# Machine Learning (CS 4/5780)

**Kilian Q. Weinberger**
**Karthik Sridharan**
**Computer Science**

# Administrative Stuff



This baby says:
BORING!

# Course Overview
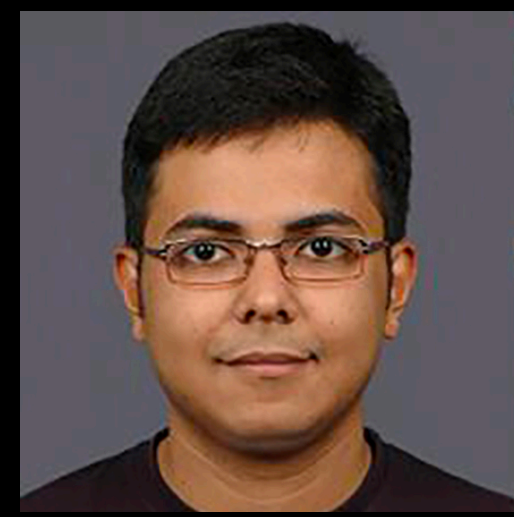
- Instructors:

  – Kilian Q. Weinberger, Karthik Sridharan

- Homepage:

  – **https://www.cs.cornell.edu/courses/cs4780/**

  – Vocareum: http://vocareum.com

  – Ed: https://edstem.org/us/courses/8451/

- TAs:

  – Many (more and better than you think)

- Office Hours / Recitations:

  - TA Office Hours: **Every day.** (Details will be posted on course webpage.)

  - Prof. Office Hours: **TBA**

- Questions:

  - Post all questions on ED (you can make them private)

  - **Do not email Professors directly (except in an emergency)**

  - Email staff at: **cs4780staff@gmail.com**

# Course resources

- Primary texts

  - Machine Learning a Probabilistic Perspective (K.P. Murphy)

  - The Elements of Statistical Learning (Hastie, Tibshirani, Friedman)

- Additional texts

  - An Introduction to Statistical Learning by James, Witten, Hastie, and Tibshirani

  - Patterns, Predictions, and Actions by Hardt and Recht

  - Fairness and Machine Learning by Barocas, Hardt, and Narayanan

- Background and programming resources on the website
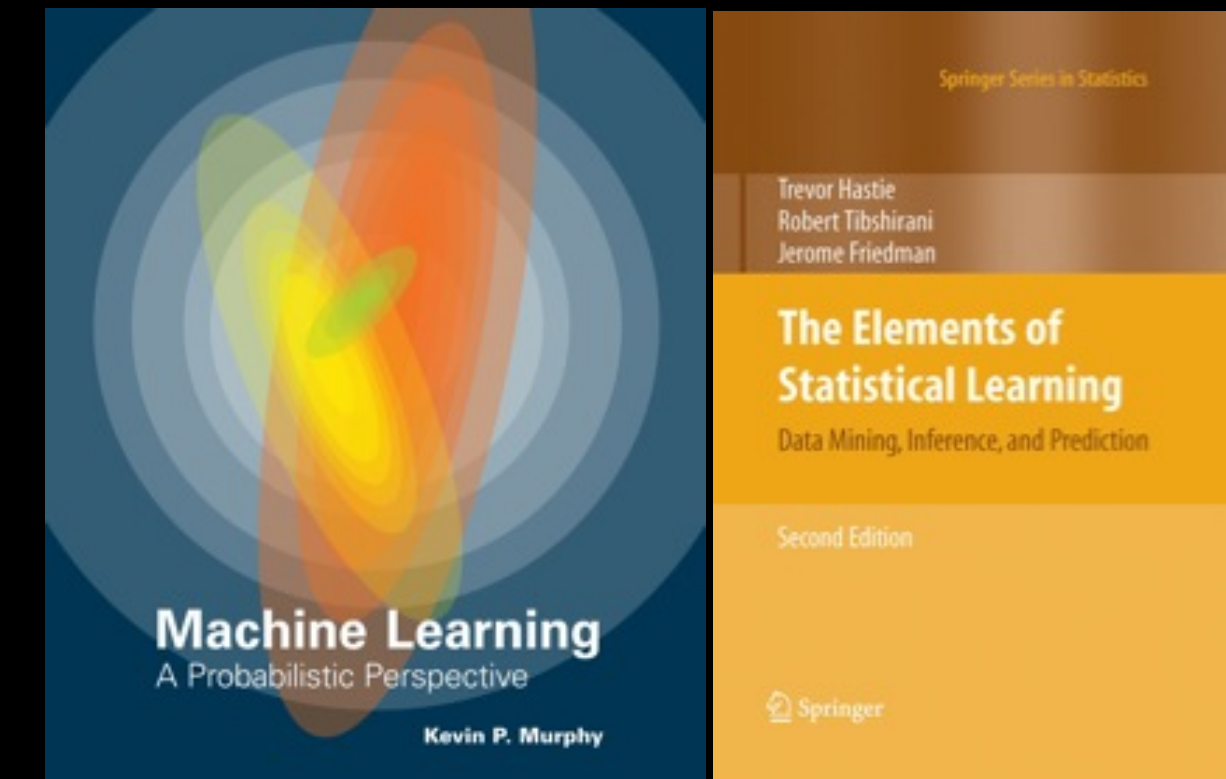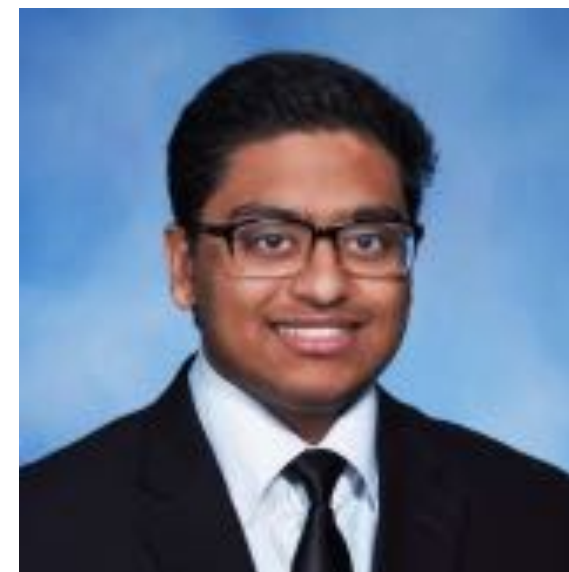
Kilian
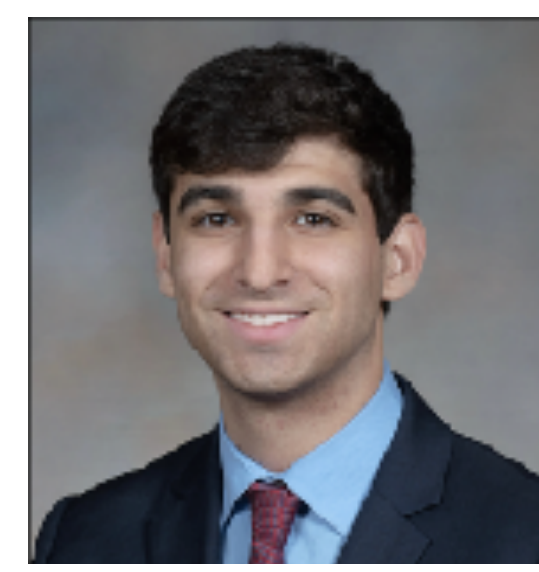"The Dark Forest"

Raphael
"Crime and Punishment"

Sally
"The Alignment Problem"

Tanvi
"The Fishermen
and the Dragon"

Mohammad
"The Queen's Thief"

Duncan
"Slaughterhouse-Five"

Jack
"Going Infinite"
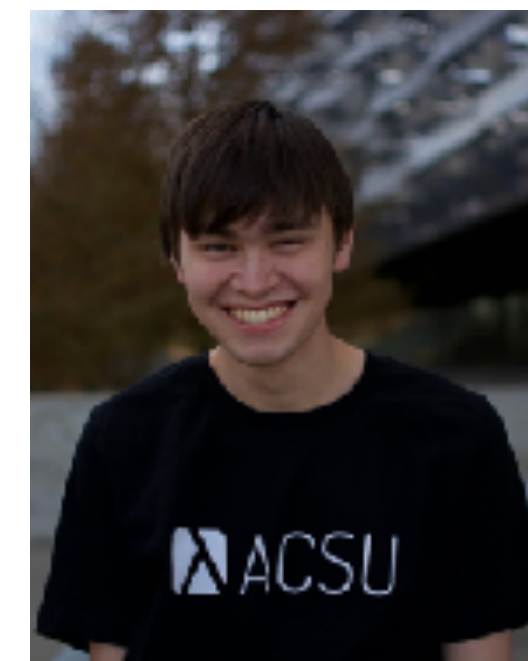
Emily W.
"Tender is the Flesh"

Prakriti
"The Girl on the Train"

Yuqing
"In The Blink of An
Eye" by Andrew Parker

Fengyu
"The Prince of Milk"

Tyler
"When Breath Becomes Air"

Joanna
"The Devotion of
Suspect X"

Afua
"Born A Crime"

Gloria G.
"The Anthropocene
Reviewed"

James
"The Sandman"

Yousef
"The Elements of
Statistical Learning"

Chris
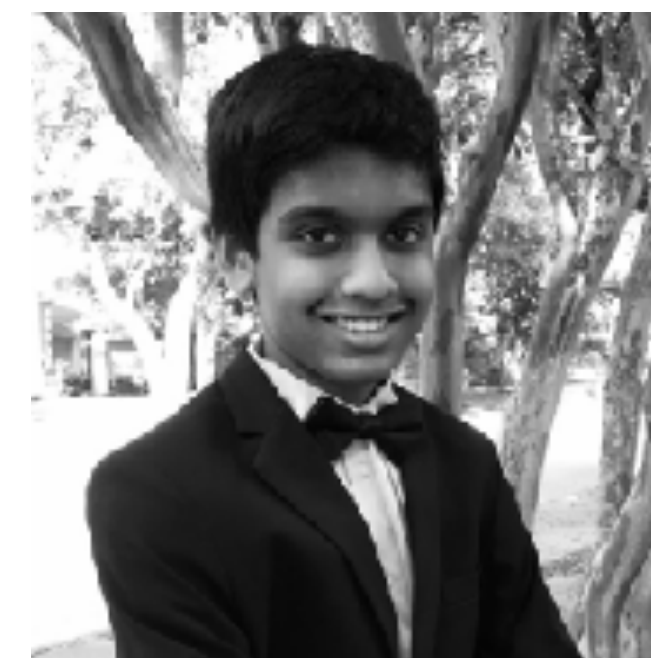"Passwords"

