize data on a 2-D coordinate plot

to win the following game:

Select three students: A, B, C

1. Student A draws a new alien and hands it to student B
2. Student B thinks about where to plot it and comes up with a 2-D coordinate, (x, y)
3. Student C looks at the coordinate and the plot (but not the drawing from A) and draws a new alien

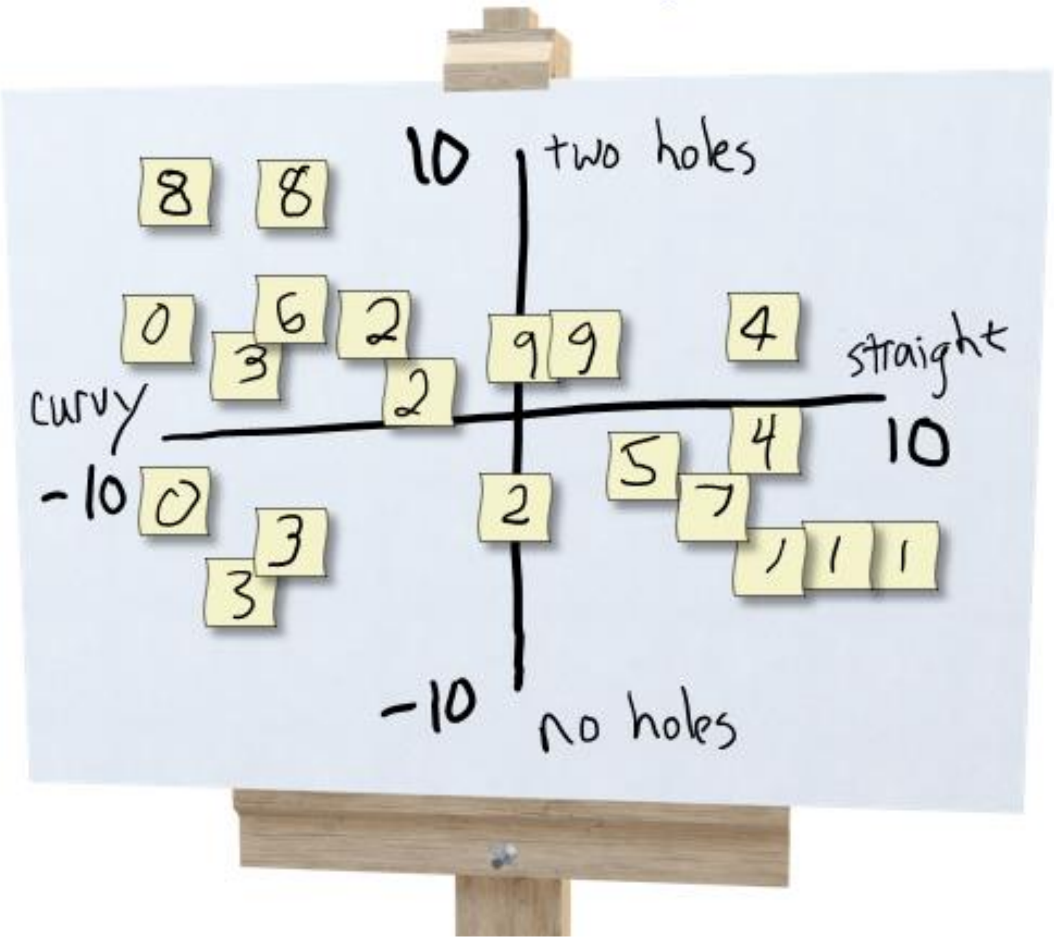

10-315
Introduction to ML

Instructor: Pat Virtue

Today

Aliens

Course team

What is AI and ML?

ML framework

Example ML models

- Autoencoder (Aliens)

Piazza and website for course info

DALL-E: "Logo of a Scotty dog with a red collar whose brain is made of circuits"

Course Information

Website: https://www.cs.cmu.edu/~10315

Communication:

piazza.com

E-mail Joshmin if piazza doesn't work:

joshminr@andrew.cmu.edu

Important posts coming soon!

- Office Hours
- Jupyter notebooks
- Recitation (Fri)
- Pre-reading (due Sun)
- HW0
- HW1

Course Team

Instructor

Pat
Virtue
pvirtue

Best way to contact Pat:

- Post on Piazza (private post as needed)
- Just grab an OH appointment slot

Education Associate

Joshmin Ray
joshminr

Email Joshmin, joshminr@andrew.cmu.edu, for any:

- Exceptions, extensions, etc
- Any course logistics

Course Team

Teaching Assistants


Alex
alextiax


Vincent
vtombari


Deep
dmpatel


Tanay
tbennur


Medha
mpalaval


Myles
msharris


Shreeya
srkhuran


Anisha
anishac

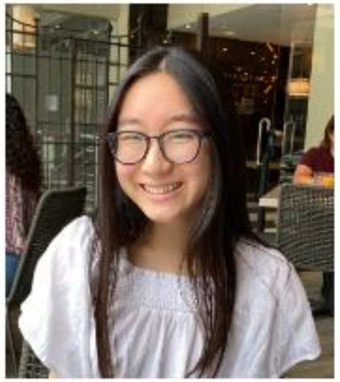
Ruthie
rylin


Ethan
ethanwan

Course Team

Students!!

"Poll"

Which of these are artificial? Neither

Which of these are more intelligent? 

AI Definition by John McCarthy

What is artificial intelligence

- It is the science and engineering of making intelligent machines, especially intelligent computer programs

What is intelligence

- Intelligence is the computational part of the ability to achieve goals in the world

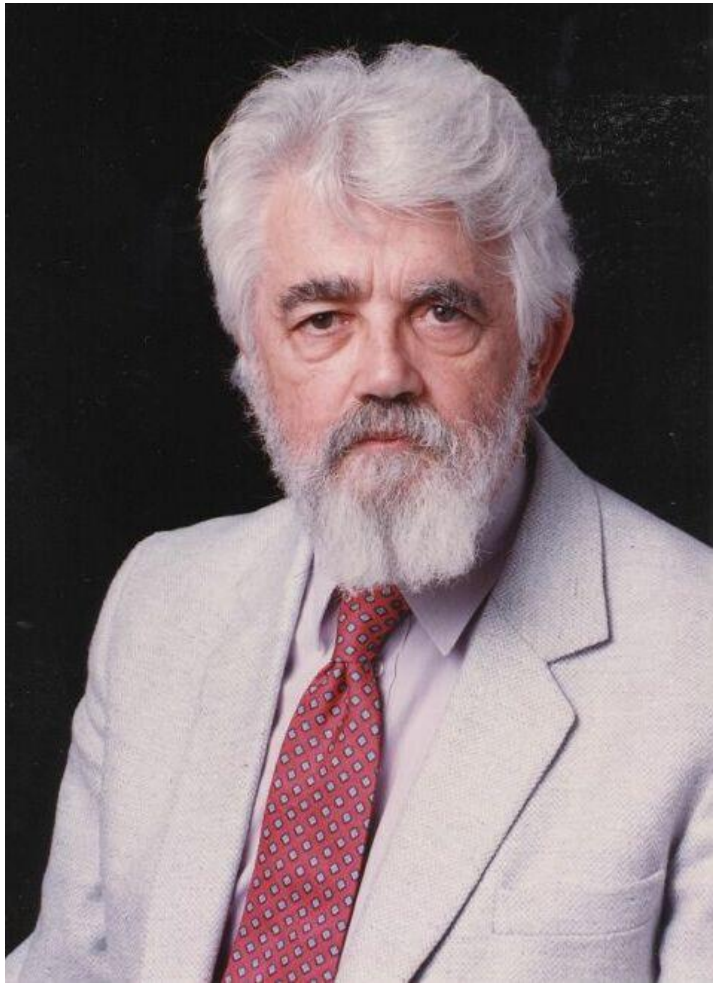

http://www-formal.stanford.edu/jmc/whatisai/whatisai.html

Attributes of Intelligence

- evoke emotion
- plan
- calc
- learn
- pat. recog.
- imititating
- extrapolating
  → adapt to new environ.
- hypothesizing
  - prediction based on exp
- testing

"Poll" Turn to your neighbor

Which of these more intelligent?

- Robot that assemble cars in factory
- Robot that fold clothes

Intelligence and Uncertainty

Uncertainty can have lots of sources, including anything we attribute to random chance

- Hidden information
  - Cards in another player's hand
- Noise
  - Sensor noise
- Way to complicated to model
  - Leaves blowing in the wind
- Infinite number of possible configurations
- More possibilities than any computer can compute in a reasonable time
  - Tic-tac-toe → Checkers → Chess

AI Definition by John McCarthy

What is artificial intelligence

- It is the science and engineering of making intelligent machines, especially intelligent computer programs

What is intelligence

- Intelligence is the computational part of the ability to achieve goals in the world

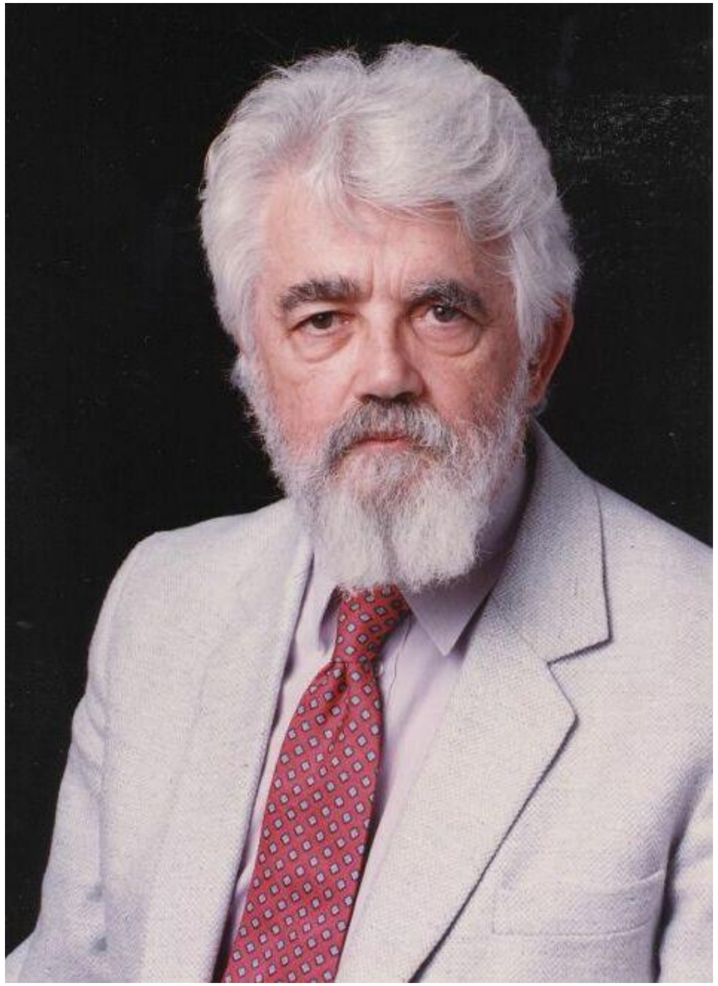

http://www-formal.stanford.edu/jmc/whatisai/whatisai.html

AI Definition (Pat's Version)

What is intelligence

- Intelligence is the ability to perform well on a task that involves uncertainty

Intelligence is not binary

- How **well** an agent performs and how much **uncertainty** is involved will determine how intelligent we consider the agent to be

Artificial Intelligence vs Machine Learning?

