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 →

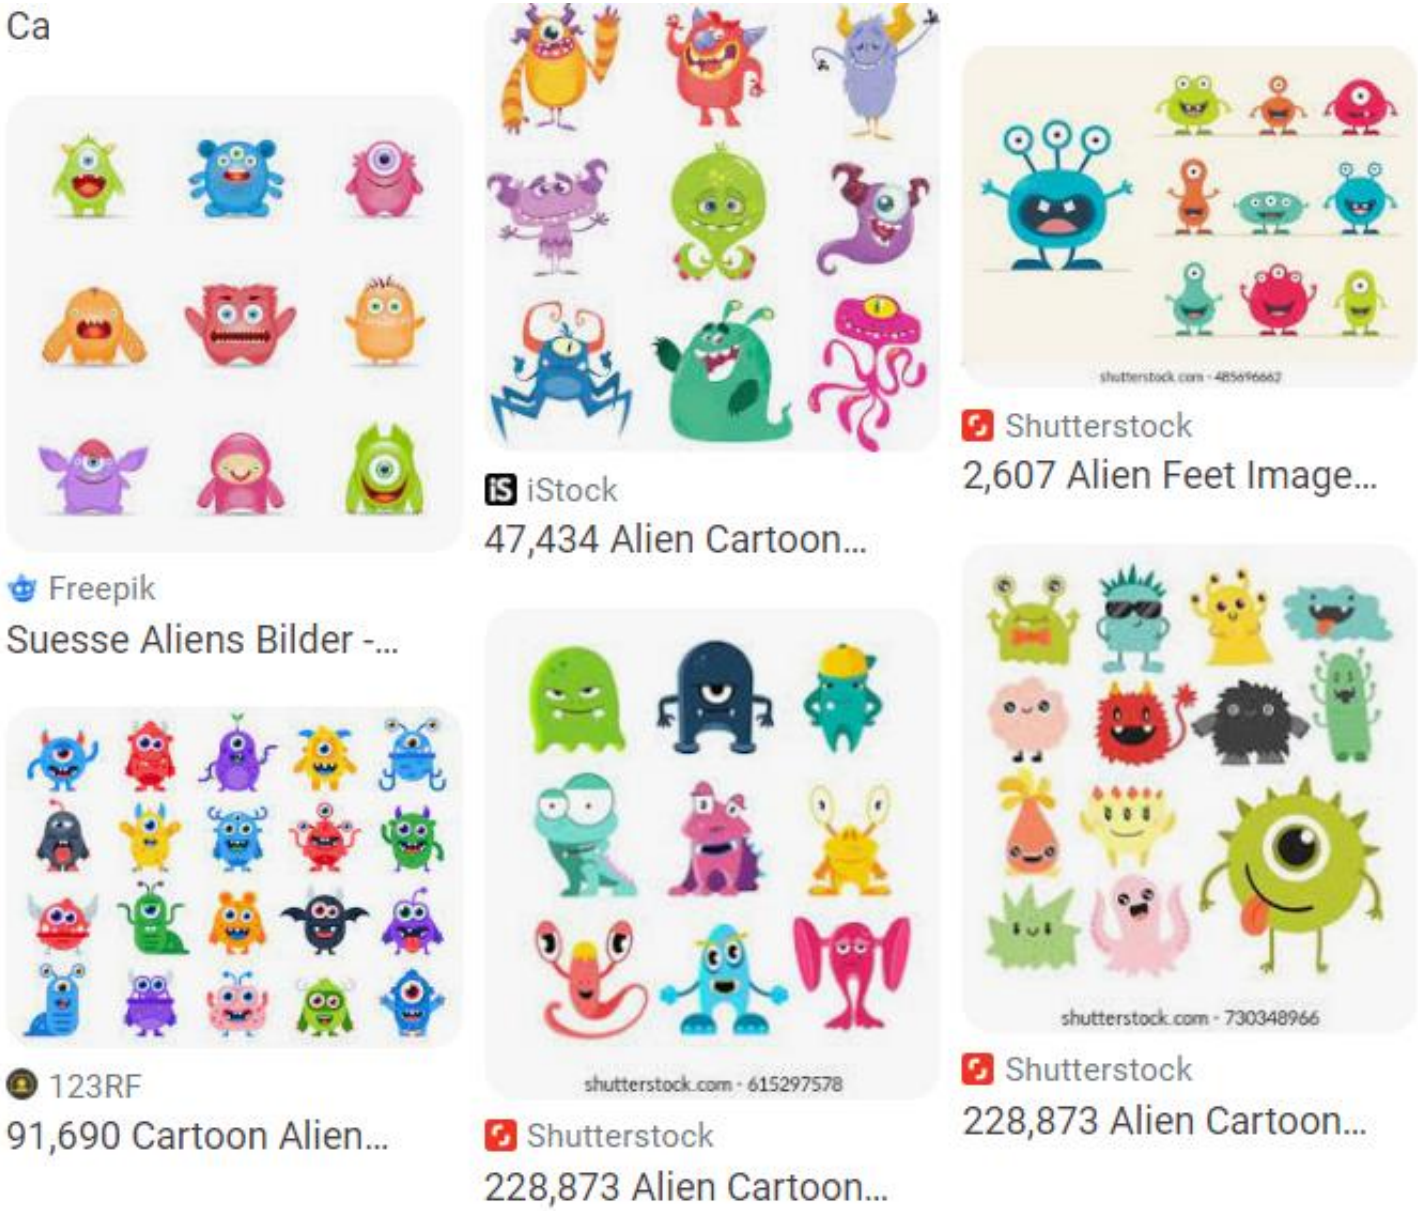

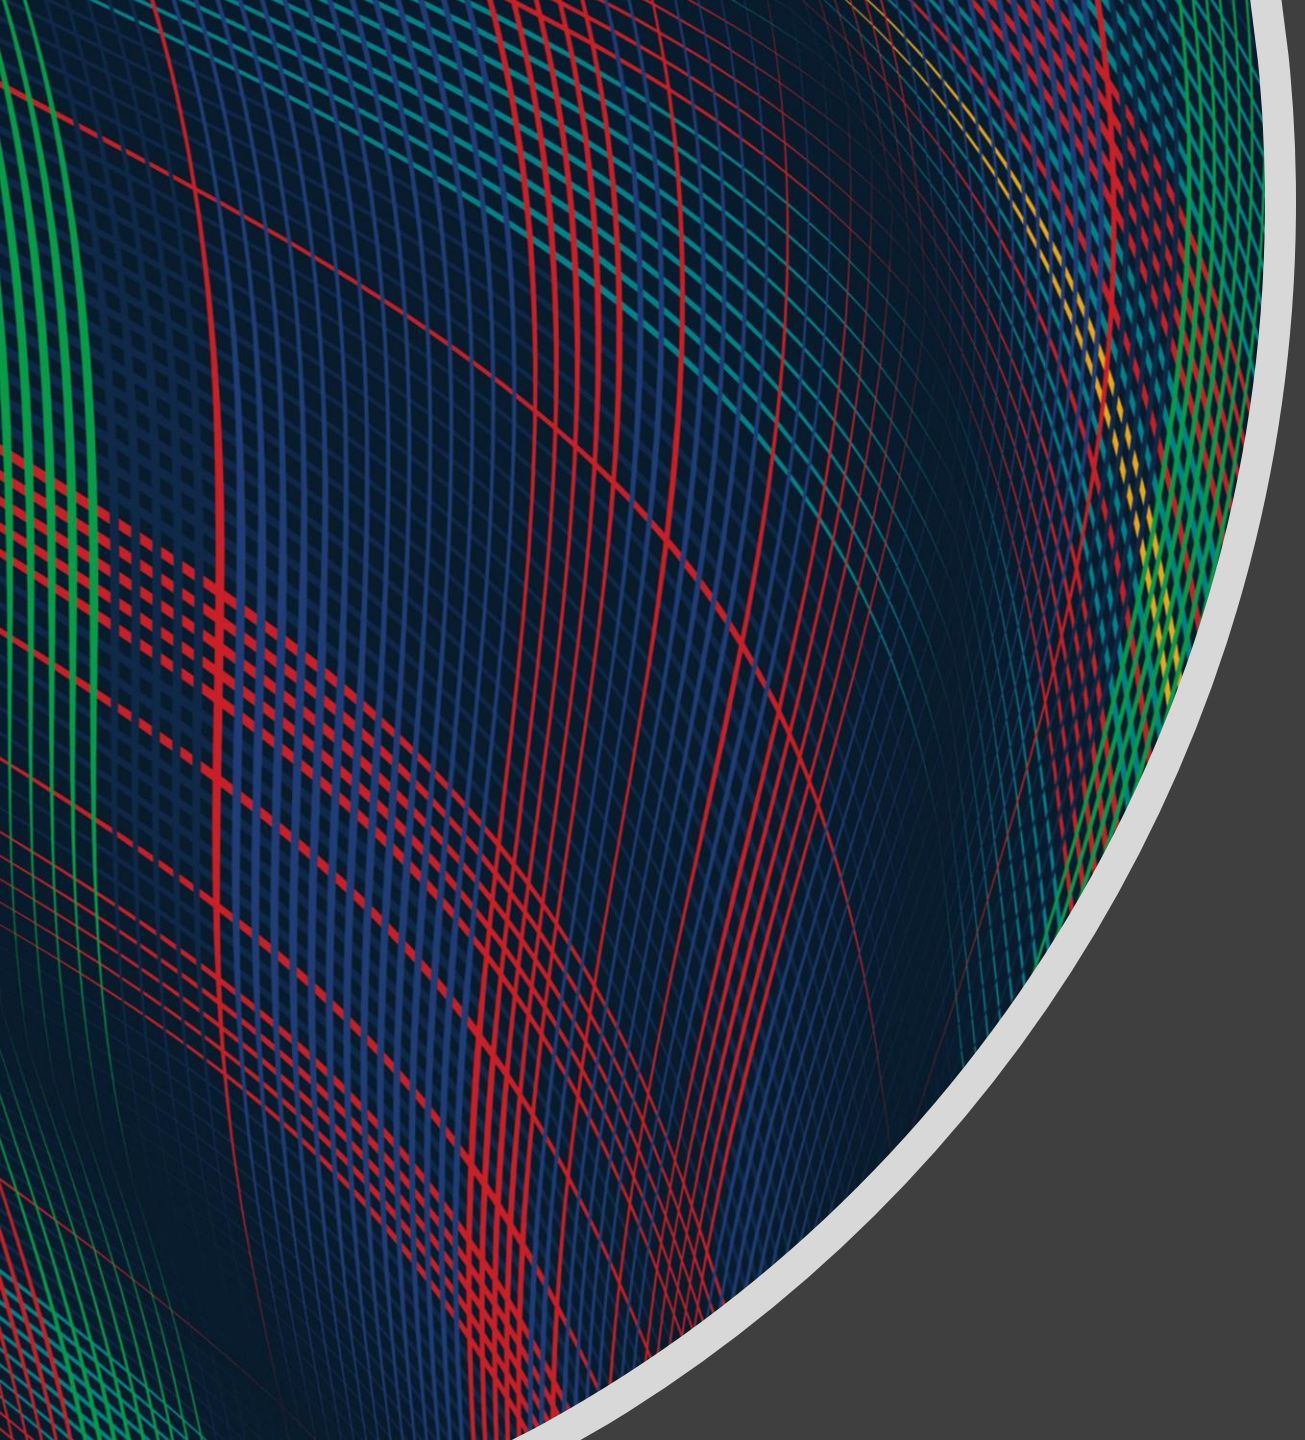

10-315 Introduction to ML

Instructor: Pat Virtue

Today

Course team

ML framework

Elephants in the room (ChatGPT, DALL-E 2, …)

Alien exercise

ML Models

Autoencoders

More course info

Announcements

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

Course Team

Instructor

Pat
Virtue
pvirtue

Education Associate

Joshmin Ray
joshminr

Course Team

Teaching Assistants

Saumya
ssgandhi

Shreeya
srkhuran

Deep
dmpatel

Devanshi
devanshg

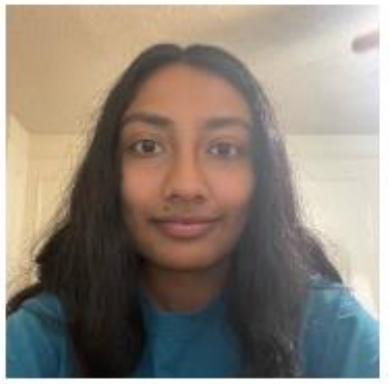

Medha
mpalaval

Alex
alextiax

Arya
aryas

Meher
mmankika

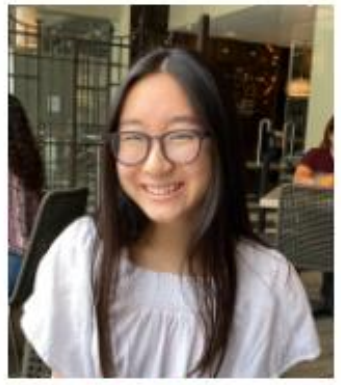

Ruthie
rylin

Saloni
salonipa

Course Team

Students!!

What is ML?

Domain of Interest

Computer Science

Machine Learning

Optimization

Statistics

Linear Algebra

Probability

Calculus

Measure Theory

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

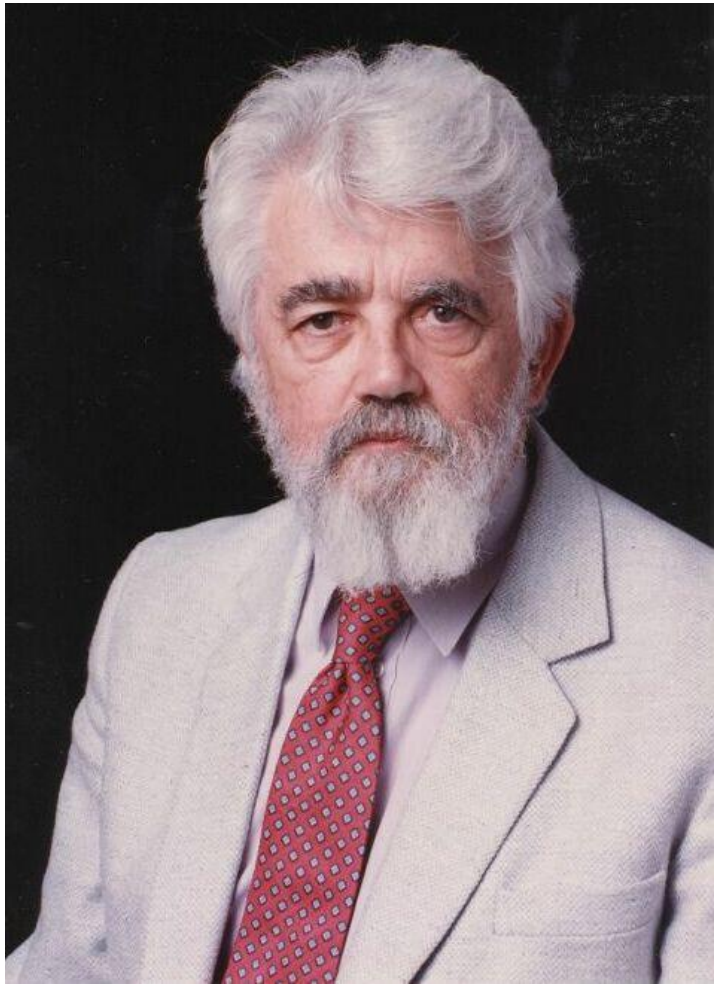

http://www-formal.stanford.edu/jmc/whatisai/whatisai.html

AI Stack for CMU AI

"AI must understand the human needs and it must make smart design decisions based on that understanding"