**Elizabeth**
"Bad Blood"

**Sydney**
"The Institute"

**Žiga**
"Mistborn Trilogy"

**Wendy**
"Surrounded by Idiots"

**Varun**
"Twelve Angry Men"

**Jessica**
"Pachinko"

**George**
"The Lord of the Rings"

**Leo**
"What If"

**Yao**
"The Midnight Library"

**Ved**
"Alone Together"

**Elida**
"A Tree Grows in
Brooklyn"
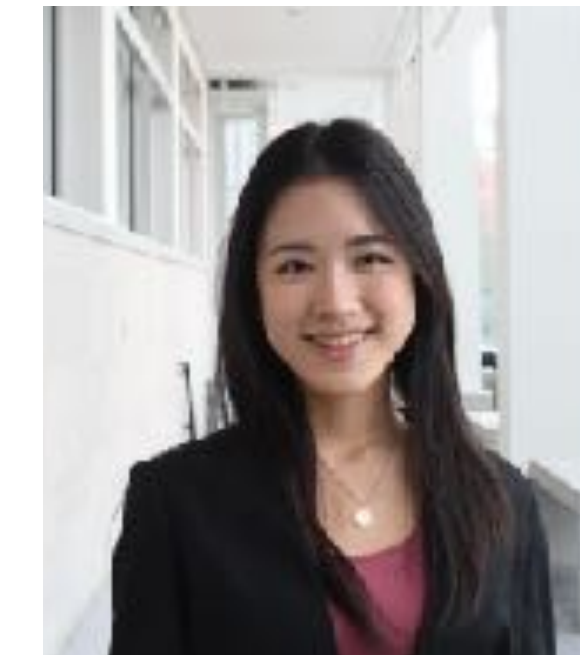
**Lingfei**
"Atonement"

**Ian**
"Fantastic Mr. Fox"

**Osama**
"Rain of Gold"
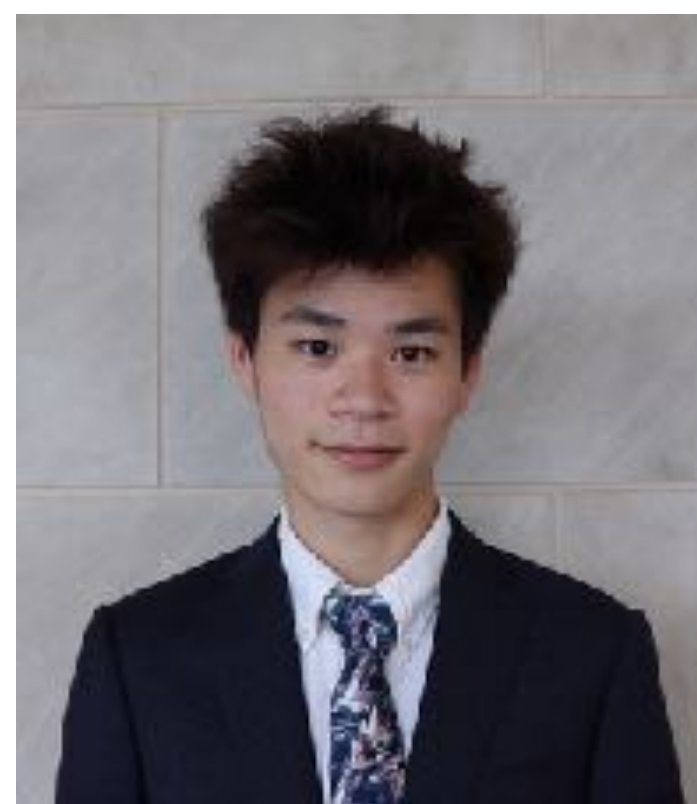
**Johann**
"The Mental Game of
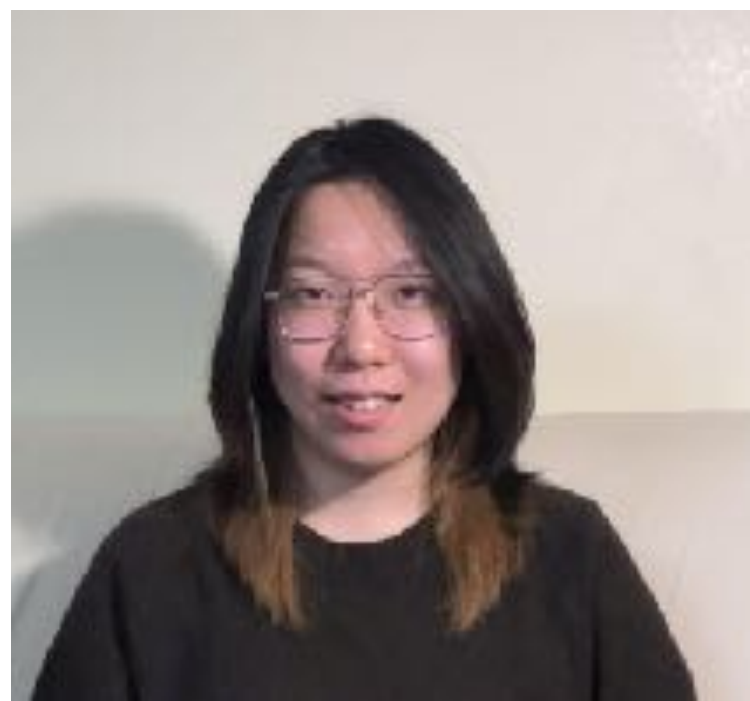Poker"
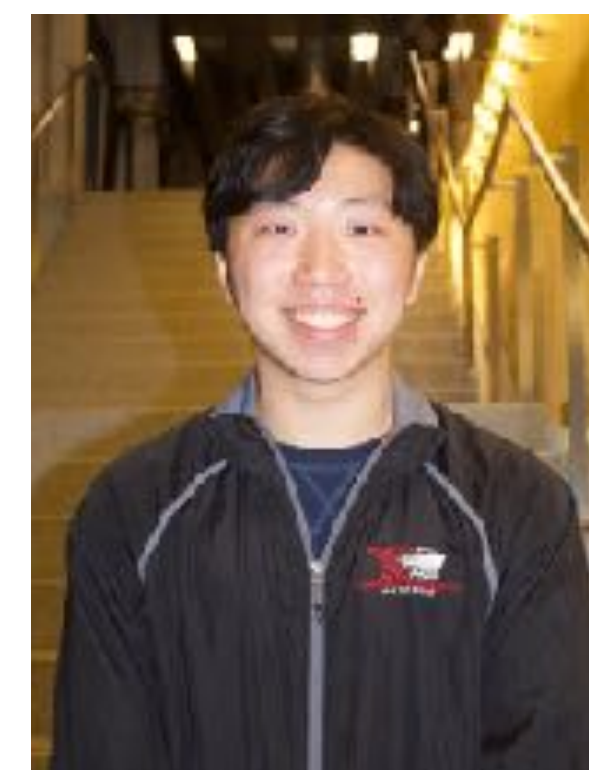
**Kevin**
"Kiss of the
Spider Woman"
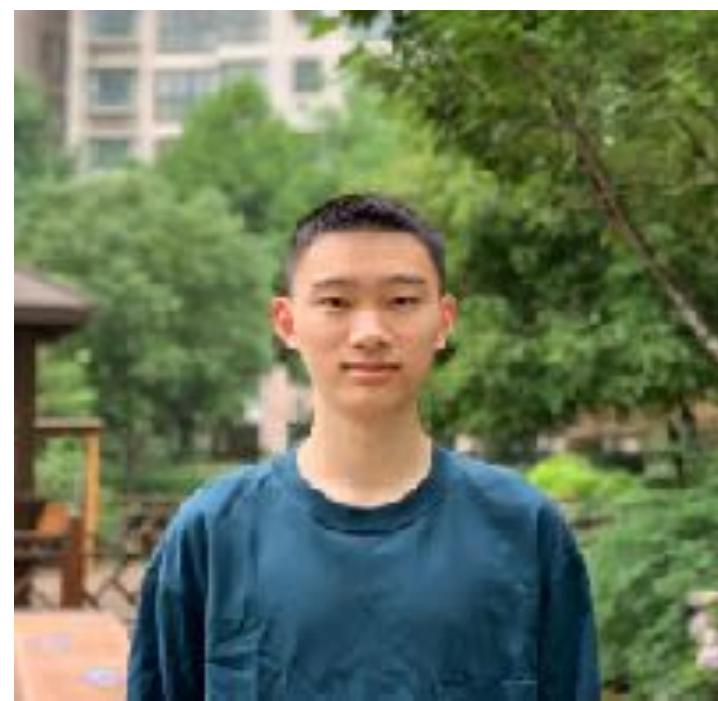
**Roberto**
"Piranesi"

**Dylan**
"Welcome to the
Monkey House"

Vivian
"The Foundation Trilogy"

Selina
"And Then There Were None"

Xilai
"One Two Three... Infinity"
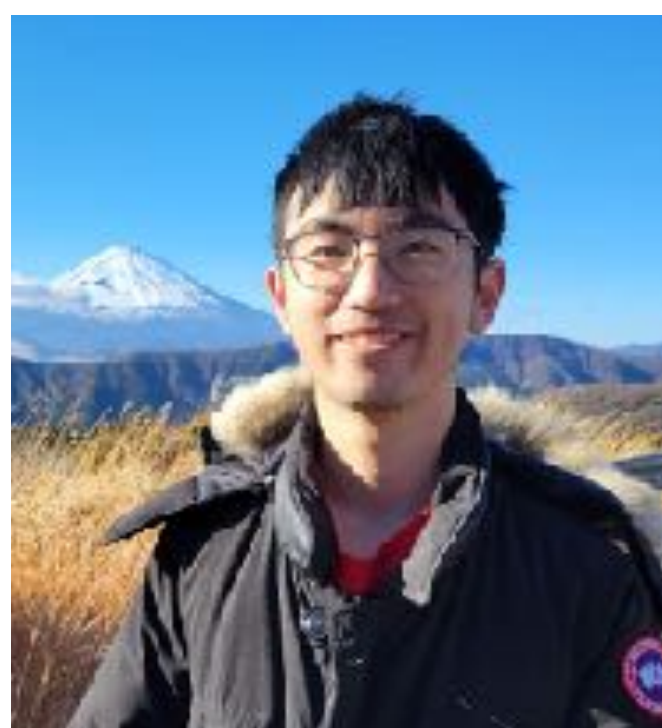
Yanny
*Angels in America*

Eugene
"The Art of War"

Taha
"Outliers"

JT
"Dune"

Daniel
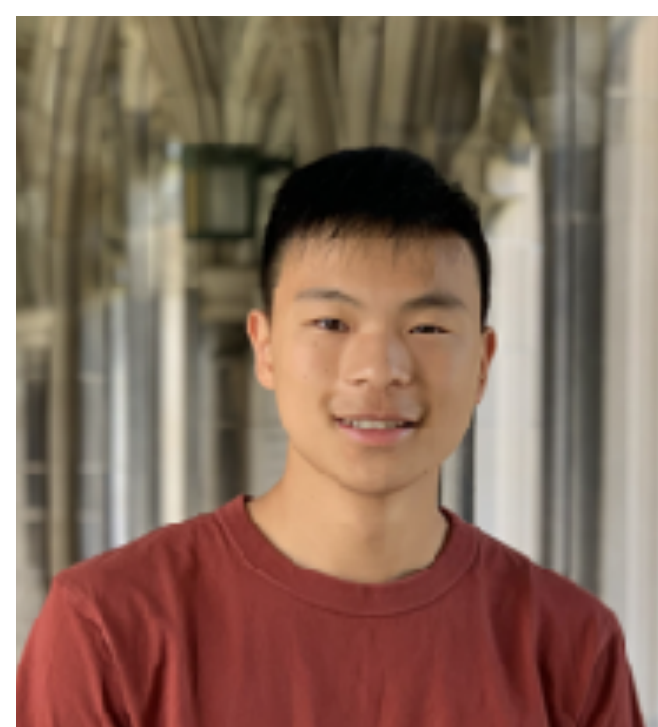"One Fish Two Fish Red Fish Blue Fish"

Andrew
"Diary of a Wimpy Kid"

Yash
"Project Hail Mary"

Lin
"The Silent Patient"

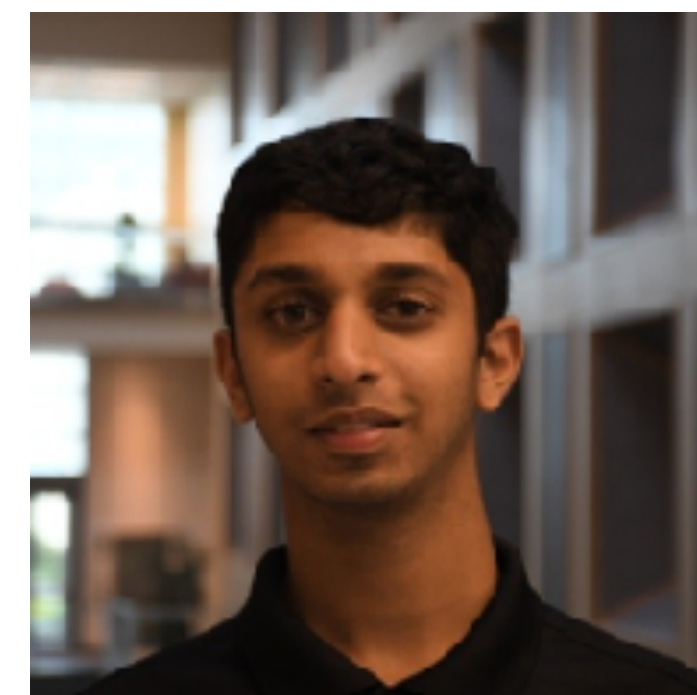Brendan
"Cat's Cradle"