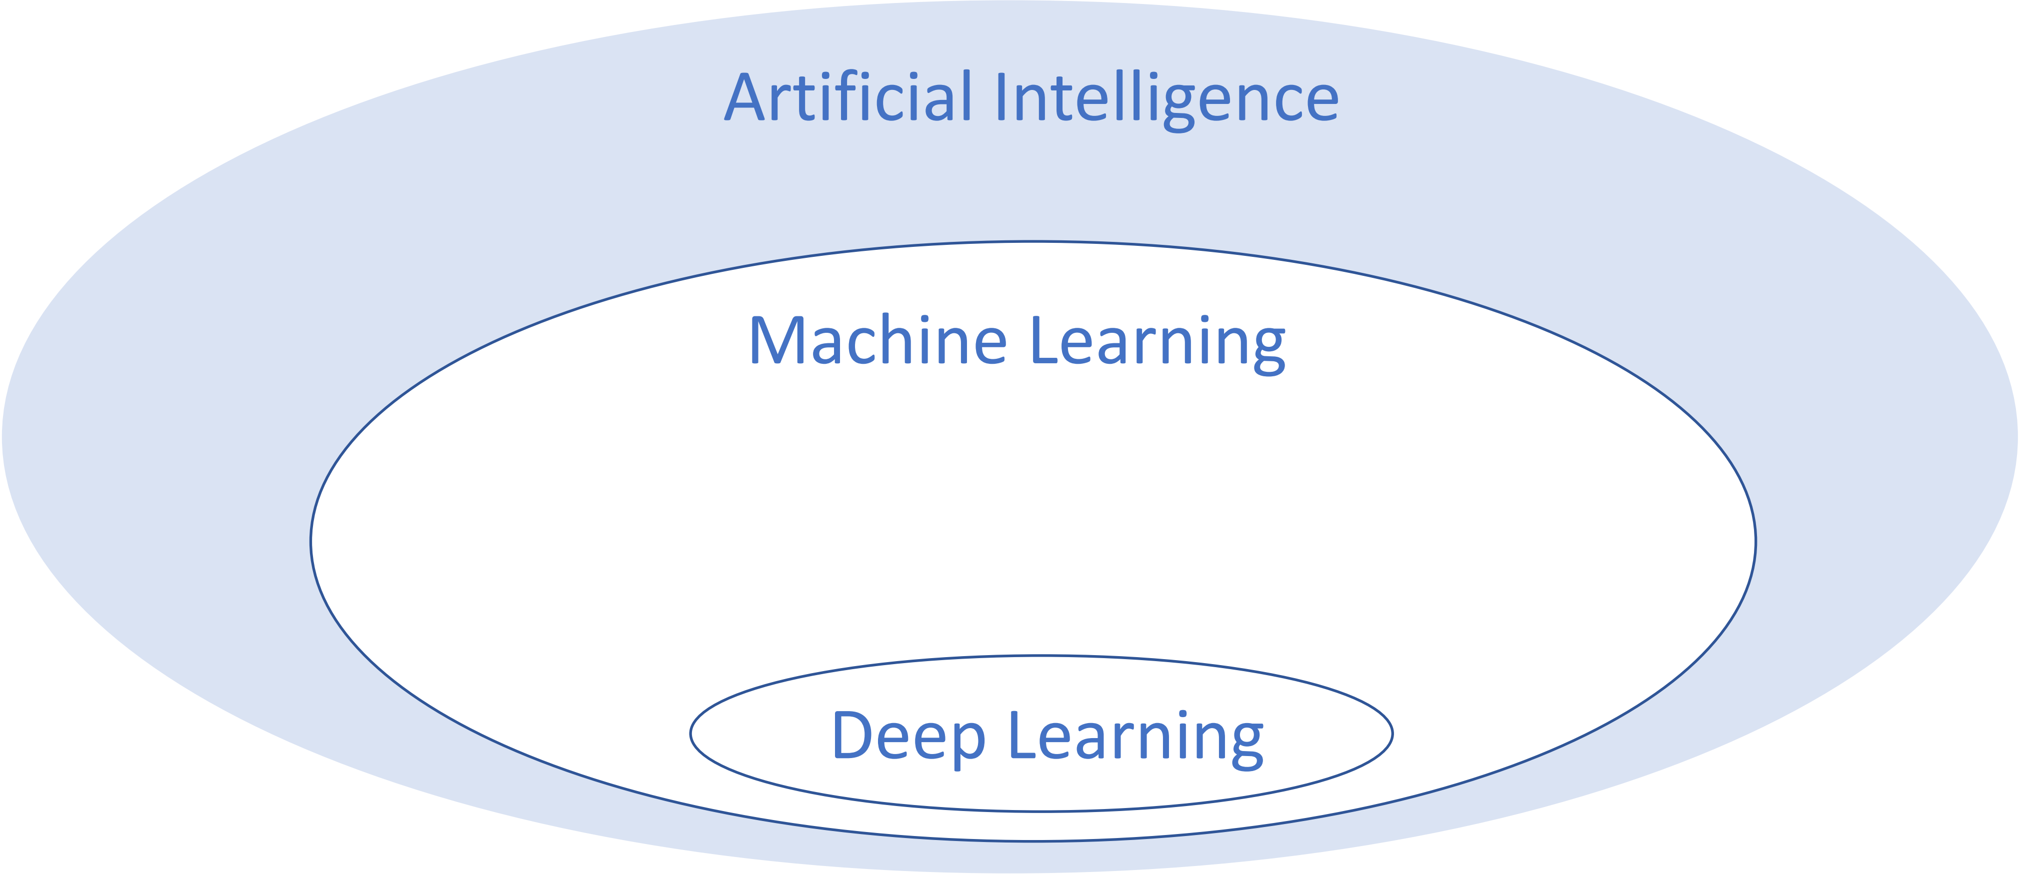

A Brief History of AI

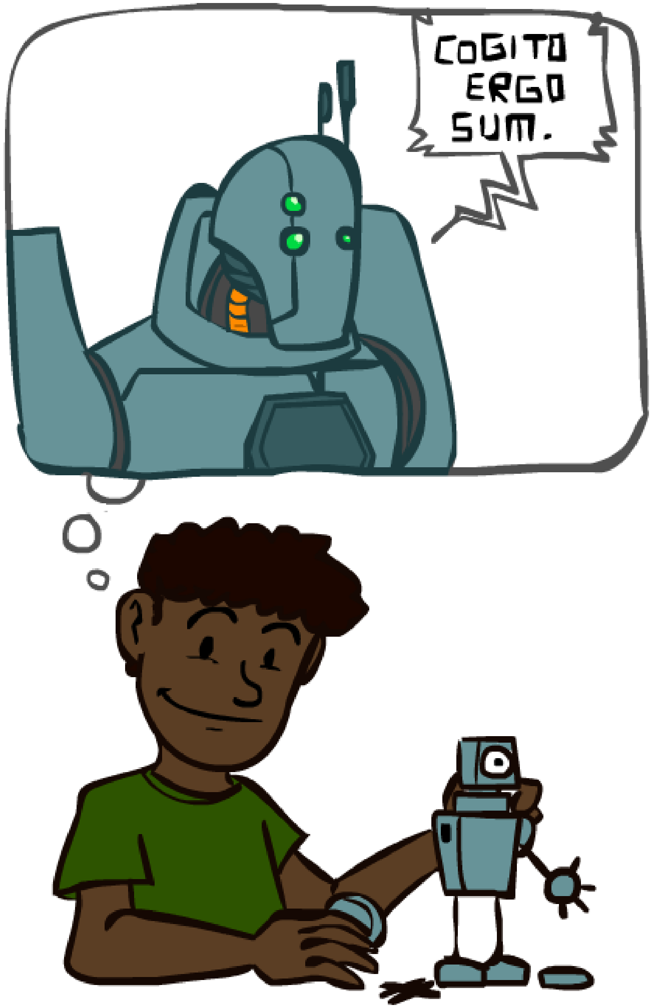

Images: ai.berkeley.edu

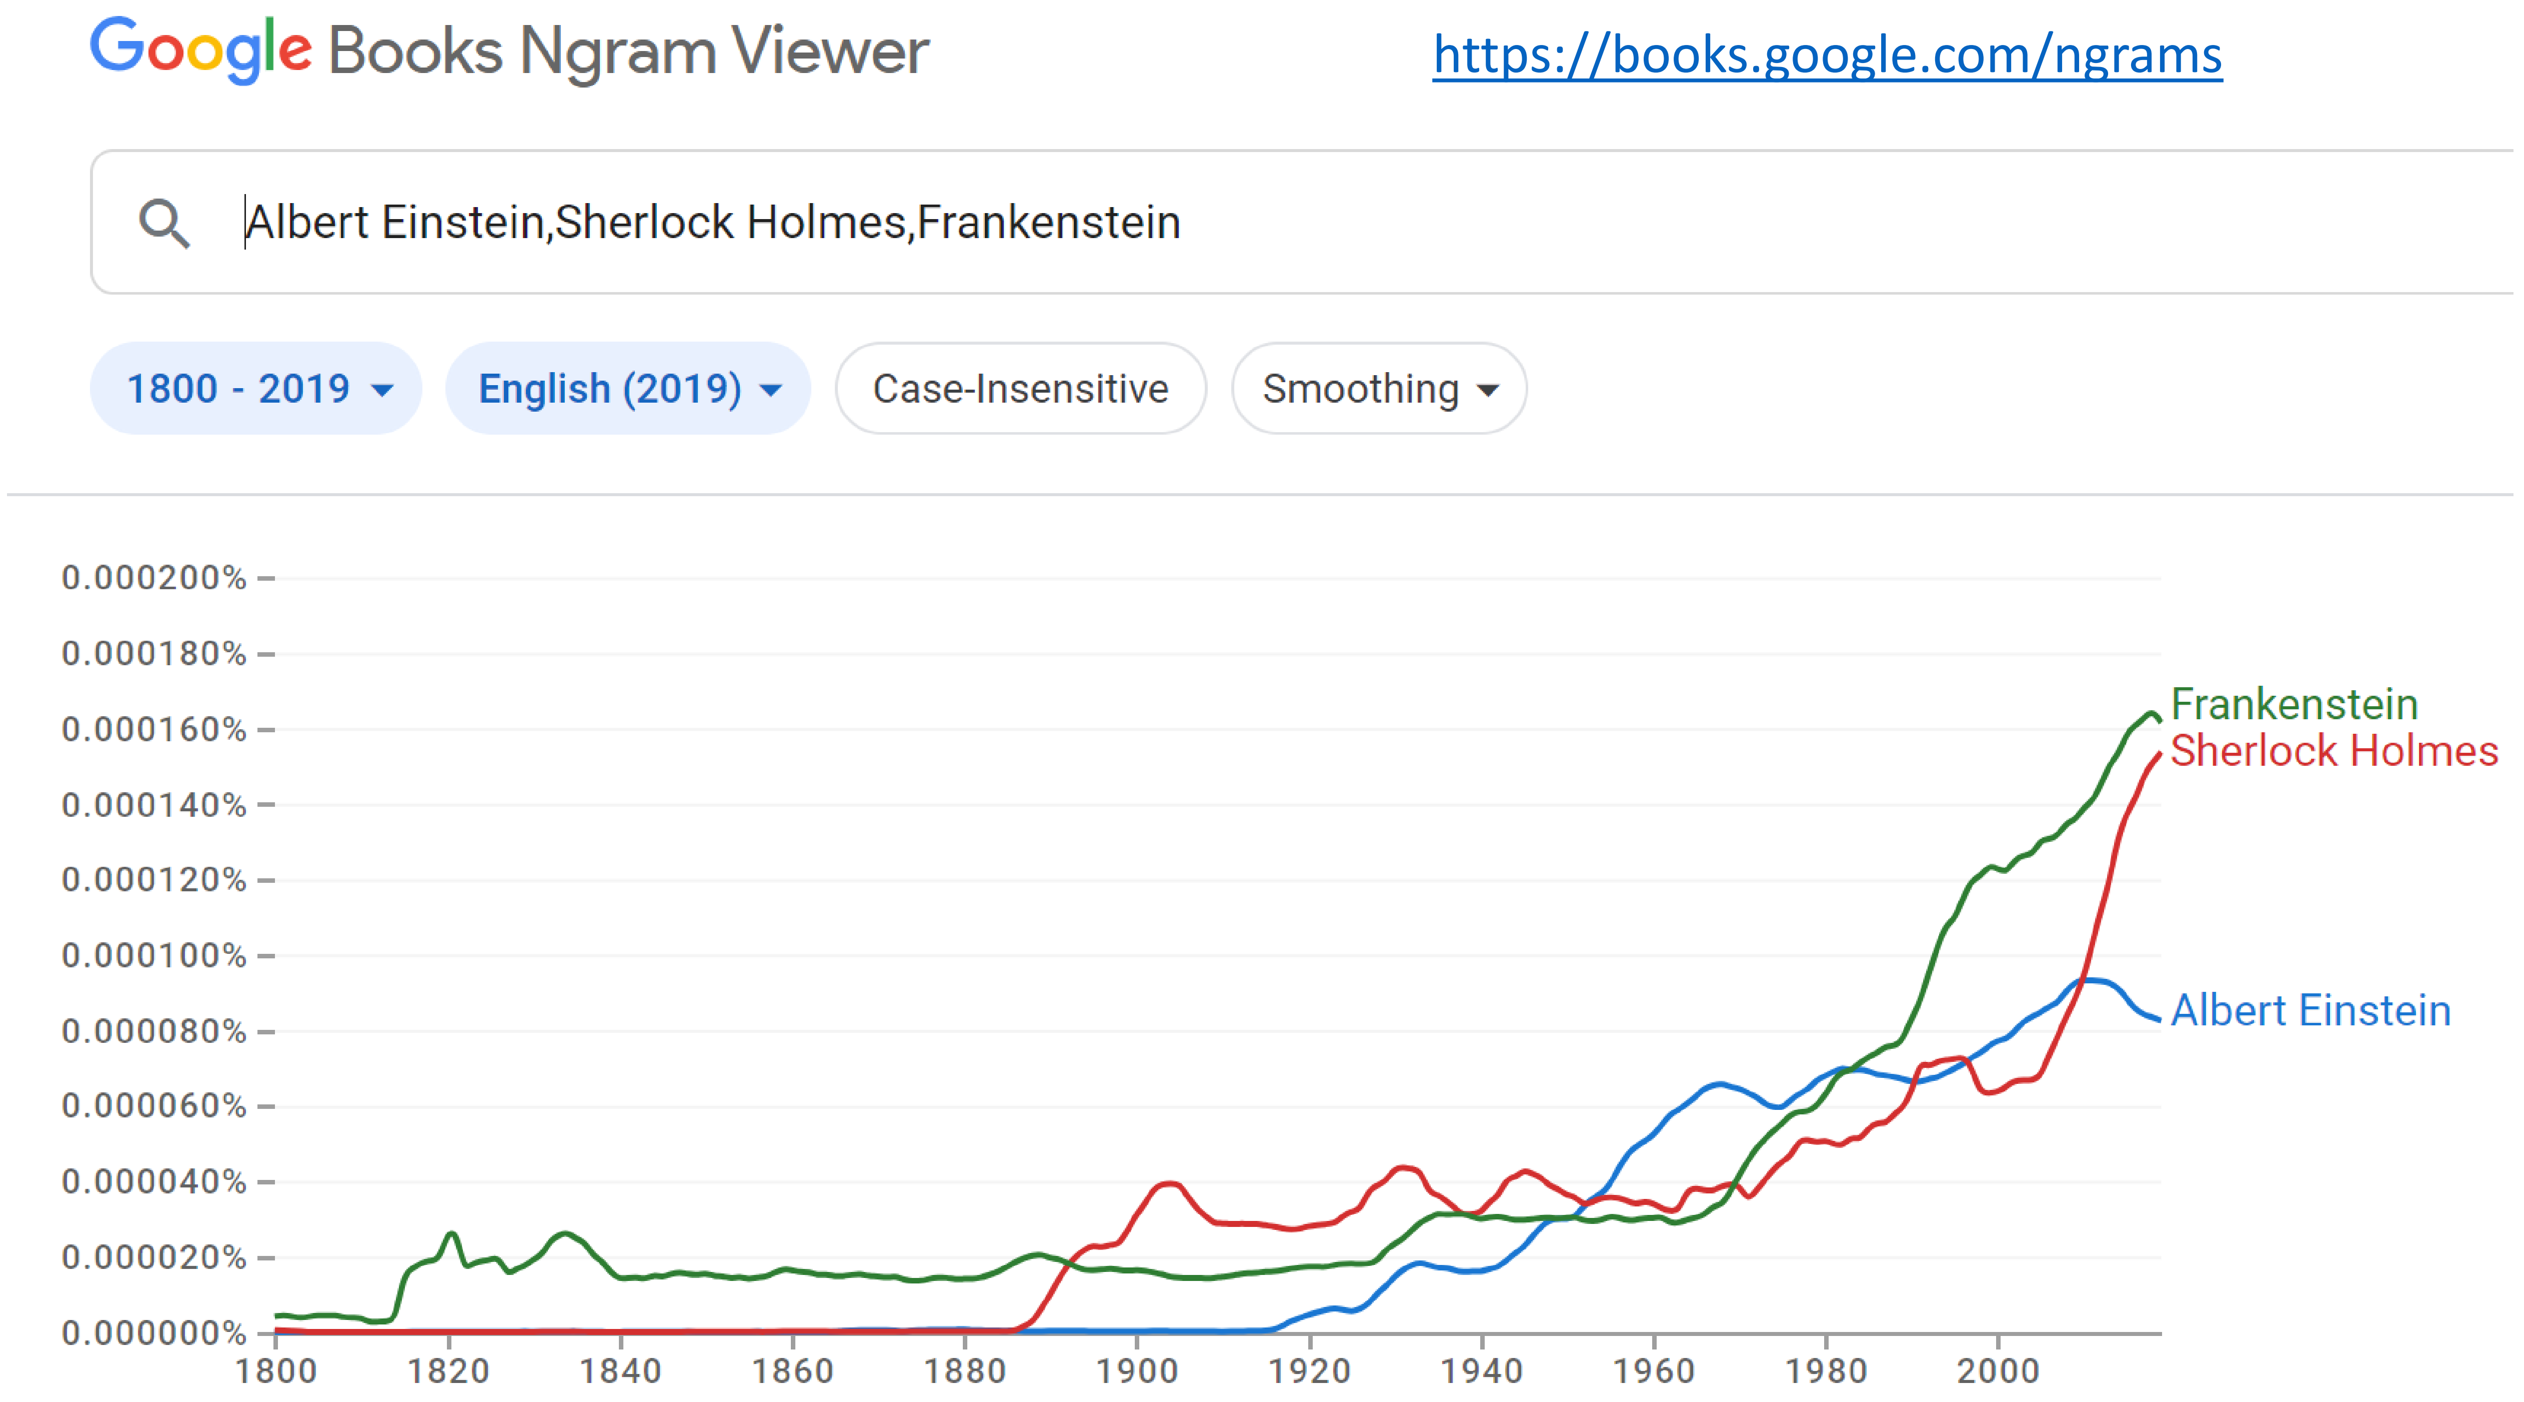
Google Books Ngram Viewer
https://books.google.com/ngrams
Albert Einstein,Sherlock Holmes,Frankenstein
1800 - 2019
English (2019)
Case-Insensitive
Smoothing
0.000200%
0.000180%
0.000160%
0.000140%
0.000120%
0.000100%
0.000080%
0.000060%
0.000040%
0.000020%
0.000000%
1800
1820
1840
1860
1880
1900
1920
1940
1960
1980
2000
Frankenstein
Sherlock Holmes
Albert Einstein