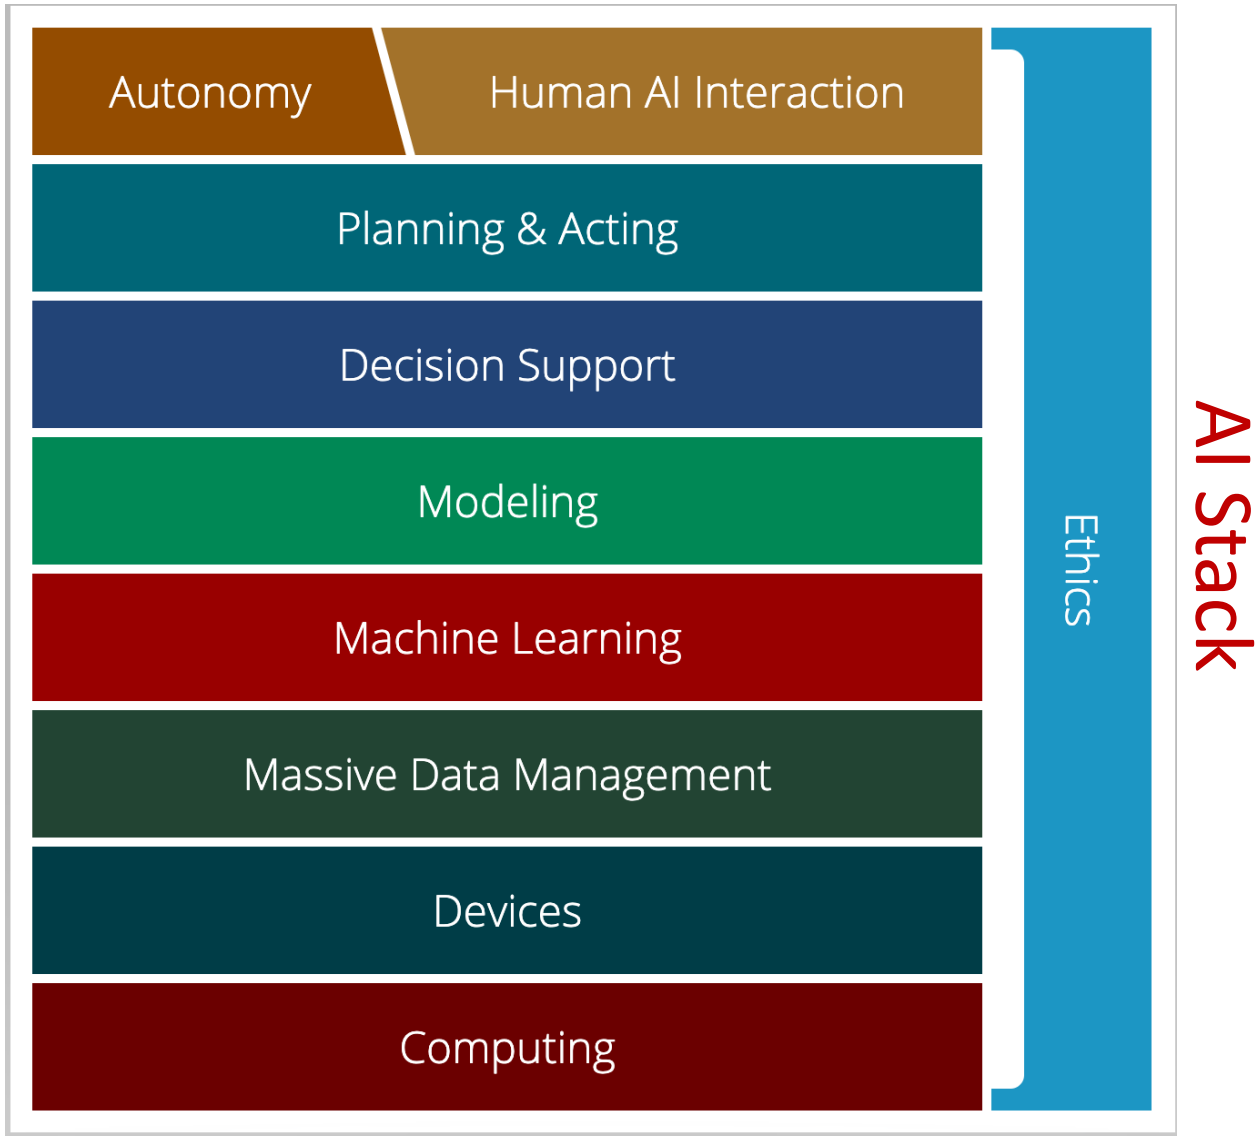

https://ai.cs.cmu.edu/about

AI Stack for CMU AI

"Machine learning focuses on creating programs that learn from experience."

"It advances computing through exposure to new scenarios, testing and adaptation, while using pattern- and trend-detection to help the computer make better decisions in similar, subsequent situations."

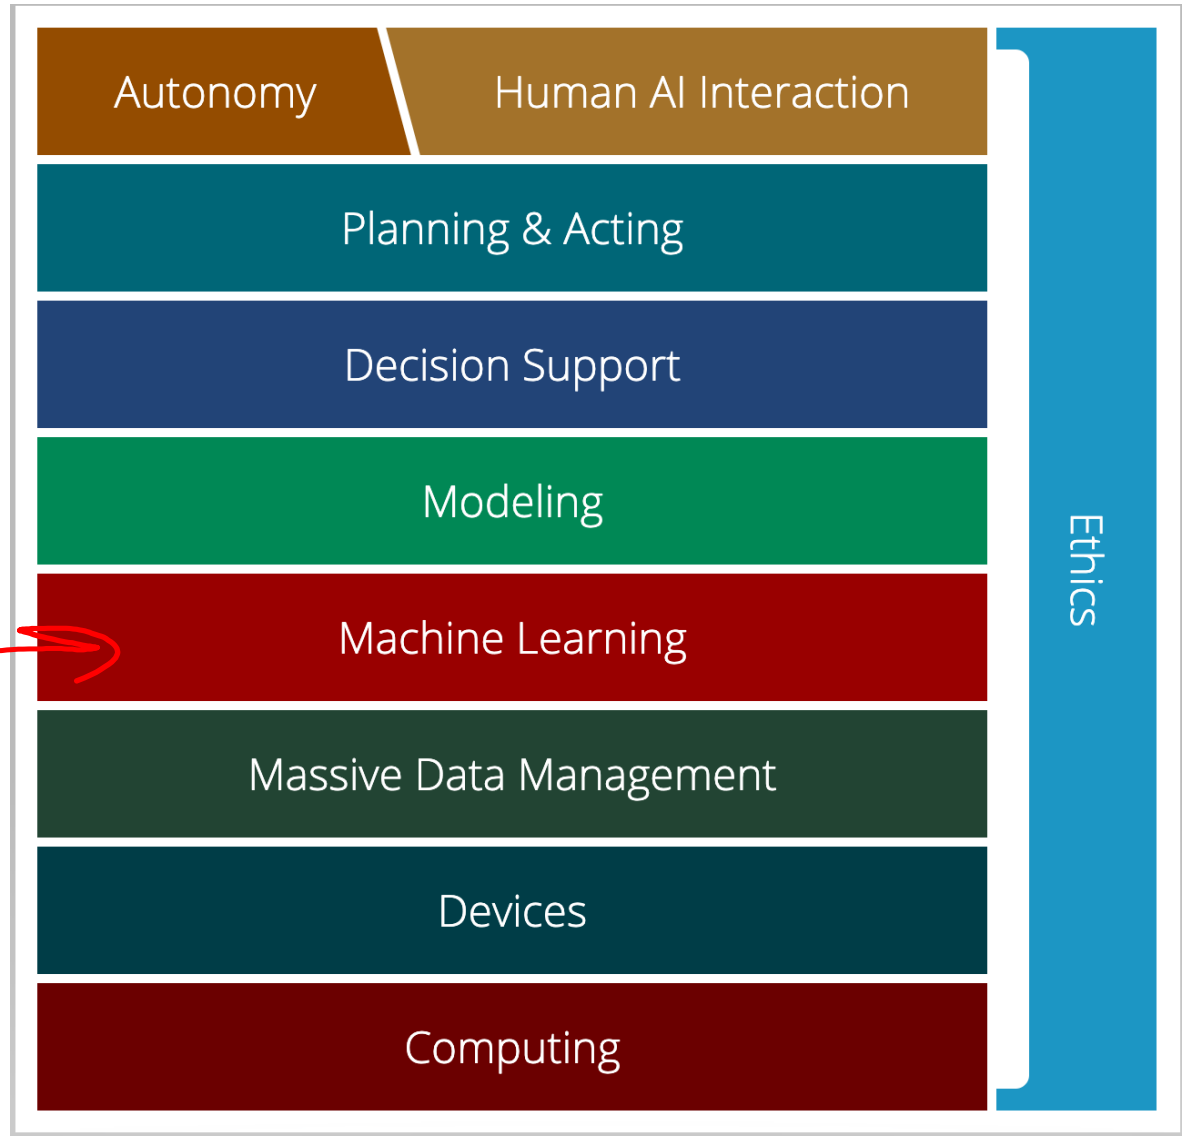

https://ai.cs.cmu.edu/about

Artificial Intelligence vs Machine Learning?