Jeffrey

Yu

Mihir
"A Mind for Numbers"

Cella
"The Sirens of Titan"

Vineet
"Never Split the Difference"

Alex
"Dark Matter"

Peter
"The Three-body Problem"



Karthik
"The Art of Programming"

# Course Breakdown 5780

- 45% **Theory: Midterm + Final**
  - Closed book
  - No cheat sheets!
  - No personal notes
- 35% **Programming Assignments**
  - Up to **2** members in each team
  - **2** days extension **per team per project**
  - Autograder (unlimited resubmissions)
  - *Extra credit if you beat my own submission*
  - *Extra credit if you win contests*
- 10% **Paper Comprehension (mandatory)**
  - *Original Research Papers in ML*
  - *Canvas Quizzes*
- **10% Homeworks**
  - Up to **4** members in each team
  - Preparation for exam



Pie chart:
- 45% Theory
- 35% Programming
- 10% Homework
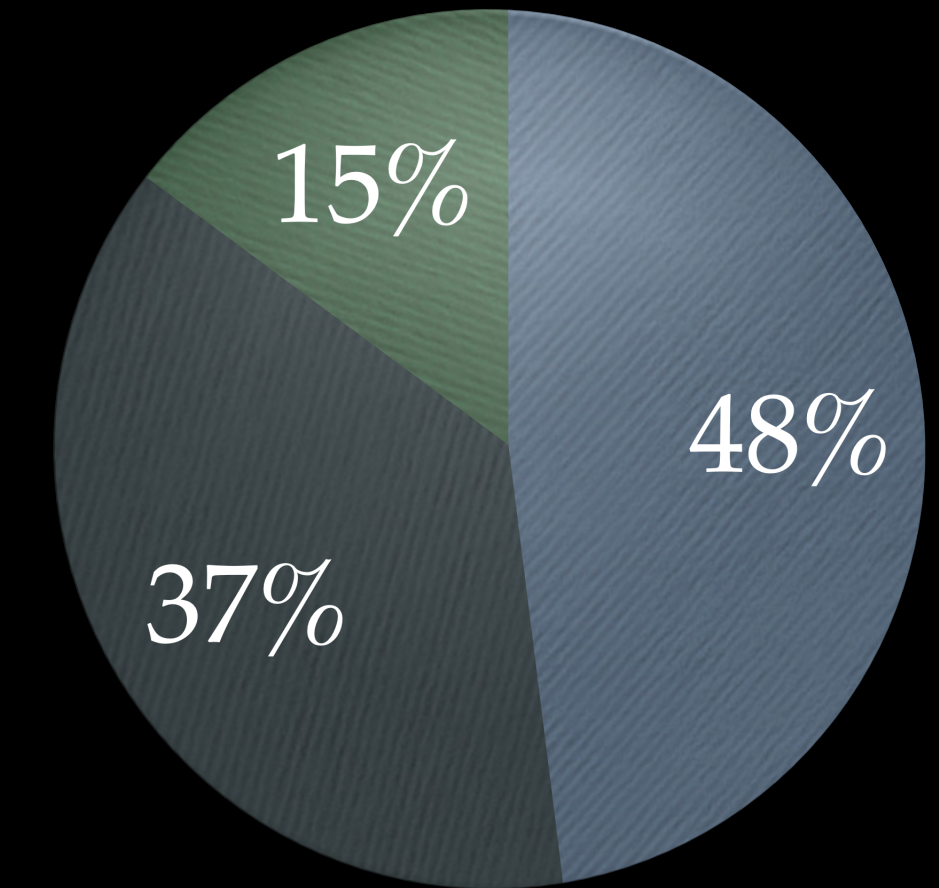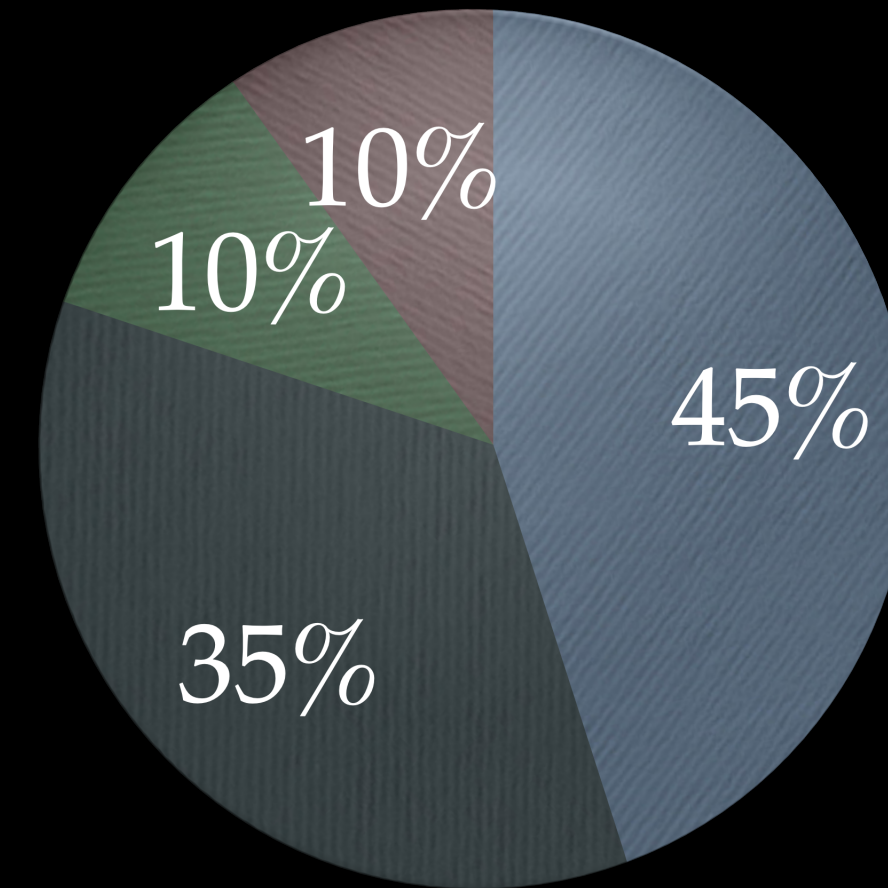- 10% Paper Comprehension
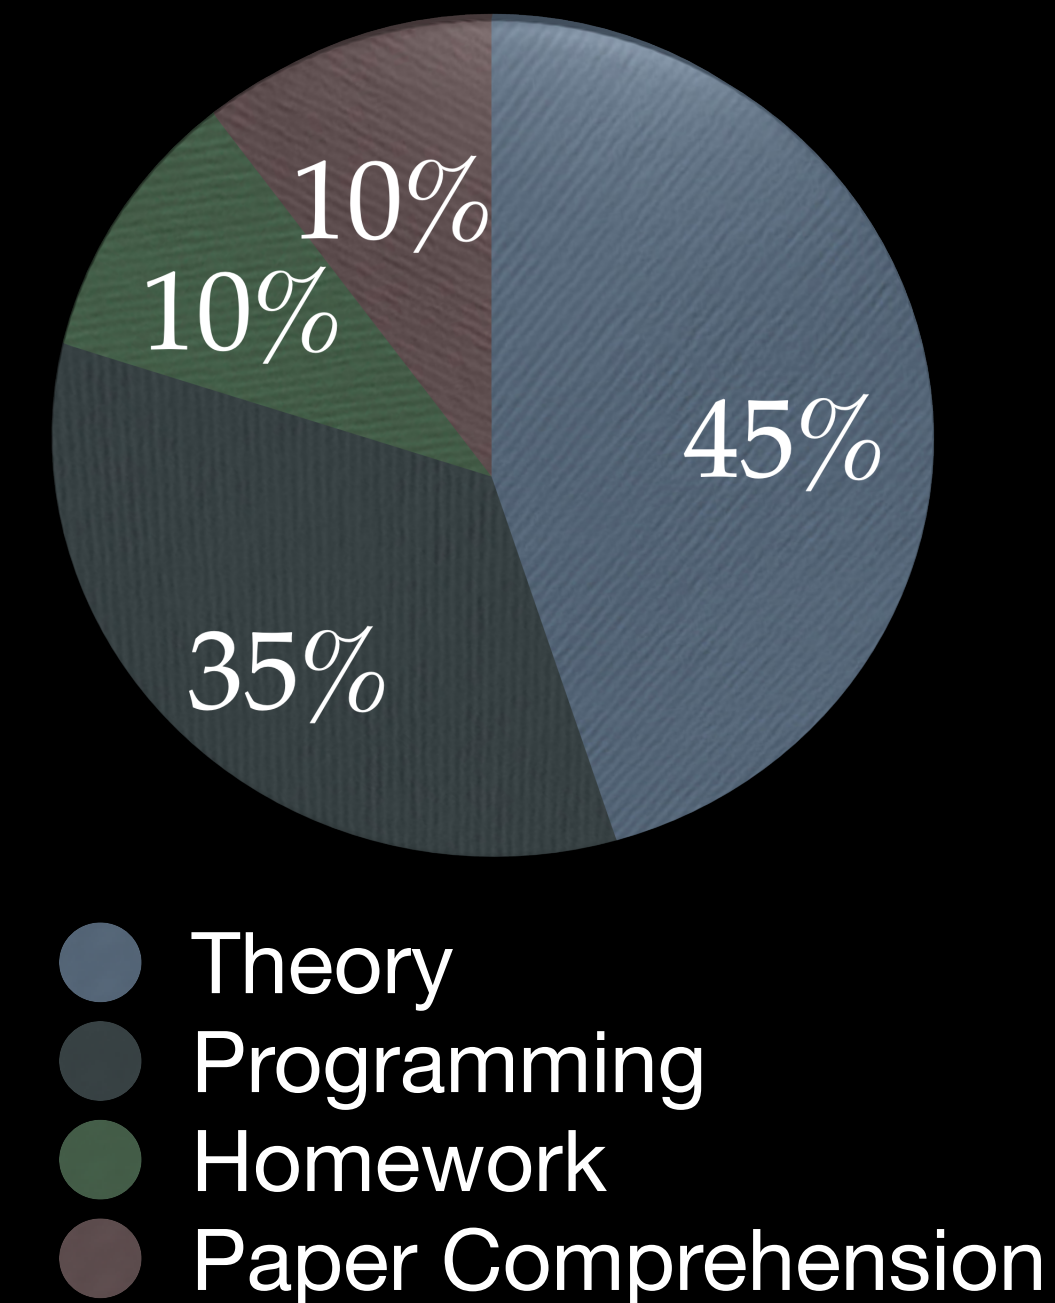
# Course Breakdown 4780

- {45% or 48%} **Theory: Midterm + Final**
  - Closed book
  - No cheat sheets!
  - No personal notes
- {35% or 37%} **Programming Assignments**
  - Up to **2** members in each team
  - **2** days extension **per team per project**
  - Autograder (unlimited resubmissions)
  - *Extra credit if you beat my own submission*
  - *Extra credit if you win contests*
- {0 or 15%} **Paper Comprehension (optional)**
  - *Original Research Papers in ML*
  - *Canvas Quizzes*
- {10% or 15% } **Homeworks**
  - Up to **4** members in each team
  - Preparation for exam



Pie chart 1: 45%, 35%, 10%, 10%

Pie chart 2: 48%, 37%, 15%

Legend:
- Exam
- Programming
- Homework
- Paper Comprehension

# Placement Exam

- Due January 31st!!!

- **<u>Canvas Page</u>**

- Get started early!!

- It is there for your own protection!

- This is how you get a Vocareum invite!