A Brief History of AI

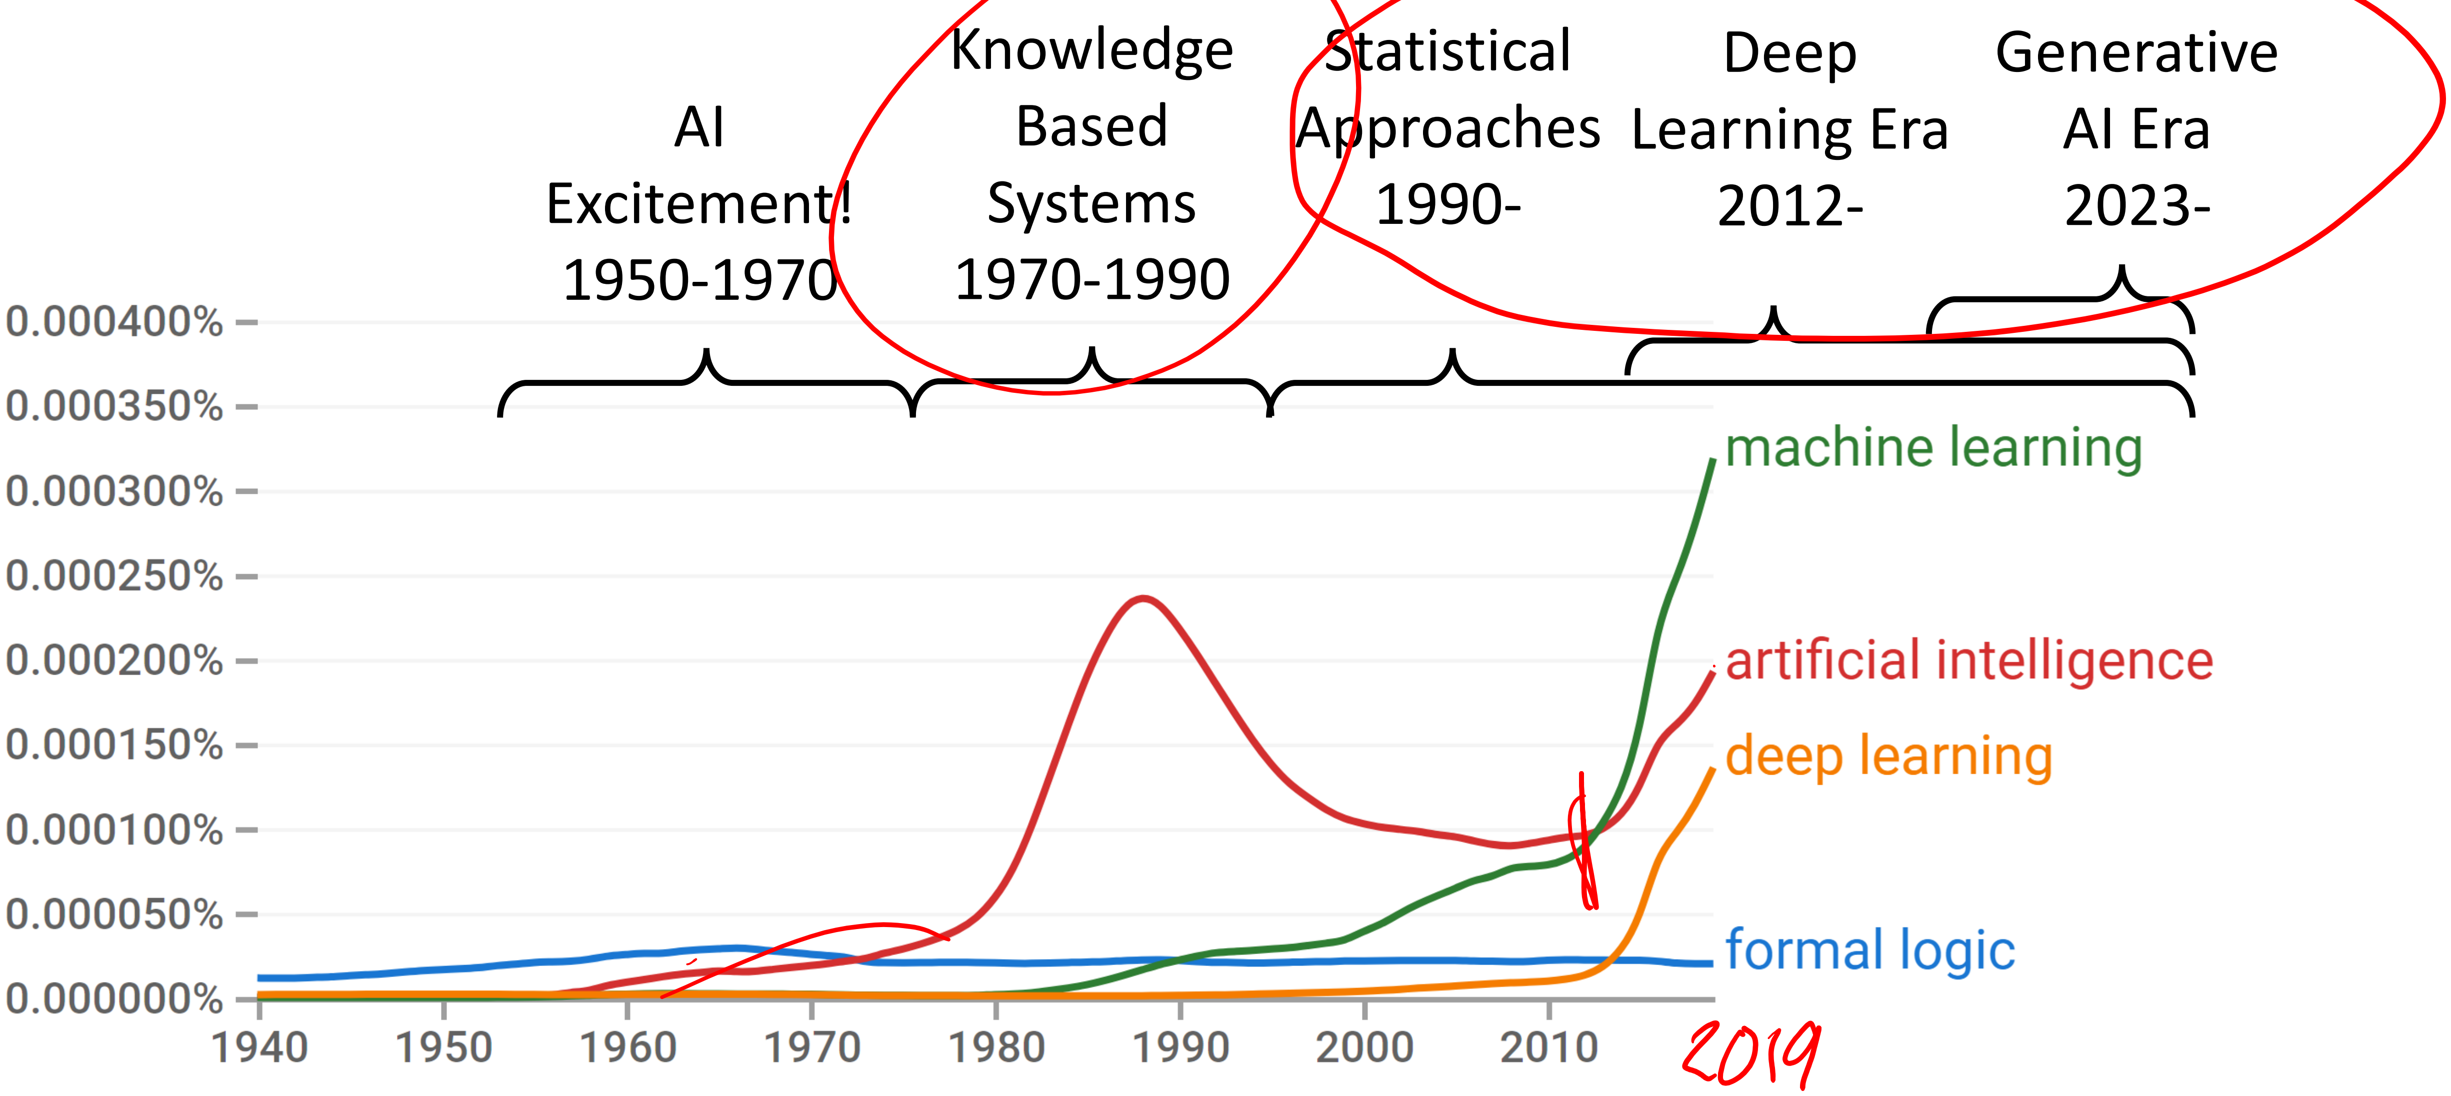

https://books.google.com/ngrams

A Brief History of AI

1940-1950: Early days

- 1943: McCulloch & Pitts: Boolean circuit model of brain
- 1950: Turing's "Computing Machinery and Intelligence"

1950—70: Excitement: Look, Ma, no hands!

- 1950s: Early AI programs, including Samuel's checkers program, Newell & Simon's Logic Theorist, Gelernter's Geometry Engine
- 1956: Dartmouth meeting: "Artificial Intelligence" adopted

1970—90: Knowledge-based approaches

- 1969—79: Early development of knowledge-based systems
- 1980—88: Expert systems industry booms
- 1988—93: Expert systems industry busts: "AI Winter"

1990—: Statistical approaches

- Resurgence of probability, focus on uncertainty
- General increase in technical depth
- Agents and learning systems… "AI Spring"?

2012—: Deep learning

- 2012: ImageNet & AlexNet

2023—: Generative AI

- 2023: ChatGPT

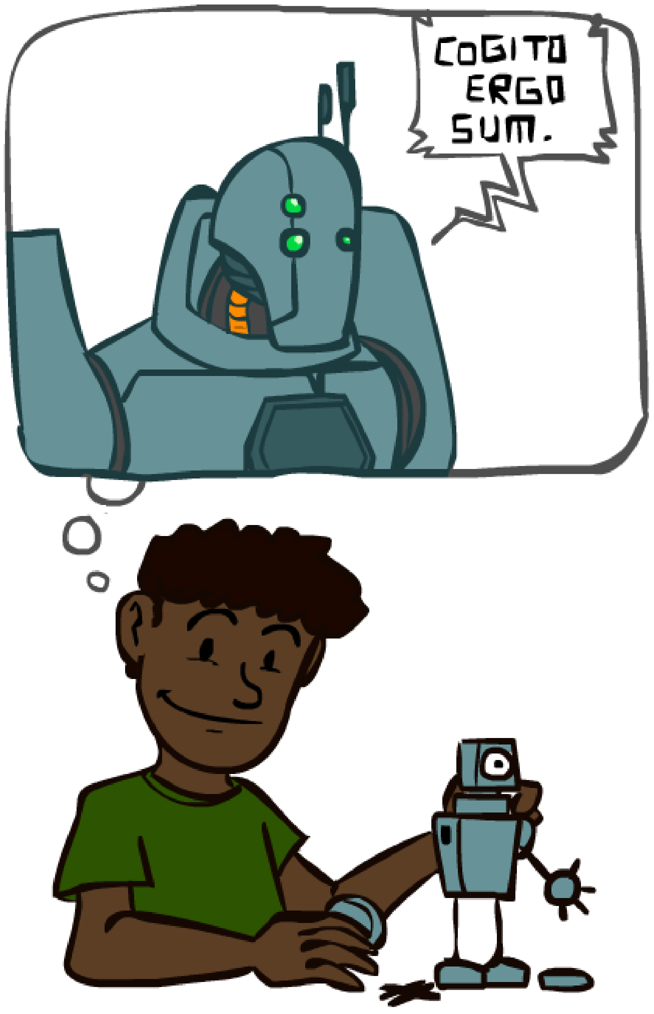

Images: ai.berkeley.edu

ML Applications?