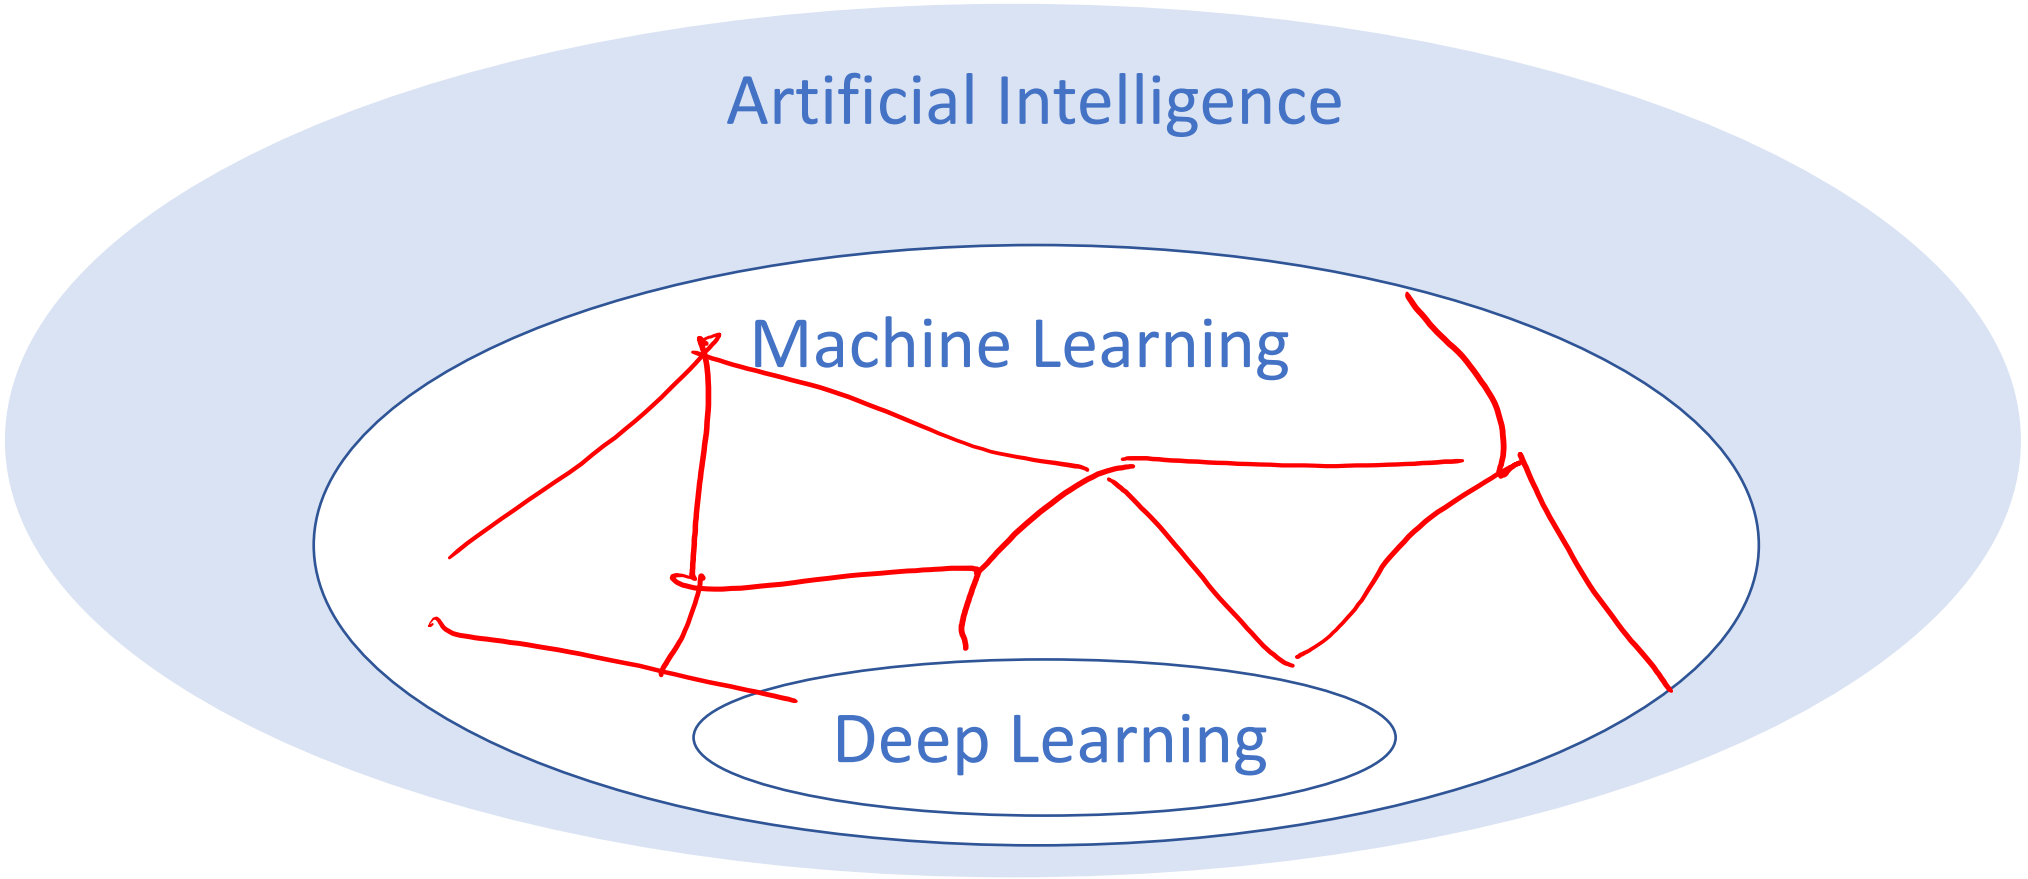

A Brief History of AI

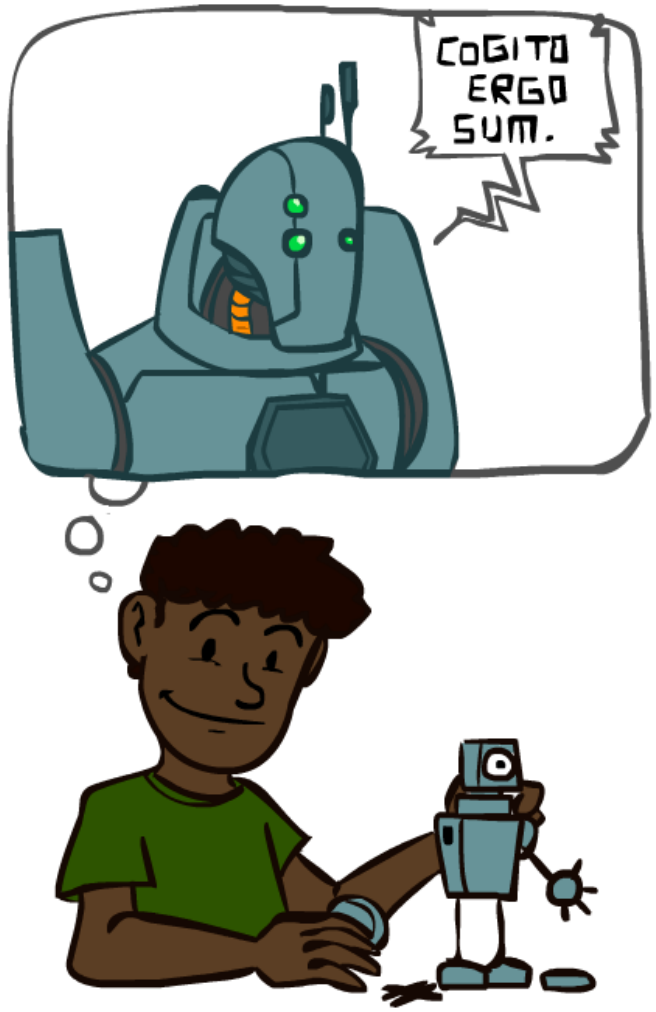

Images: ai.berkeley.edu

A Brief History of AI

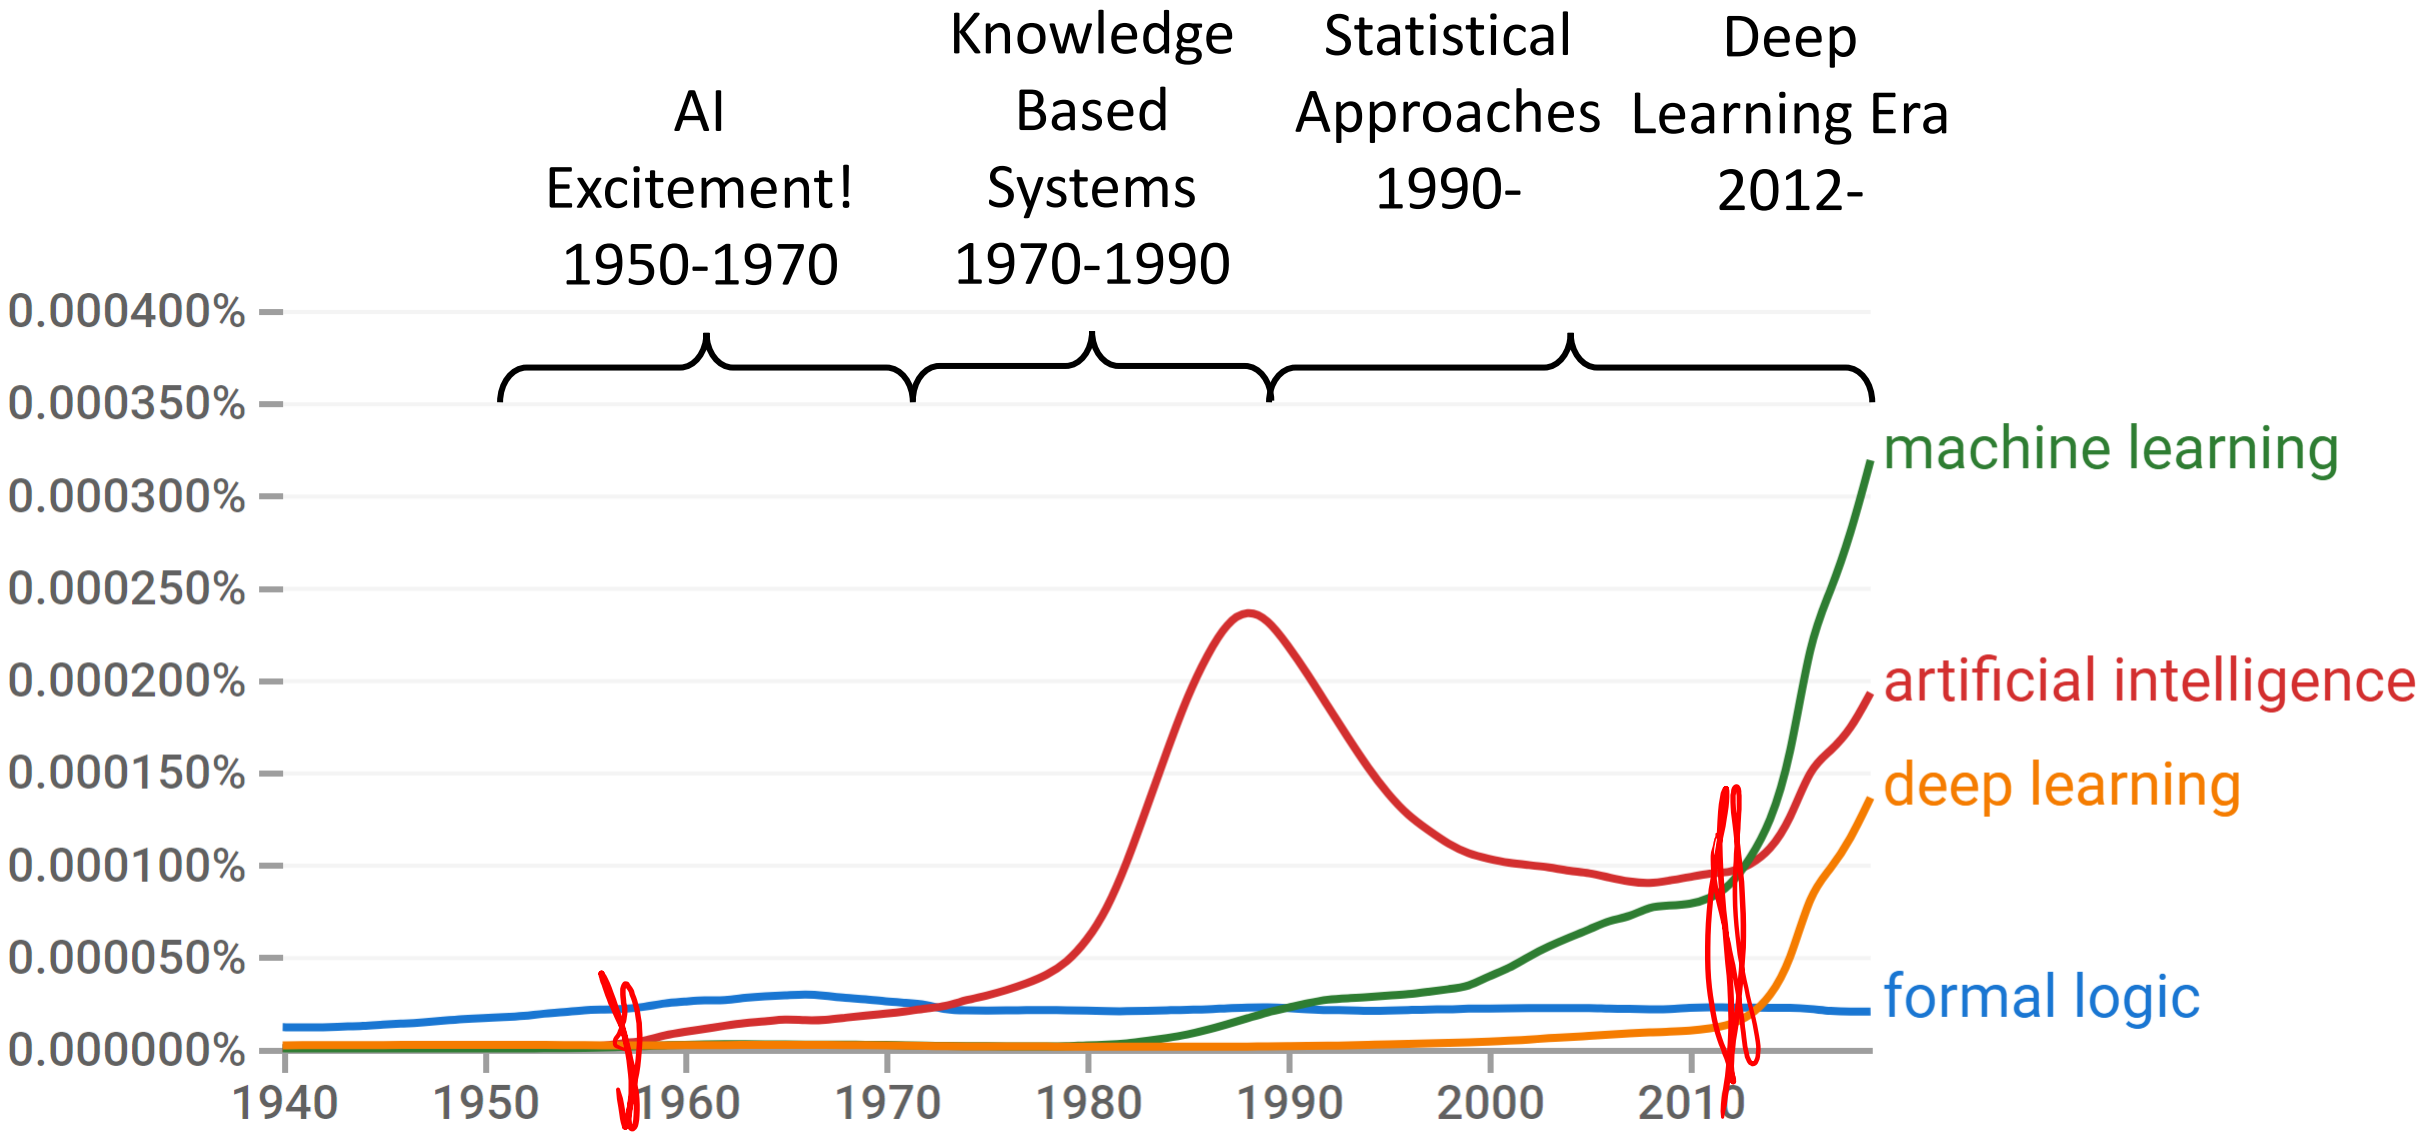

https://books.google.com/ngrams

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

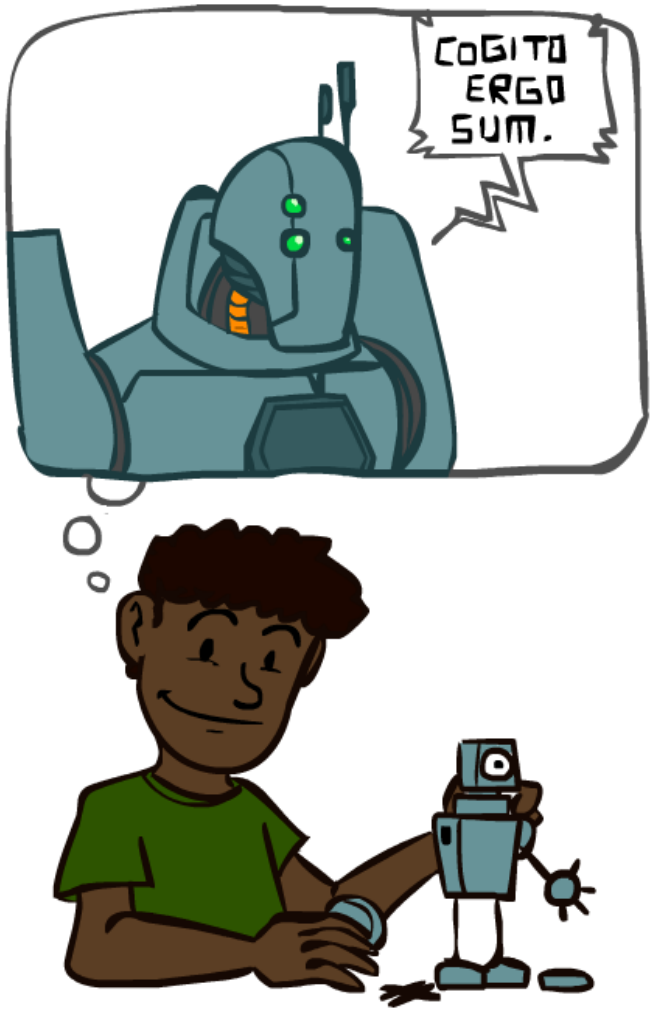

Images: ai.berkeley.edu

ML Applications?

Speech Recognition

1. Learning to recognize spoken words

THEN	NOW
"…the SPHINX system (e.g. Lee 1989) learns speaker-specific strategies for recognizing the primitive sounds (phonemes) and words from the observed speech signal…neural network methods…hidden Markov models…" (Mitchell, 1997)	

Source: https://www.stonetemple.com/great-knowledge-box-showdown/#VoiceStudyResults

8

Robotics

2. Learning to drive an autonomous vehicle

THEN	NOW
"…the ALVINN system (Pomerleau 1989) has used its learned strategies to drive unassisted at 70 miles per hour for 90 miles on public highways among other cars…" (Mitchell, 1997)	 waymo.com

9

Games / Reasoning

3. Learning to beat the masters at board games

THEN	NOW
"…the world's top computer program for backgammon, TD-GAMMON (Tesauro, 1992, 1995), learned its strategy by playing over one million practice games against itself…" (Mitchell, 1997)	

11

Computer Vision

4. Learning to recognize images

THEN	NOW
"…The recognizer is a convolution network that can be spatially replicated. From the network output, a hidden Markov model produces word scores. The entire system is globally trained to minimize word-level errors…." (LeCun et al., 1995)	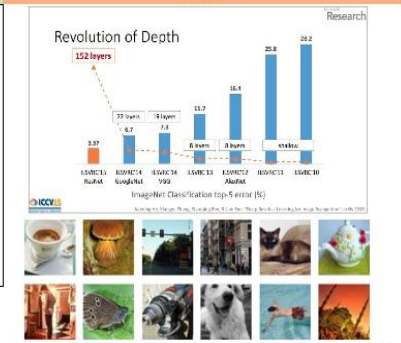

Images from https://blog.openai.com/generative-models/

12

Learning Theory

- 5. In what cases and how well can we learn?

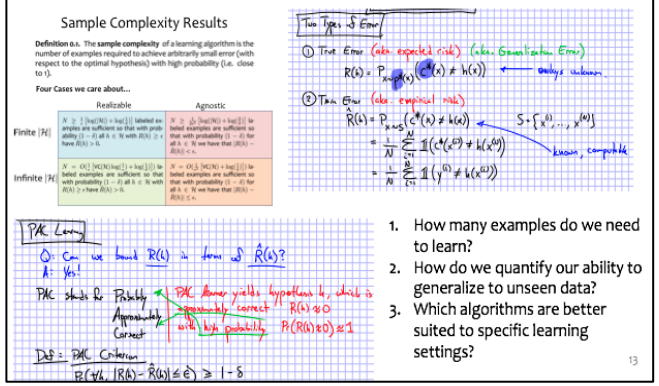

1. How many examples do we need to learn?
2. How do we quantify our ability to generalize to unseen data?
3. Which algorithms are better suited to specific learning settings?

13

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
  - Not much awareness of computational complexity

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

Machine Learning and Statistics

From the mid 1990s:

The two fields have effectively merged

- Carnegie Mellon has led the way!