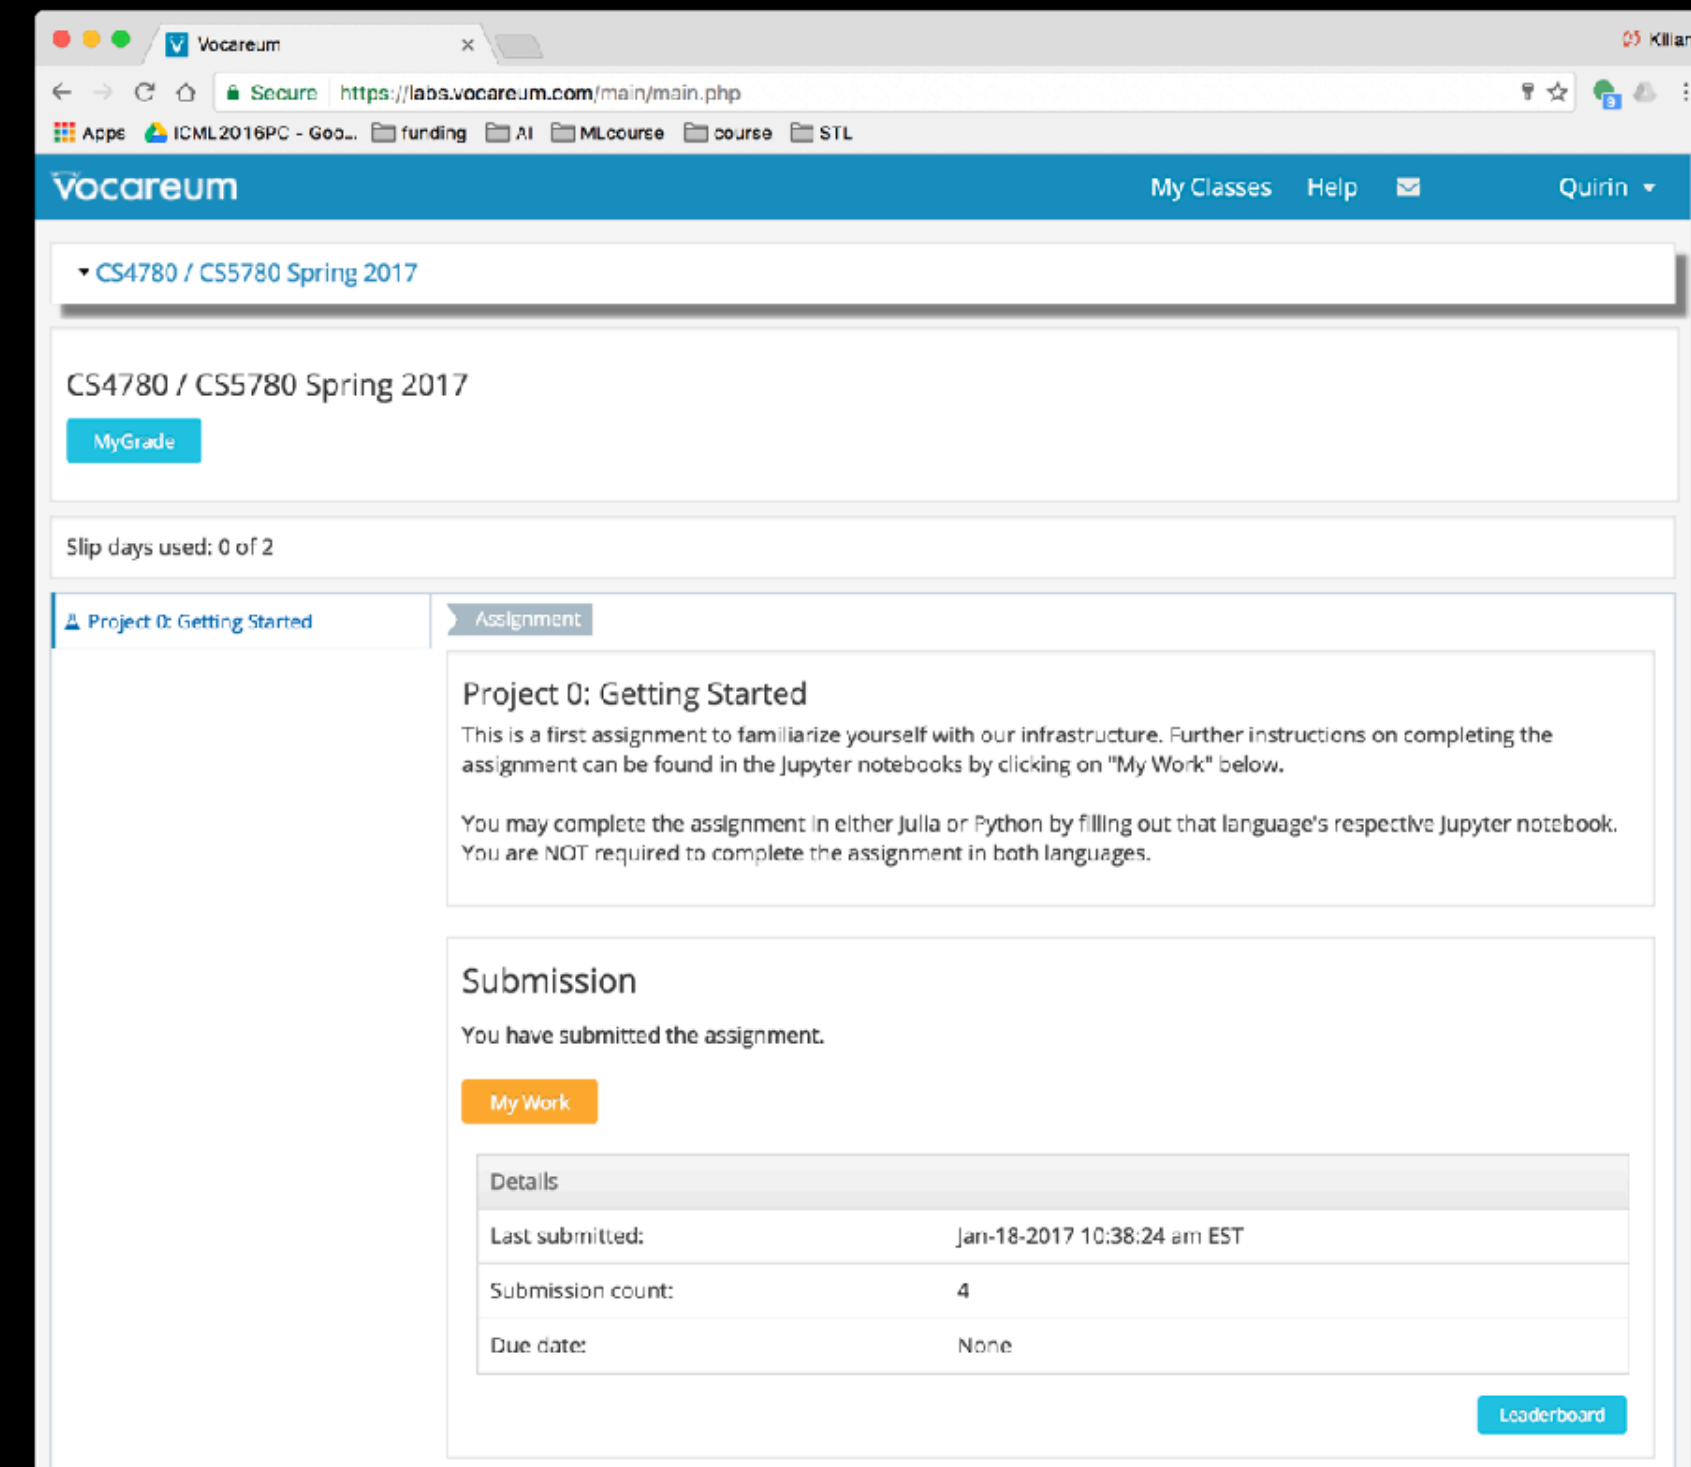- Take it even if you are on the **waitlist**!
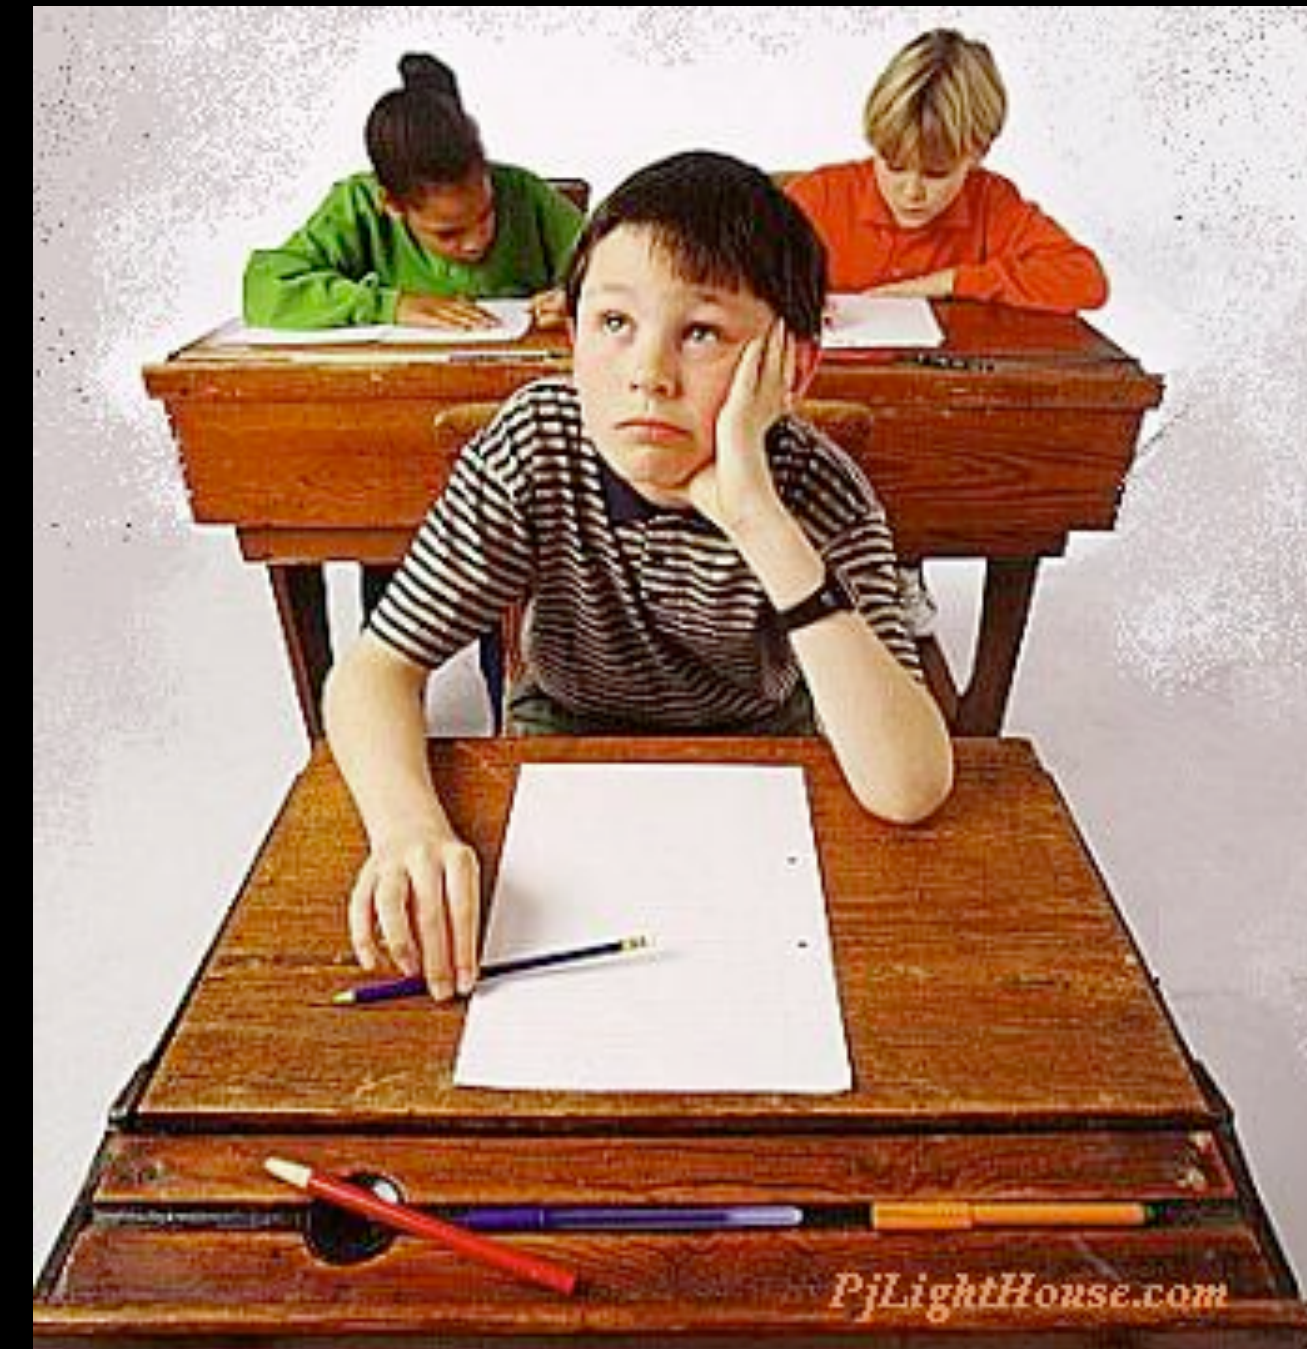
# Study Groups (2-4 people)

- You **must** join a study group by the end of next week.