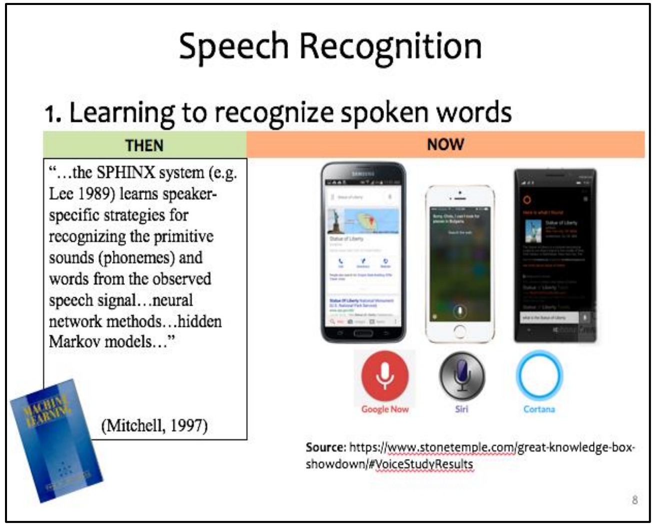

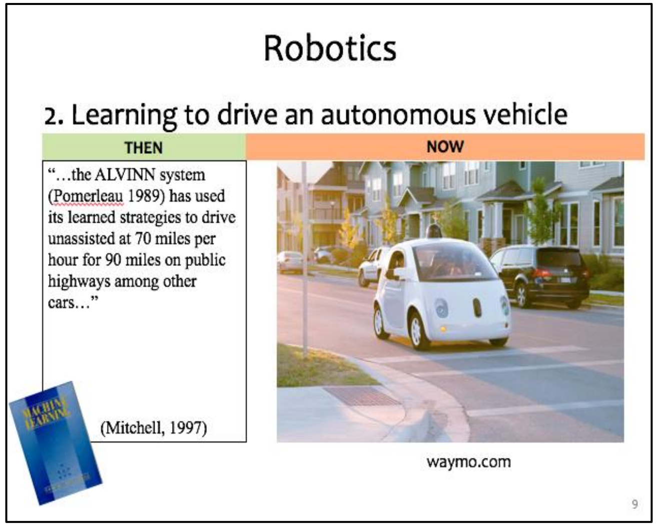

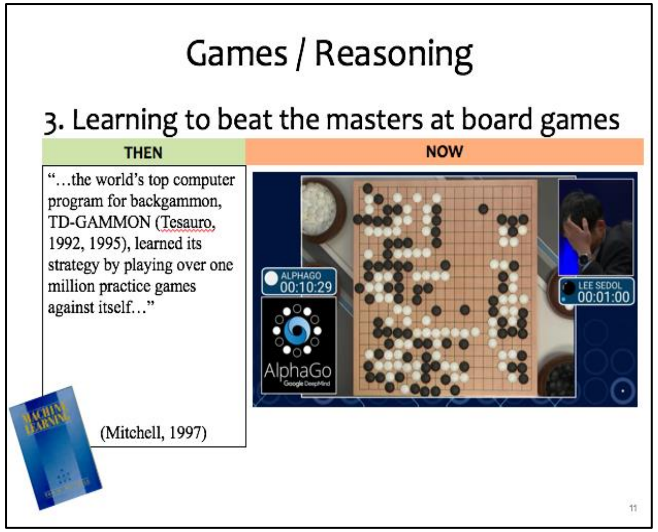

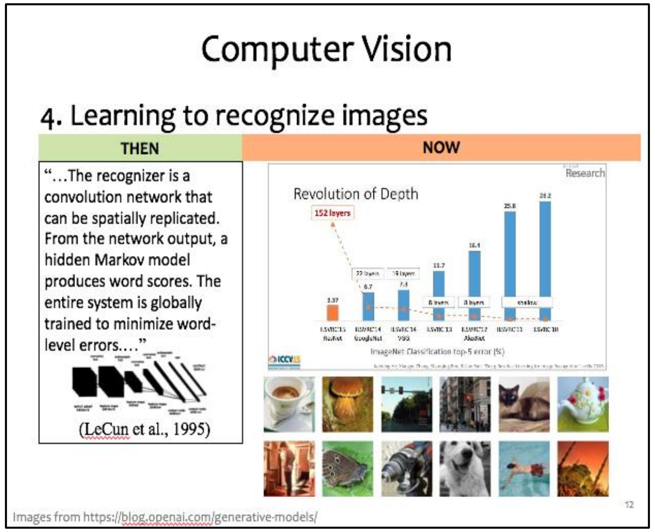

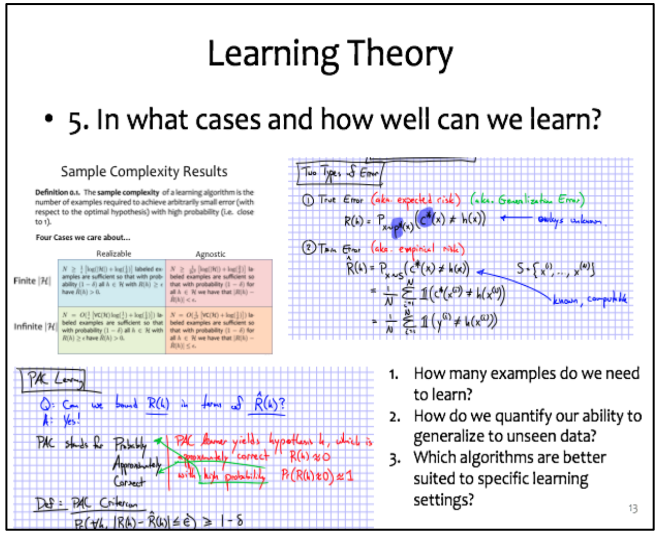

Machine Learning and Statistics

Statistics is also about learning from data

Statistics has been around from much longer!

What's the difference?

Until the mid 1990s:

Statistics:

- A branch of mathematics
- Emphasized rigor, correctness, provable properties ("is it correct?")
- Was not very concerned with scaling

Machine Learning:

- A branch of Computer Science / AI
- Focus on heuristics, making things work in practice ("does it work?")
- Not much awareness of statistical theory

Slide: CMU ML, Roni Rosenfeld

As you walk in

Welcome!

1) Help draw some aliens for our dataset today!
   - See table up front
   - Just stick figures, nothing quite this fancy →

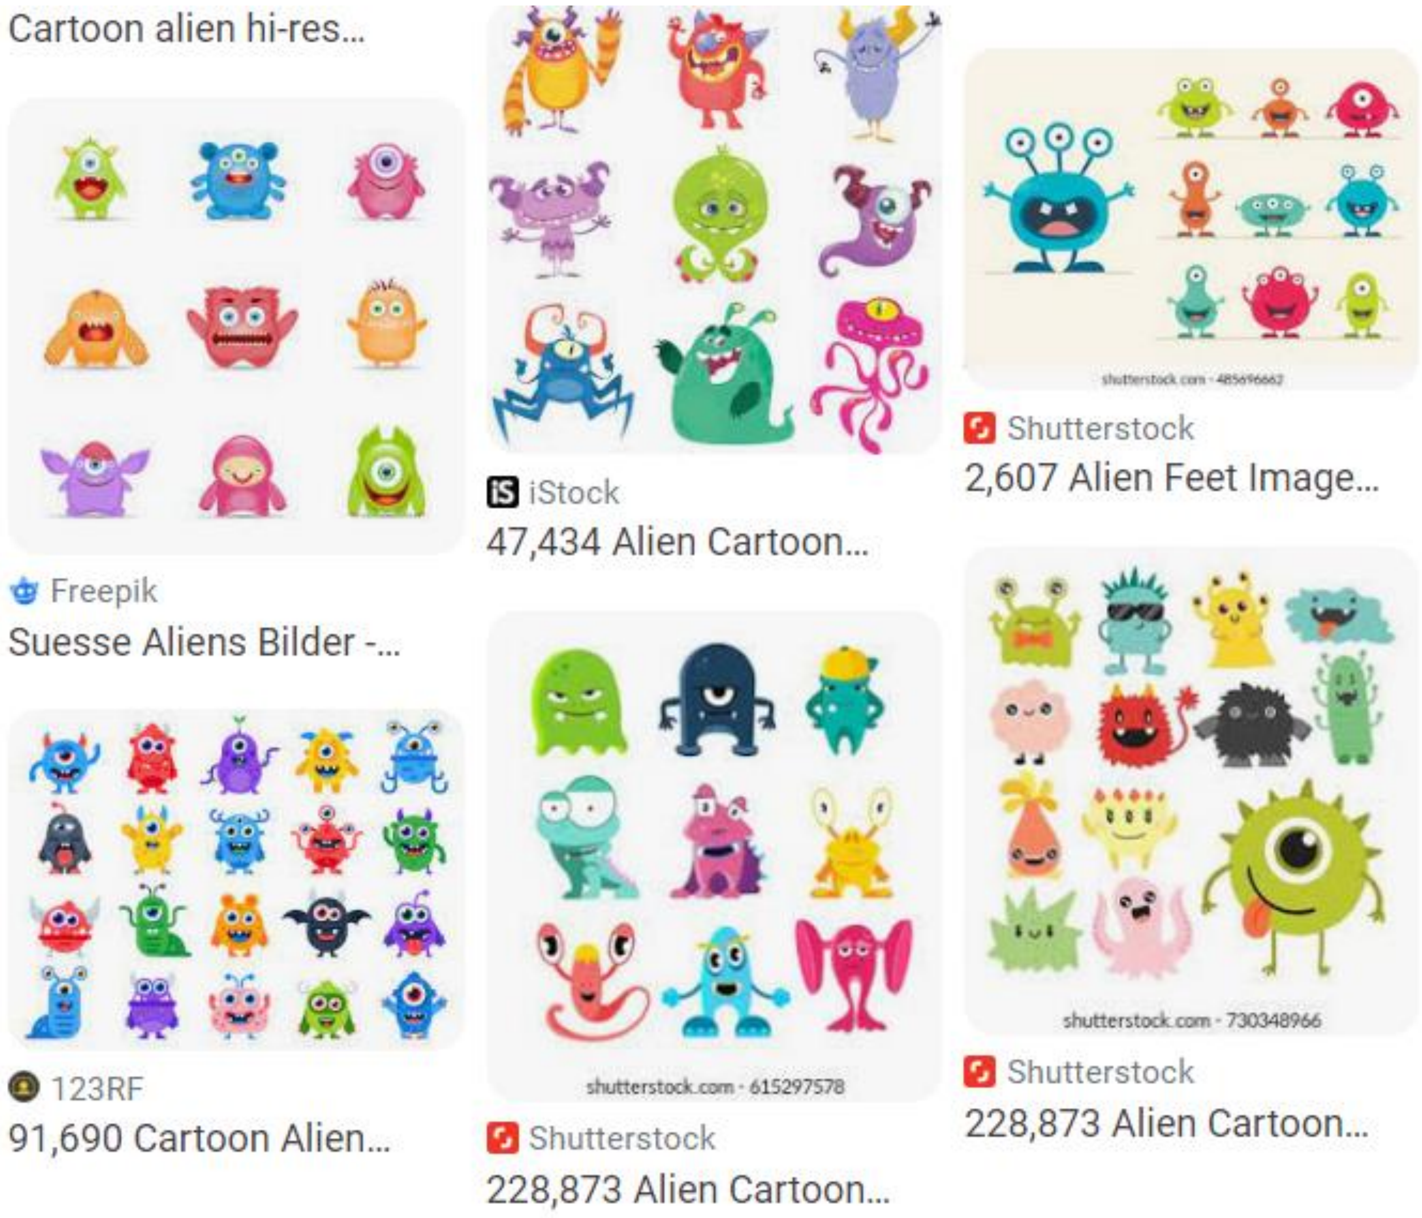

Exercise: Human-defined Feature Space

Try to organize data on a 2-D coordinate plot

to win the following game:

Select three students: A, B, C

1. Student A draws a new alien and hands it to student B
2. Student B thinks about where to plot it and comes up with a 2-D coordinate, (x, y)
3. Student C looks at the coordinate and the plot (but not the drawing from A) and draws a new alien