ML is now often called “Statistical Machine Learning”

- There is very little non-statistical ML today

Slide: CMU ML, Roni Rosenfeld

Machine Learning Systems

Task

Experience

Performance measure

Machine Learning Framework

Formalize the task as a mapping from input to output

- Task examples will usually be pairs: (input, correct_output)

Formalize performance as an error measure

- or more generally, as an objective function (aka Loss function)

Examples:

- Medical Diagnosis
  - mapping input to one of several classes/categories (aka classification)
- Predict tomorrow's Temperature
  - mapping input to a number (aka regression)
- Chance of Survival: From patient data to p(survive >= 5 years)
  - mapping input to probability (aka logistic regression)
- Driving recommendation
  - mapping input into a plan (aka Planning)

Slide: CMU ML, Roni Rosenfeld

Input/Output Tasks

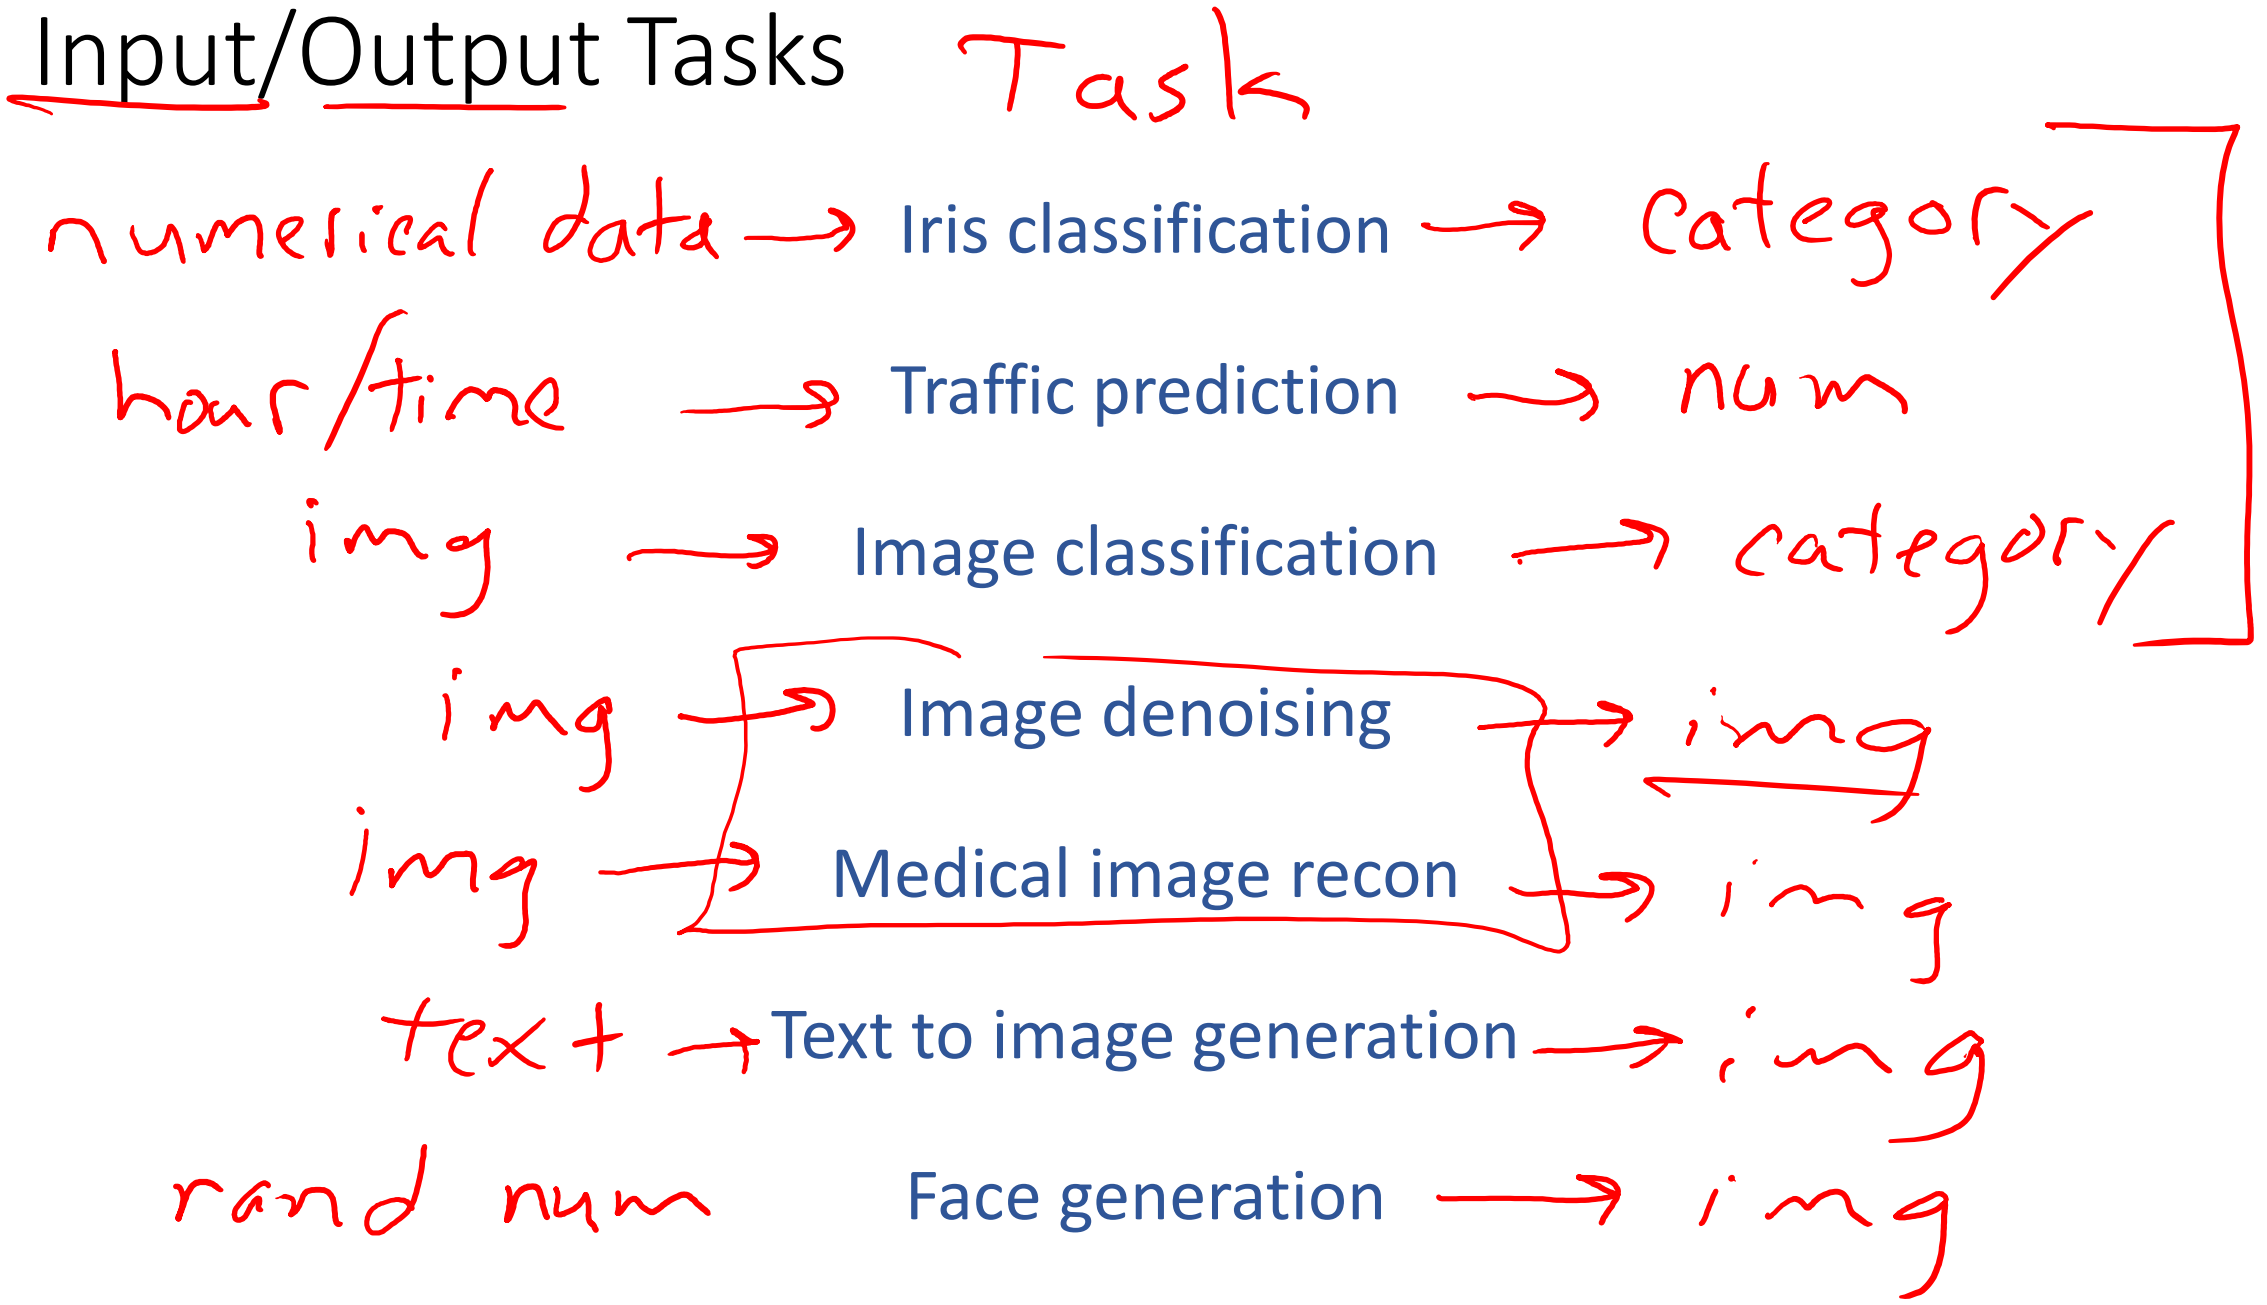

Elephant in the room

Image: Zootopia

Elephant in the room

Craiyon (Formerly DALL-E Mini)

Free online AI image generator from text

Image: Zootopia

Elephant in the room

Craiyon (Formerly DALL-E Mini)

Free online AI image generator from text

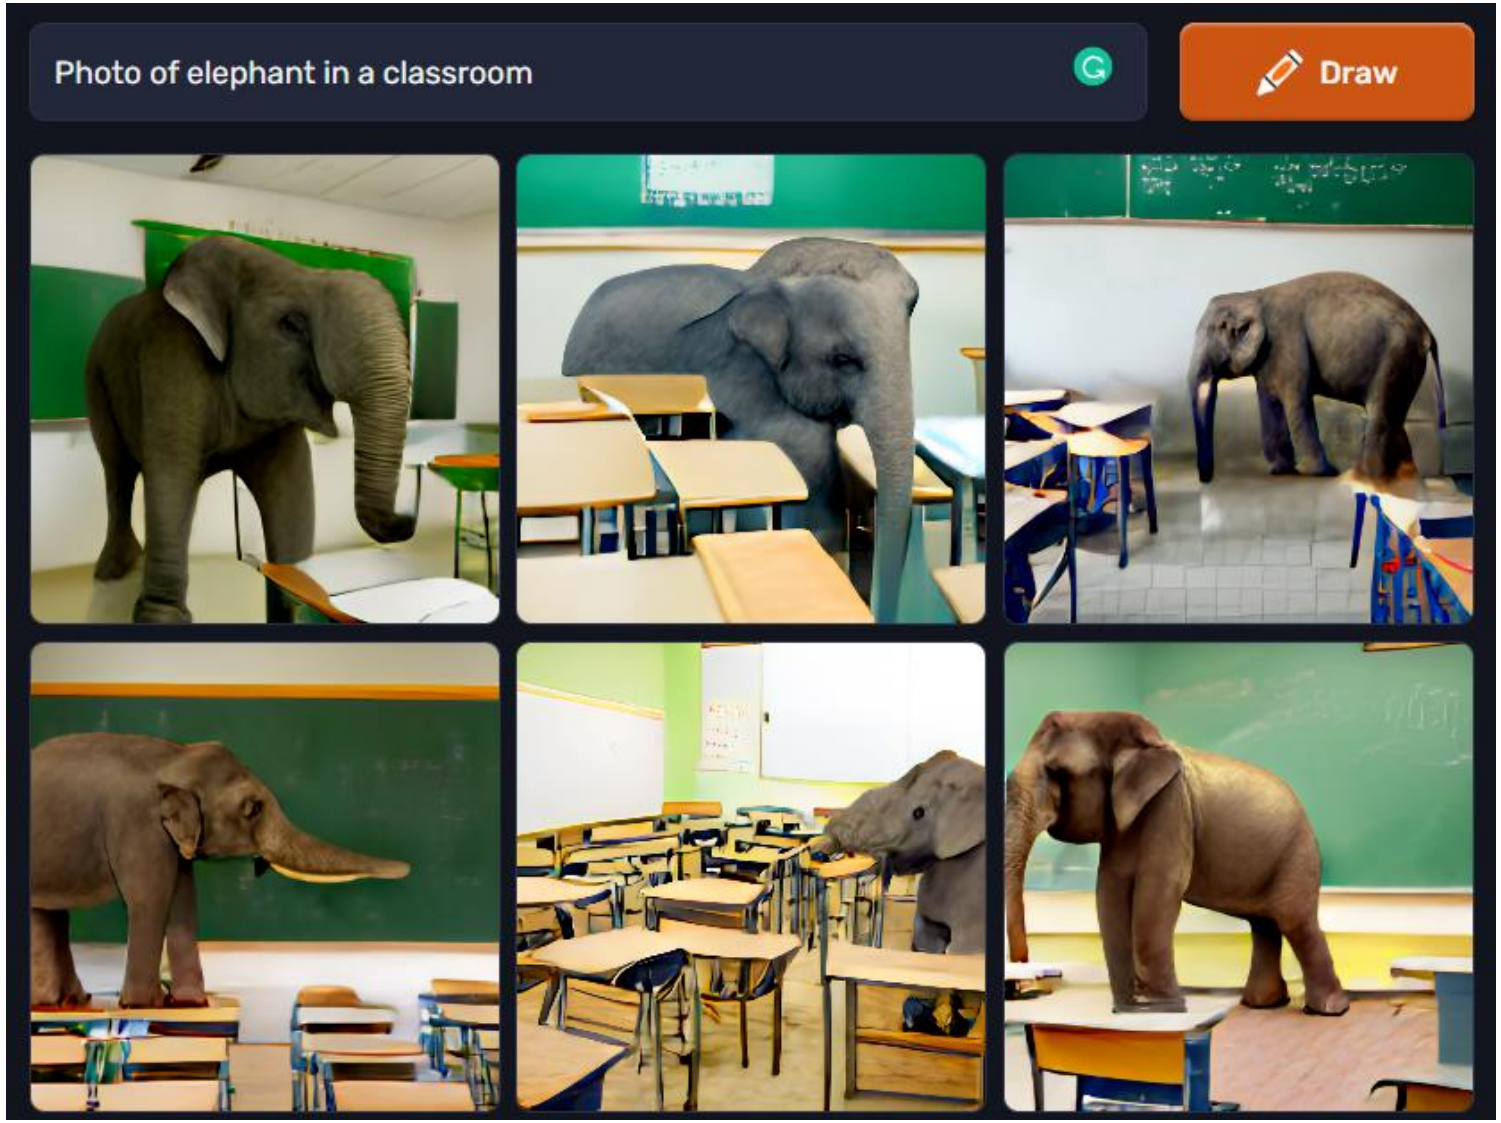

Image: Zootopia

Elephant in the room

Axios

AI could someday make medical decisions instead of your doctor

ChatGPT recently passed all three parts of the U.S. Medical Licensing Examination.

4 hours ago

Ad Age

What generative AI means for brands—a marketing guide to ChatGPT, DALL-E and other artificial intelligence

Chatbots and image generators are the latest technology piquing marketers' interest.

7 hours ago

Business Insider

ChatGPT is a 'game changer,' says Coursera CEO

Coursera CEO Jeff Maggioncalda said that he uses ChatGPT daily and plans to integrate the AI into his company's coursework despite its early...

1 day ago

Image: Google News: “chatgpt”, 1/18/2023

Dimensionality Reduction

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

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 →

Cartoon alien hi-res...

Freepik

Suesse Aliens Bilder -...

123RF

91,690 Cartoon Alien...

iStock

47,434 Alien Cartoon...

Shutterstock

228,873 Alien Cartoon...

Shutterstock

2,607 Alien Feet Image...

Shutterstock

228,873 Alien Cartoon...

