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 →

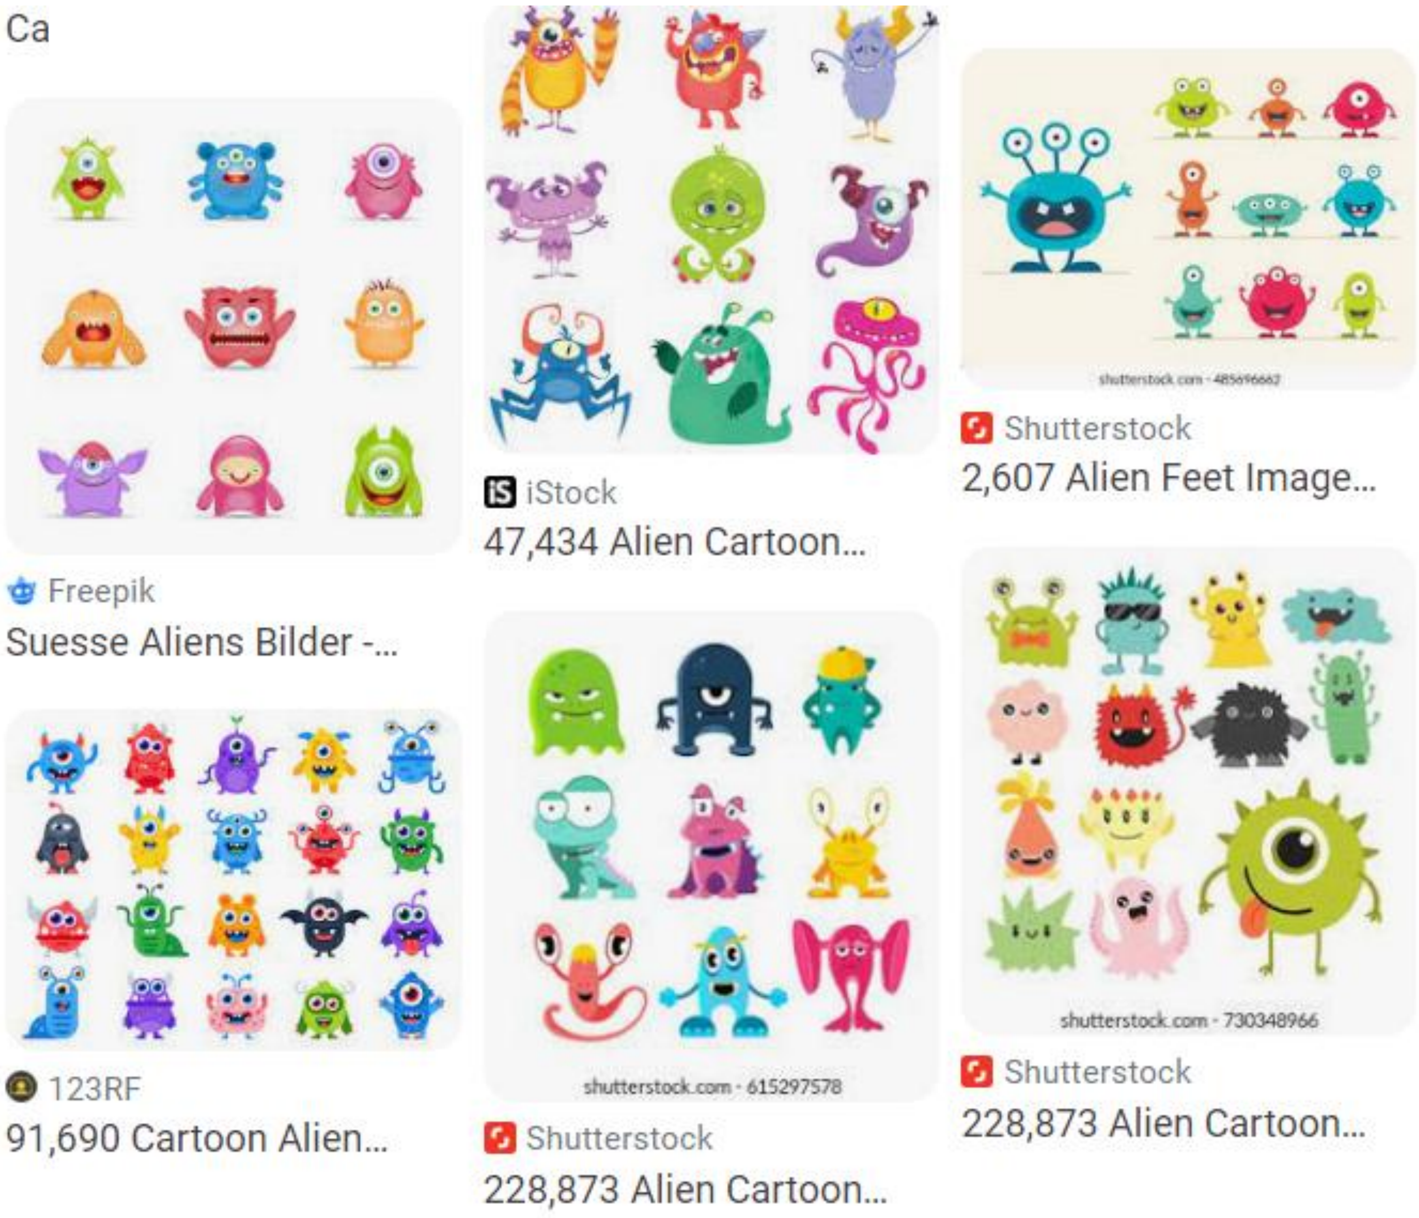

10-315
Introduction to ML

Instructor: Pat Virtue

Today

Course team

Quick course info

What is AI and ML?

ML: Input/Ouput Tasks

Example Input/Output Tasks

- Autoencoder (Aliens)

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

Course Team

Instructor

Pat Virtue
pvirtue

Best way to contact Pat:

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

(Not email ☹)

Education Associate

Nichelle Phillips
nichellp

Email Nichelle, nichellp@andrew.cmu.edu:

- Exceptions, extensions, etc
  - (may replace with exceptions form)
- Any course logistics

Course Team

Teaching Assistants

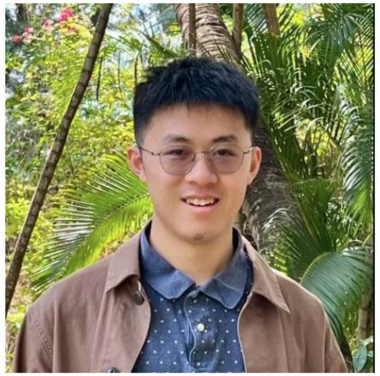

Ethan
ethanwan

AJ
ajseo

Margaret
mche

Derek
dereky

Shreya
ssridha4

Gaurika
gsawhney

Jerick
junkais

Course Team

Students!!

Course Information

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

Communication:

https://edstem.org/

E-mail Nichelle if Ed doesn't work:

nichellp@andrew.cmu.edu

Important posts coming soon!

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

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

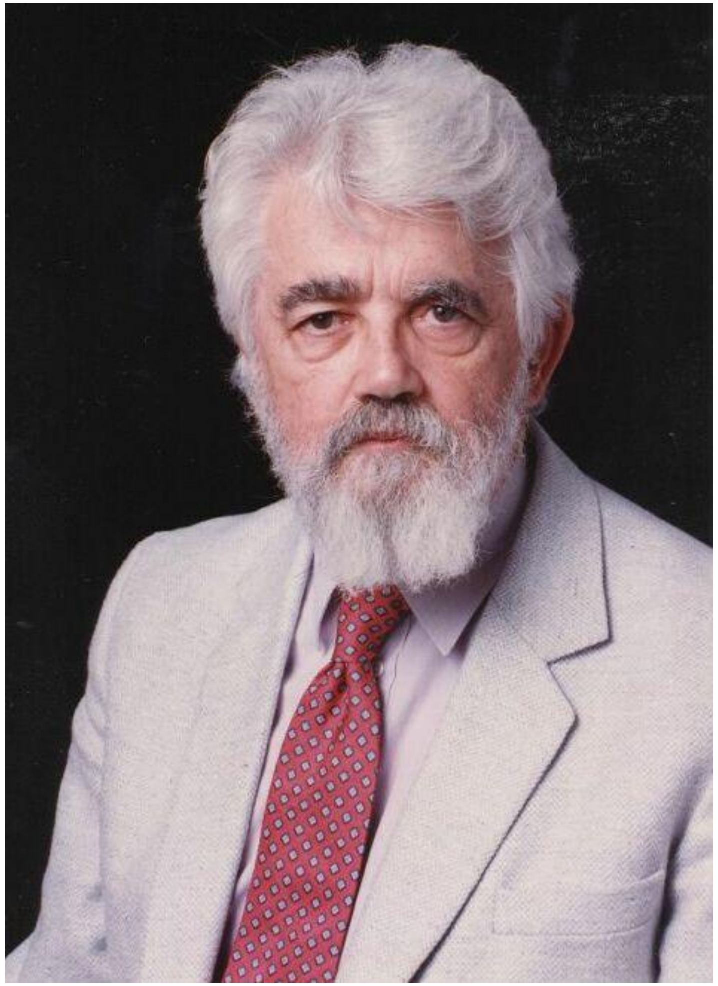

http://www-formal.stanford.edu/jmc/whatisai/whatisai.html

What is artificial Intelligence?

Why is folding clothes a research problem while robots are actively being used to assemble cars in factories?

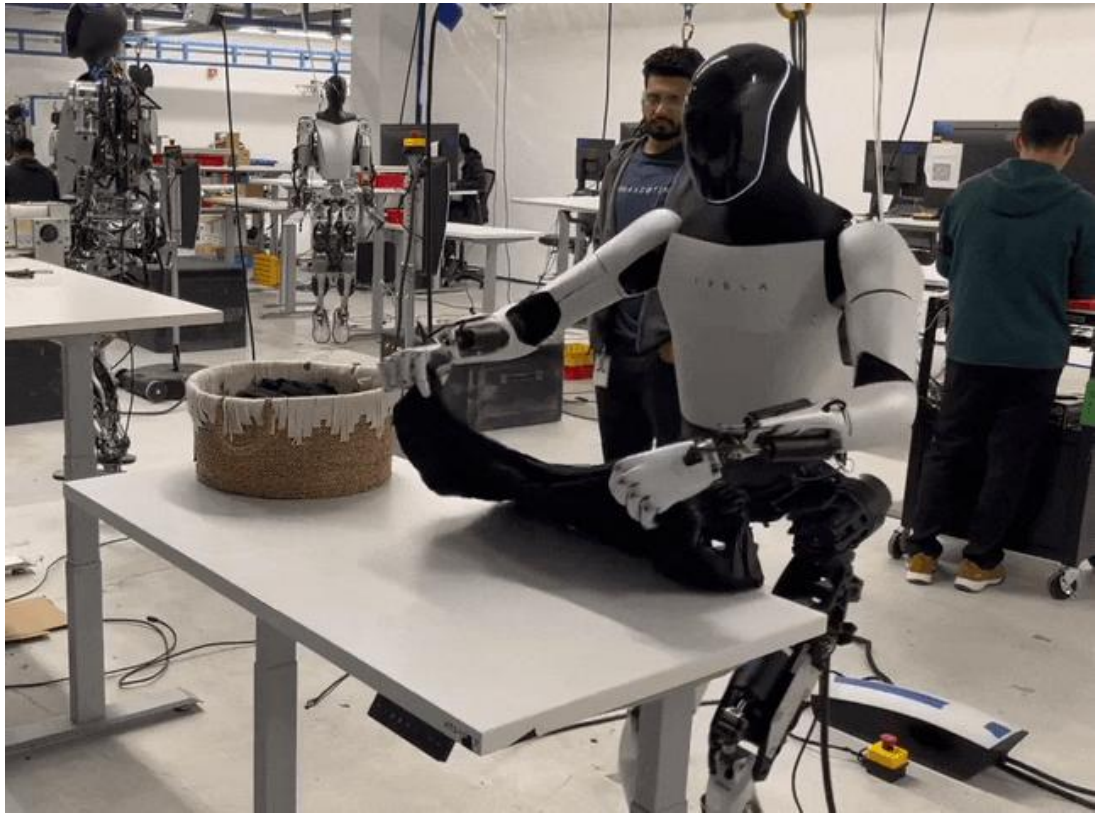

AI Definition (Pat's Version)

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?