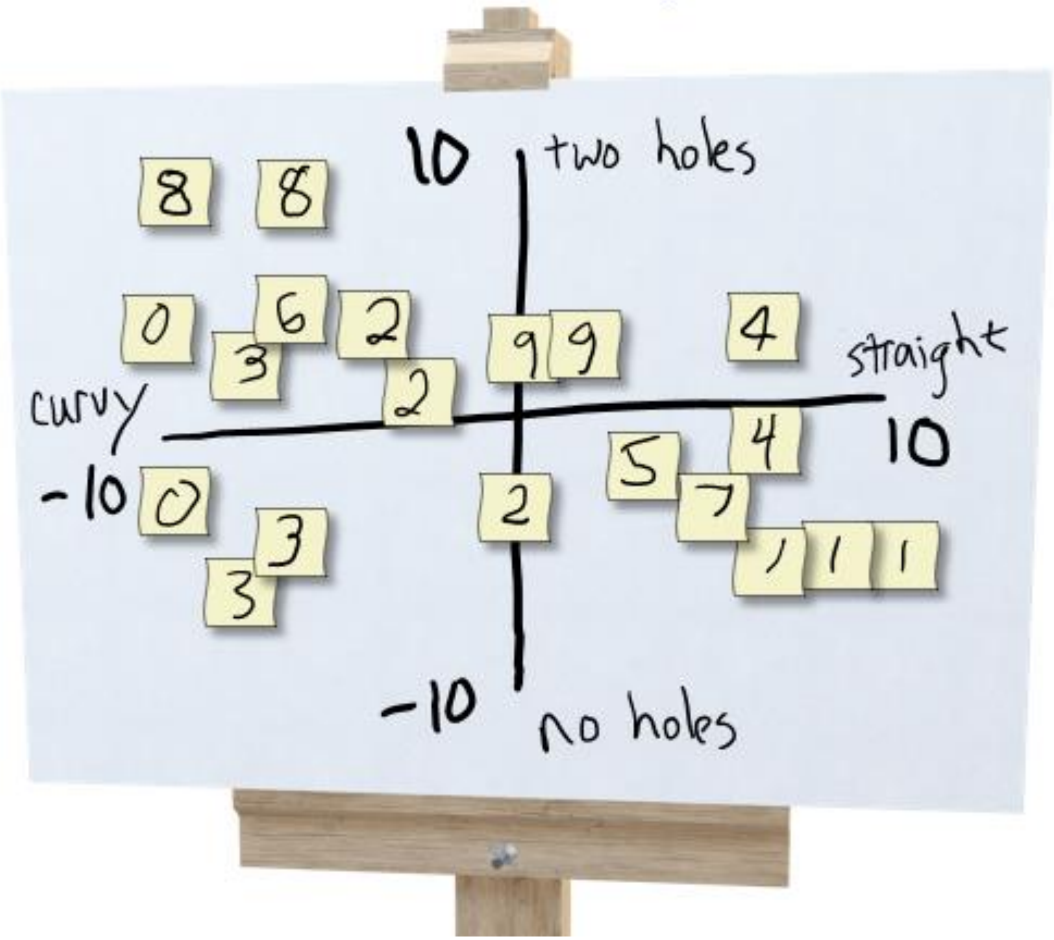

Exercise: Human-defined Feature Space

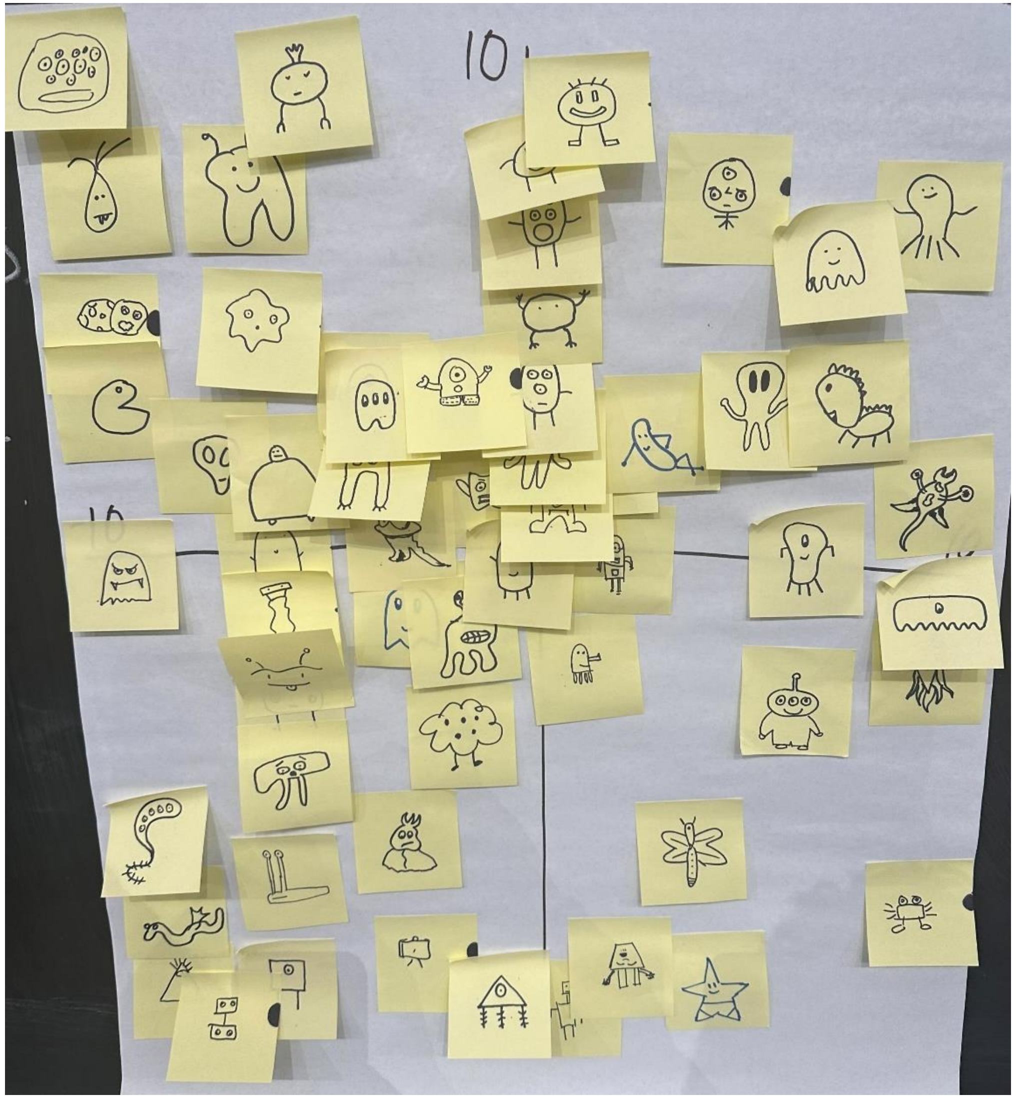

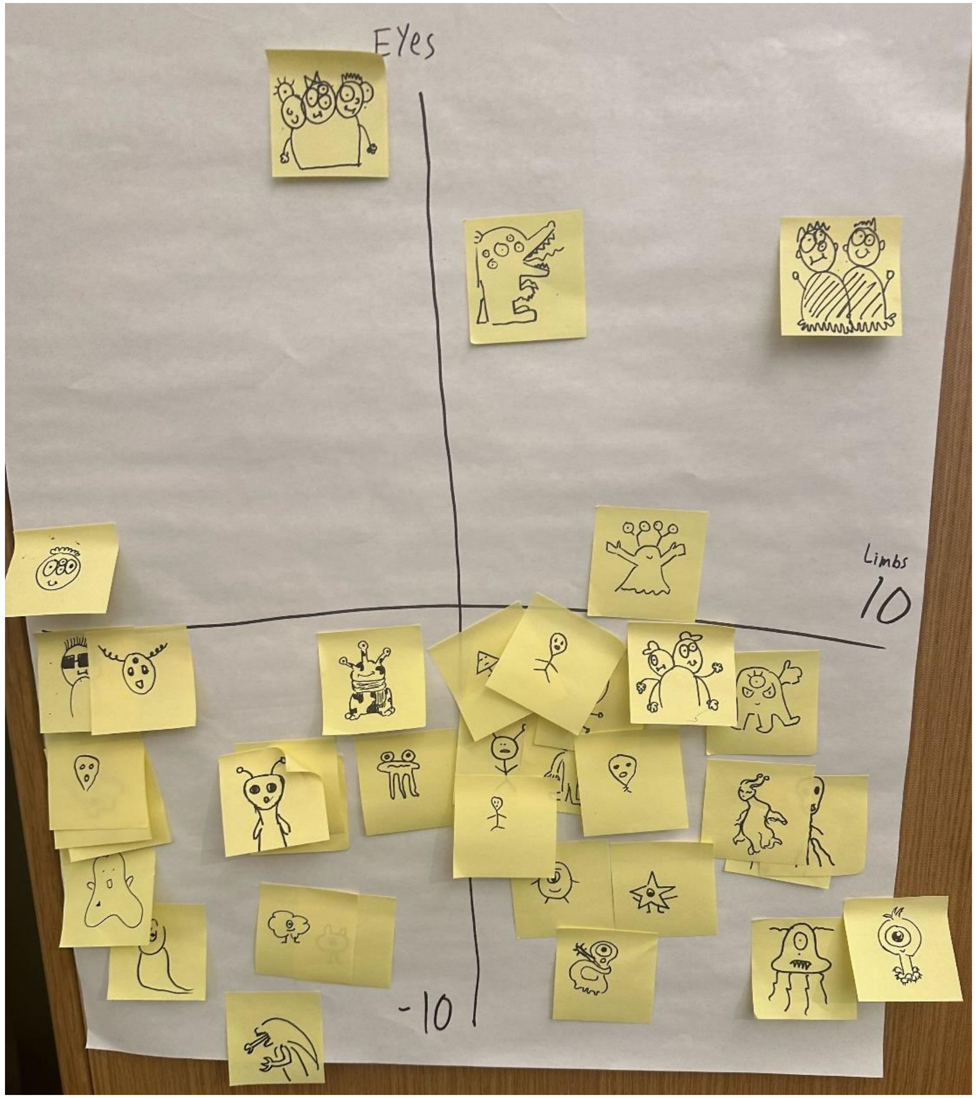

Exercise: Human-defined Feature Space

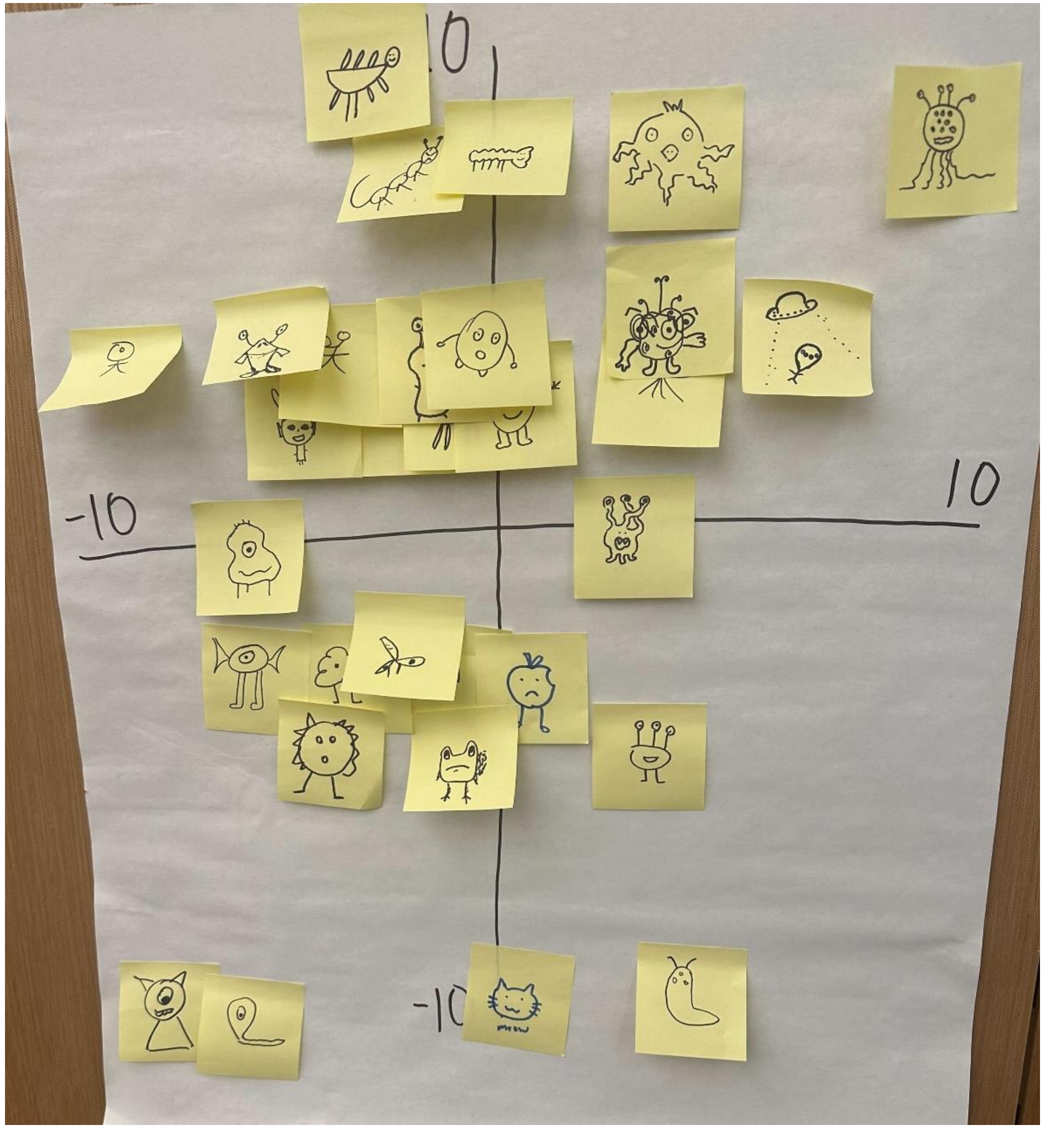

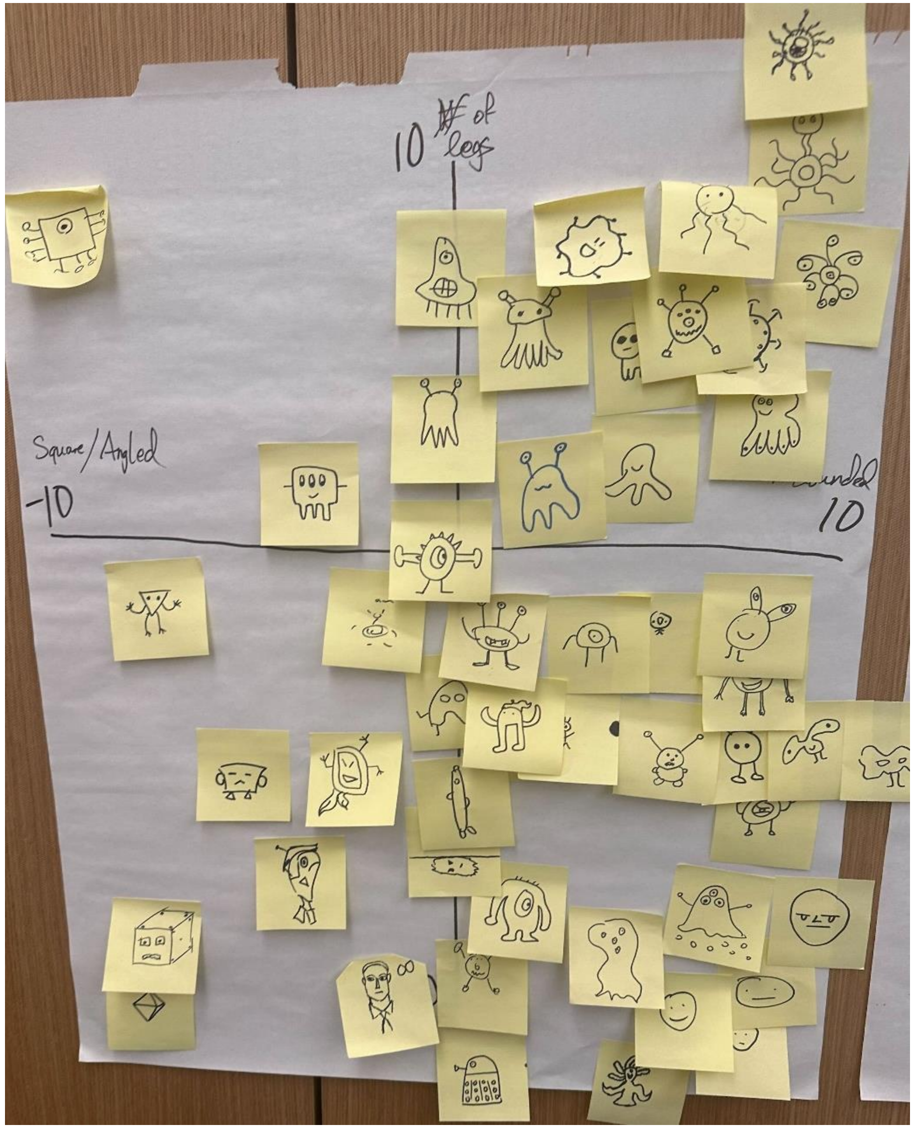