Exercise: Human-defined Feature Space

Step1: Write a bunch of digits 0-9 on post-it notes

Exercise: Human-defined Feature Space

Step2: In groups, students try to organize digits on a 2-D coordinate plot

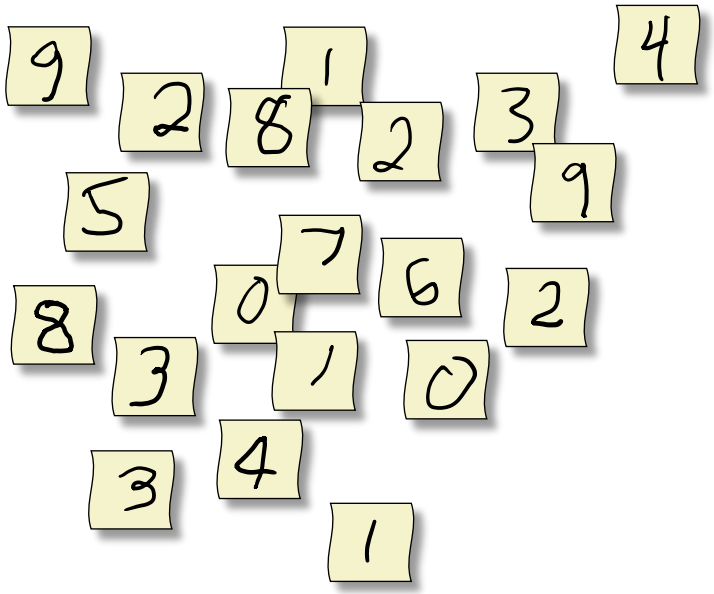

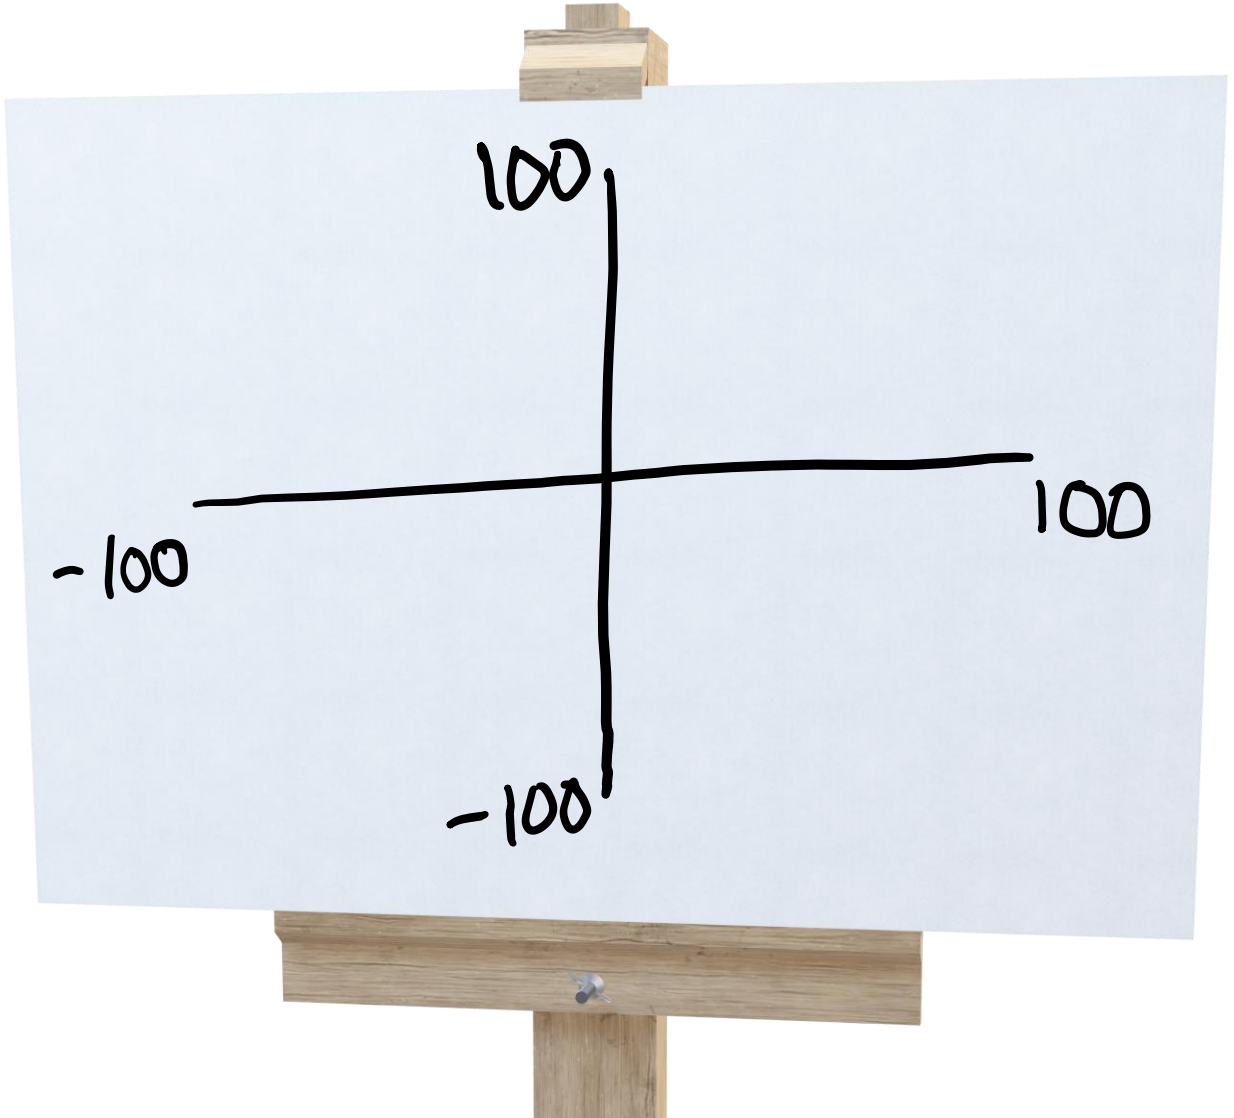

Optional: label the extreme ends of both coordinate axes

Exercise: Human-defined Feature Space

Step2: In groups, try to organize digits on a 2-D coordinate plot

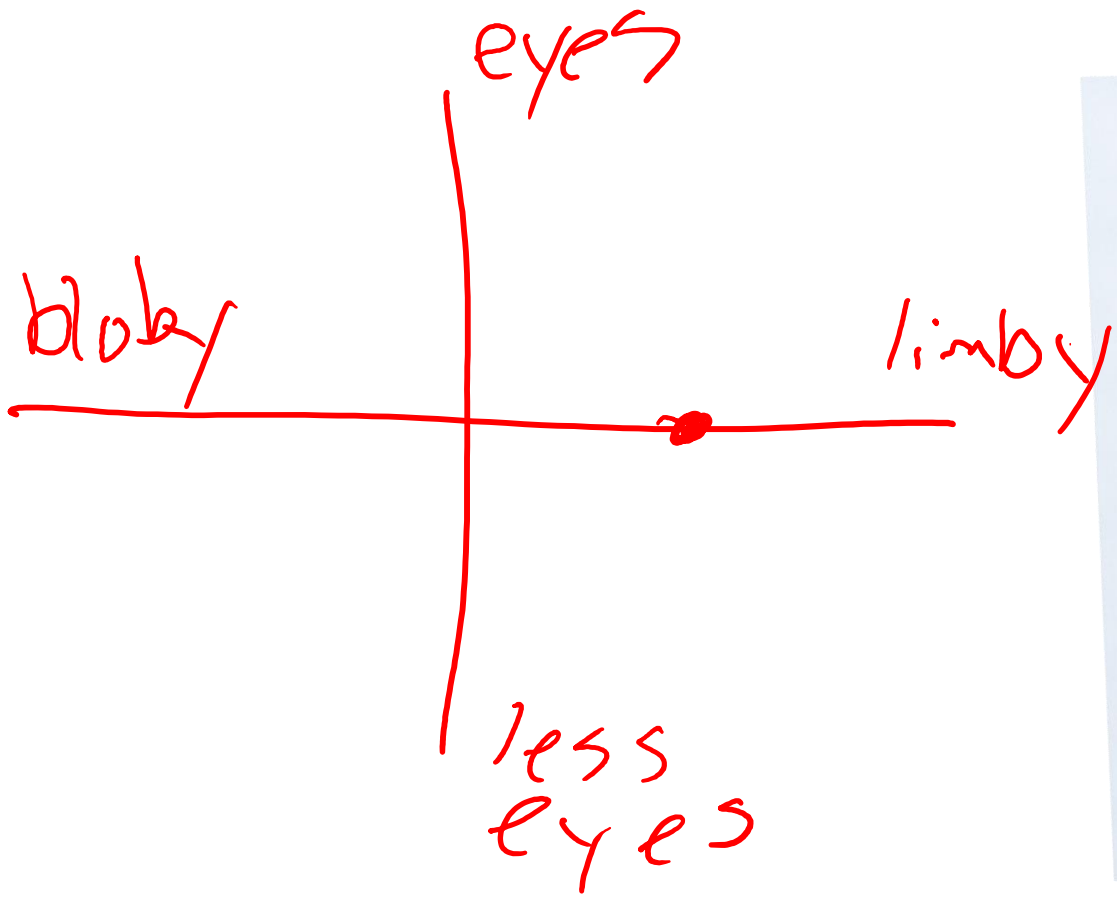

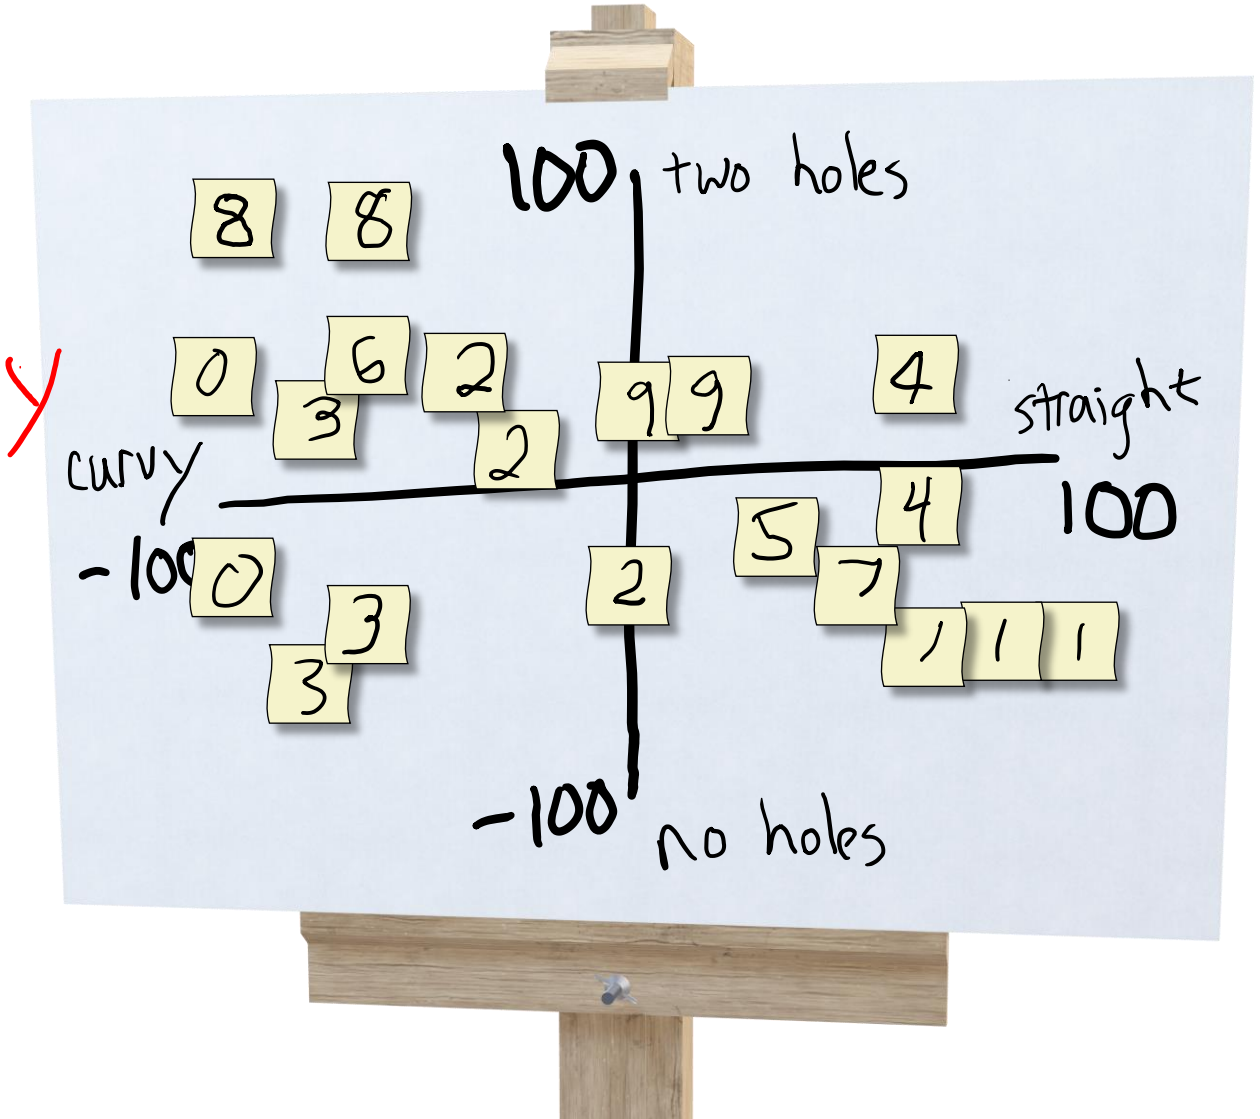

Optional: label the extreme ends of both coordinate axes

Exercise: Human-defined Feature Space

Step 3: Prediction!

1. Select three students: A,B,C
2. Student A draws a new digit and hands it to student B
3. Student B thinks about where to plot it and comes up with a 2-D coordinate, (x, y)
4. Student C looks at the coordinate and the plot (but not the drawing from A) and predicts the digit, 0-9

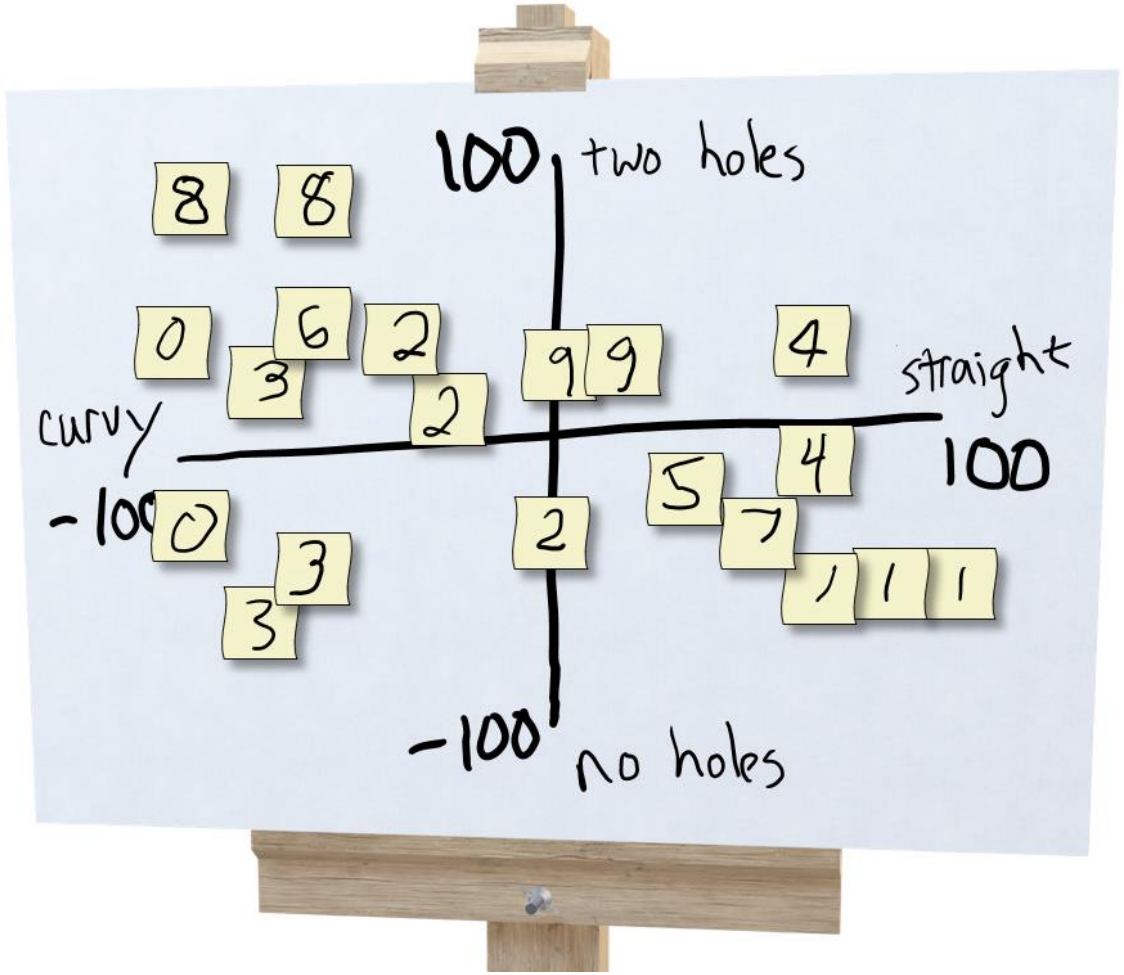

Exercise: Human-defined Feature Space

Step 4: Creation!