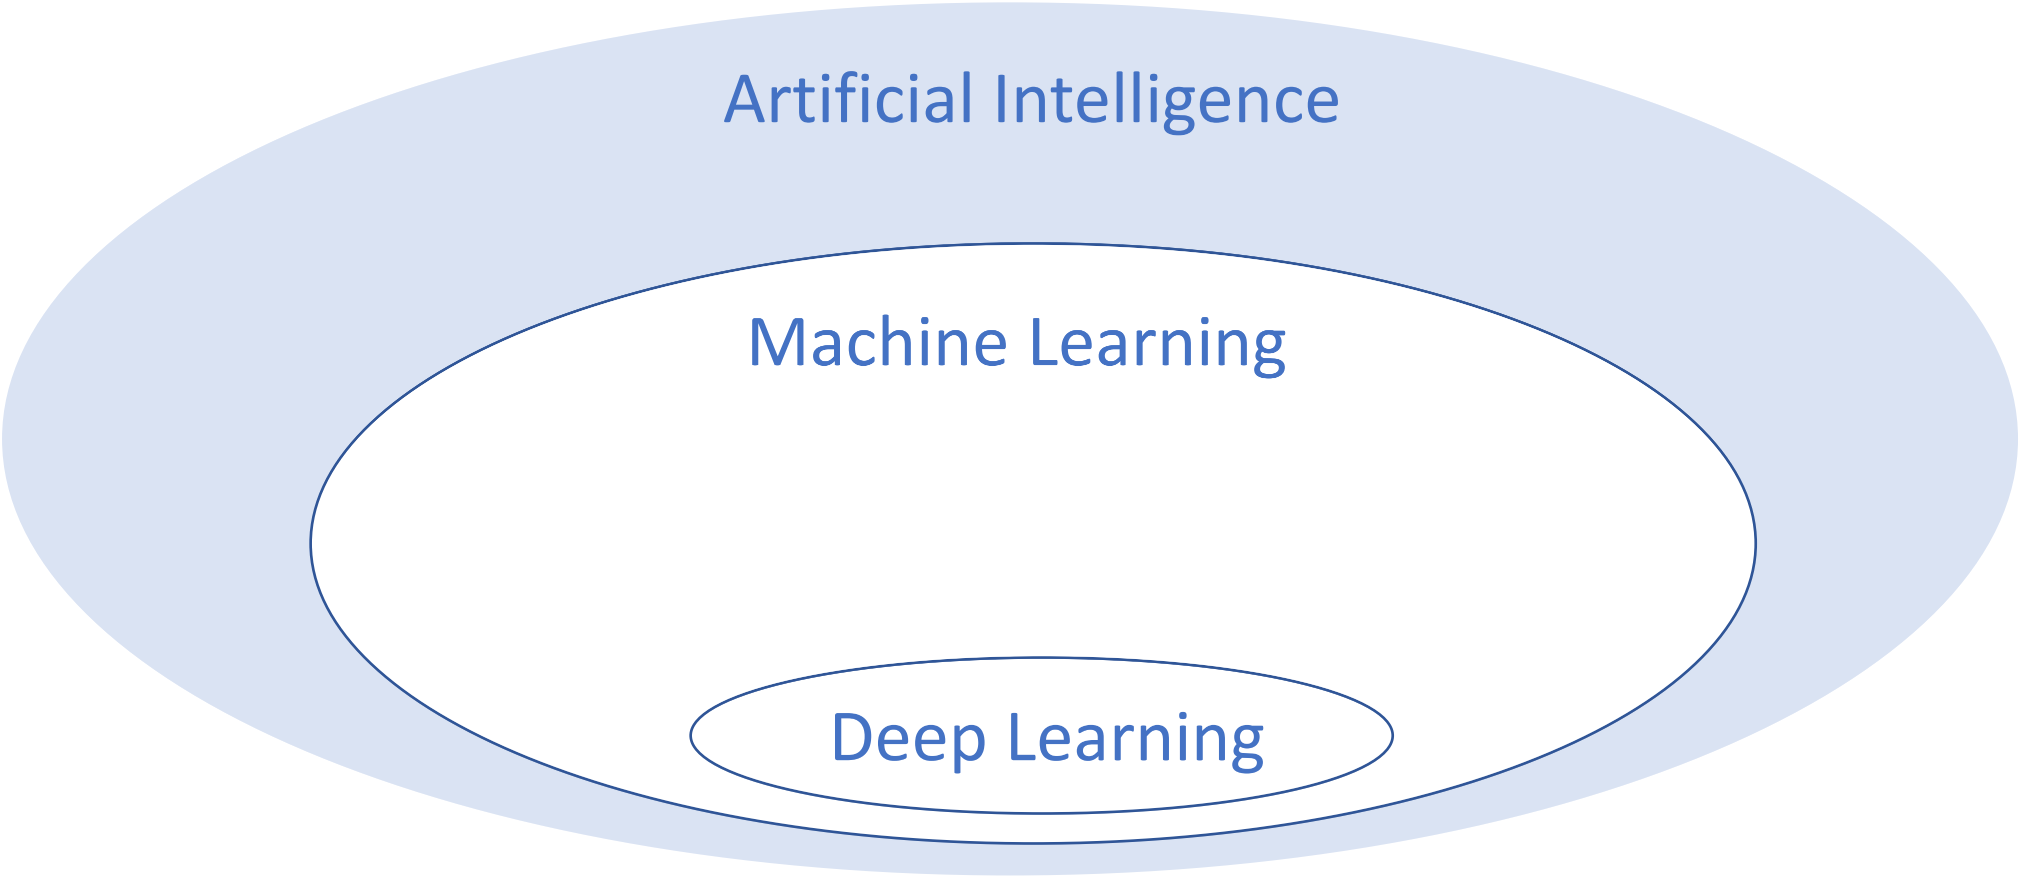

A Brief History of AI

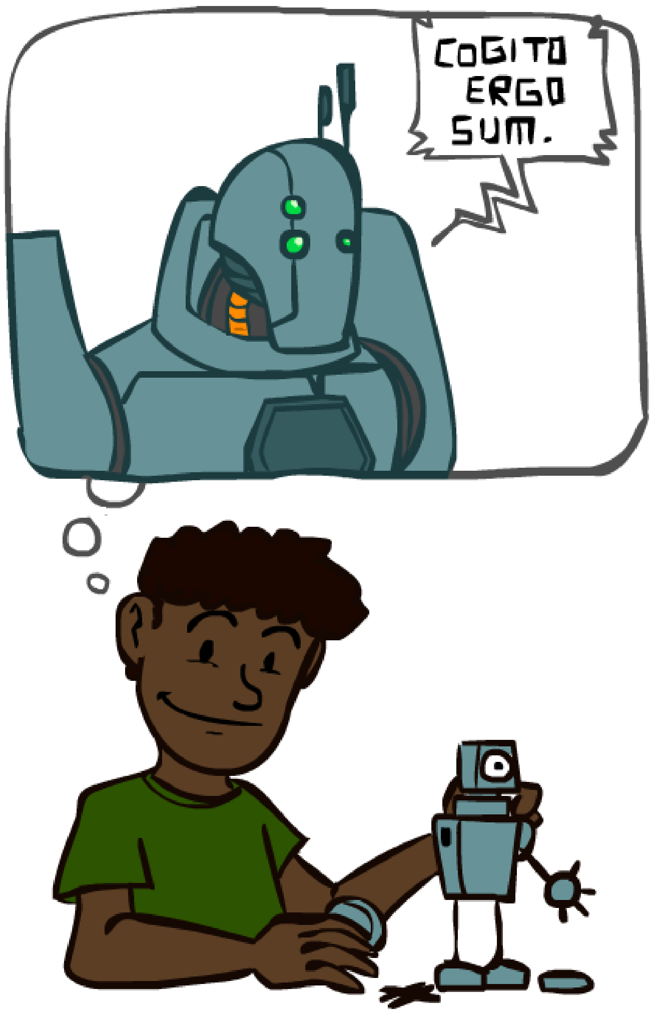

Images: ai.berkeley.edu

Google Books Ngram Viewer https://books.google.com/ngrams

Albert Einstein,Sherlock Holmes,Frankenstein

1800 - 2019 | English (2019) | Case-Insensitive | Smoothing