- Find people on Ed

- This course will cover non-trivial material, learning in a group makes it easier and more fun!

# http://vocareum.com

- Pass placement exam -> get account
- There will be 8(+2) projects
- You have (roughly) **2 weeks** for each project
- Unlimited submits until deadline
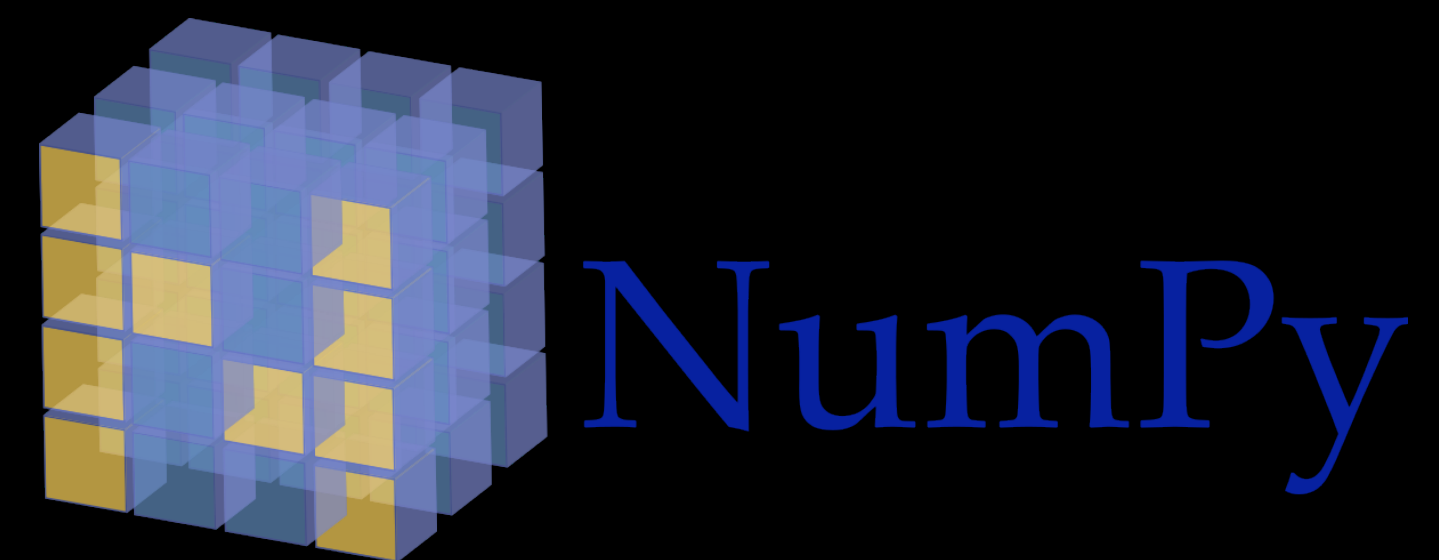- Costs around $30 :-(

# Placement Exam II (Project -1)

- NUMPY proficiency test

- Will turn into your own cheat-sheet

- Please take it seriously - this is for your own good

- We will go through the solutions in a dedicated section

# Autograder

- Important notes:
  - Setup a (secret) **screen name** for your team for the leaderboard
    - (top right corner - click on your login)
  - Only text with **#<GRADED>** and **#</GRADED>** will be graded
  - **!!!!You MUST form teams before you get started!!!!!!**

# Course Topics

- We will cover:

  - Parametric / Non-parametric learning

  - Empirical Risk Minimization

  - Unsupervised Learning

  - Bias/Variance Trade-off

  - Boosting

  - Support Vector Machines

  - Deep Learning

- We will **not** cover:

  - Graphical Models -> [CS4700]

  - Reinforcement Learning -> [CS4789]

  - Learning Theory  -> [CS4783]

  - Genetic Programming

# AI Path @ Cornell

## Level 1

Intro to AI [CS4700]

Intro to ML [CS4780]

Fund. of ML [ECE4200]

Math Found. [CS4850]

**unlock**

## Level 2
### (Some L1 courses as prerequisite)

Deep Learning [CS4782]

Principles of Large Sc. ML. [CS 4787]

Reinforcement Learning [CS 4789]

Math. Found. Of ML [CS 4783]

Computer Vision [CS 4670]

Natural Language Processing [CS 4740]

Found. of Robotics [CS 4750]

Comp. Genetics [CS 4775]

Many more, cross listed from other parts of Cornell!