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

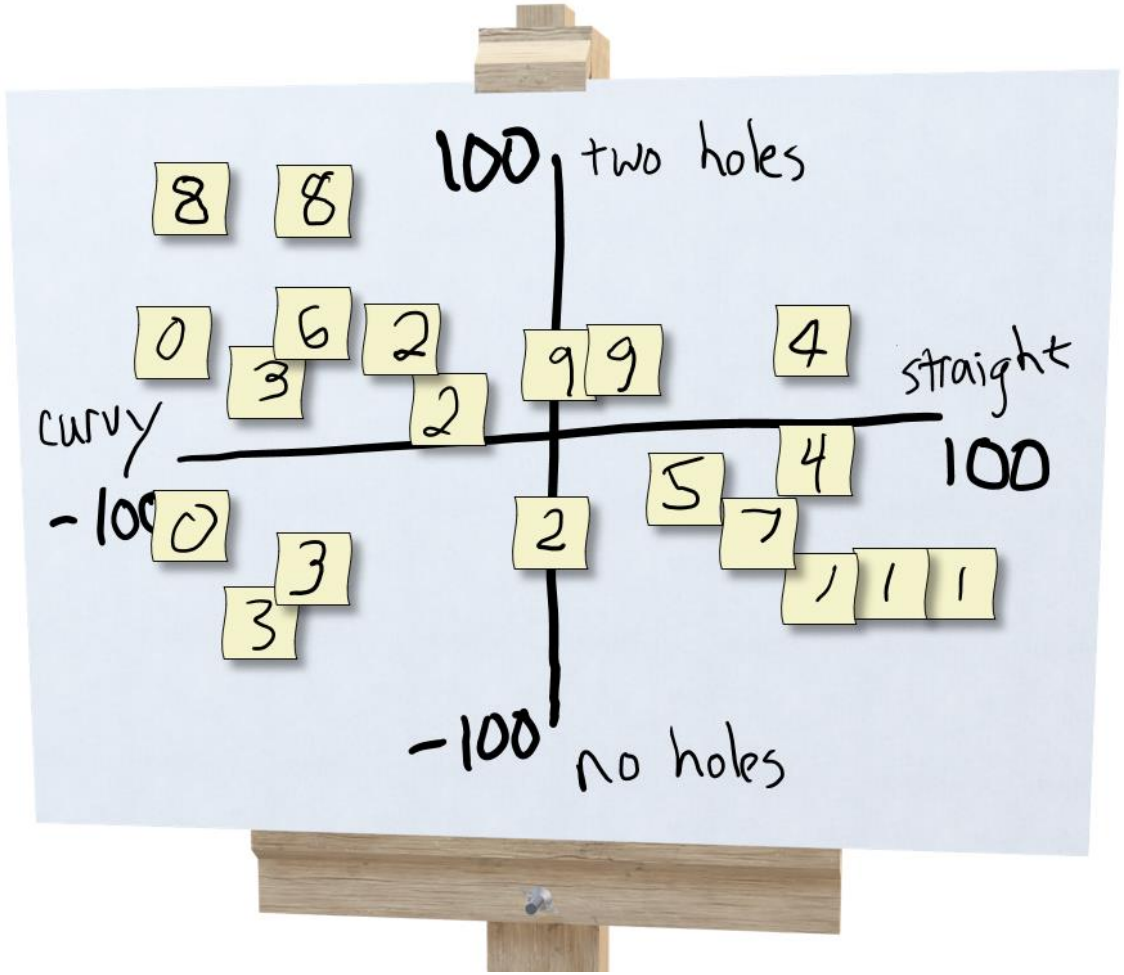

Exercise: Human-defined Feature Space

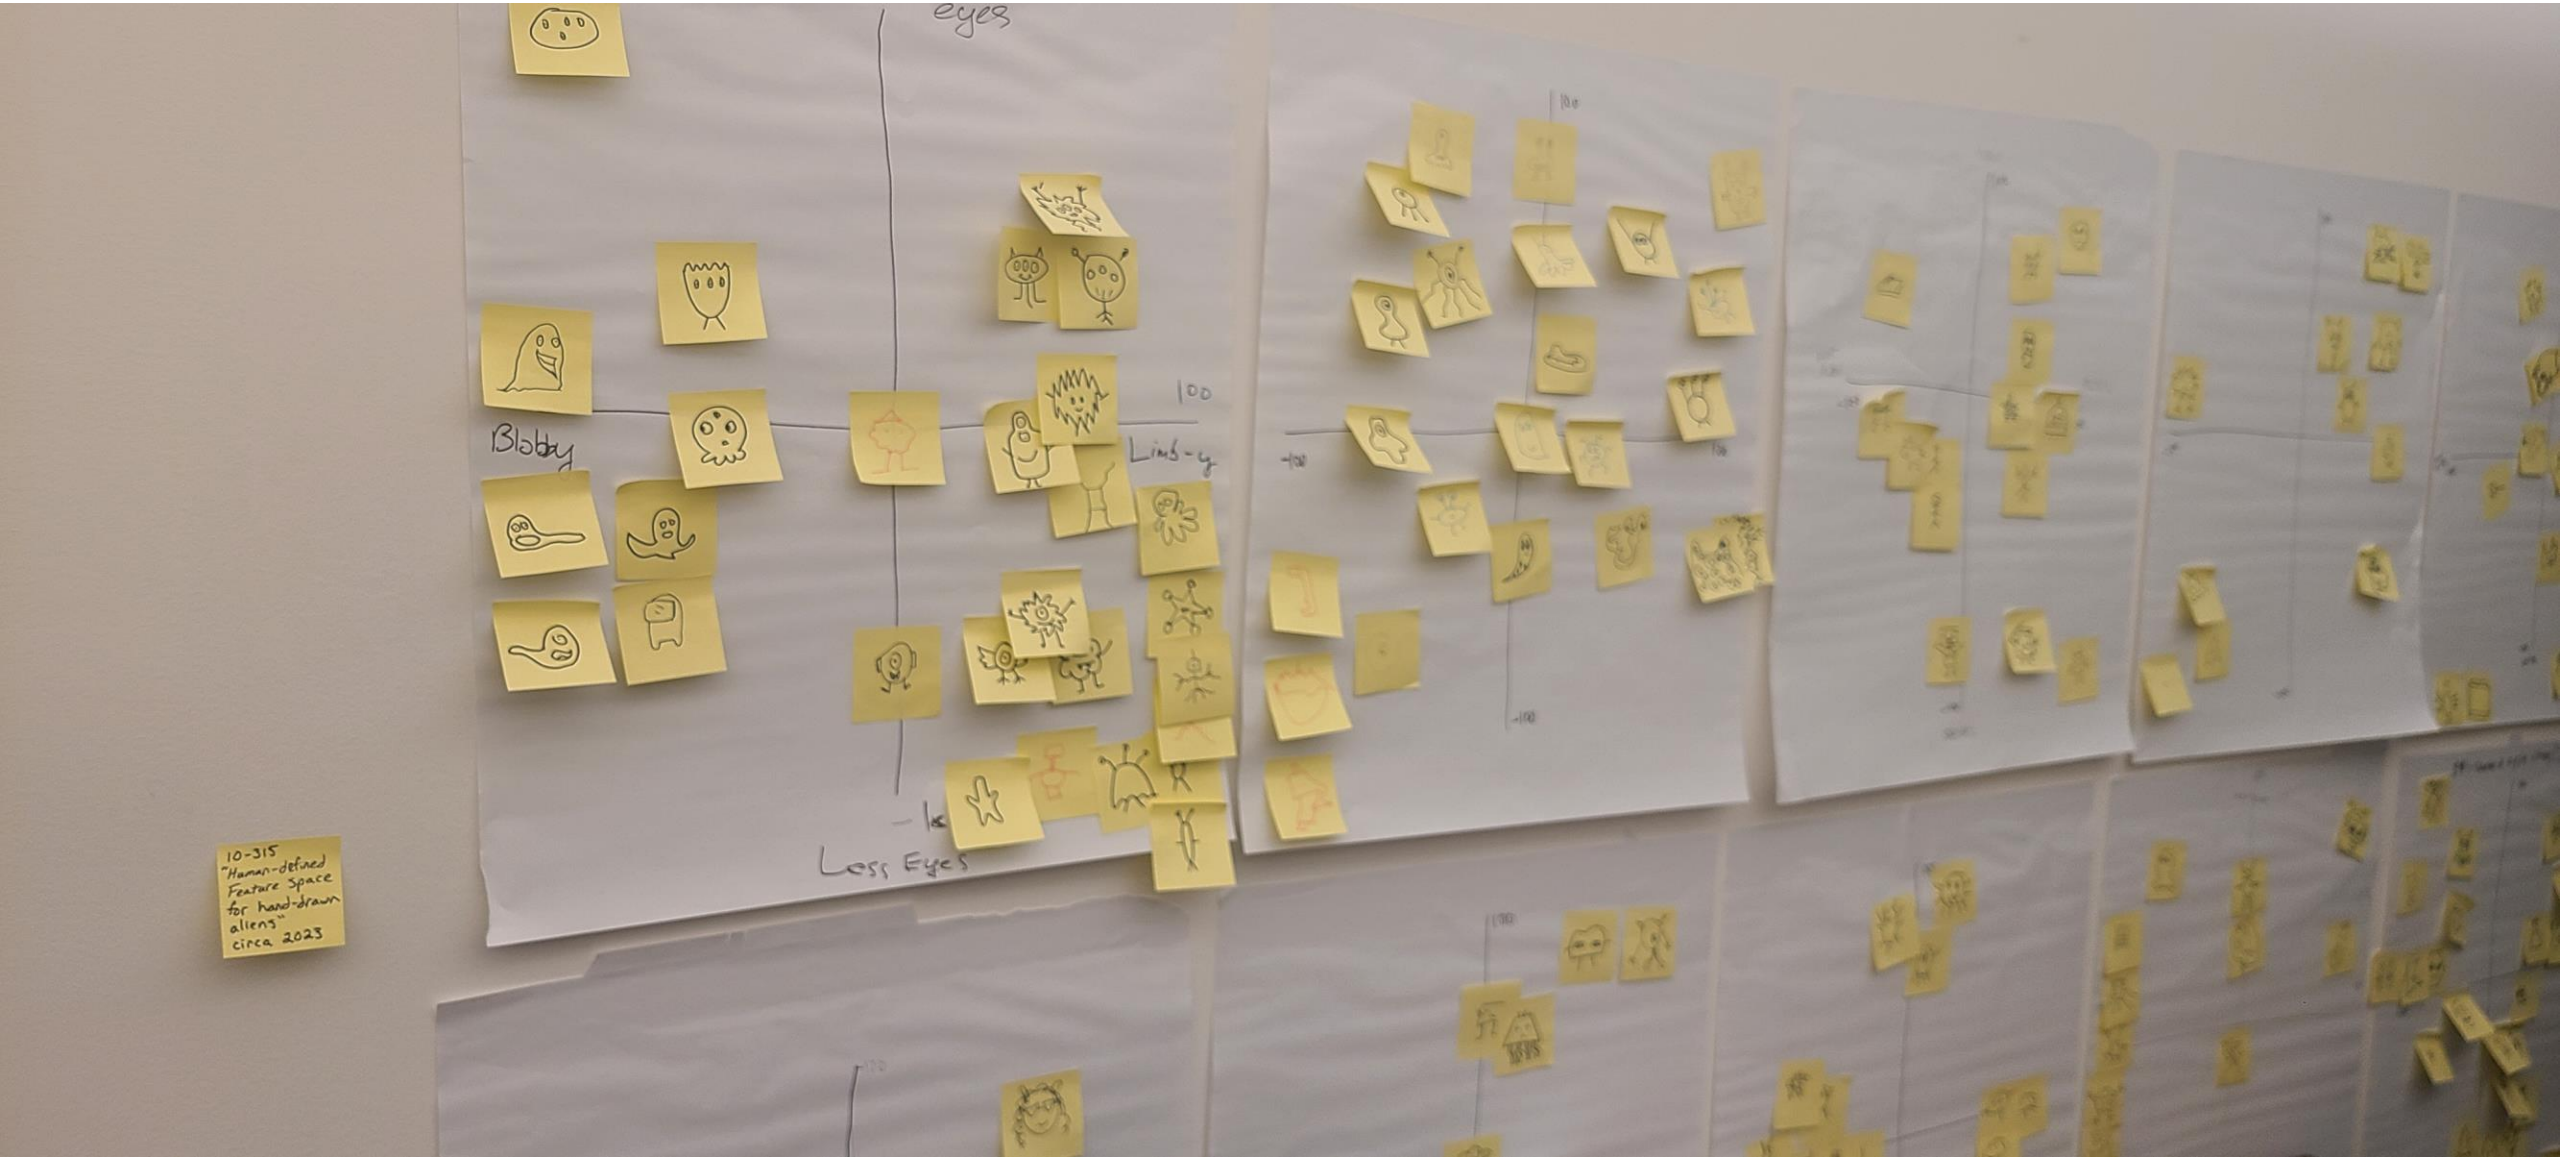

Learning to Organize Data

Neural networks can learn to organization too!