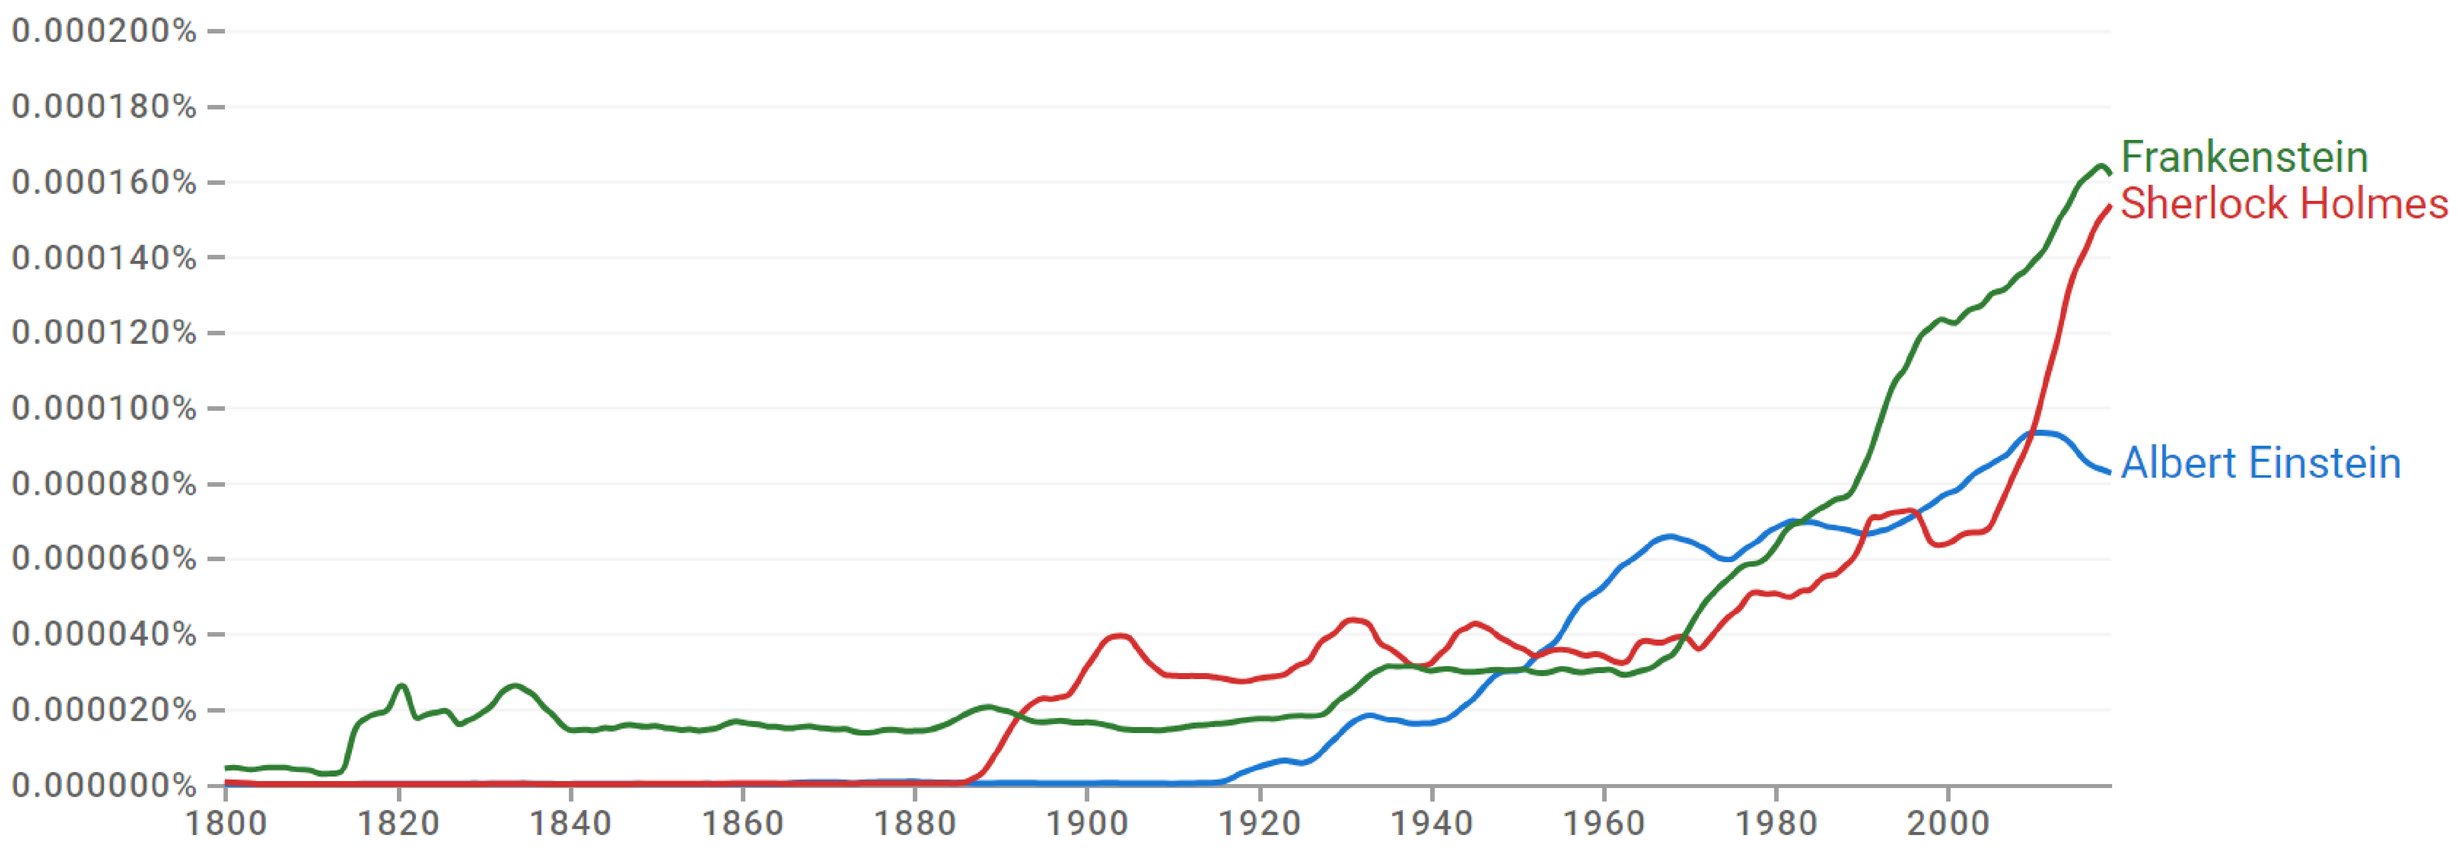

A Brief History of AI

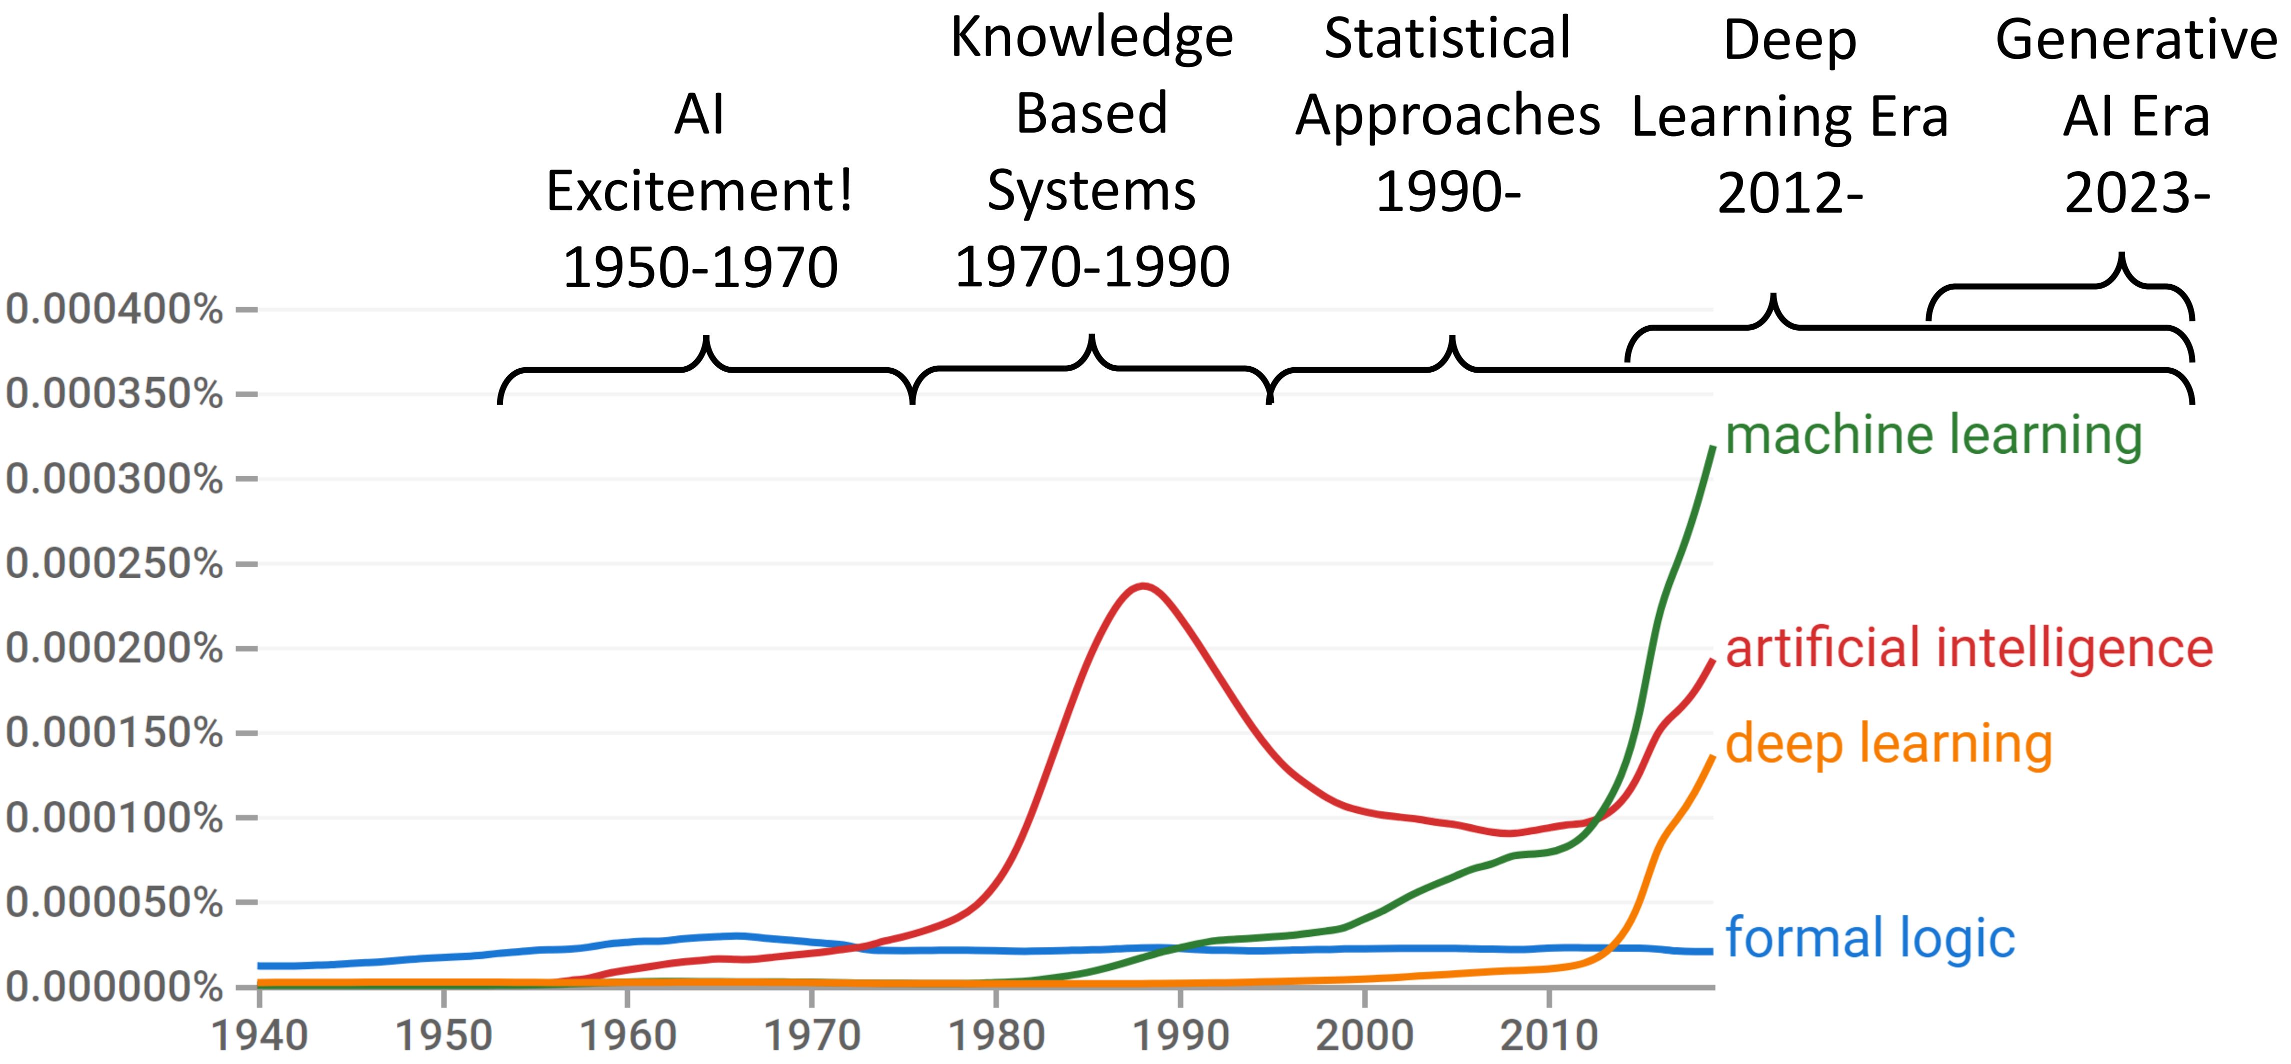

https://books.google.com/ngrams

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

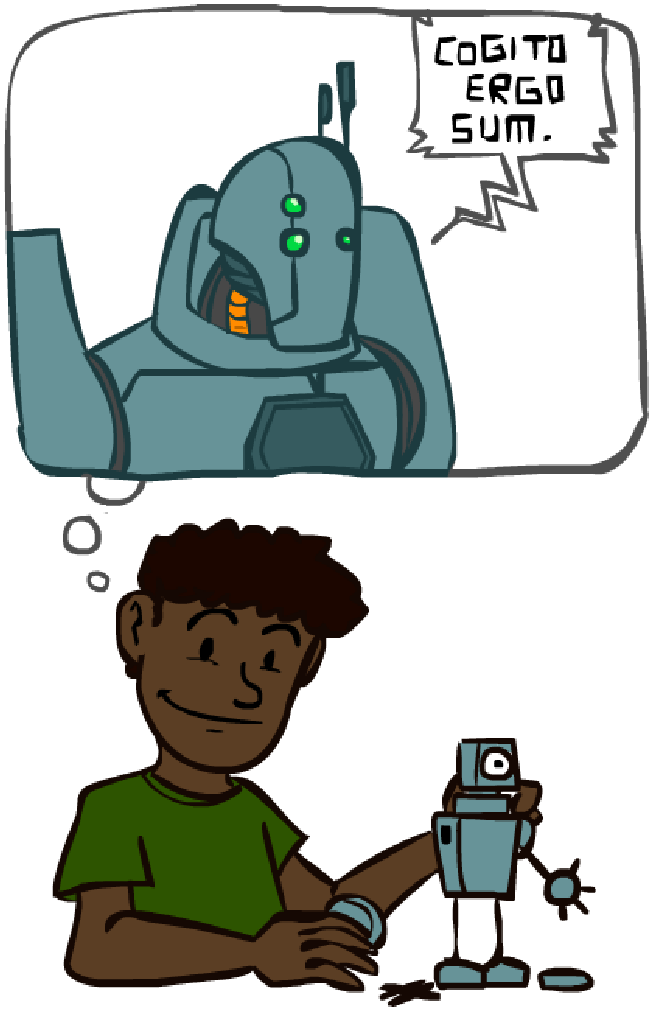

Images: ai.berkeley.edu

ML Applications?