Exercise: Human-defined Feature Space

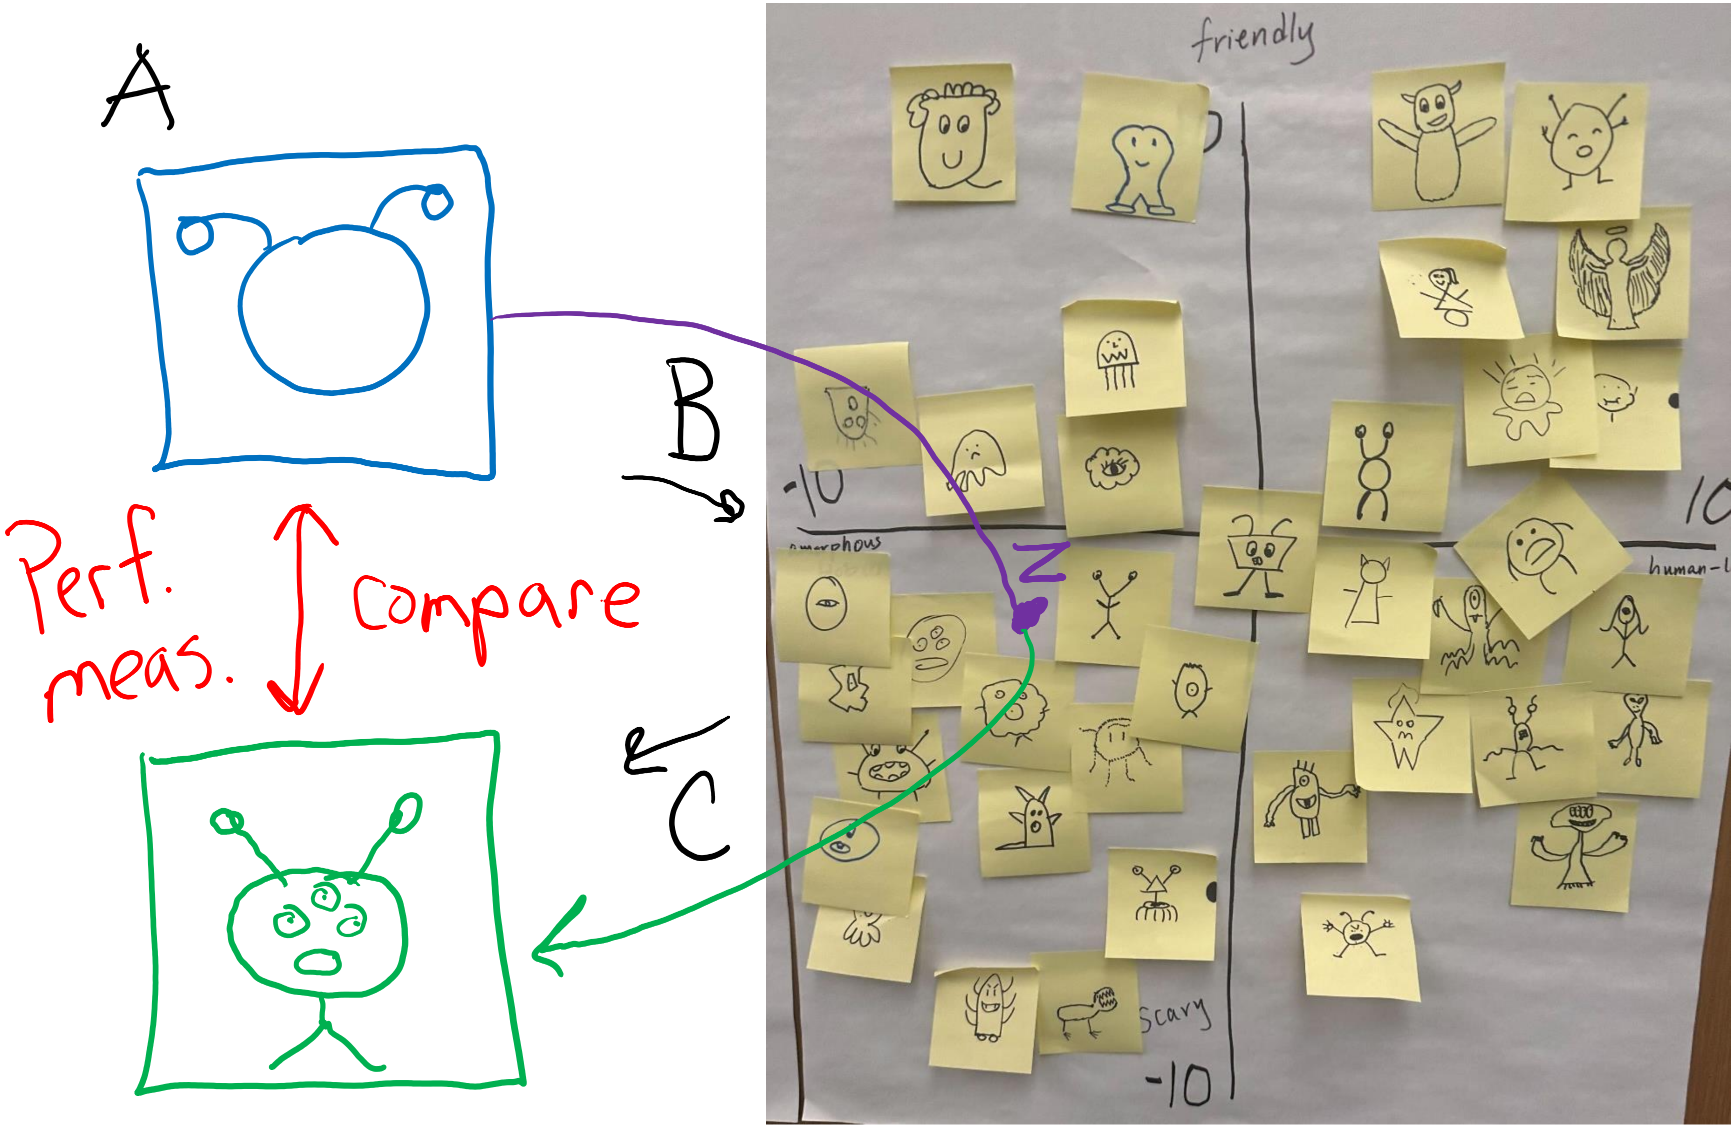

$z = (-5, -2)$

Poll 1

What features did you use?

Poll 1

What features did you use?

Learning to Organize Data

$\min \|\vec{x} - \hat{y}\|_2^2$

Neural networks can learn to organization too!