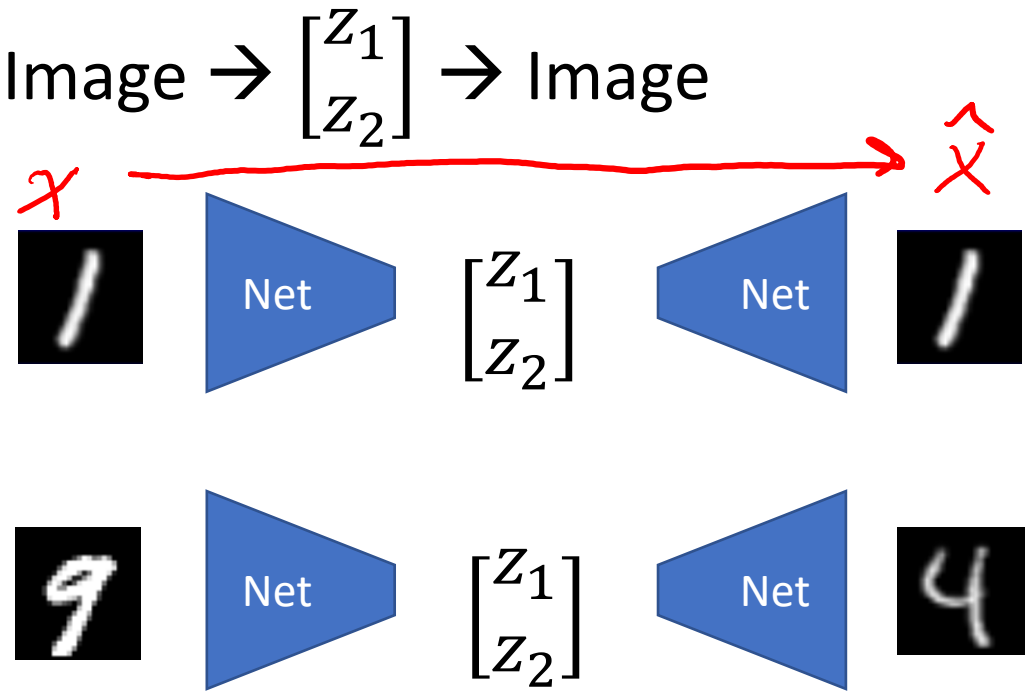

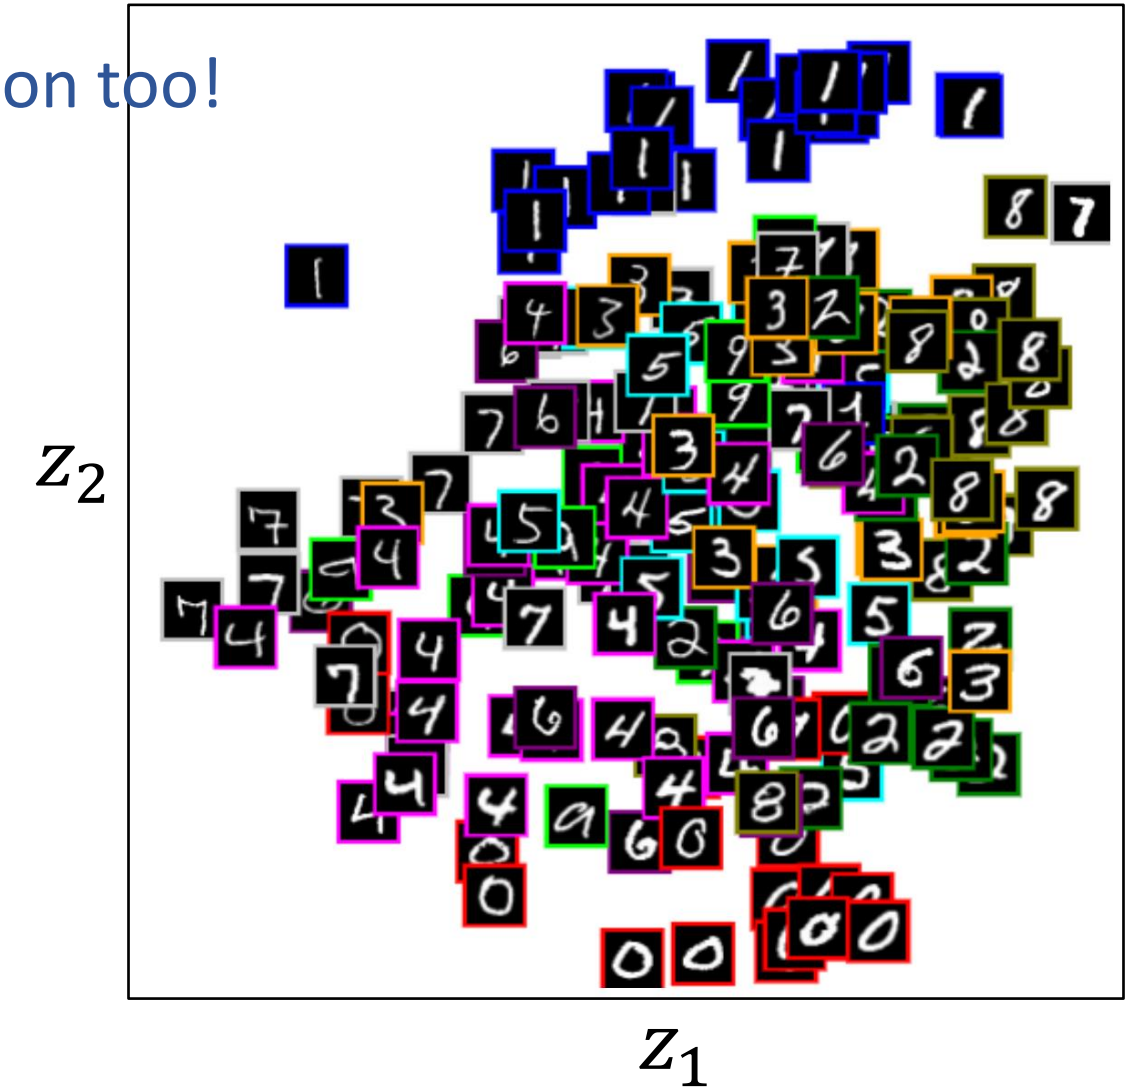

https://cs.stanford.edu/people/karpathy/convnetjs/demo/autoencoder.html

Digit Autoencoder

Demo: Using a learned feature space

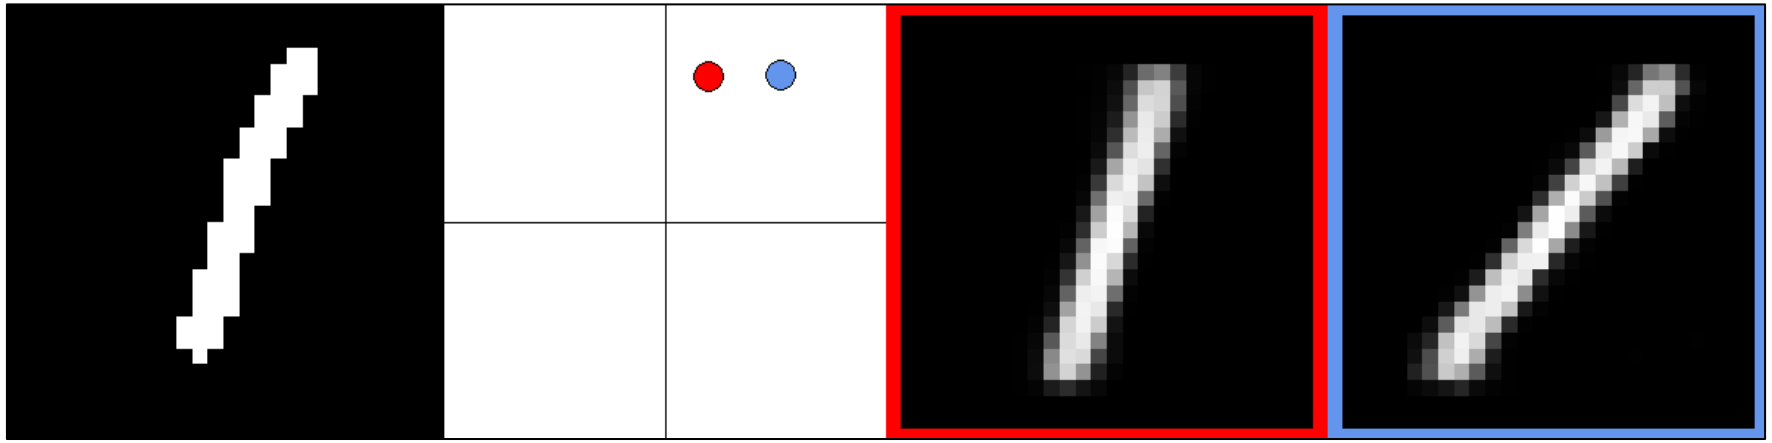

$\vec{x}$

$\vec{z}$

$\text{encoder}(\vec{x}) \rightarrow \begin{bmatrix} z_1 \\ z_2 \end{bmatrix} \rightarrow \text{decode}(\vec{z}) \rightarrow \hat{\vec{x}}$