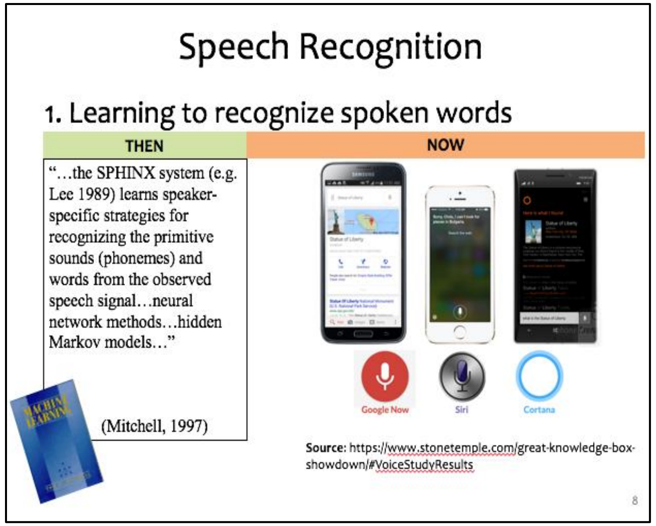

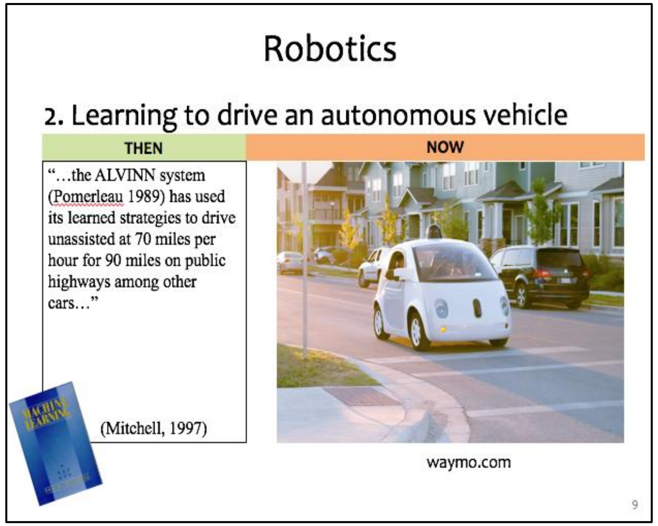

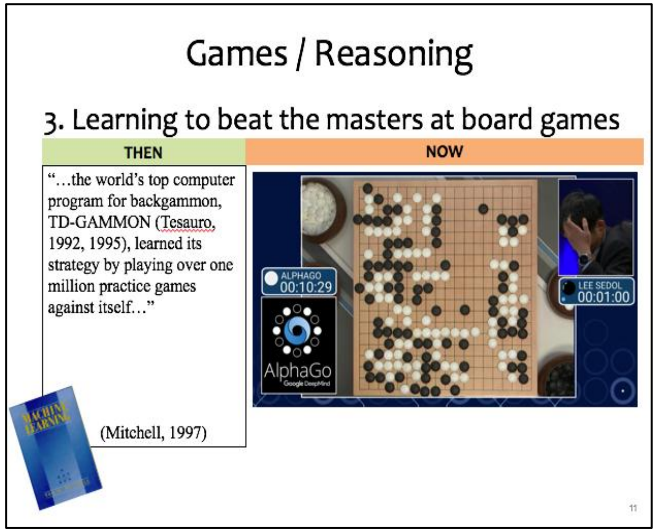

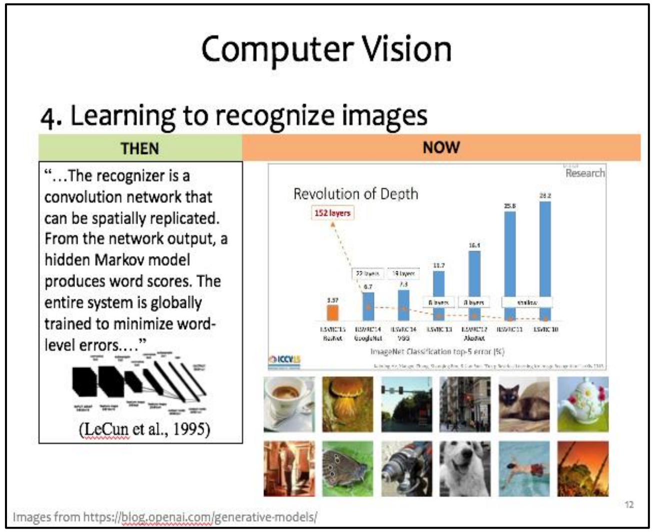

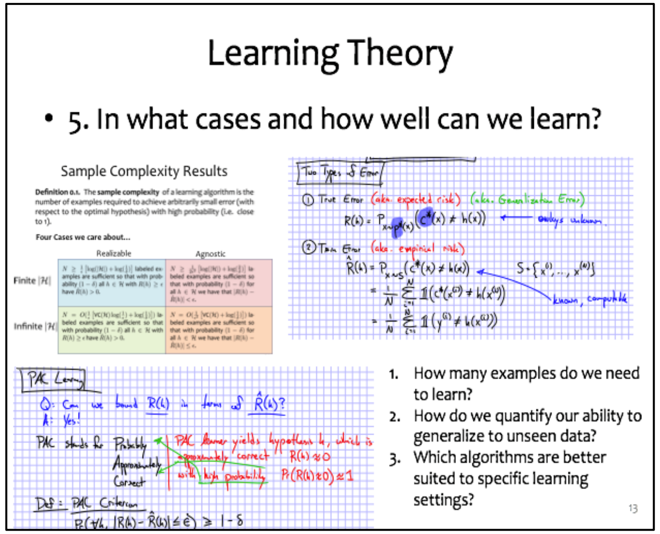

Machine Learning and Statistics

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

ML: Input/Output Tasks

Agents

An **agent** is an entity that *perceives* and *acts*.

Actions can have an effect on the *environment.*

The specific *sensors* and *actuators* affect what the agent is capable of *perceiving* and what *actions* it is capable of taking

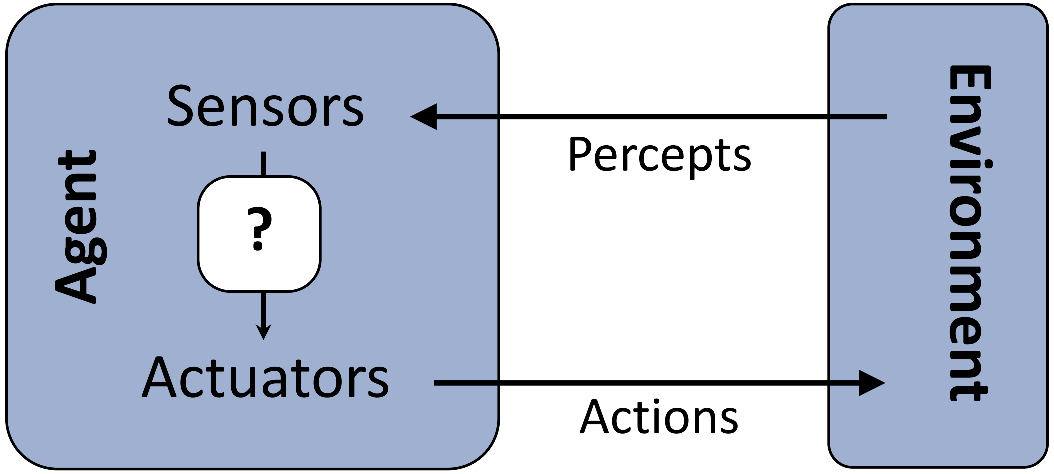

Agent: Simple Input/Output Task

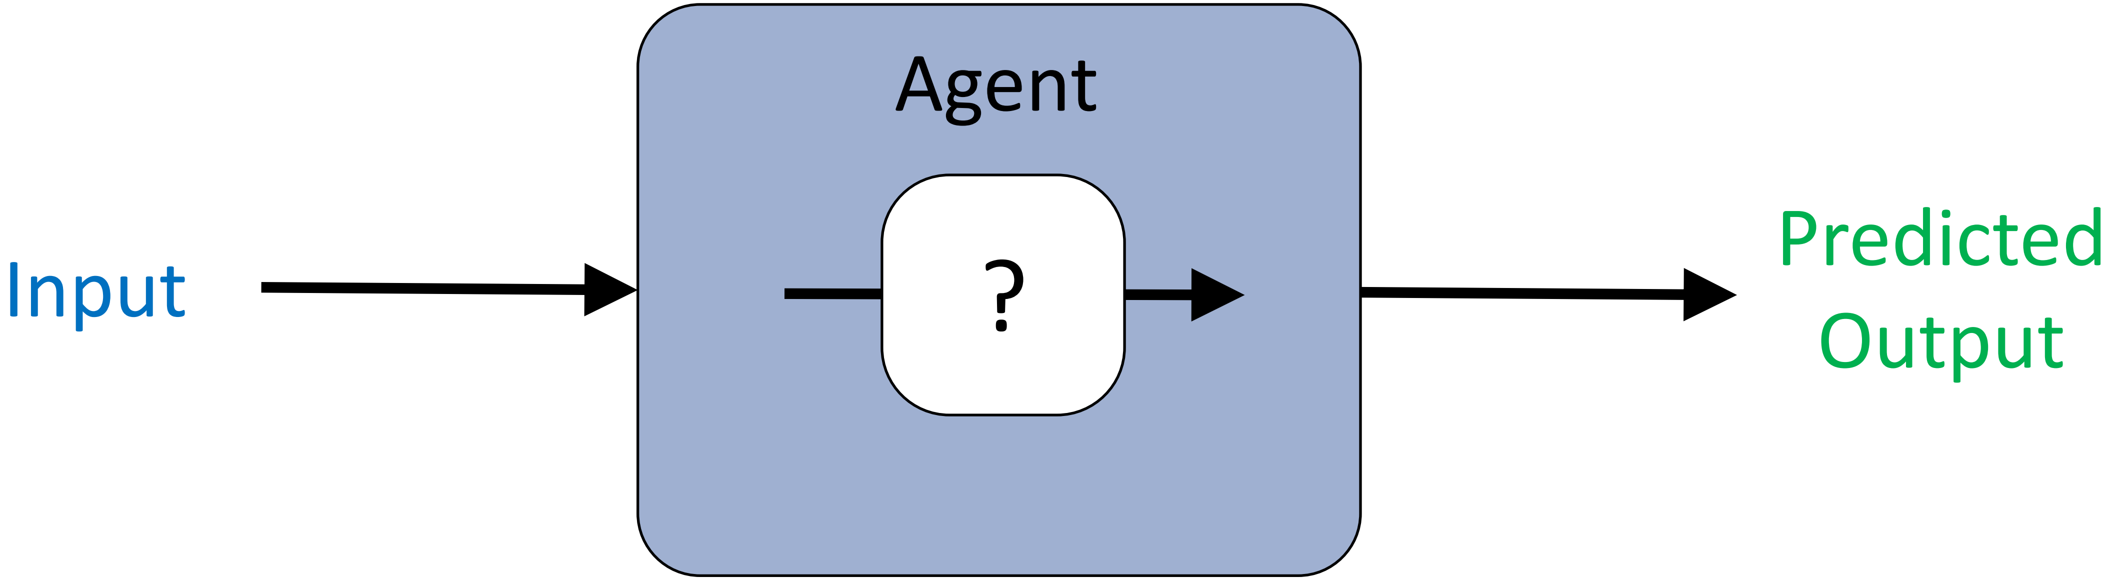

Task: Image Classification

What is this?