# Prerequisites
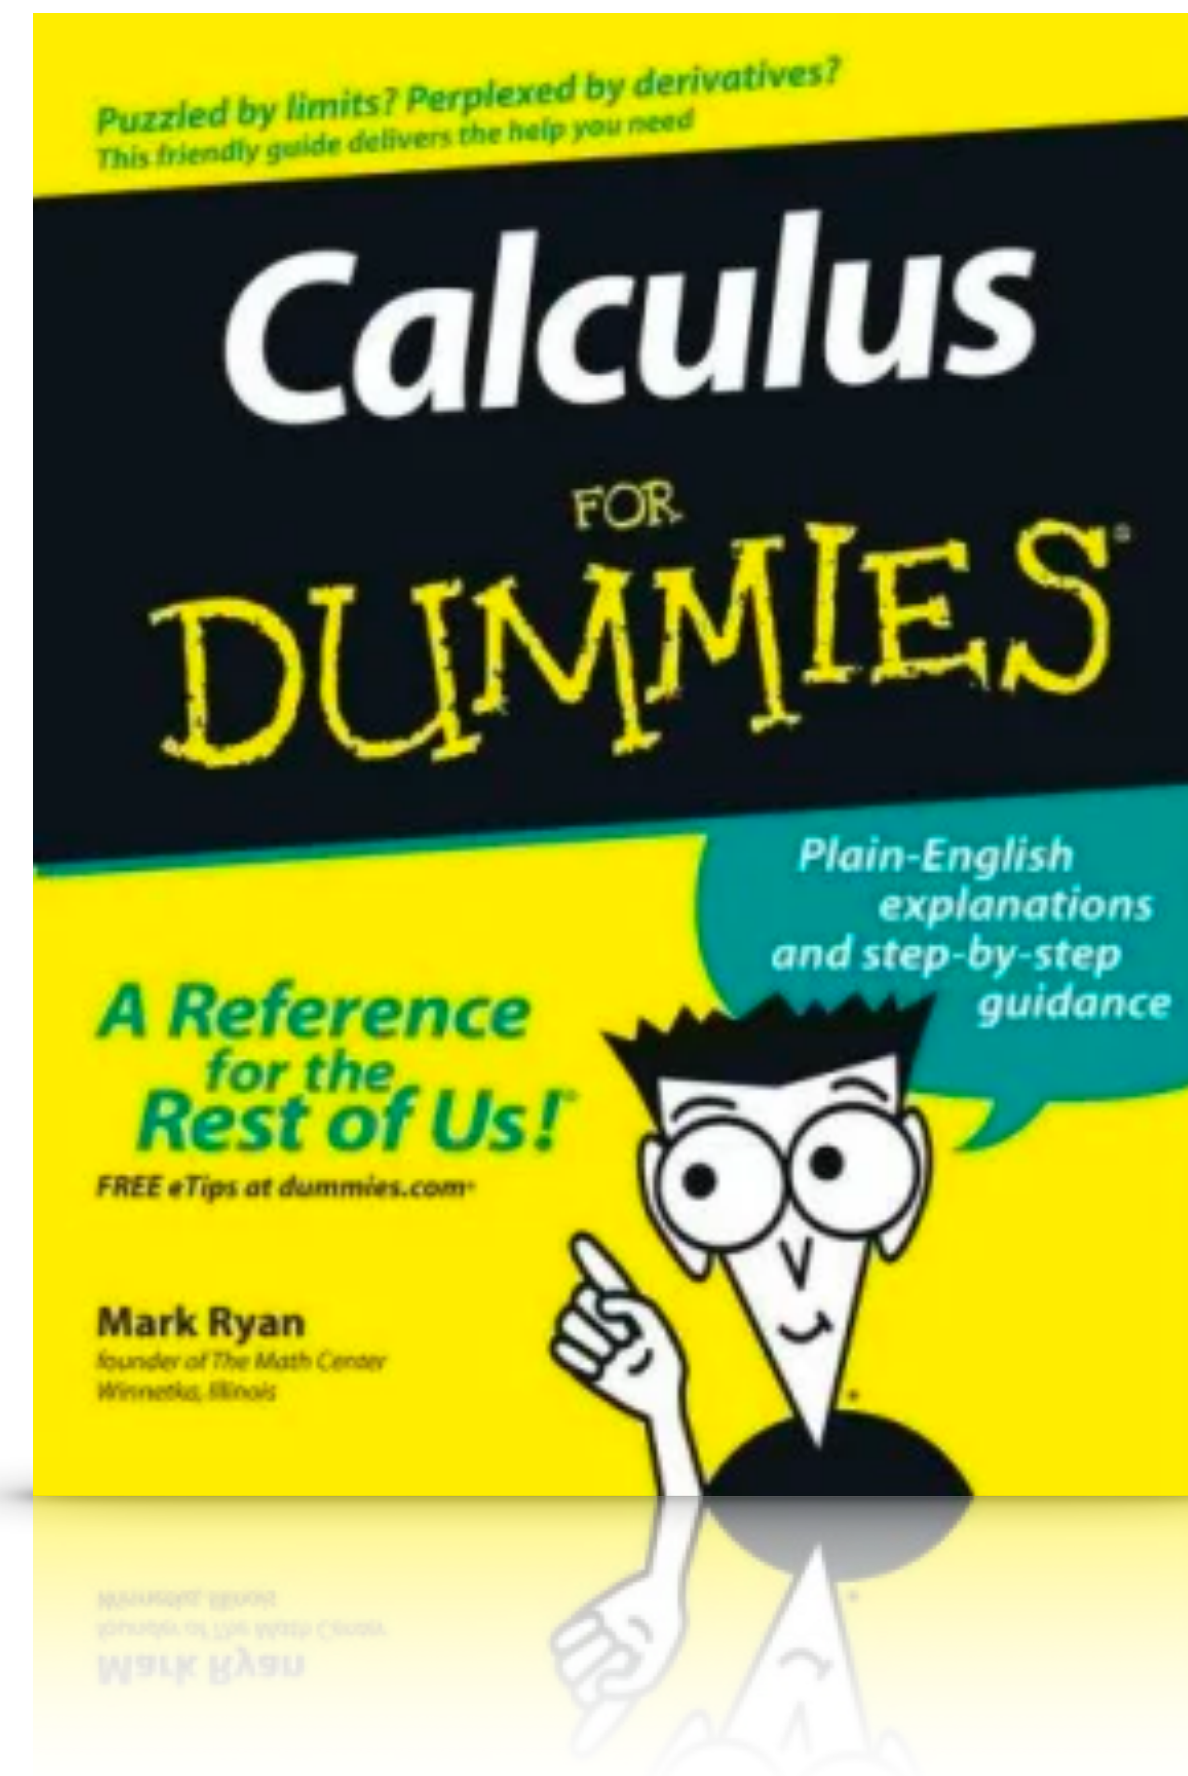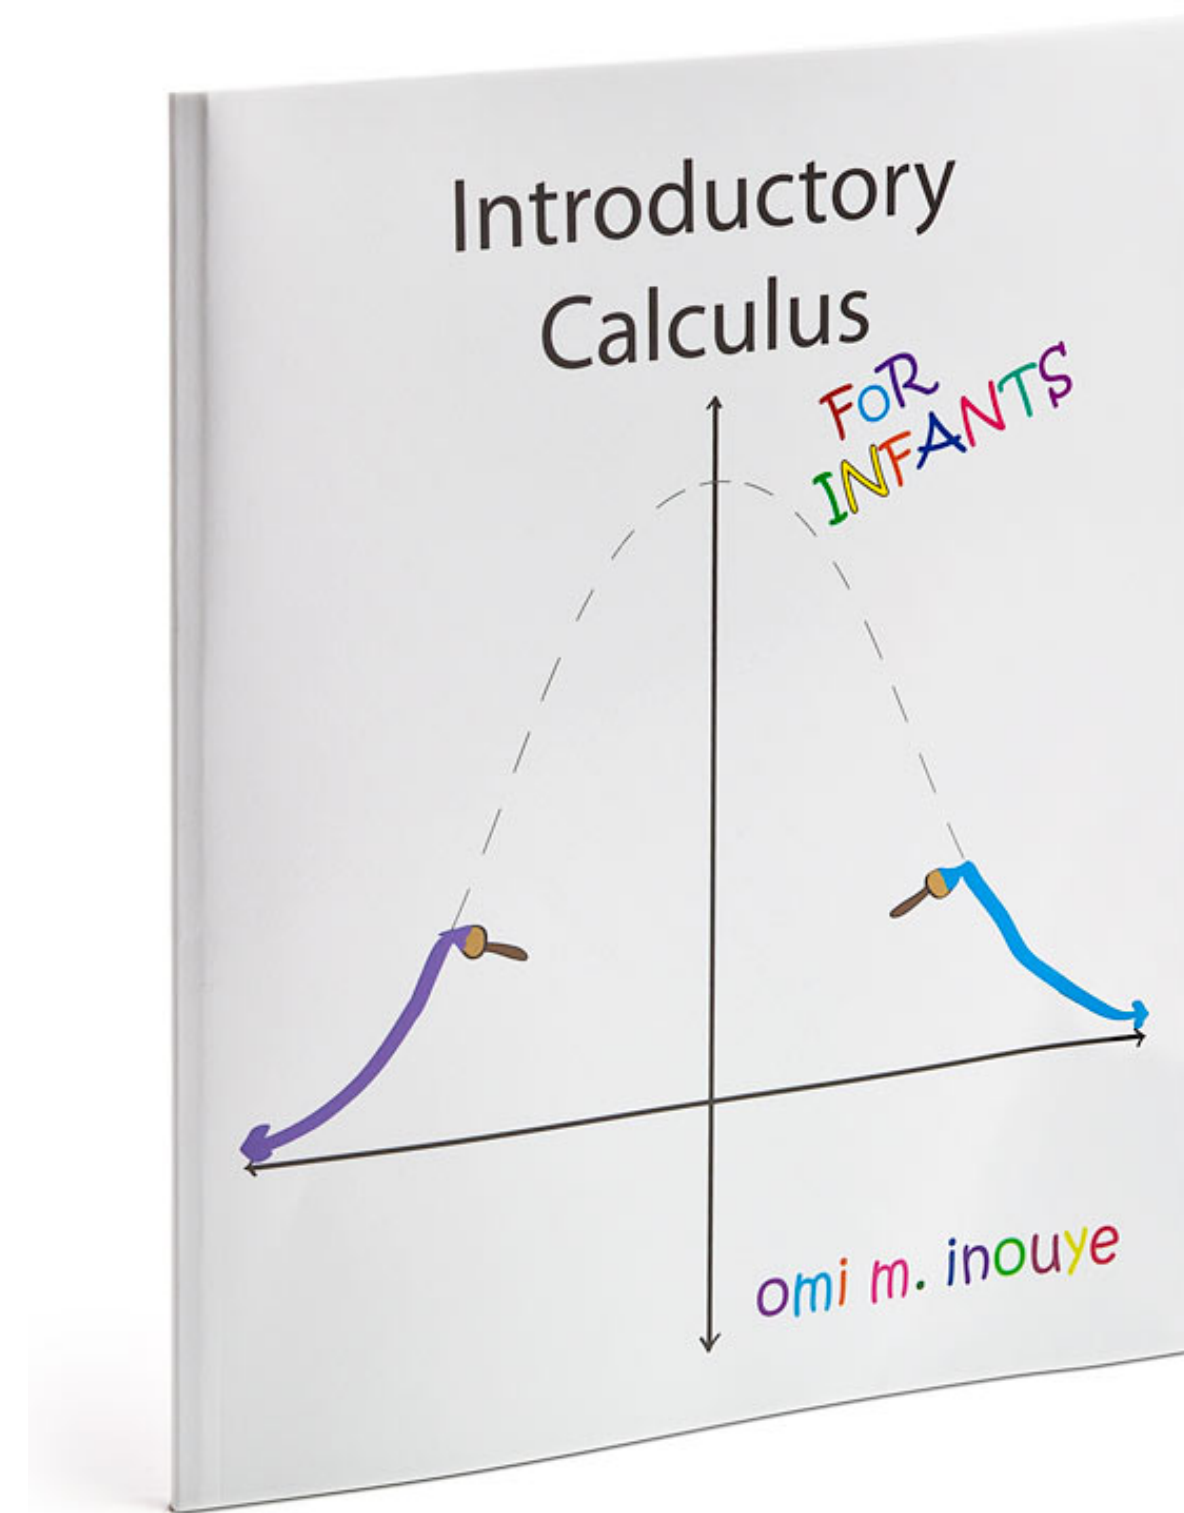
# Prerequisites

- Three pillars of ML:

  - **Statistics / Probability**

  - **Linear Algebra**

  - **Multivariate Calculus**

- Should be confident in at least 1/3, ideally 2/3.

- TAs might be able to give recitations on some topics if needed. (But don't rely on it.)
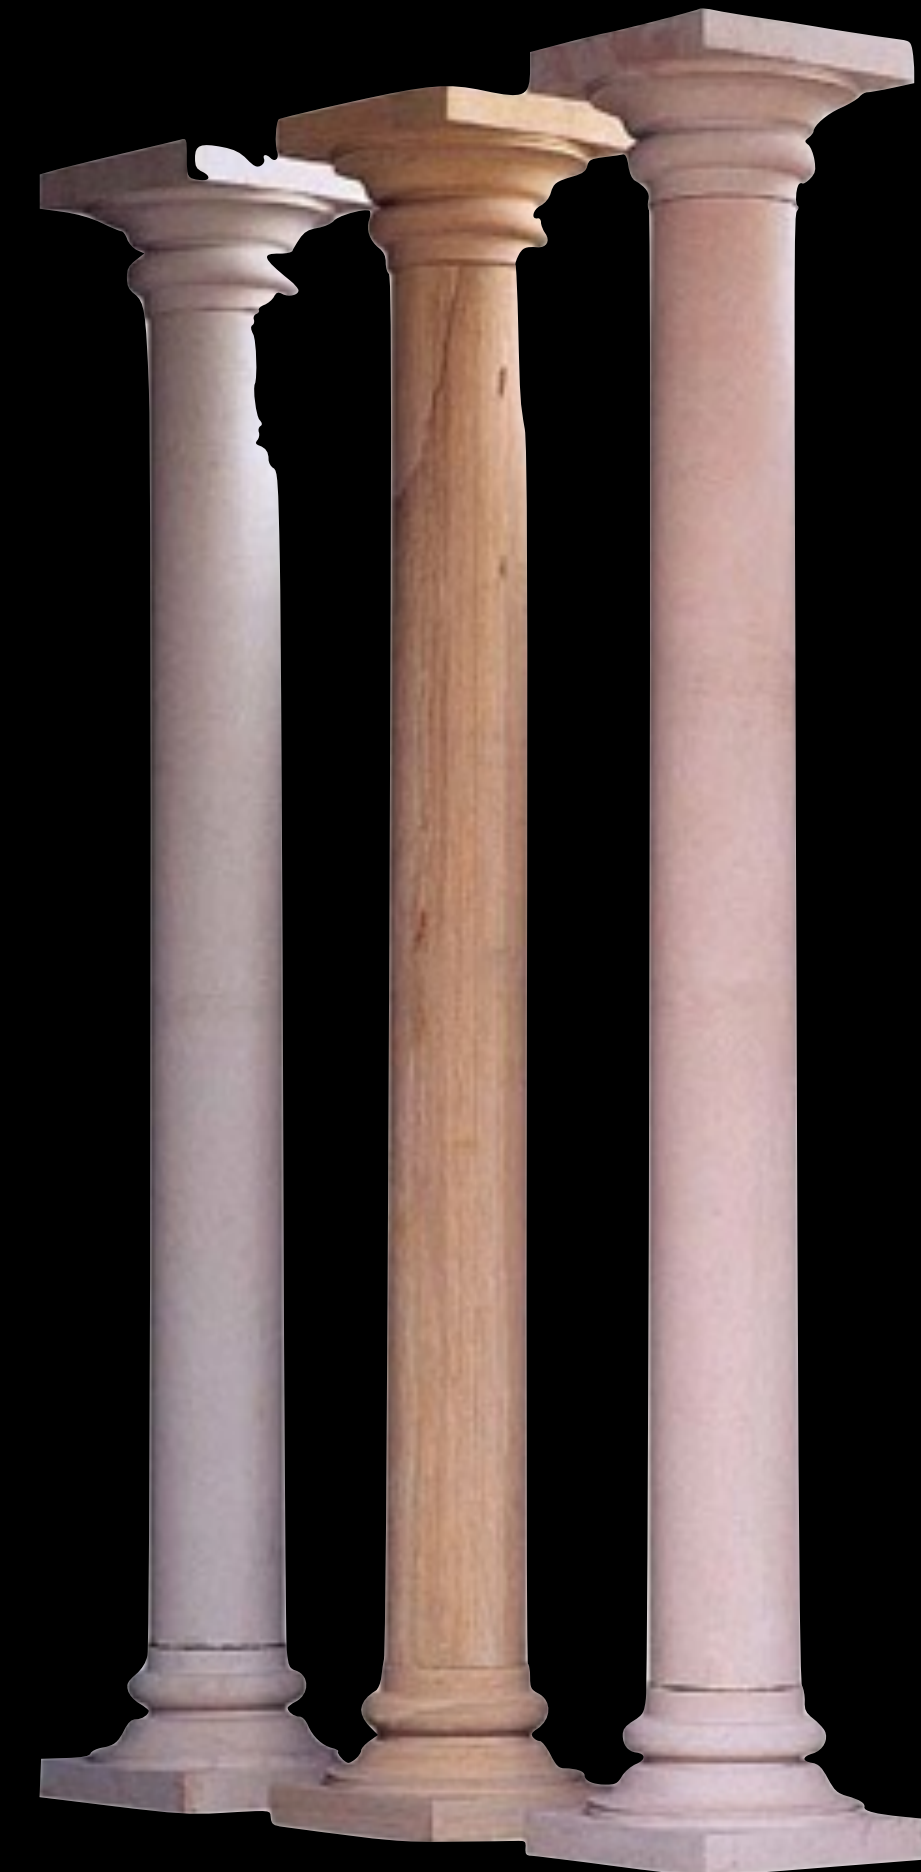
# About this course

- Take this course if …

  - you are interested in Machine Learning

  - you are comfortable with a decent amount of mathematics

  - you are not scared of programming

- Don't take this course if …

  - matrices scare you

  - you don't remember how to take derivatives

  - you want an easy A

- You cannot take this course if you fail the placement exam.

  - In that case, take appropriate prerequisites and come back next year.

# Student comments

- "[…] Requires a good knowledge in math and derivatives."

- "A TON of work, but mostly worth it for a very valuable skill."

- "great course, but prepare to work your butt off."

- "The topics were pretty complicated and difficult to understand quickly. I would have preferred a slightly slower pace."

- "It's mostly a math class"

- "Huge work load, excessive at times, but that's just the nature of the course"

# Academic Integrity

- We **actively** look for academic conduct violations

- The autograder checks for plagiarism

- (Zero tolerance policy: all occurrences will be reported.)