Listen Learner

Chris Harrison, CMU

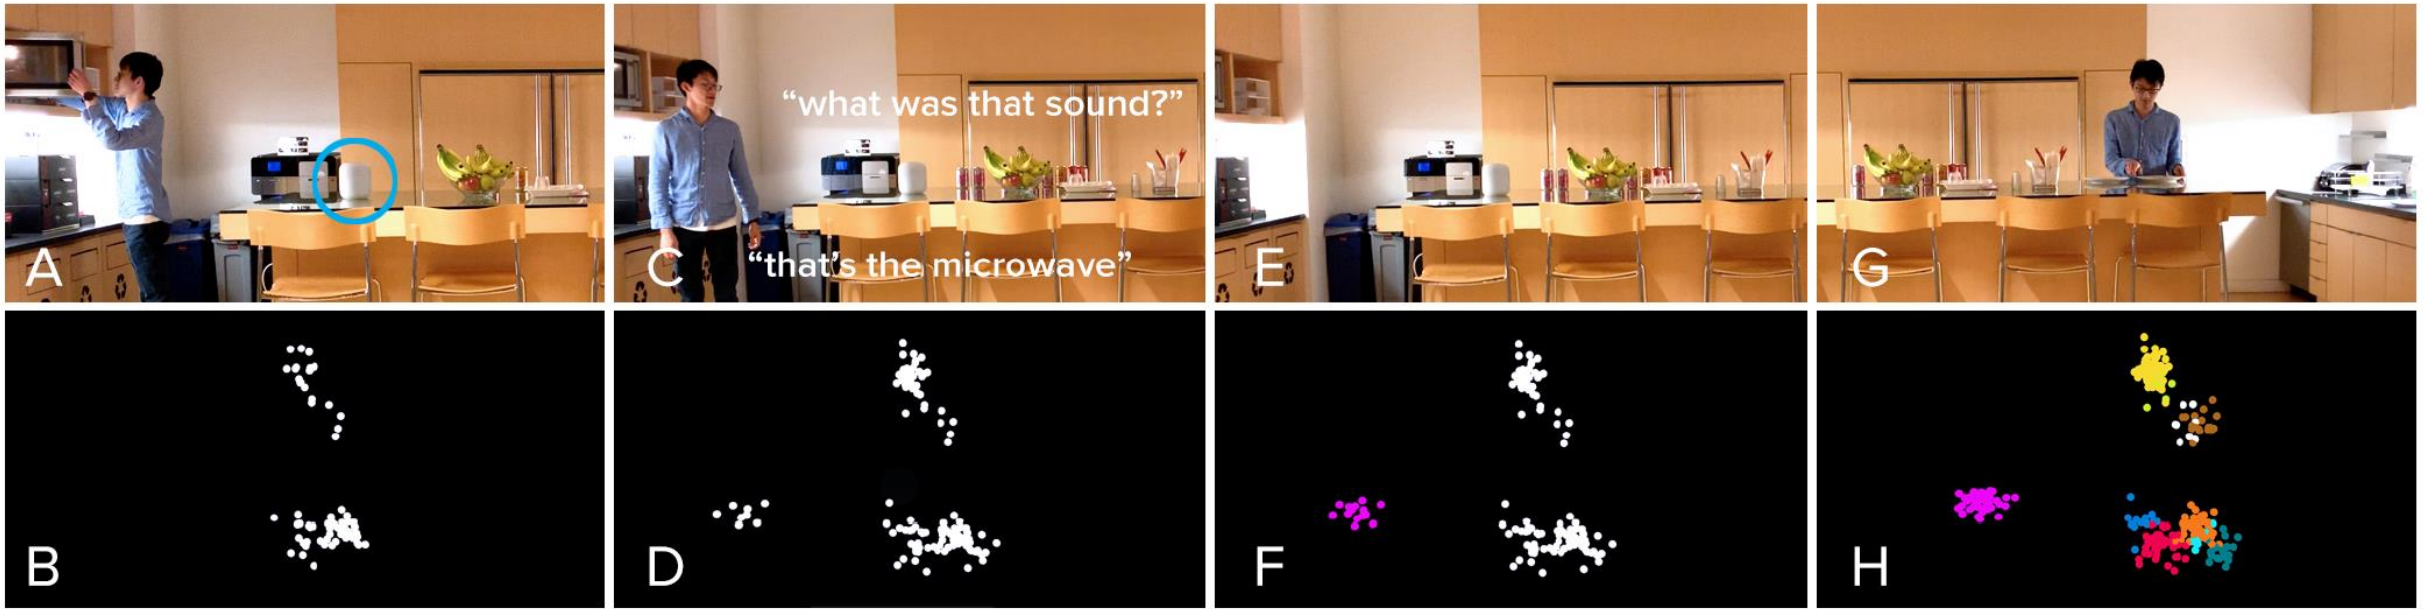

https://chrisharrison.net/index.php/Research/ListenLearner

Machine Learning

Using (training) data to learn a model that we'll later use for prediction

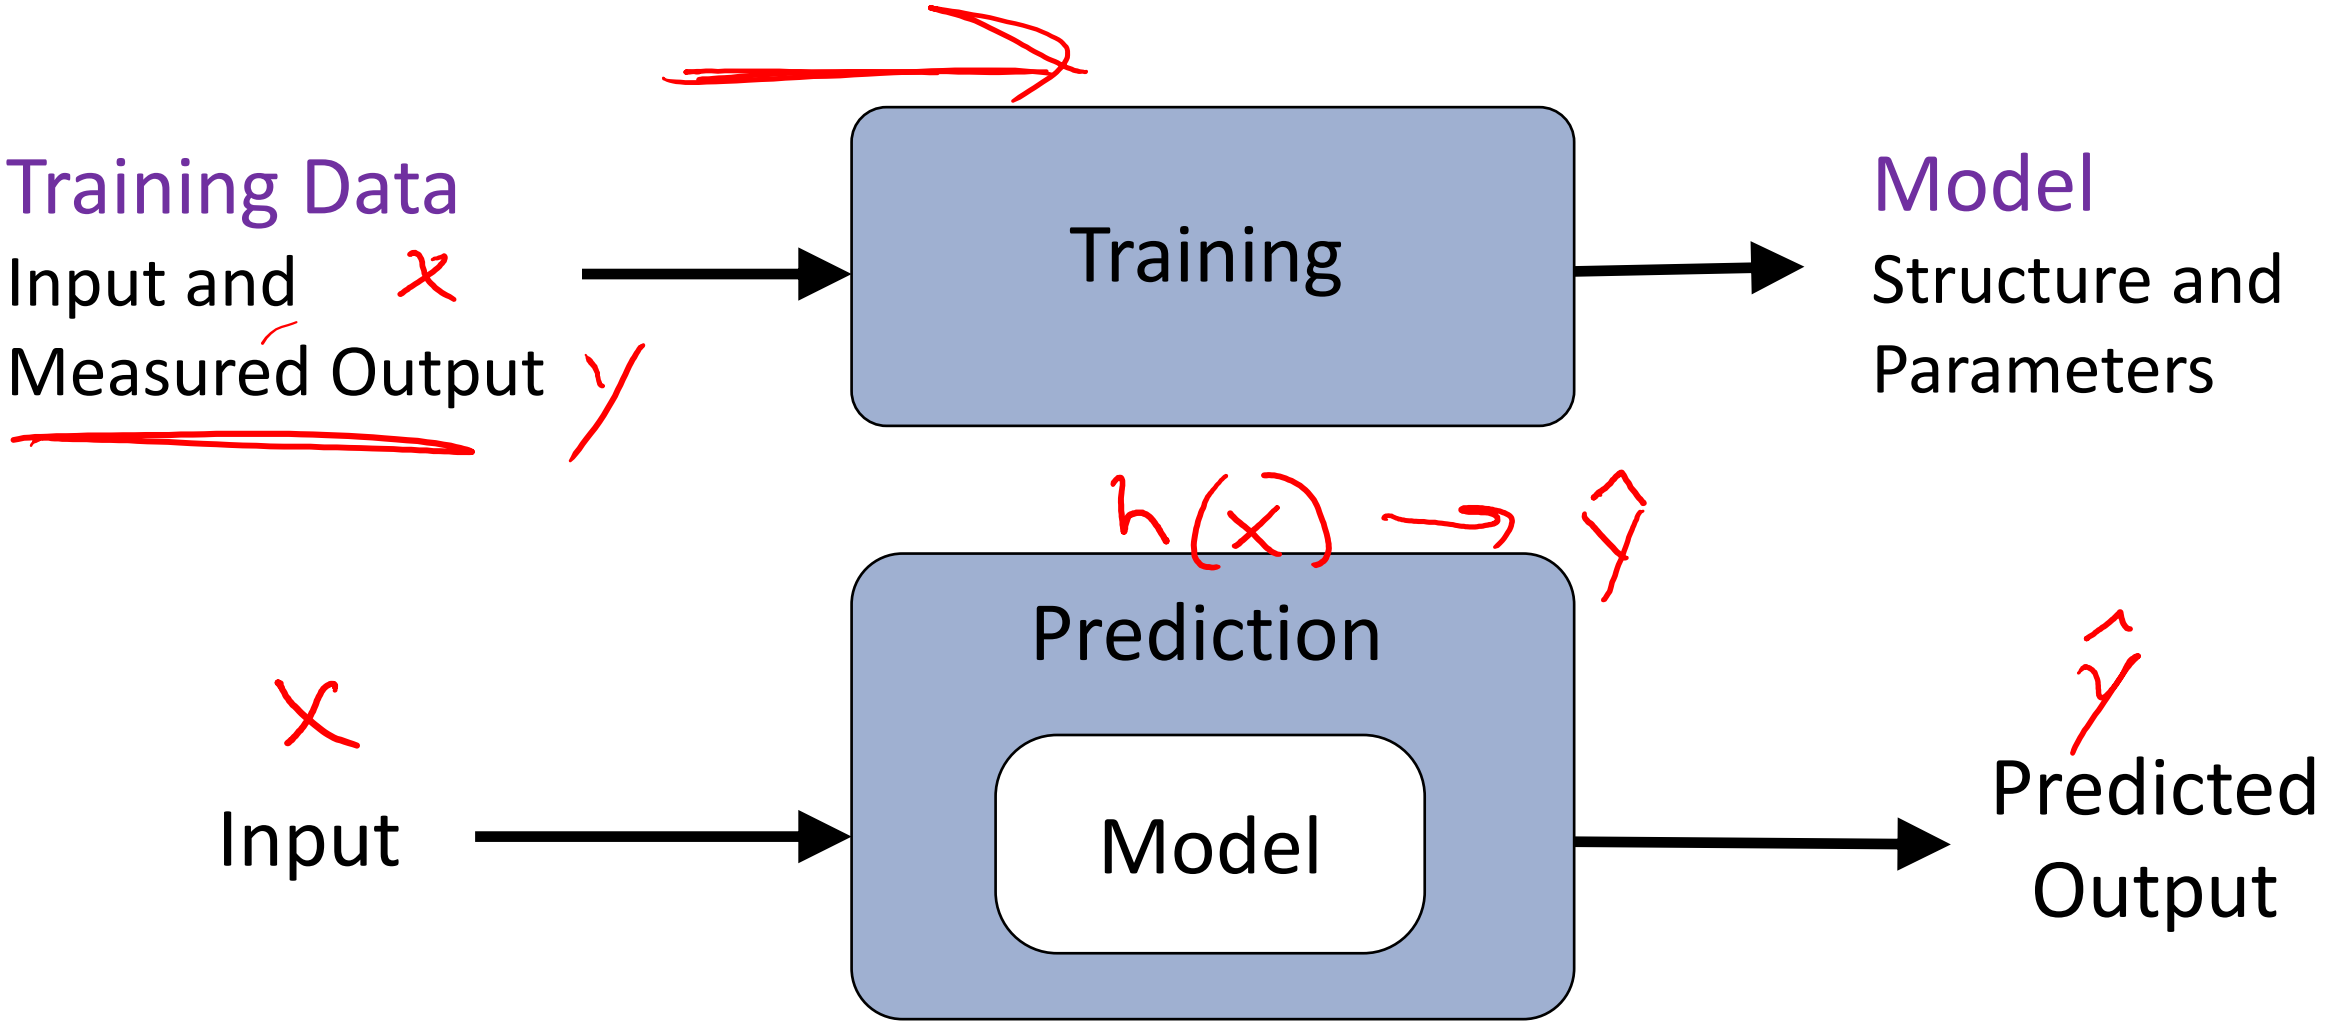

Autoencoder Model

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.

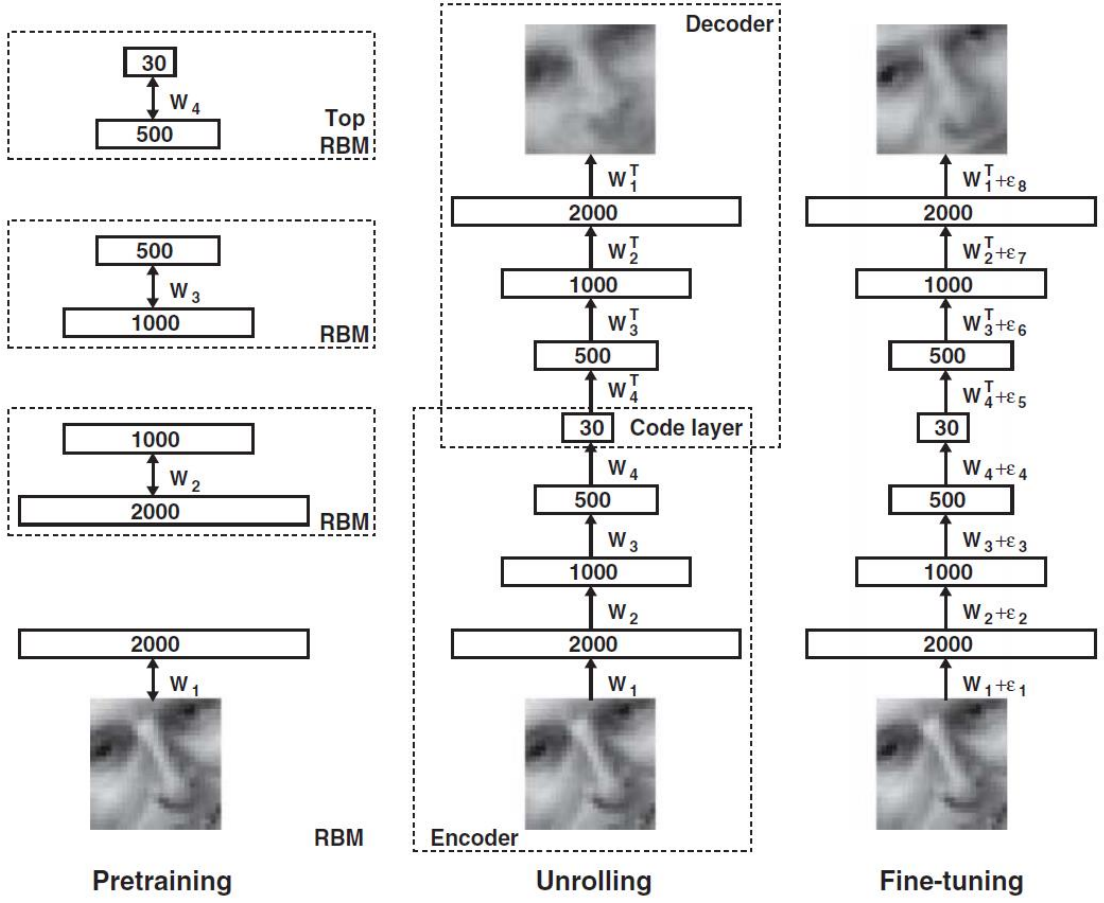

Digit Autoencoder

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

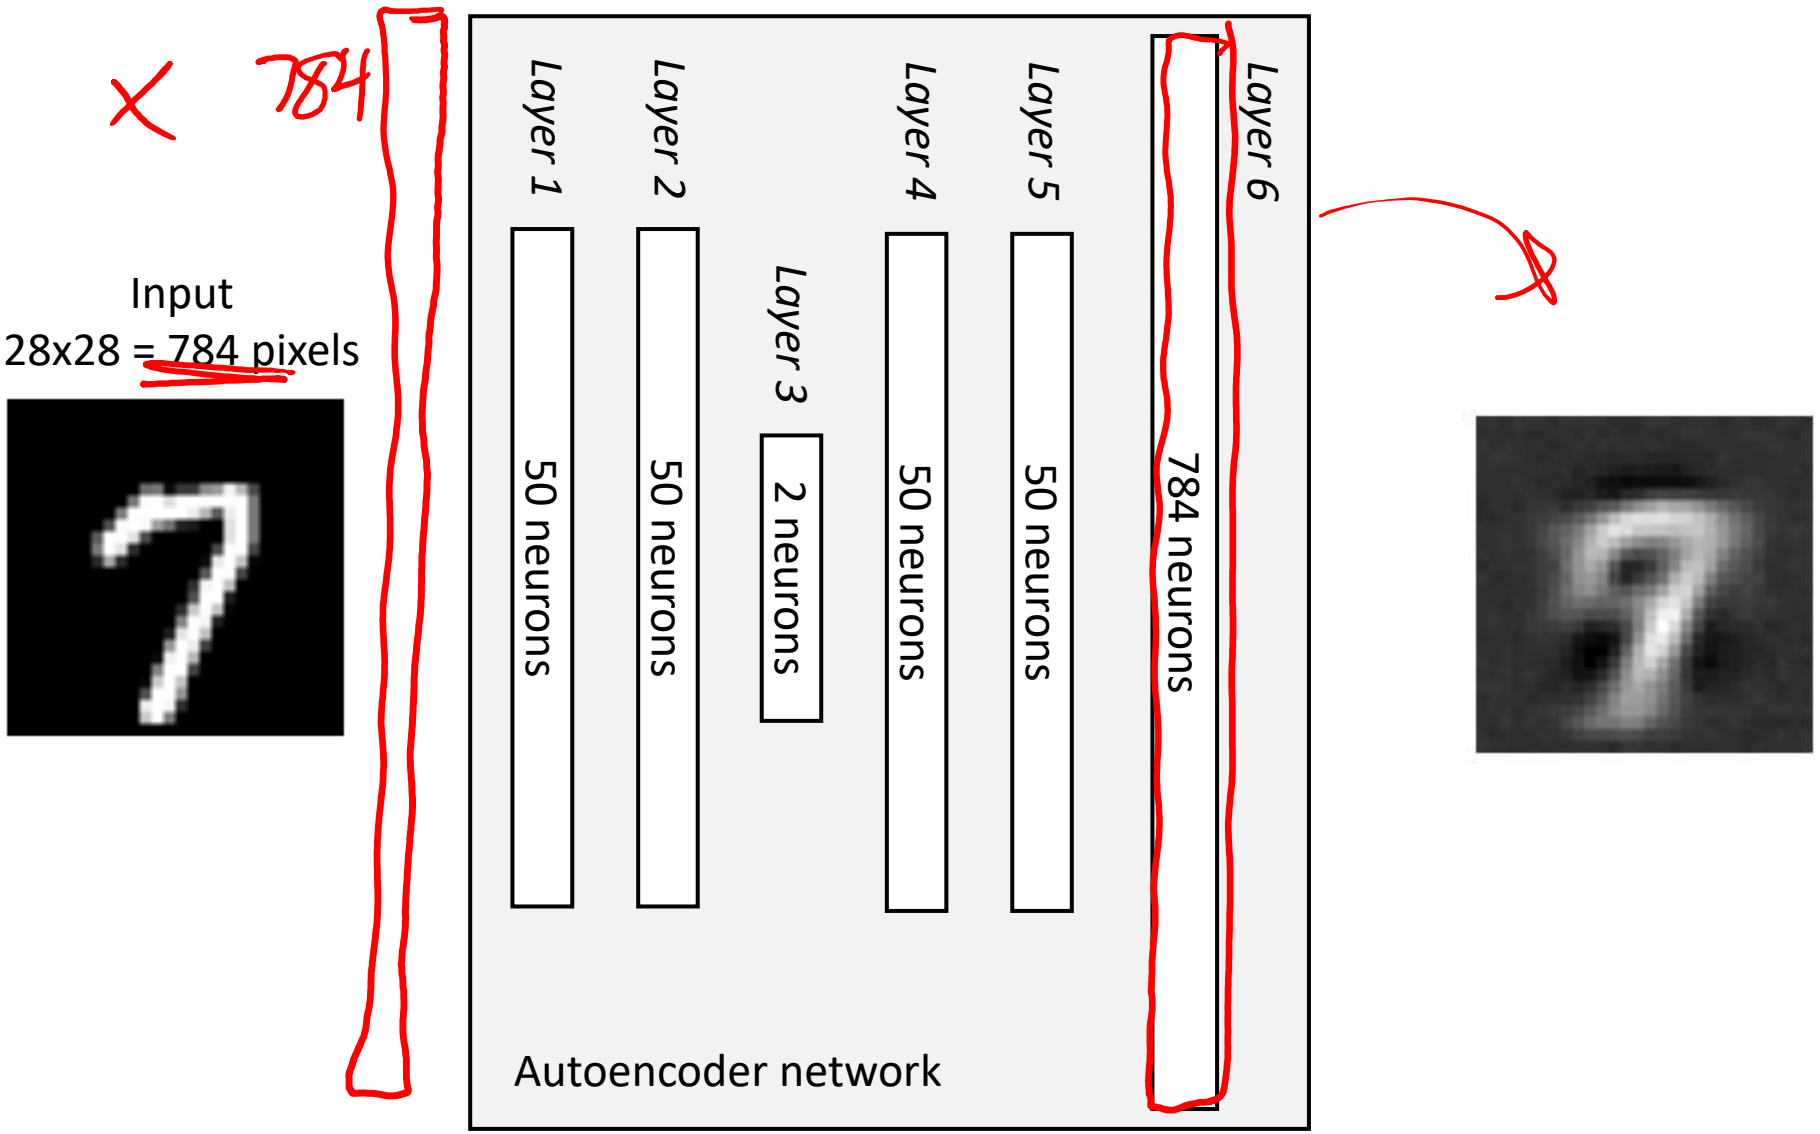

Digit Autoencoder

Math for autoencoder model

Linear layer $\vec{x}, W, \vec{b}$ $W\vec{x} + \vec{b}$

ReLU layer

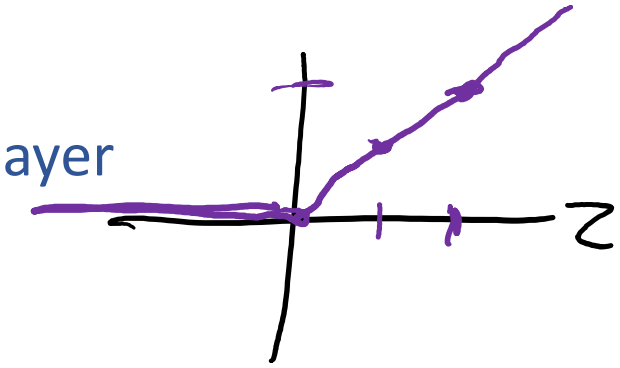

$\max(z, 0)$ $\max(\vec{z}, 0)$

Tanh layer

$\max\left(\begin{bmatrix} z_1 \\ z_2 \\ z_3 \end{bmatrix}, \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}\right)$

Course Information

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

Canvas: canvas.cmu.edu

Gradescope: gradescope.com

Communication:

piazza.com

E-mail (if piazza doesn't work):

joshminr@andrew.cmu.edu

pvirtue@andrew.cmu.edu

Course Information

Lectures

- Pre-reading before week of lecture
- Lectures are recorded
  - Shared with our course and ML course staff only
- Participation points earned by answering Piazza polls in lecture
- Slides will be posted

Recitations

- Recommended attendance
- Not recorded, no participation points in recitation
- Recitation materials are in-scope for quizzes and exams

Course Information

Office Hours

- OH calendar on course website
- OH-by-appointment requests are certainly welcome

Mental Health

Announcements

Updates for this week

- HW1
  - Out this evening
  - Due Sat 1/28, 11:59 pm
- Pick with pre-reading and check-point to get ready for next week
- Policies up on website by tomorrow morning
- Piazza polls will be practice today
  - Use this to work out any kinks