Image → $\begin{bmatrix} z_1 \\ z_2 \end{bmatrix}$ → Image

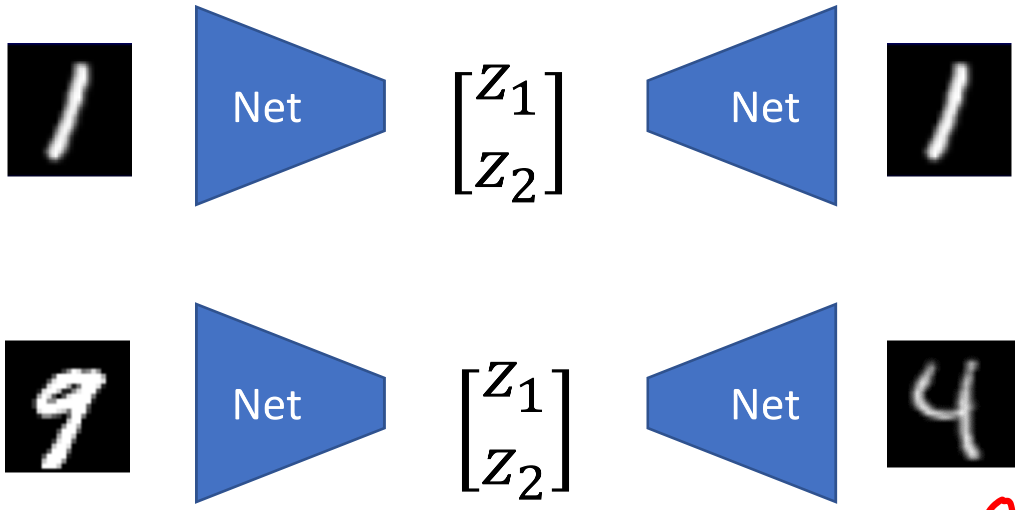

$\vec{x}$ $\vec{z}$ $\hat{y}$

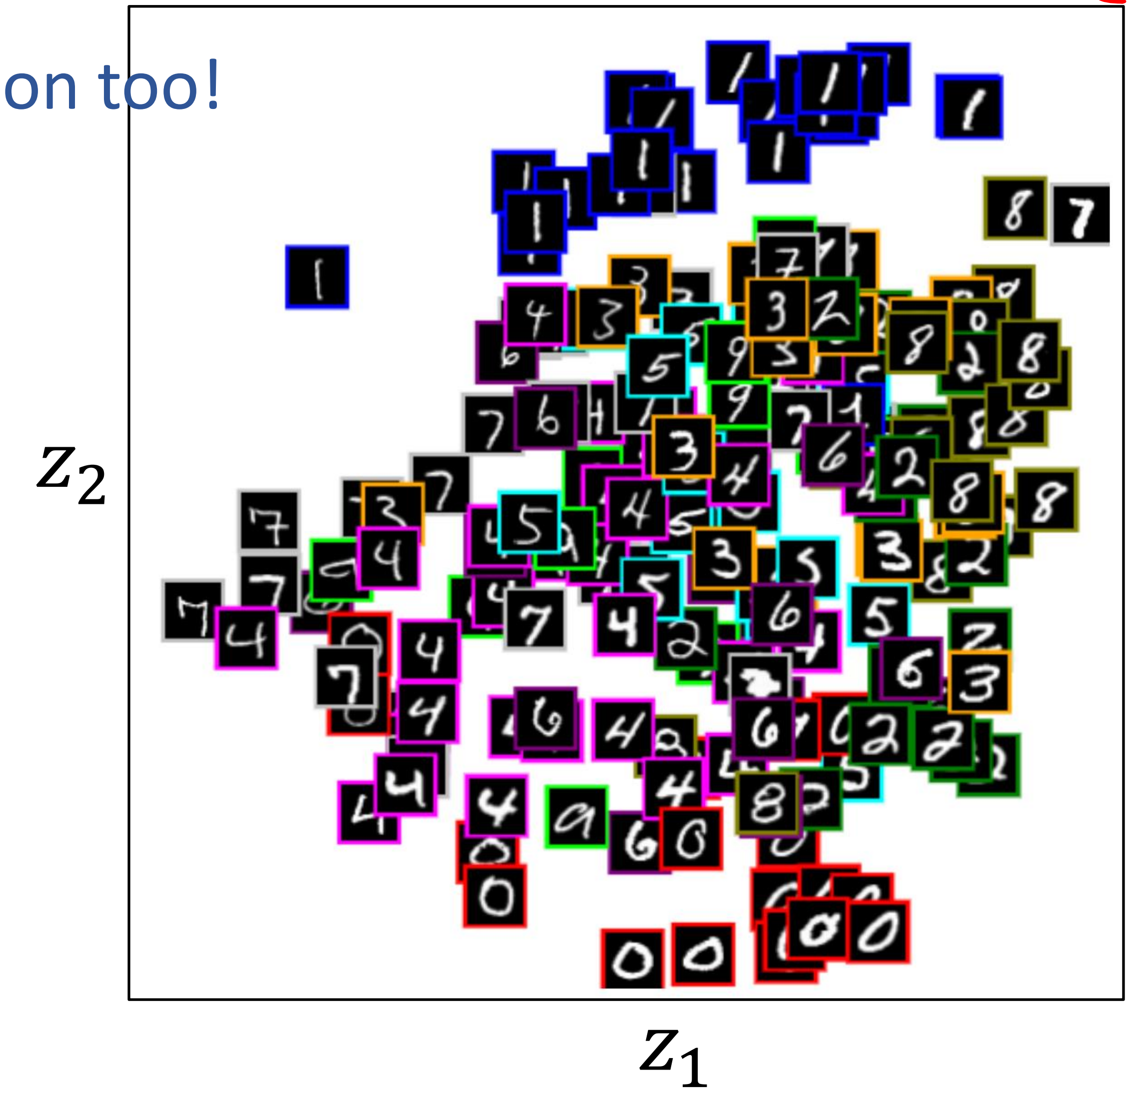

https://cs.stanford.edu/people/karpathy/convnetjs/demo/autoencoder.html

Digit Autoencoder

Demo: Using a learned feature space

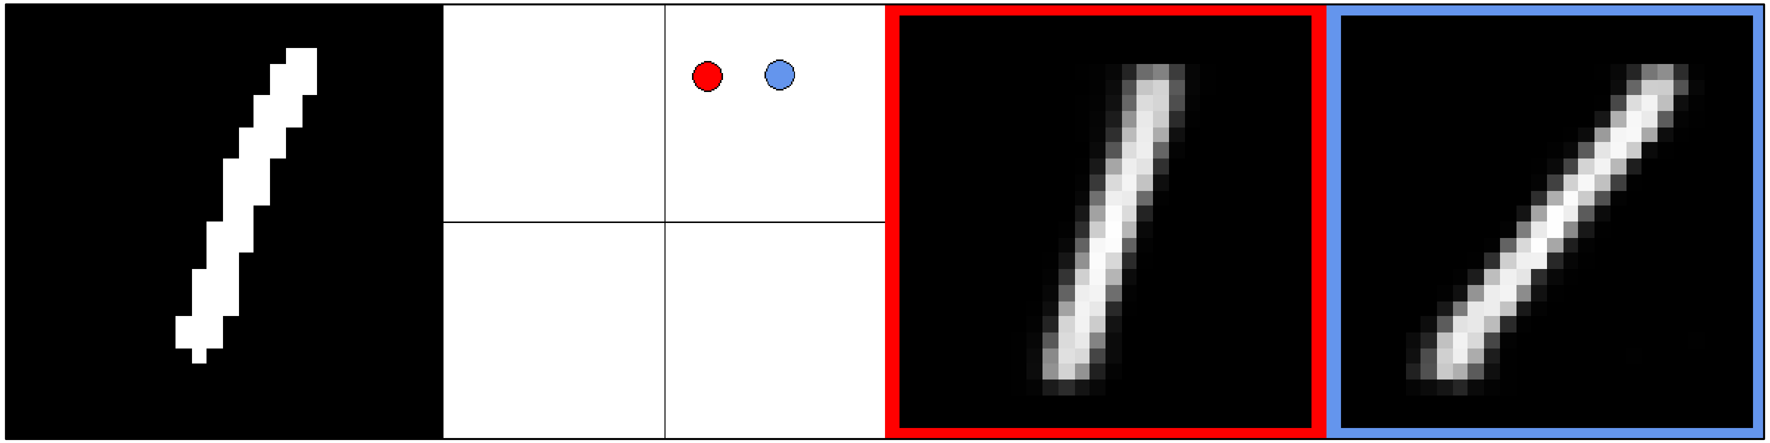

Task: Face Generation

https://thispersondoesnotexist.com/

Task: Face Generation

https://thispersondoesnotexist.com/

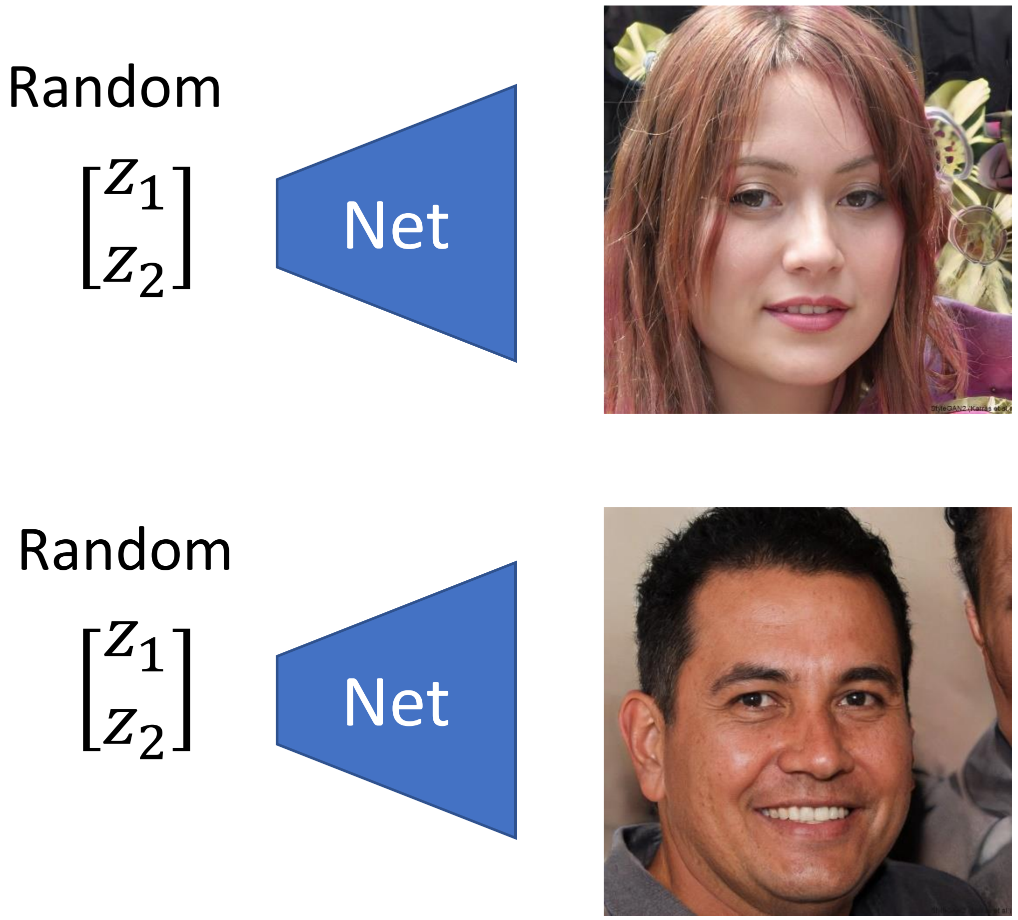

Face GAN Slider Demo

https://github.com/genforce/interfacegan

https://colab.research.google.com/github/genforce/interfacegan/blob/master/docs/InterFaceGAN.ipynb

z

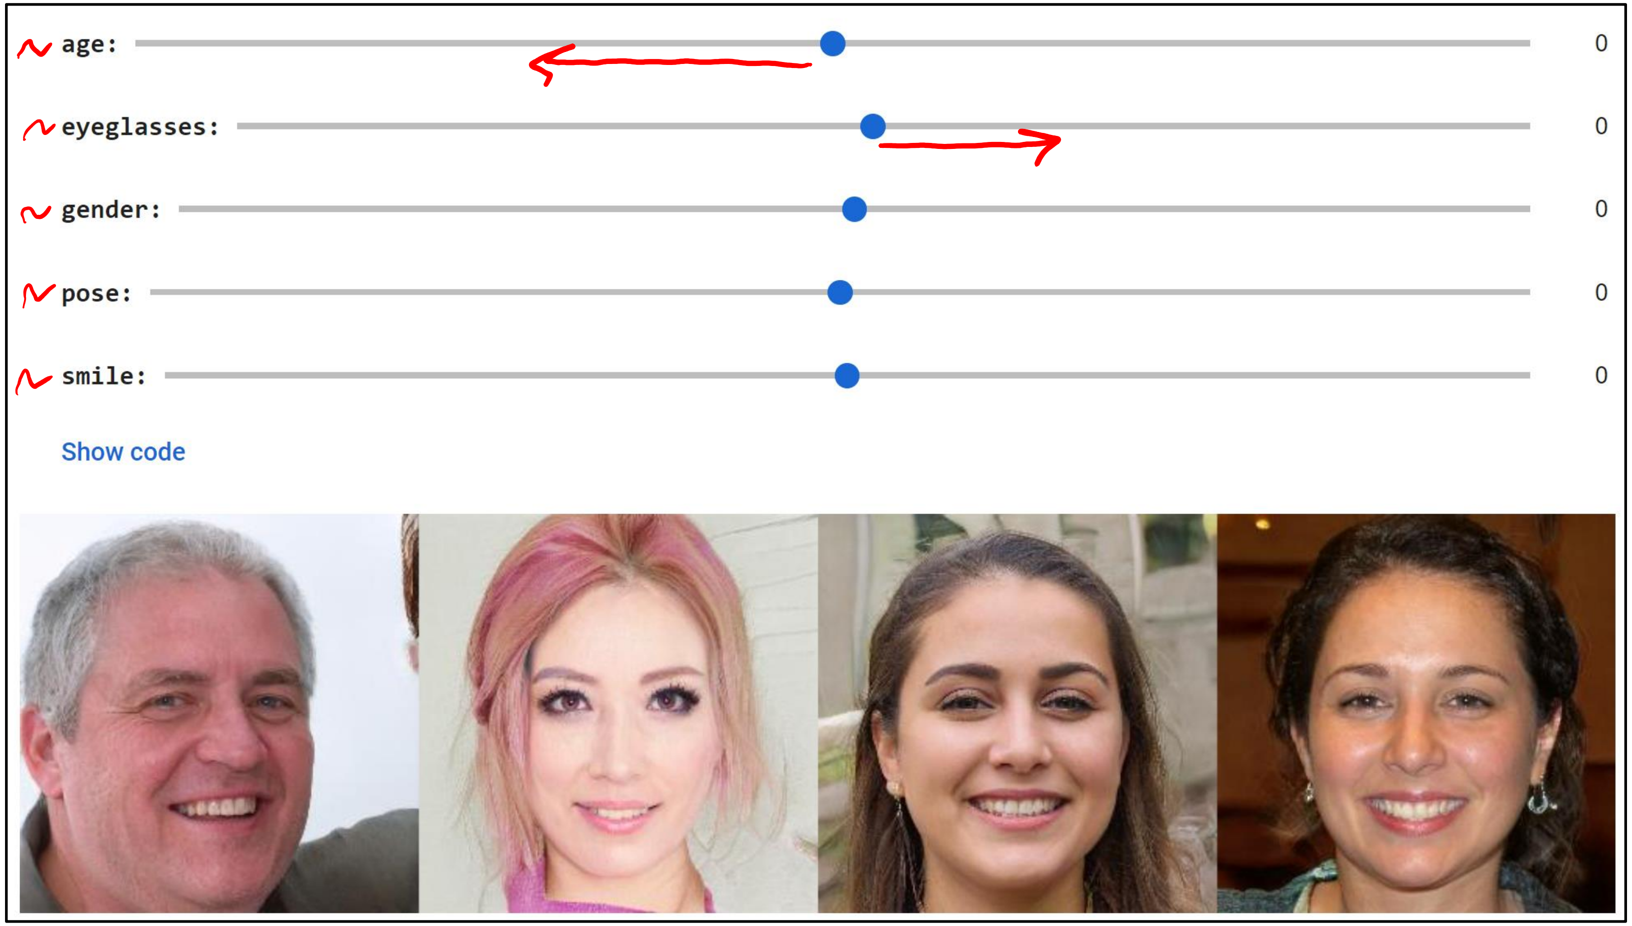

Face GAN Slider Demo

https://github.com/genforce/interfacegan

https://colab.research.google.com/github/genforce/interfacegan/blob/master/docs/InterFaceGAN.ipynb