Input:
Image (pixels)

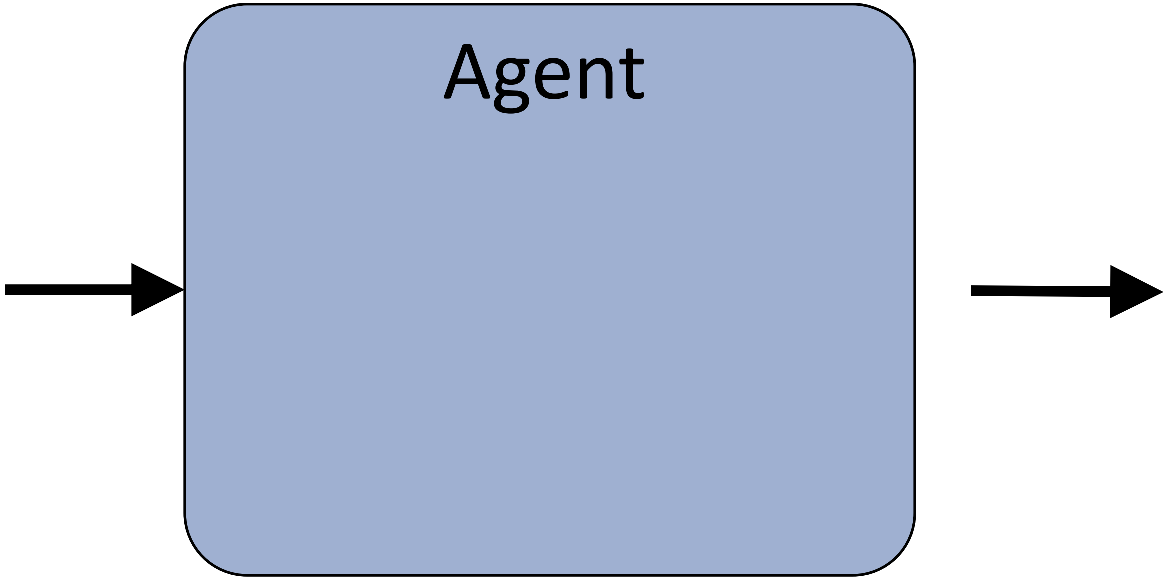

Predicted Output:
Category

Task Input and Output

Input	Task	Output
Petal measurements	Iris classification	Category
Time of day	Traffic prediction	Traffic Volume
Image	Image classification	Category
Image	Image denoising	Image
Text	Text to image generation	Image
???	Face generation	Image

Task: Face Generation

https://thispersondoesnotexist.com/

Task: Classification from Features

What is this?