# Academic Integrity

- *Examples:*

  - Most common: Students steal from outside source

  - Students post to RentACoder.com or other page

  - Students post solutions on the web

  - Students use solutions from last year's course

# But: Using LLMs is OK :-)

- You **can** use LLMs (e.g. ChatGPT, Bard) for your assignments, but you **must** hand in a **detailed description** of how you used it (including all prompts and outputs)

# 5K Run
## Early May

- Belle Sherman 5K

- Karthik and Kilian will run!

- Usually after last class
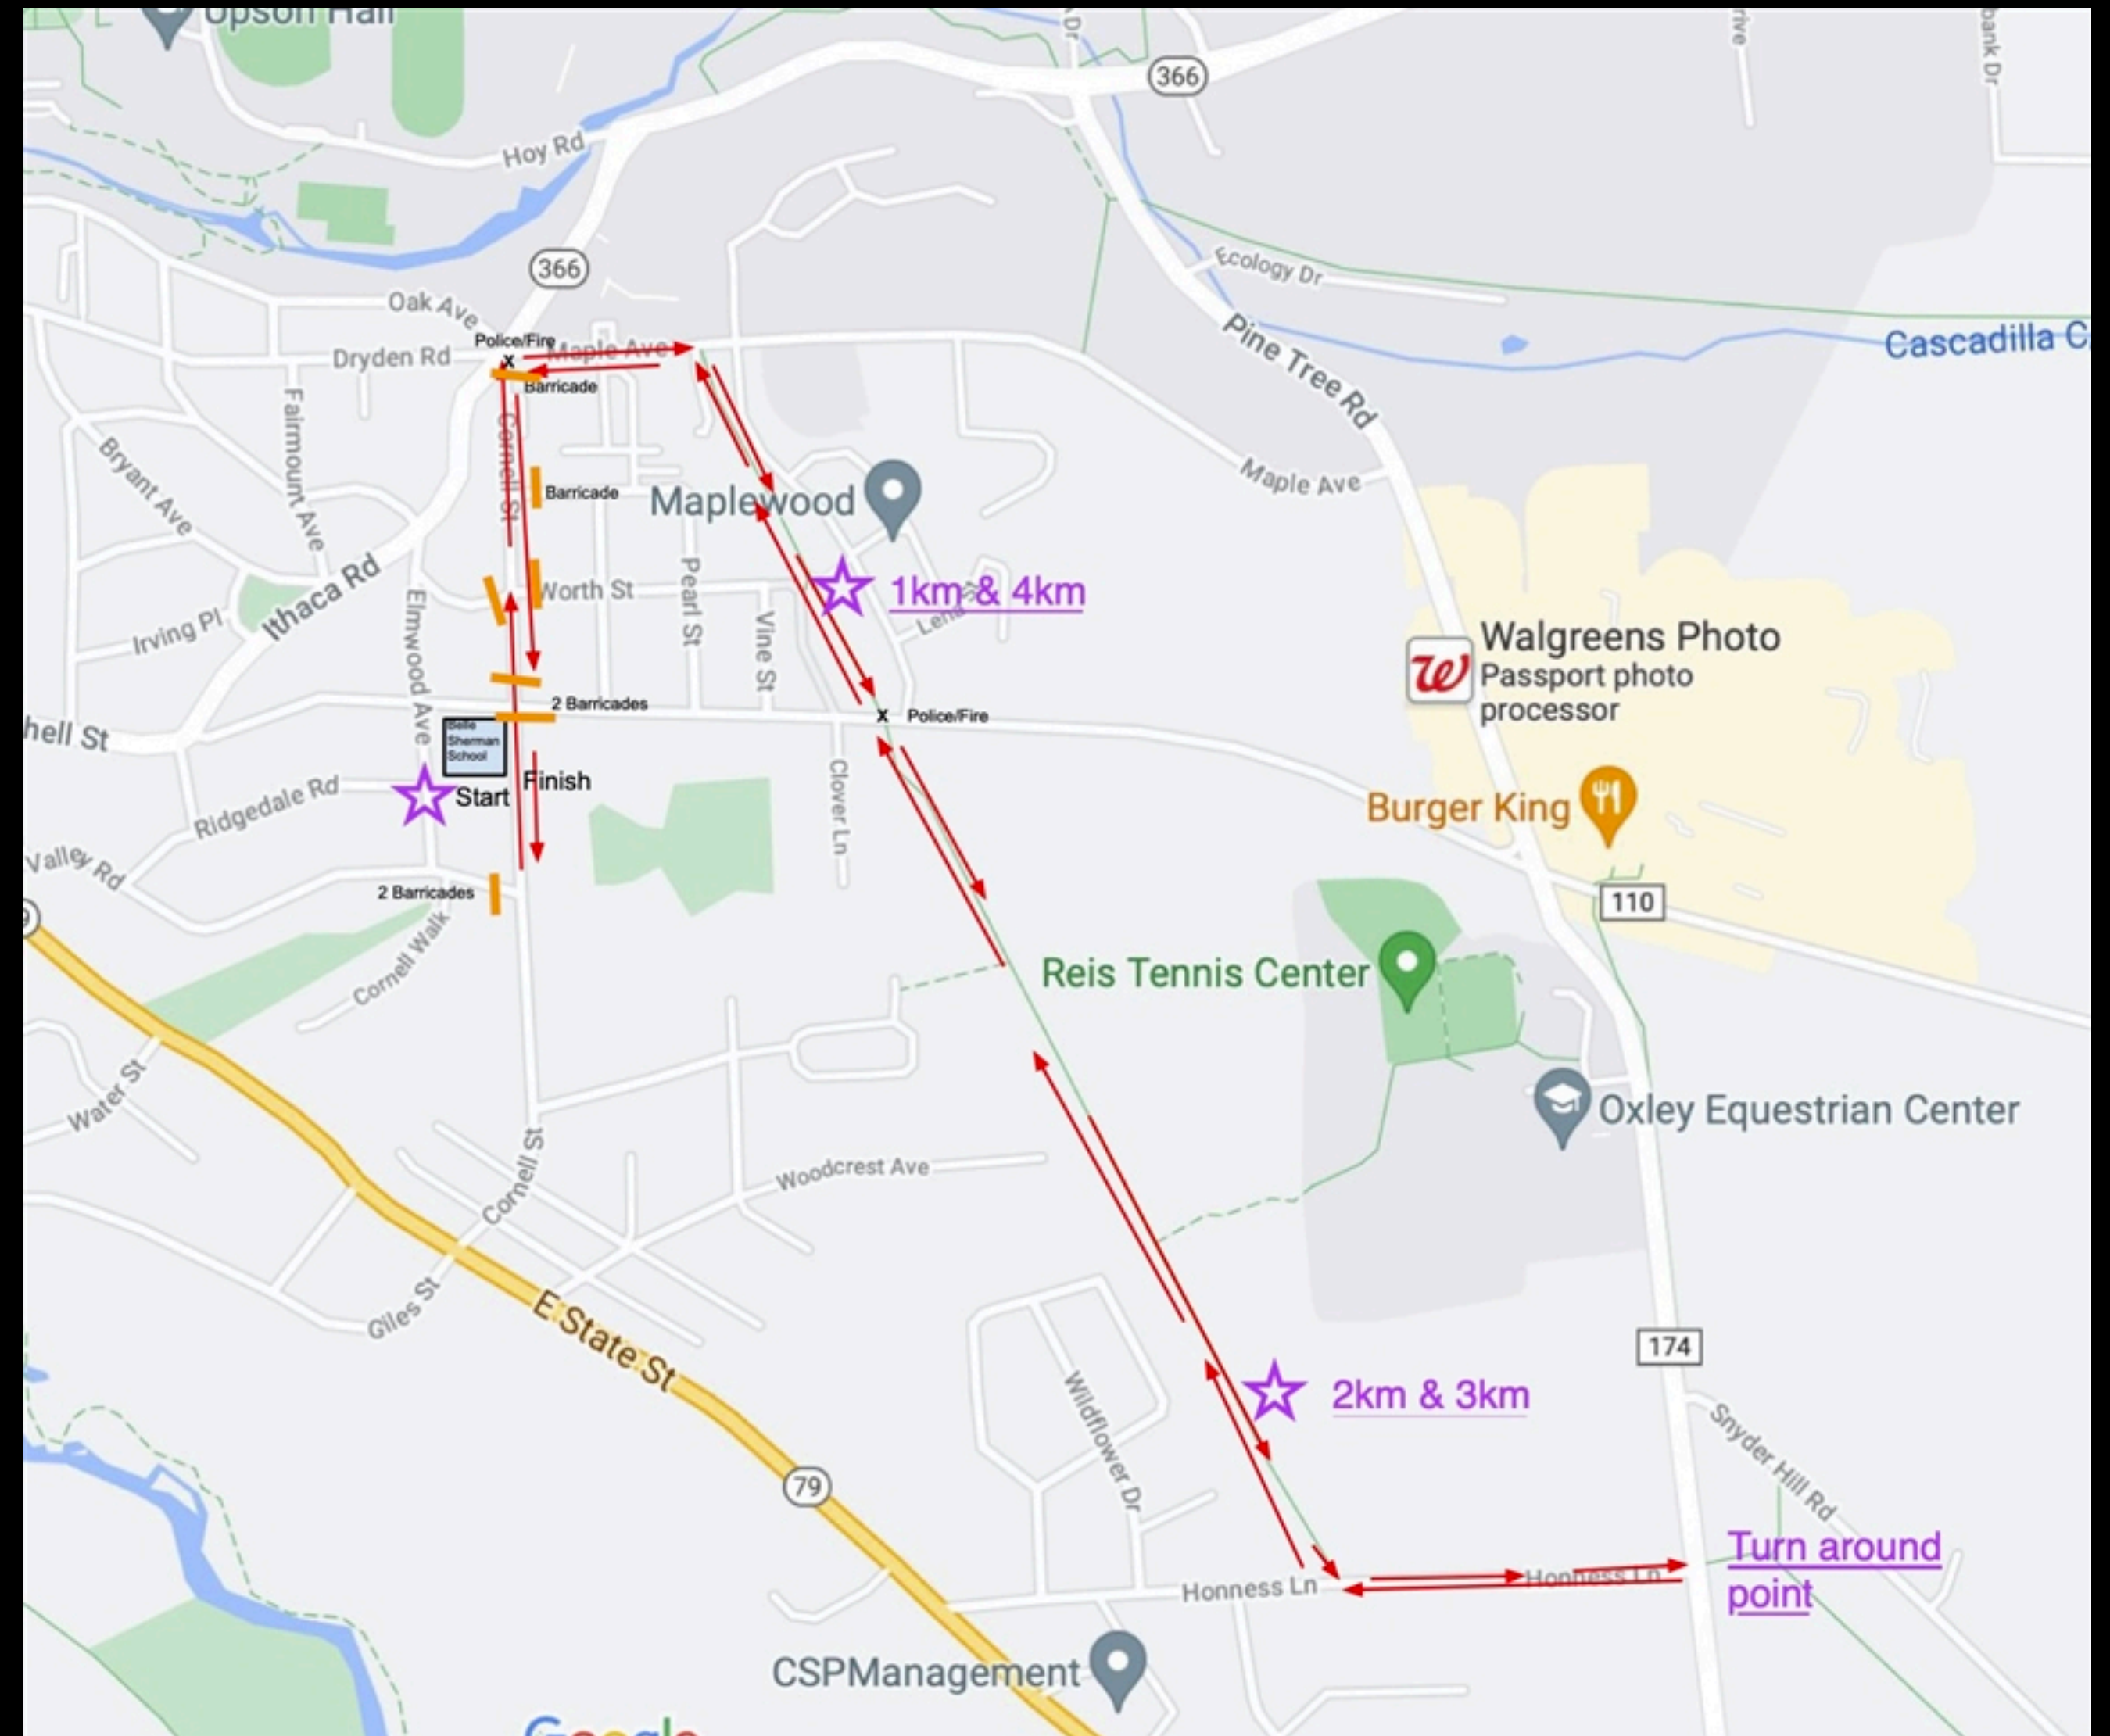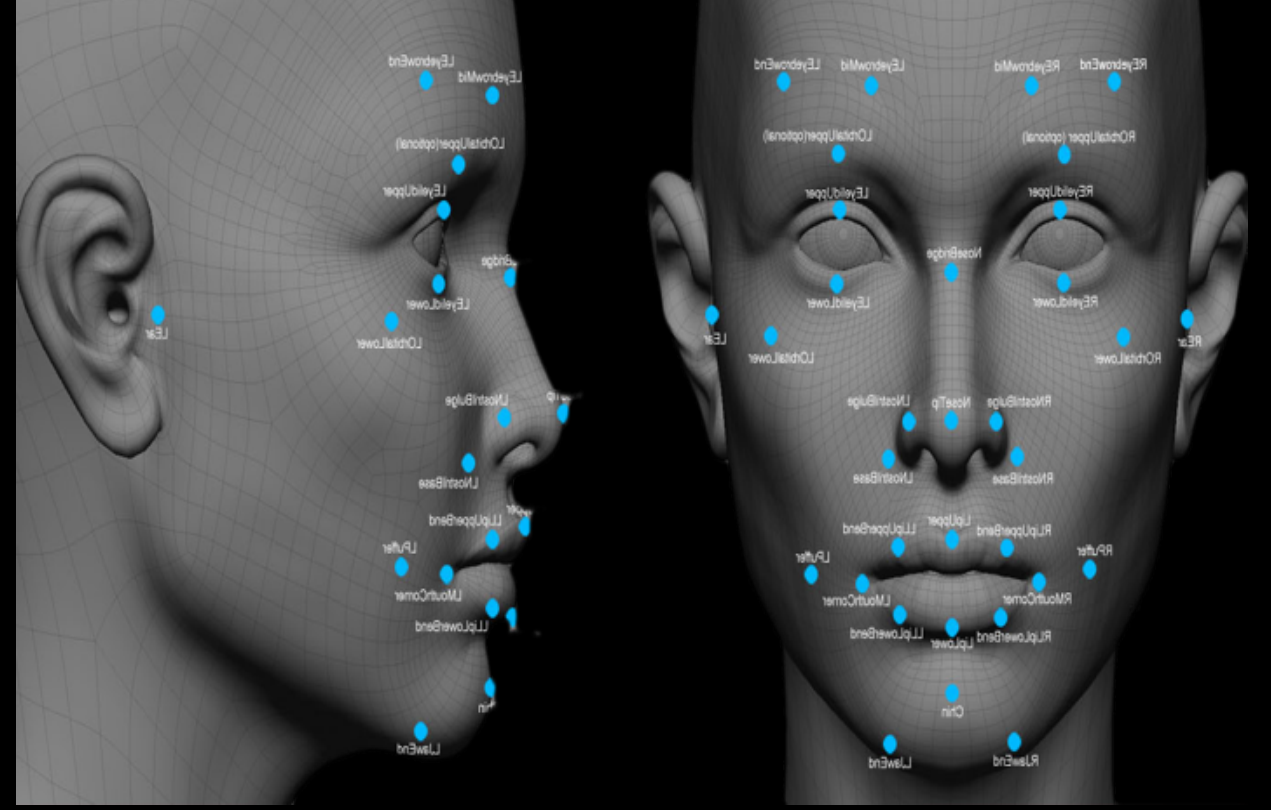
# Machine Learning (ML)