$\tilde{z}$

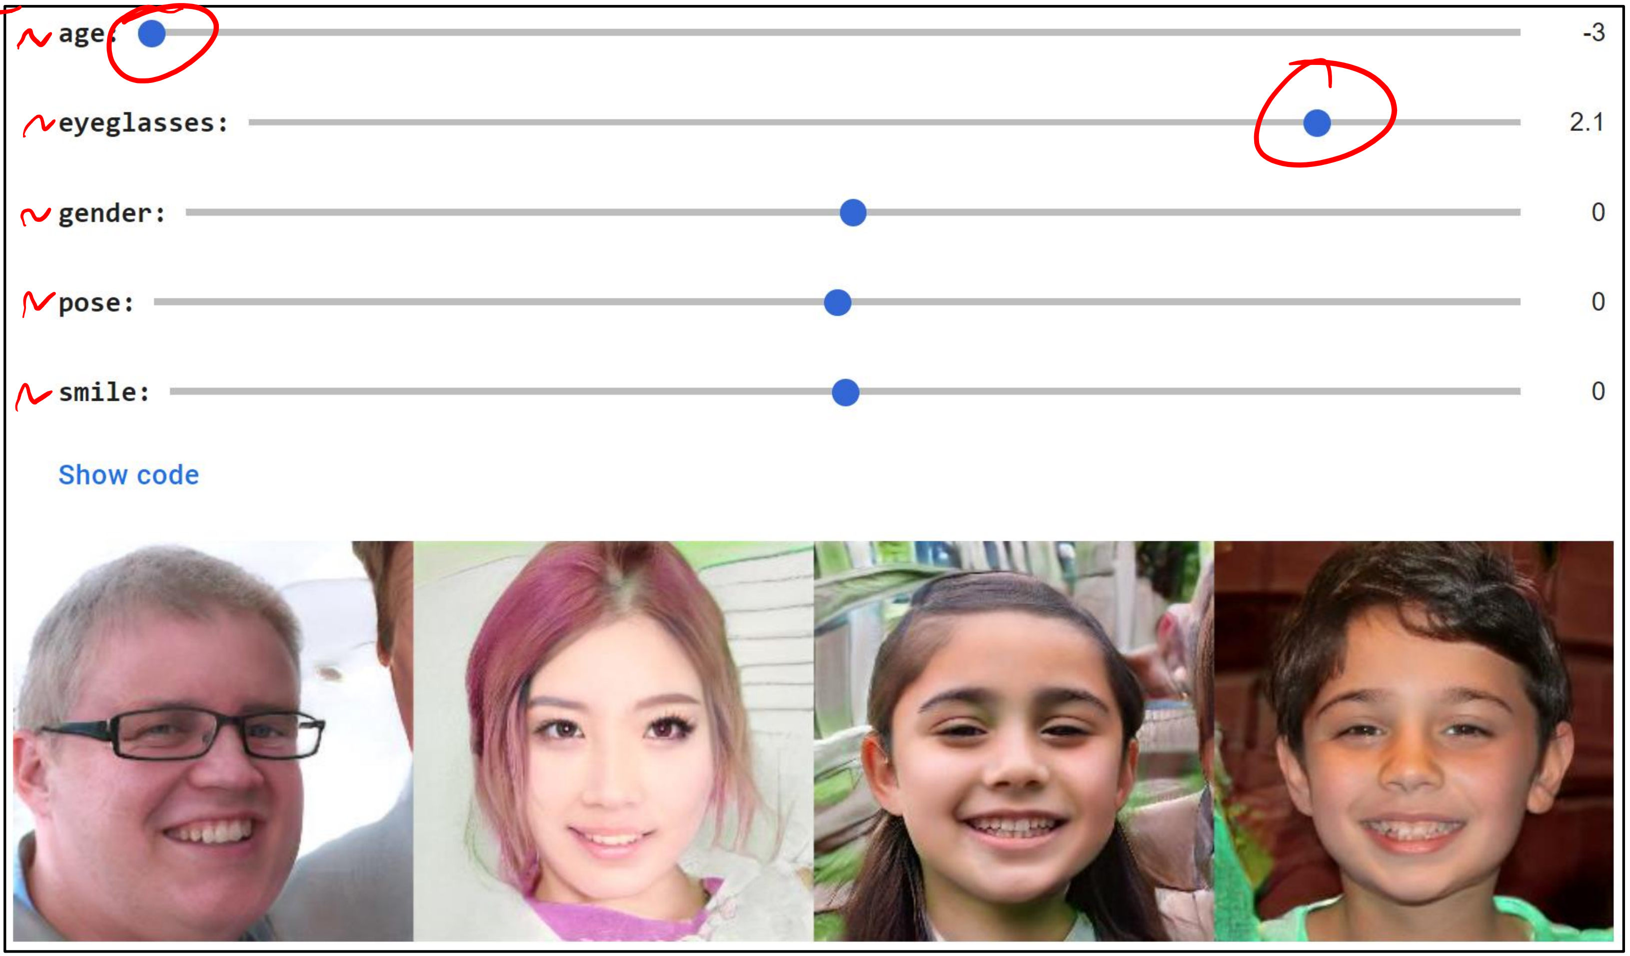

Listen Learner

$\text{audio} \longrightarrow \boxed{} \longrightarrow \text{category}$

Chris Harrison, CMU

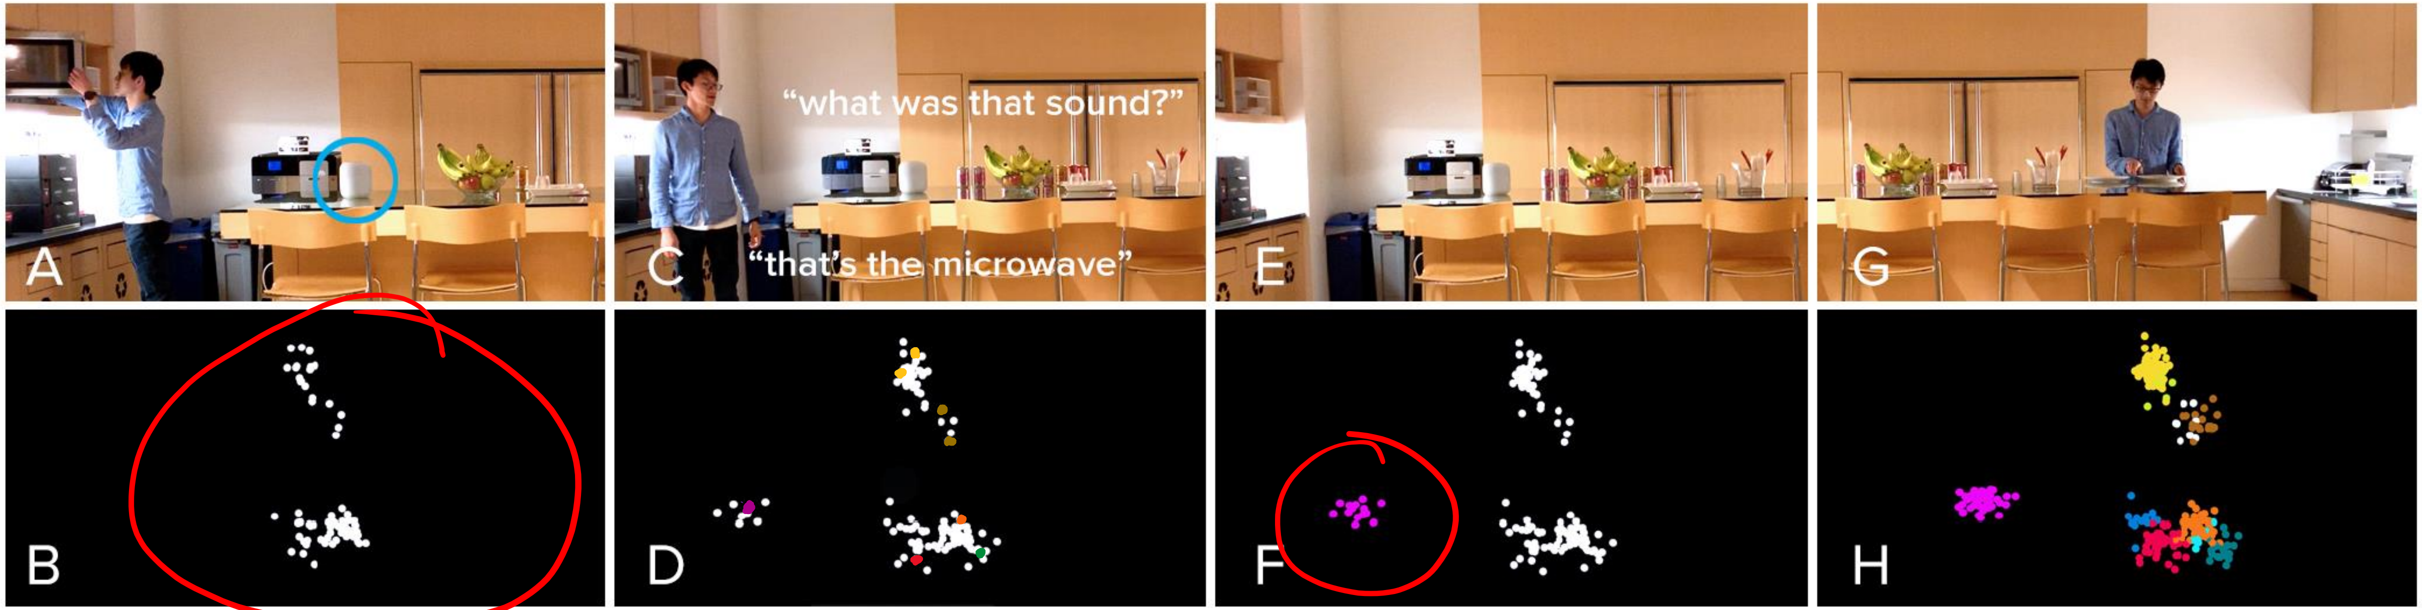

https://chrisharrison.net/index.php/Research/ListenLearner

$$\underline{\text{audio}} \longrightarrow \begin{bmatrix} z_1 \\ z_2 \end{bmatrix} \longrightarrow \underline{\text{audio}}$$

Dimensionality Reduction with Deep Learning

Hinton, Geoffrey E., and Ruslan R. Salakhutdinov.

"Reducing the dimensionality of data with neural networks."

Science 313.5786 (2006): 504-507.

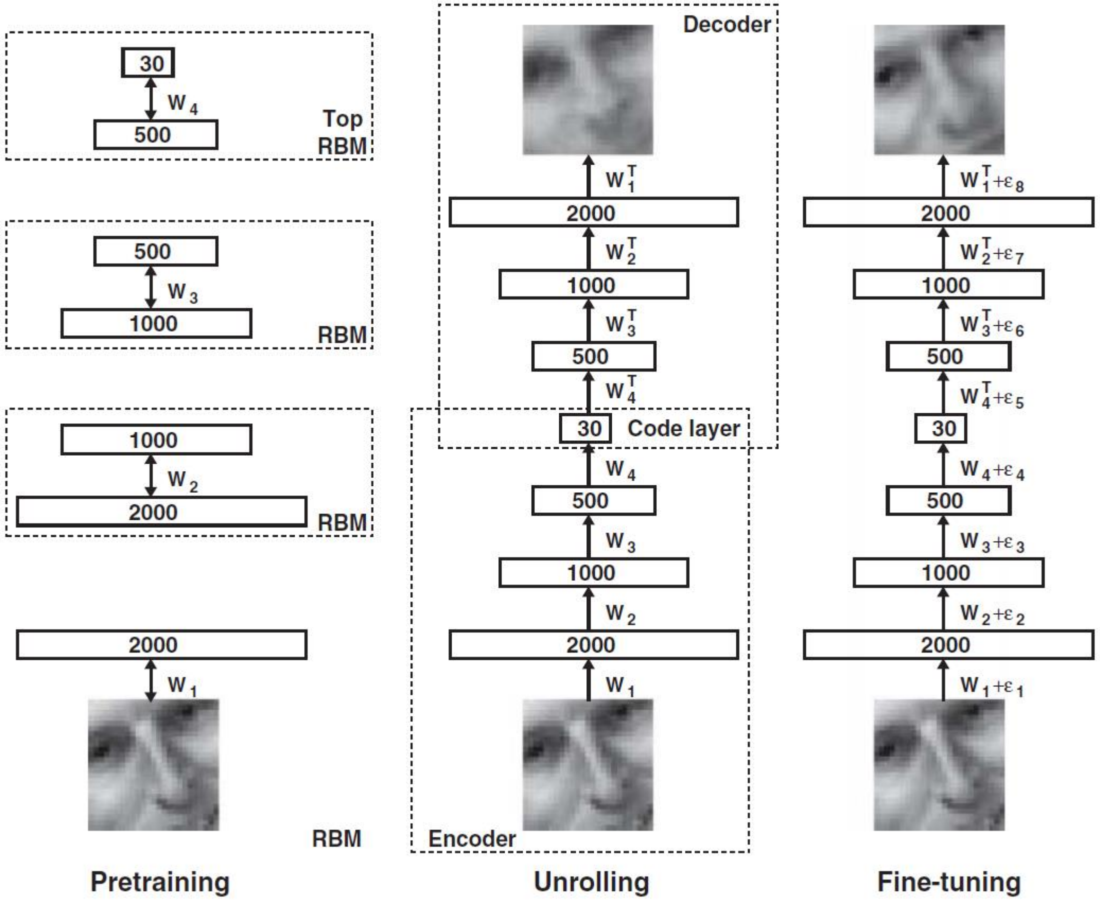

Digit Autoencoder

https://cs.stanford.edu/people/karpathy/convnetjs/demo/autoencoder.html

Input
28x28 = 784 pixels

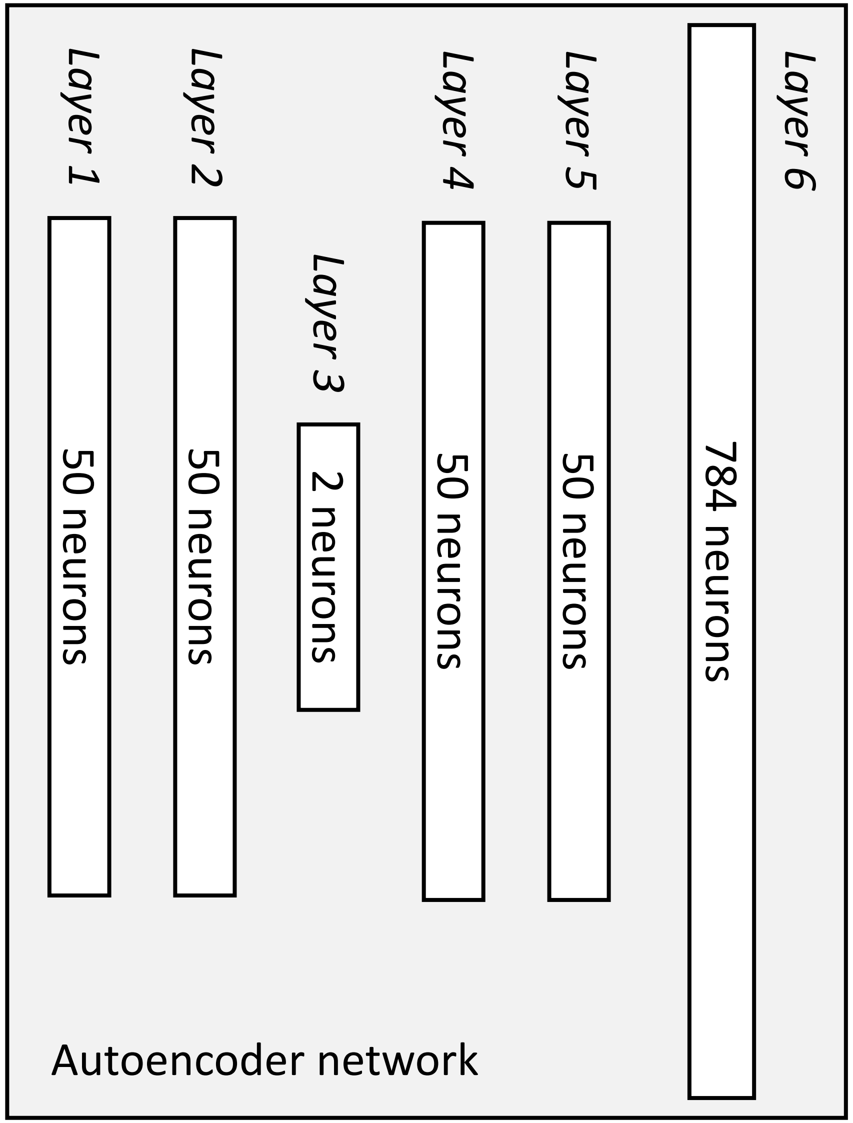

Digit Autoencoder

Math for autoencoder model

Linear layer

$$W\vec{x} + b$$

ReLU layer

$$\max(\vec{z}, 0) \qquad \begin{bmatrix} \max(z_1, 0) \\ \max(z_2, 0) \\ \vdots \\ \max(z_K, 0) \end{bmatrix}$$

Tanh layer

$$\tanh(\vec{z}) \qquad \begin{bmatrix} \tanh(z_1) \\ \tanh(z_2) \\ \vdots \\ \tanh(z_K) \end{bmatrix}$$