Input:
Image

Predicted Output:
Category

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 →

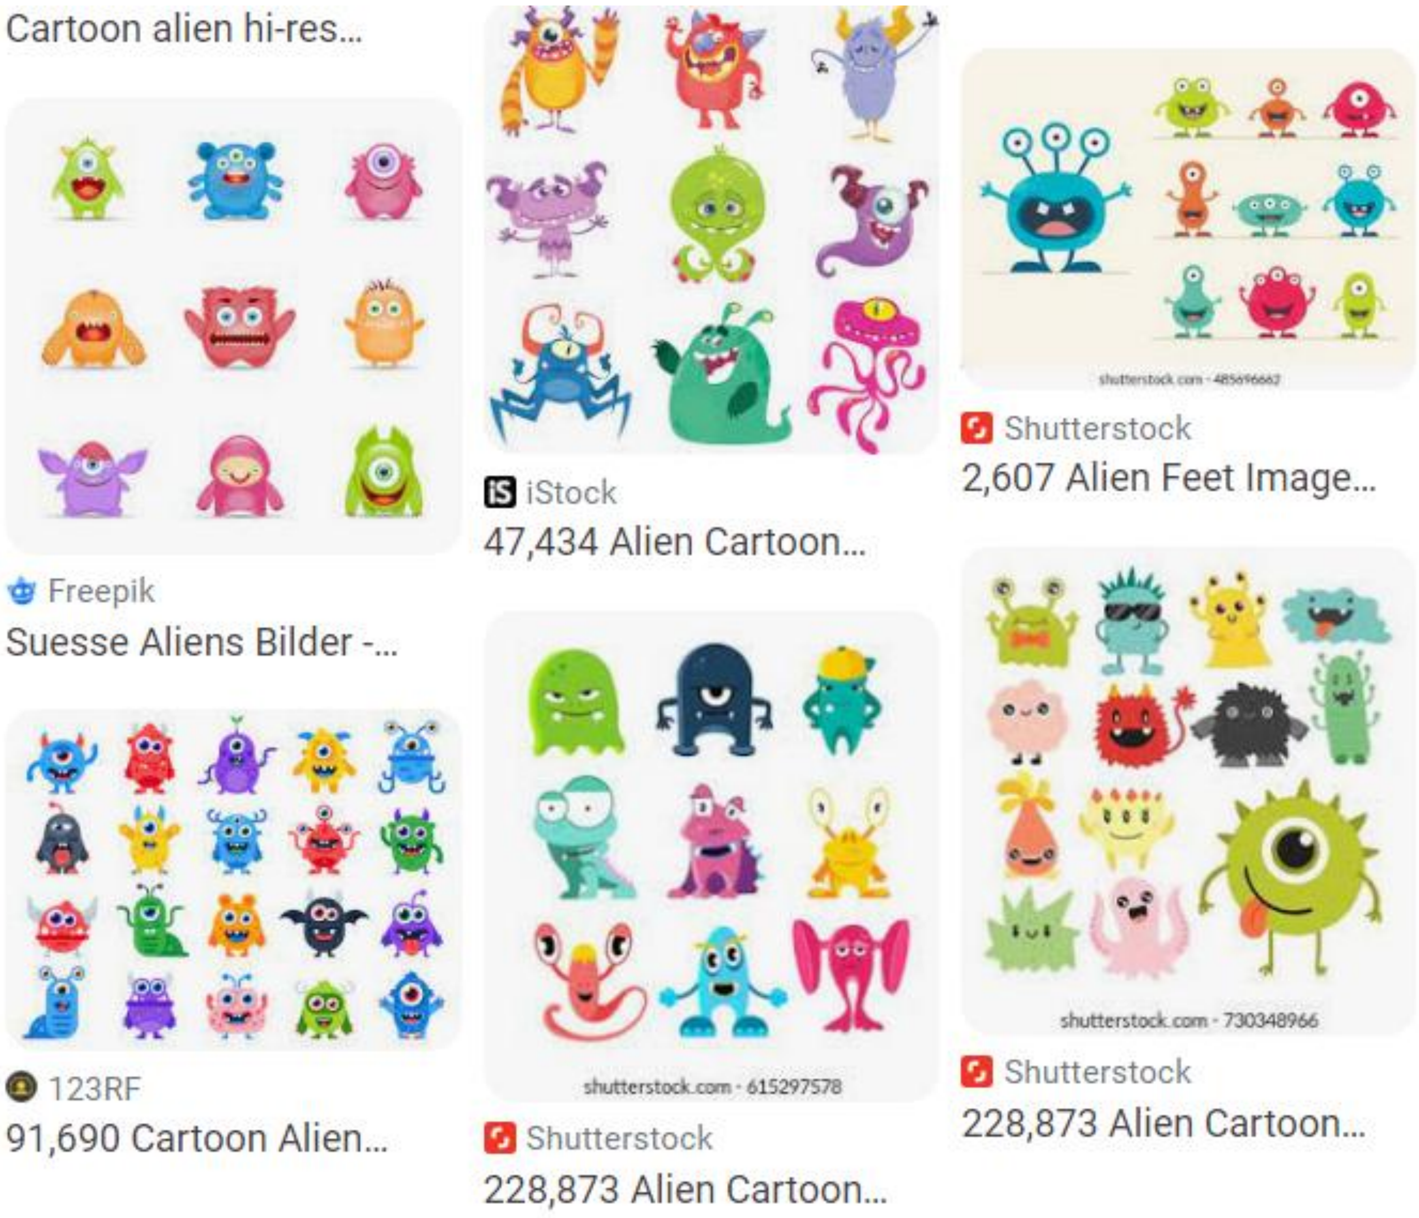

Exercise: Human-defined Feature Space

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

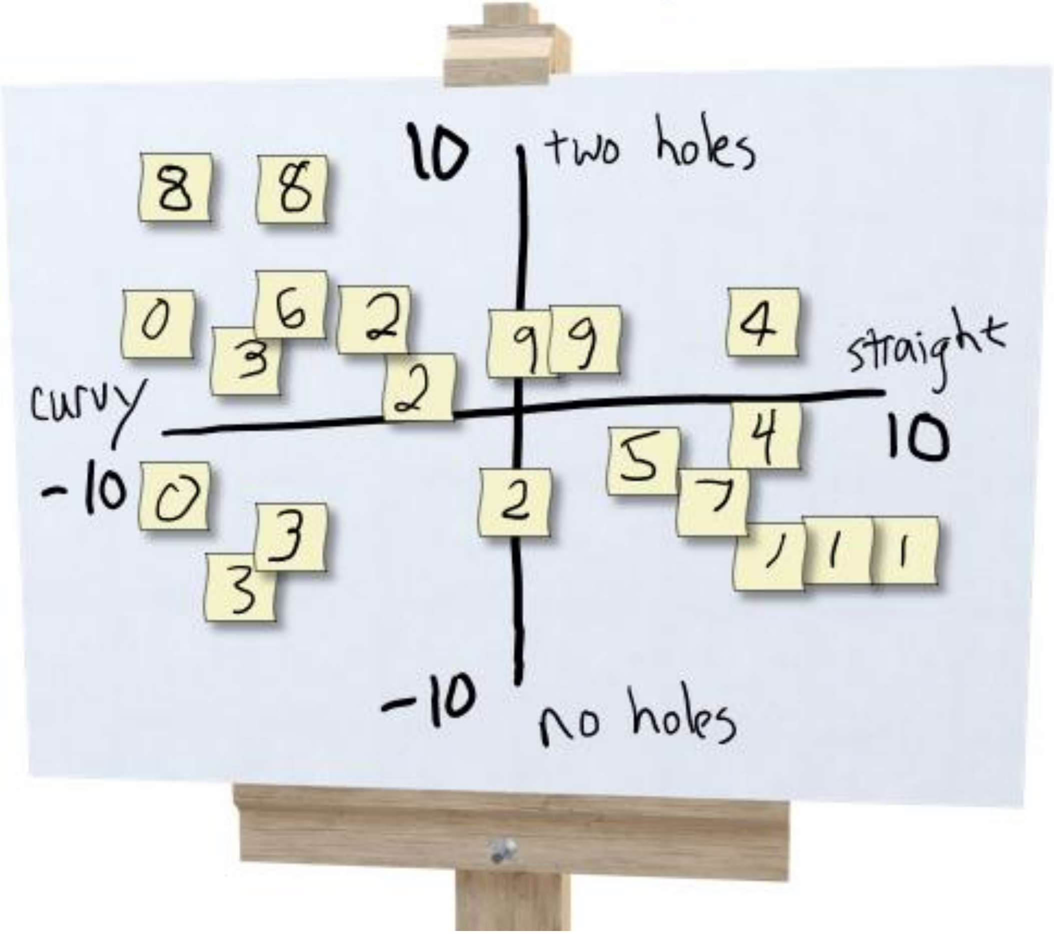

Exercise: Human-defined Feature Space

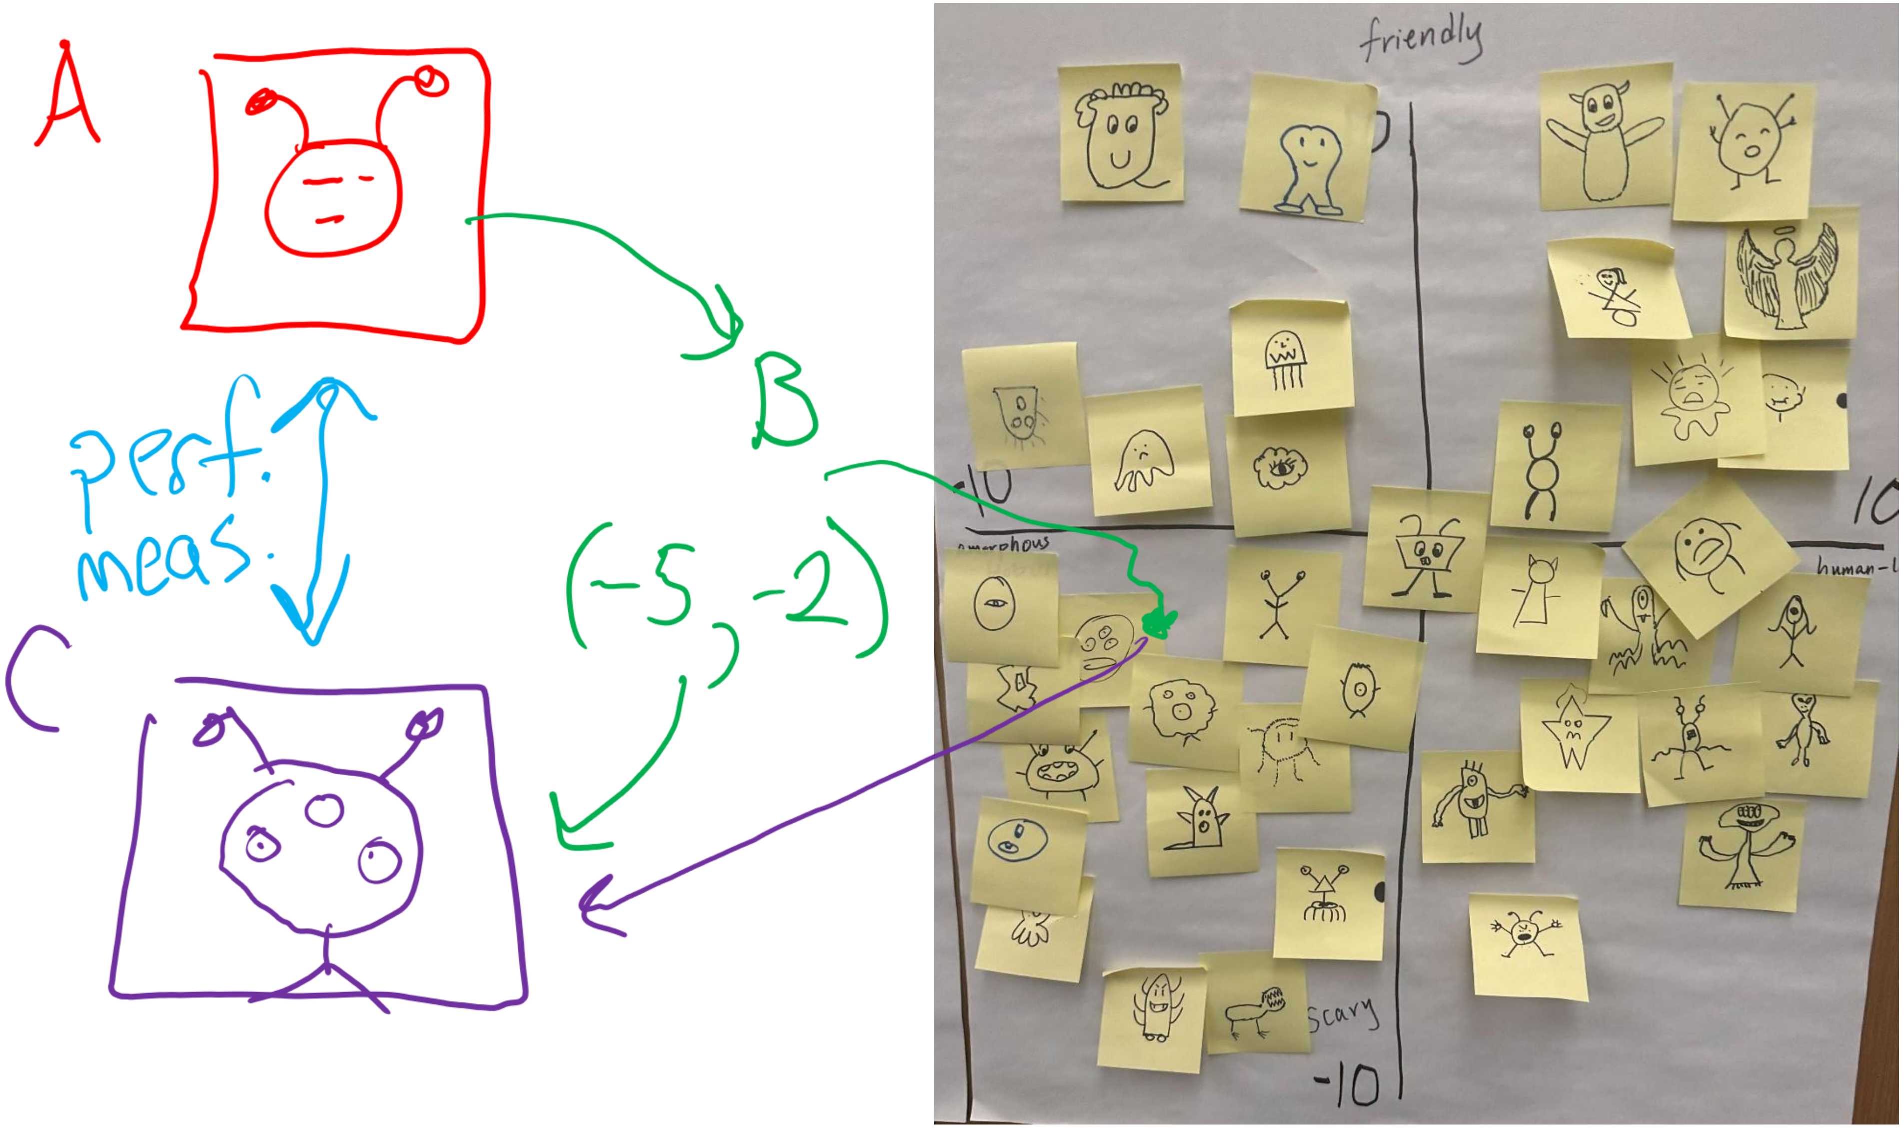

Poll 1

What features did you use?

Learning to Organize Data

Neural networks can learn to organization too!

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

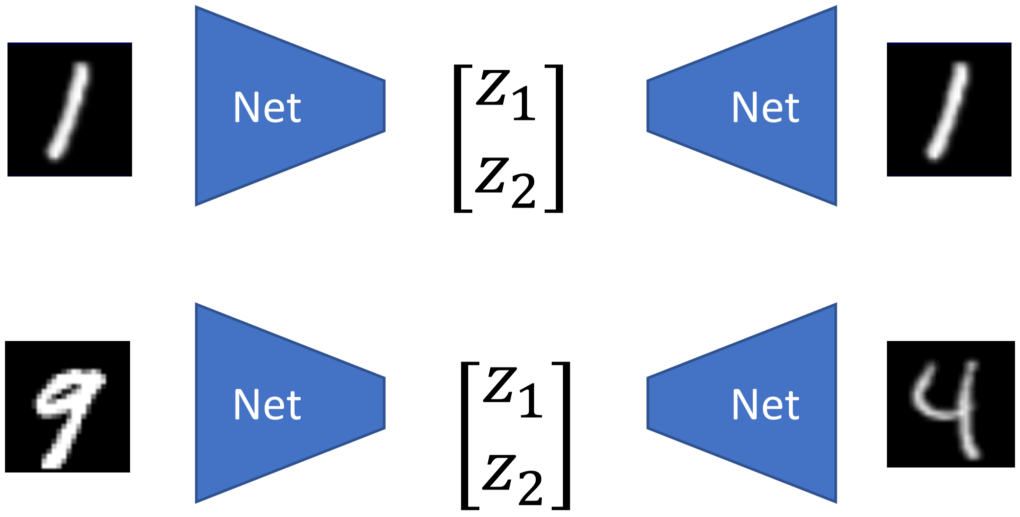

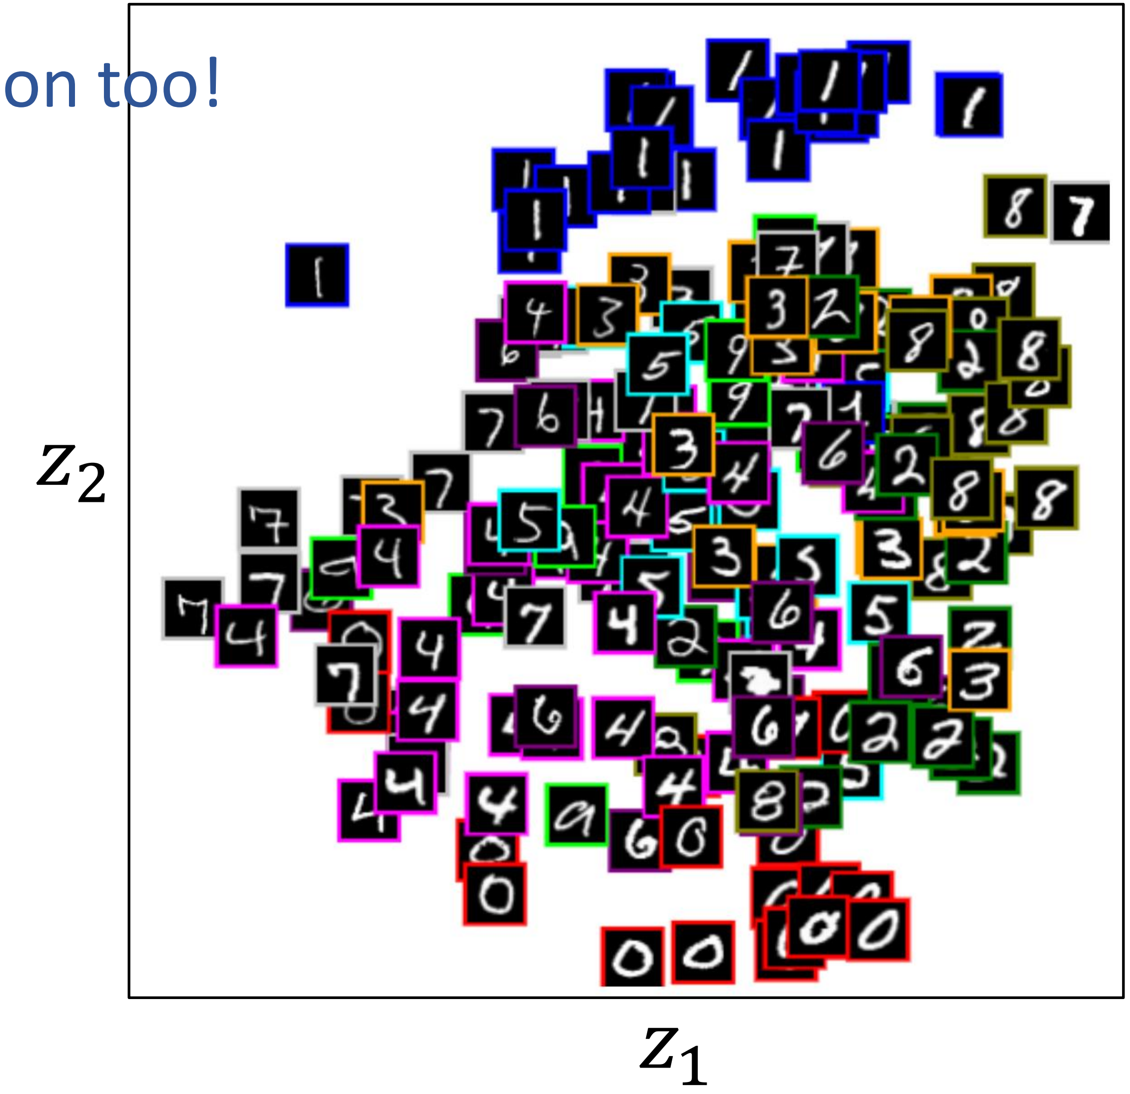

https://cs.stanford.edu/people/karpathy/convnetjs/demo/autoencoder.html

Digit Autoencoder

Demo: Using a learned feature space

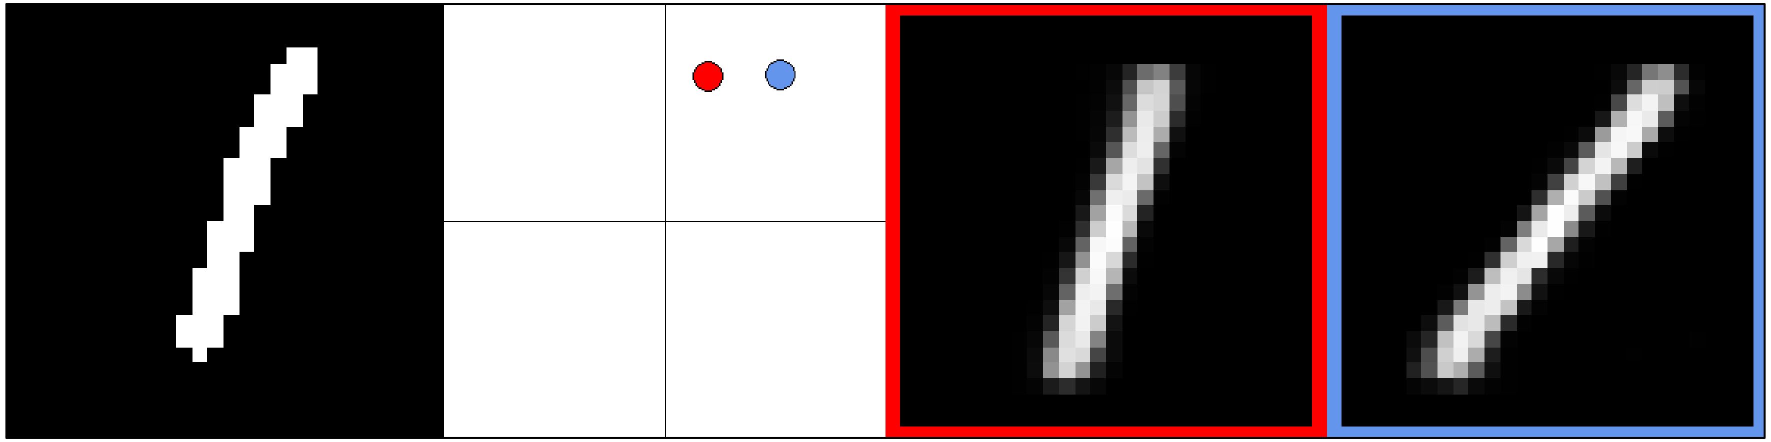

Task: Face Generation

https://thispersondoesnotexist.com/





Task: Face Generation

https://thispersondoesnotexist.com/

Random

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

Net

Random

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

Net

Face GAN Slider Demo

https://github.com/genforce/interfacegan

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

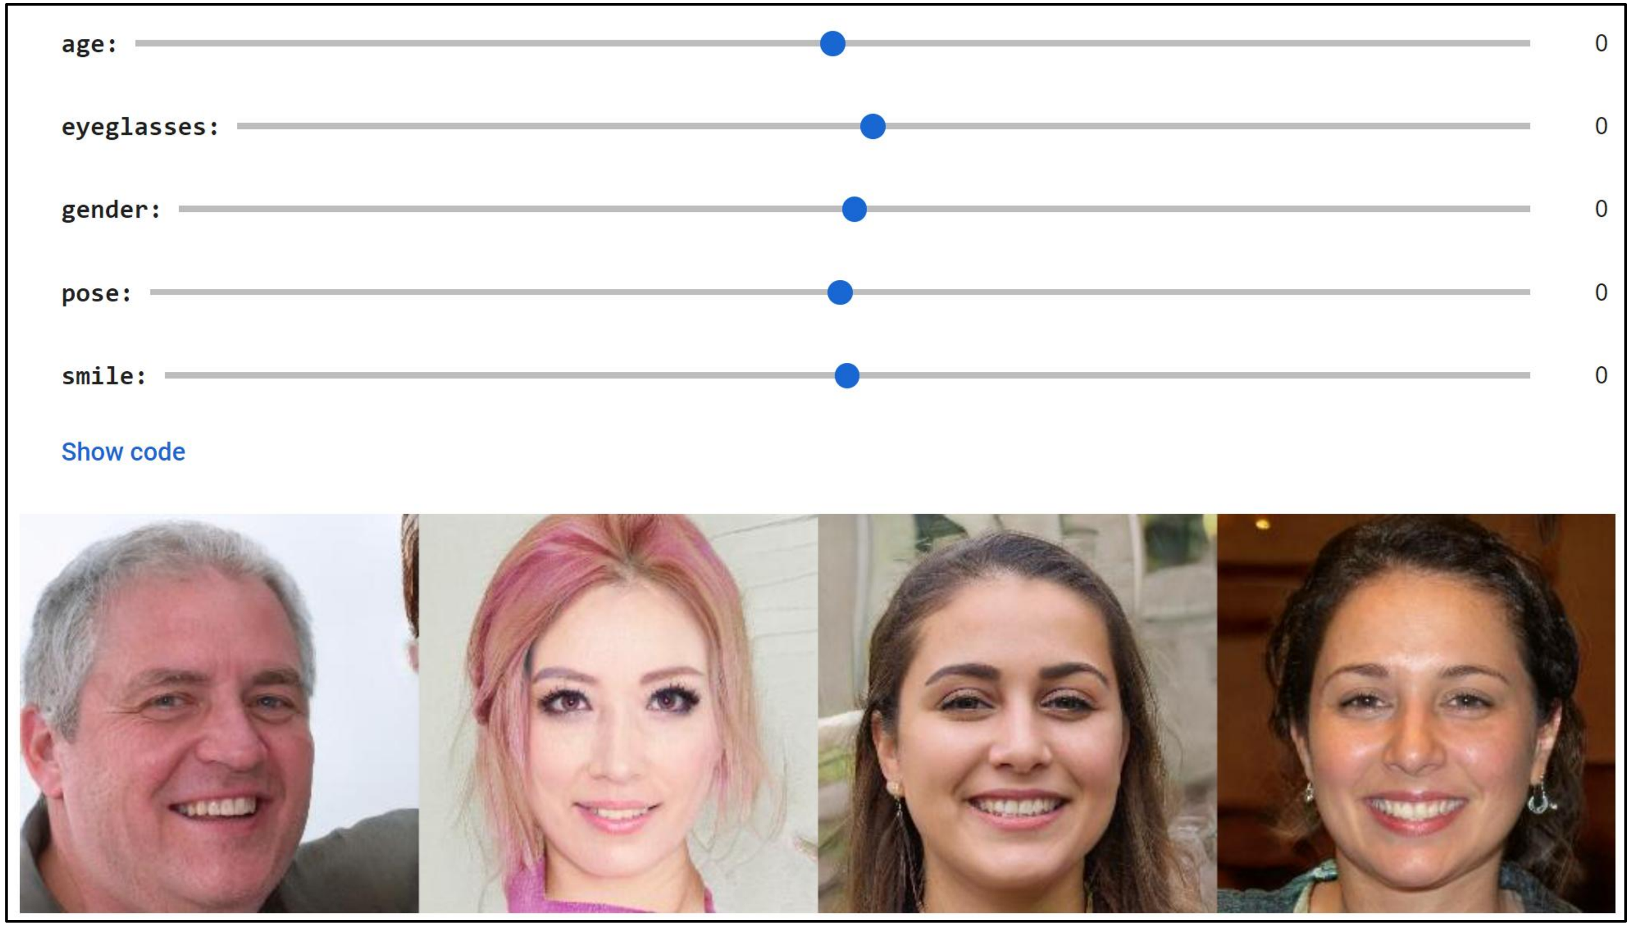

Face GAN Slider Demo

https://github.com/genforce/interfacegan

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

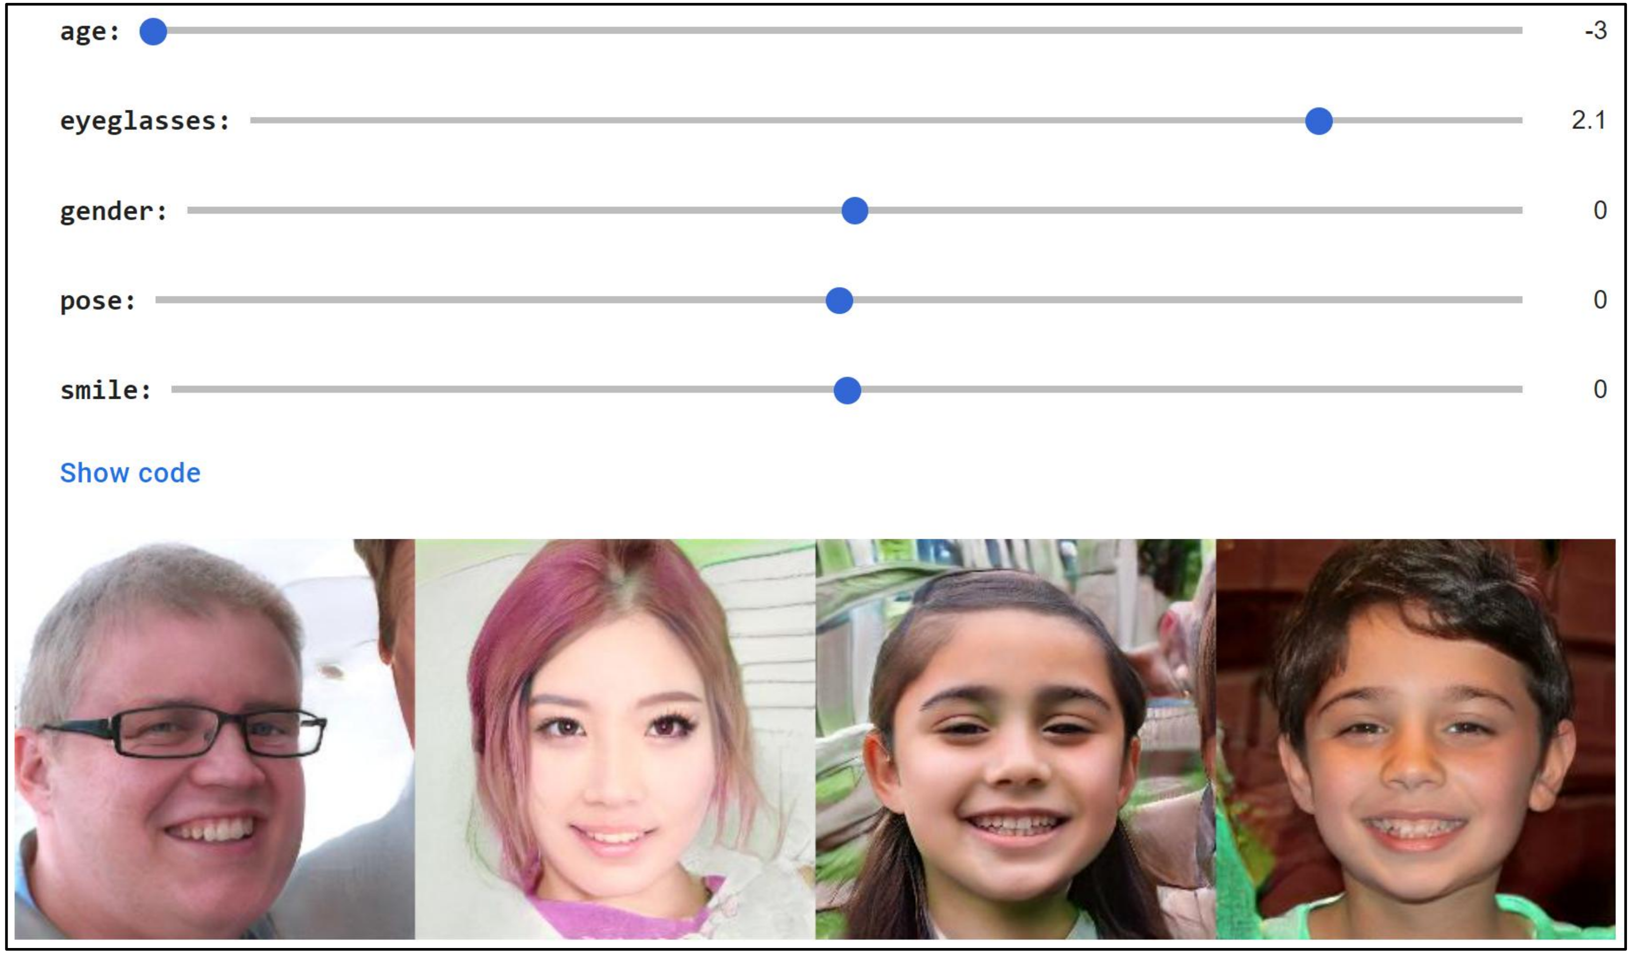

Listen Learner

Chris Harrison, CMU

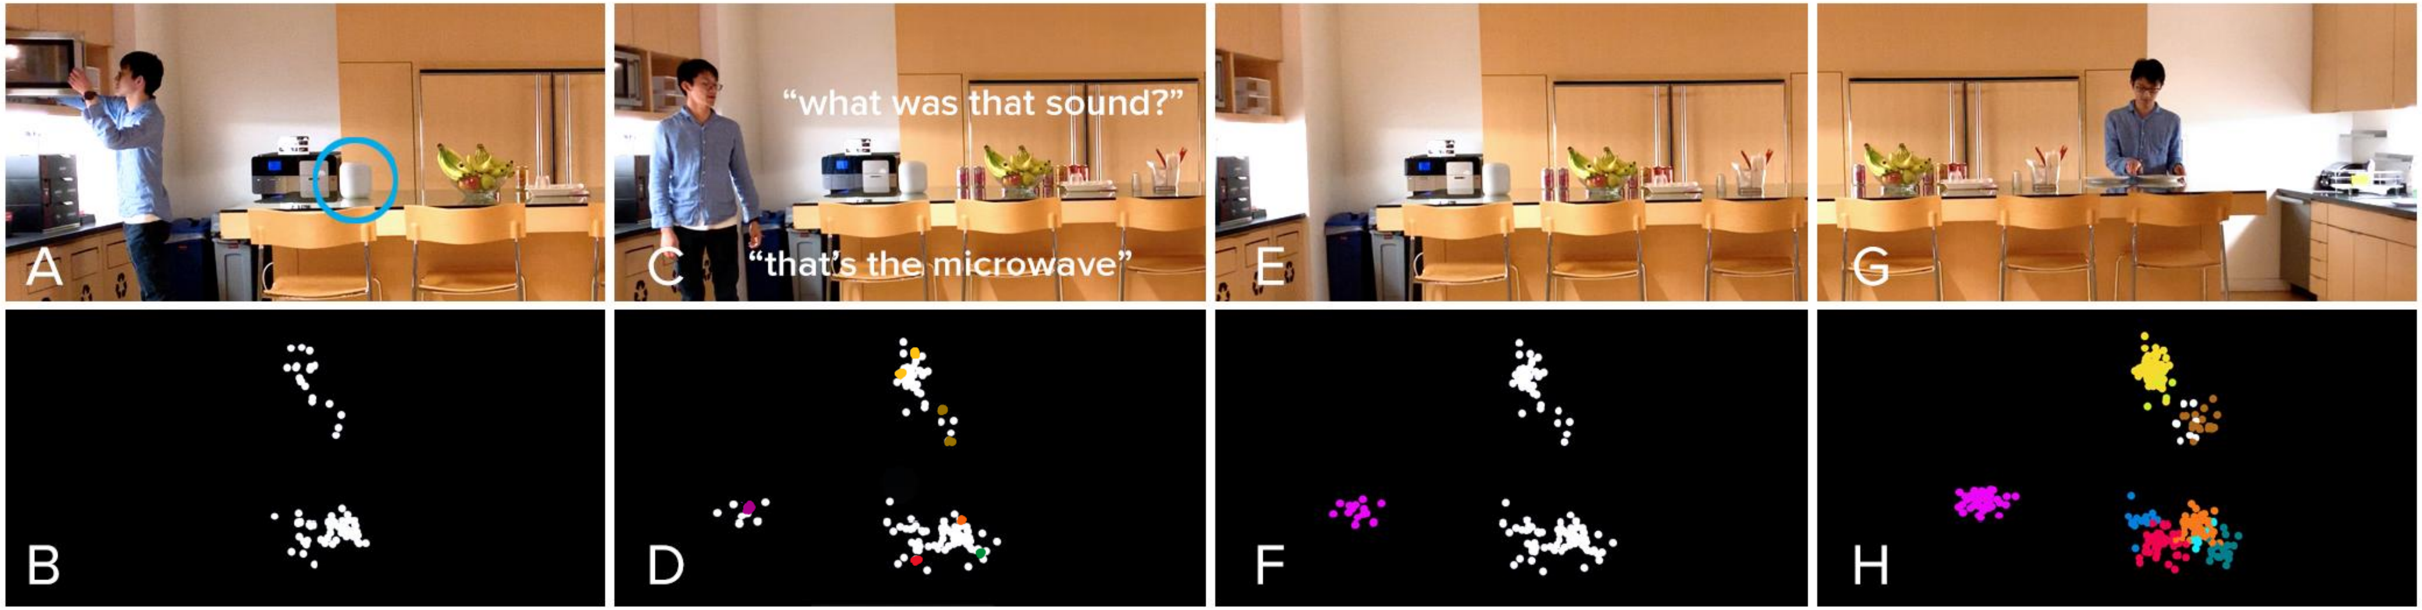

https://chrisharrison.net/index.php/Research/ListenLearner

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.

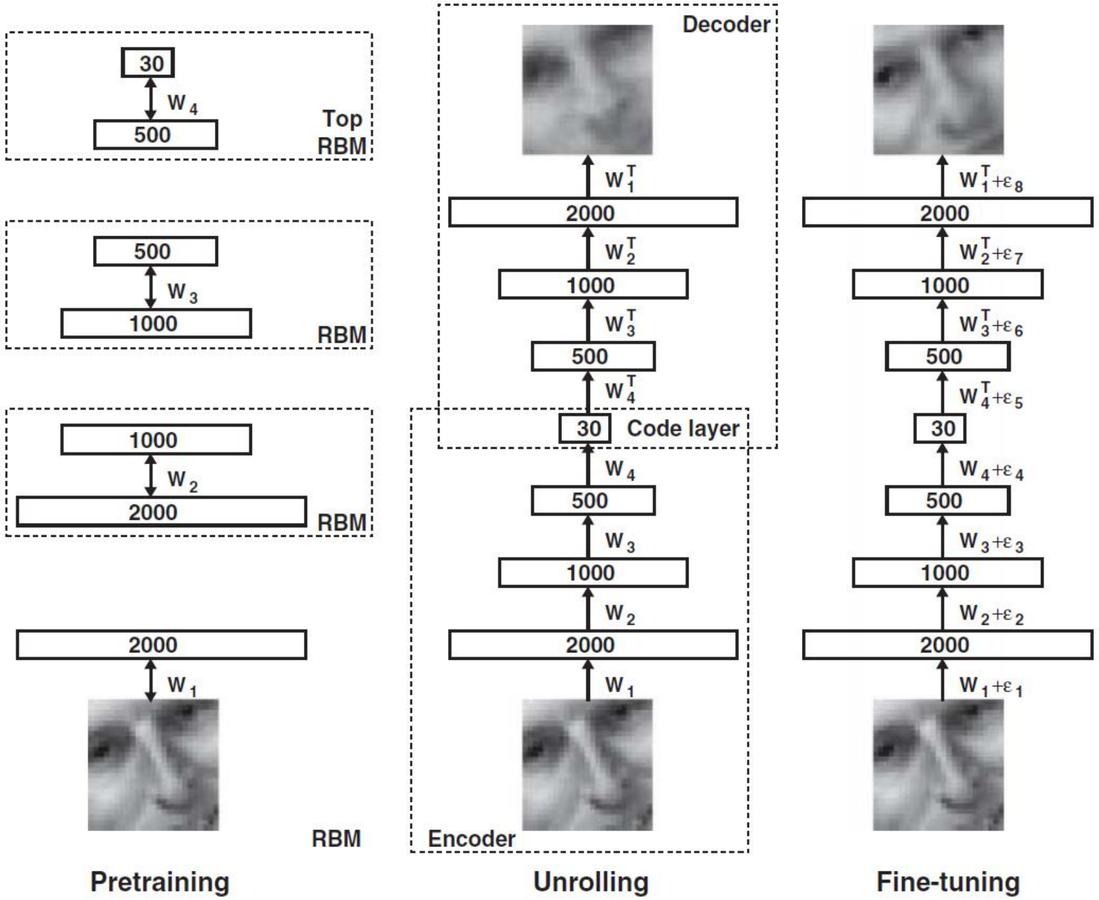

Digit Autoencoder

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

Input
28x28 = 784 pixels

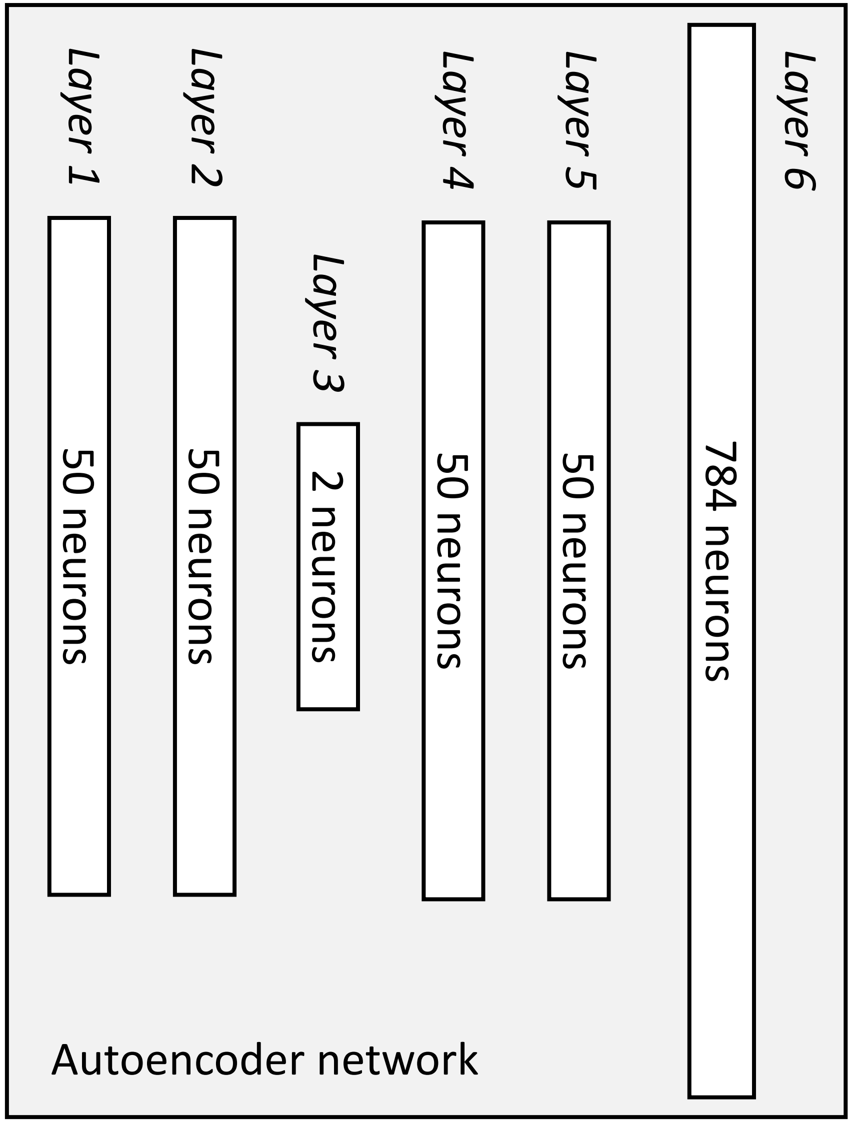

Digit Autoencoder

Math for autoencoder model

Linear layer

ReLU layer

Tanh layer