Programs that improve with experience.
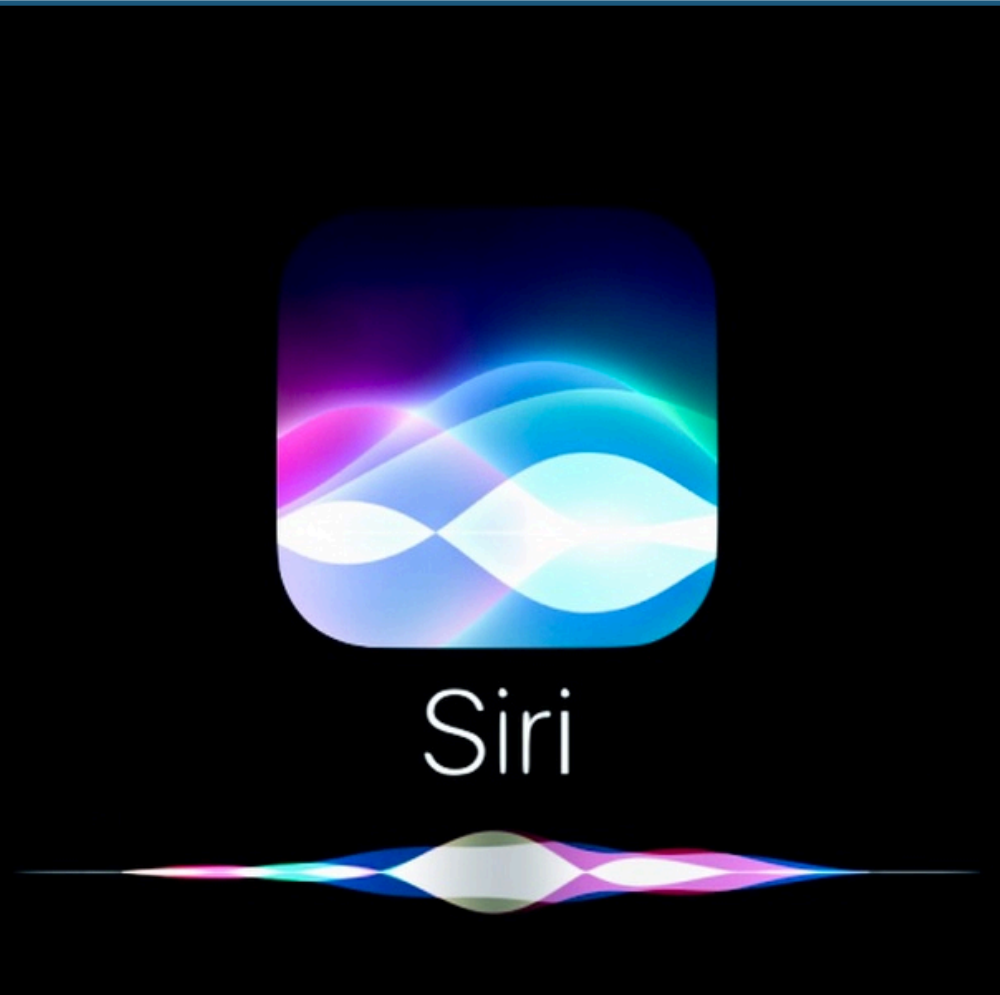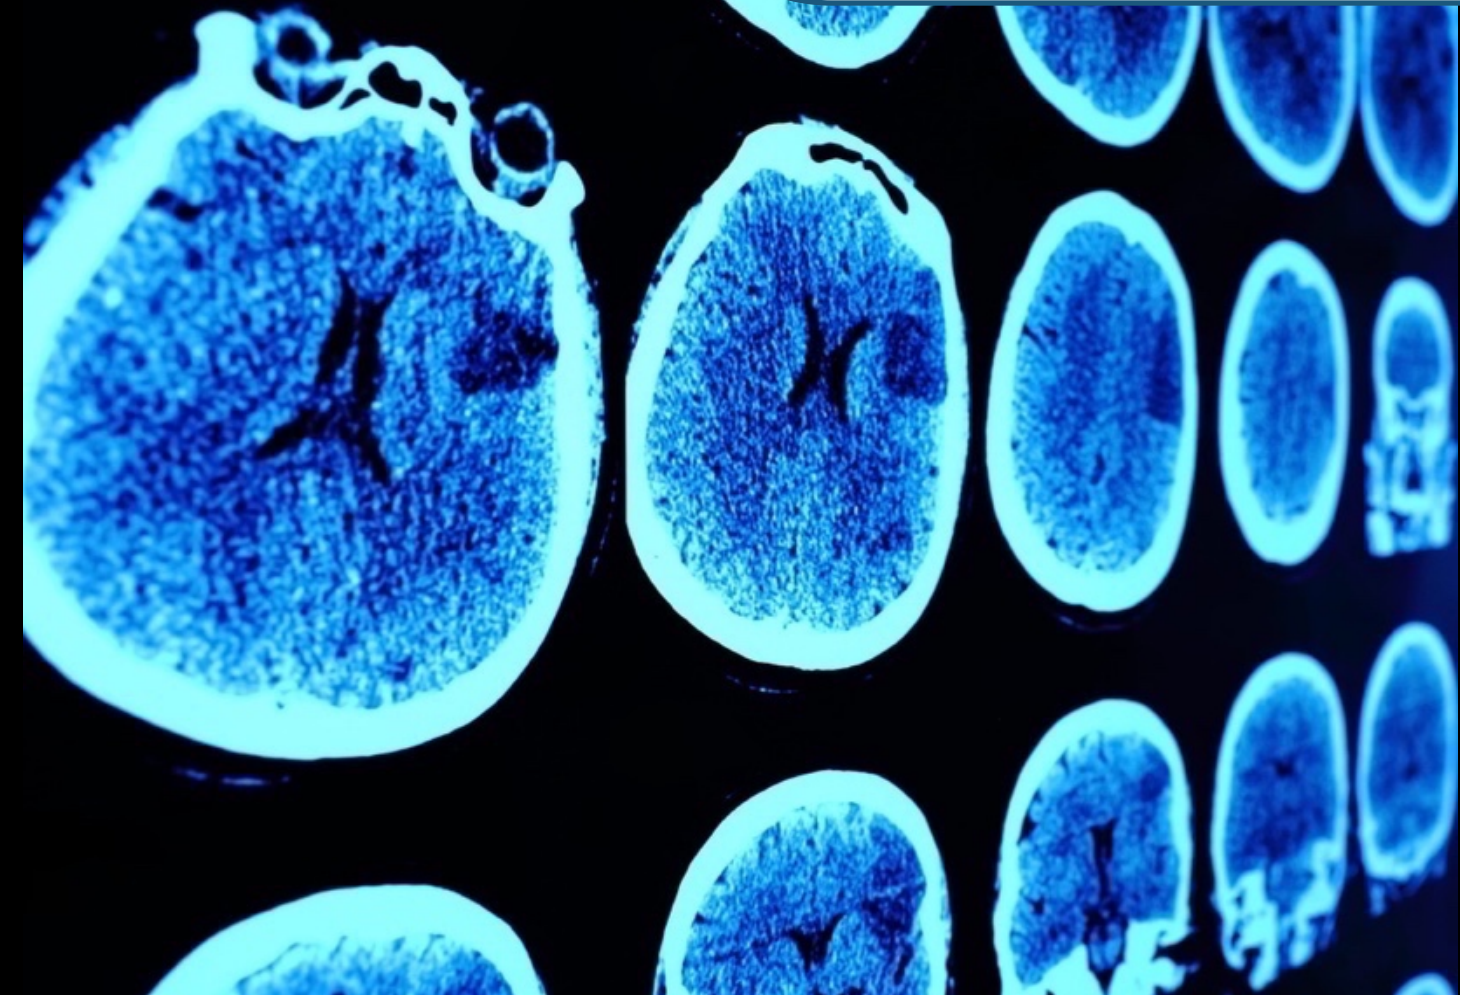
# Traditional Computer Science

**Traditional CS:**
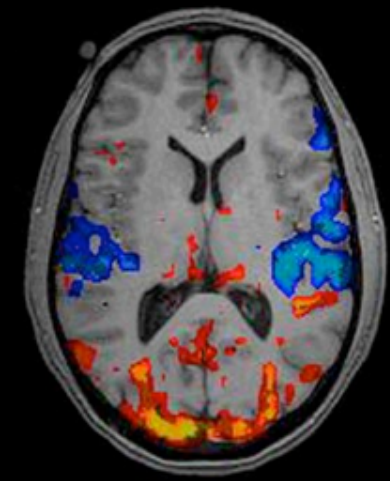


Data →

Program →
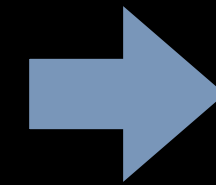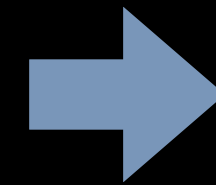
Computer
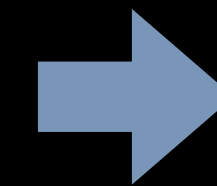
→ Output

# Machine Learning

**Traditional CS:**

Data →

Program →

Computer

→ Output

**Machine Learning:**

Data →

Output →

Computer
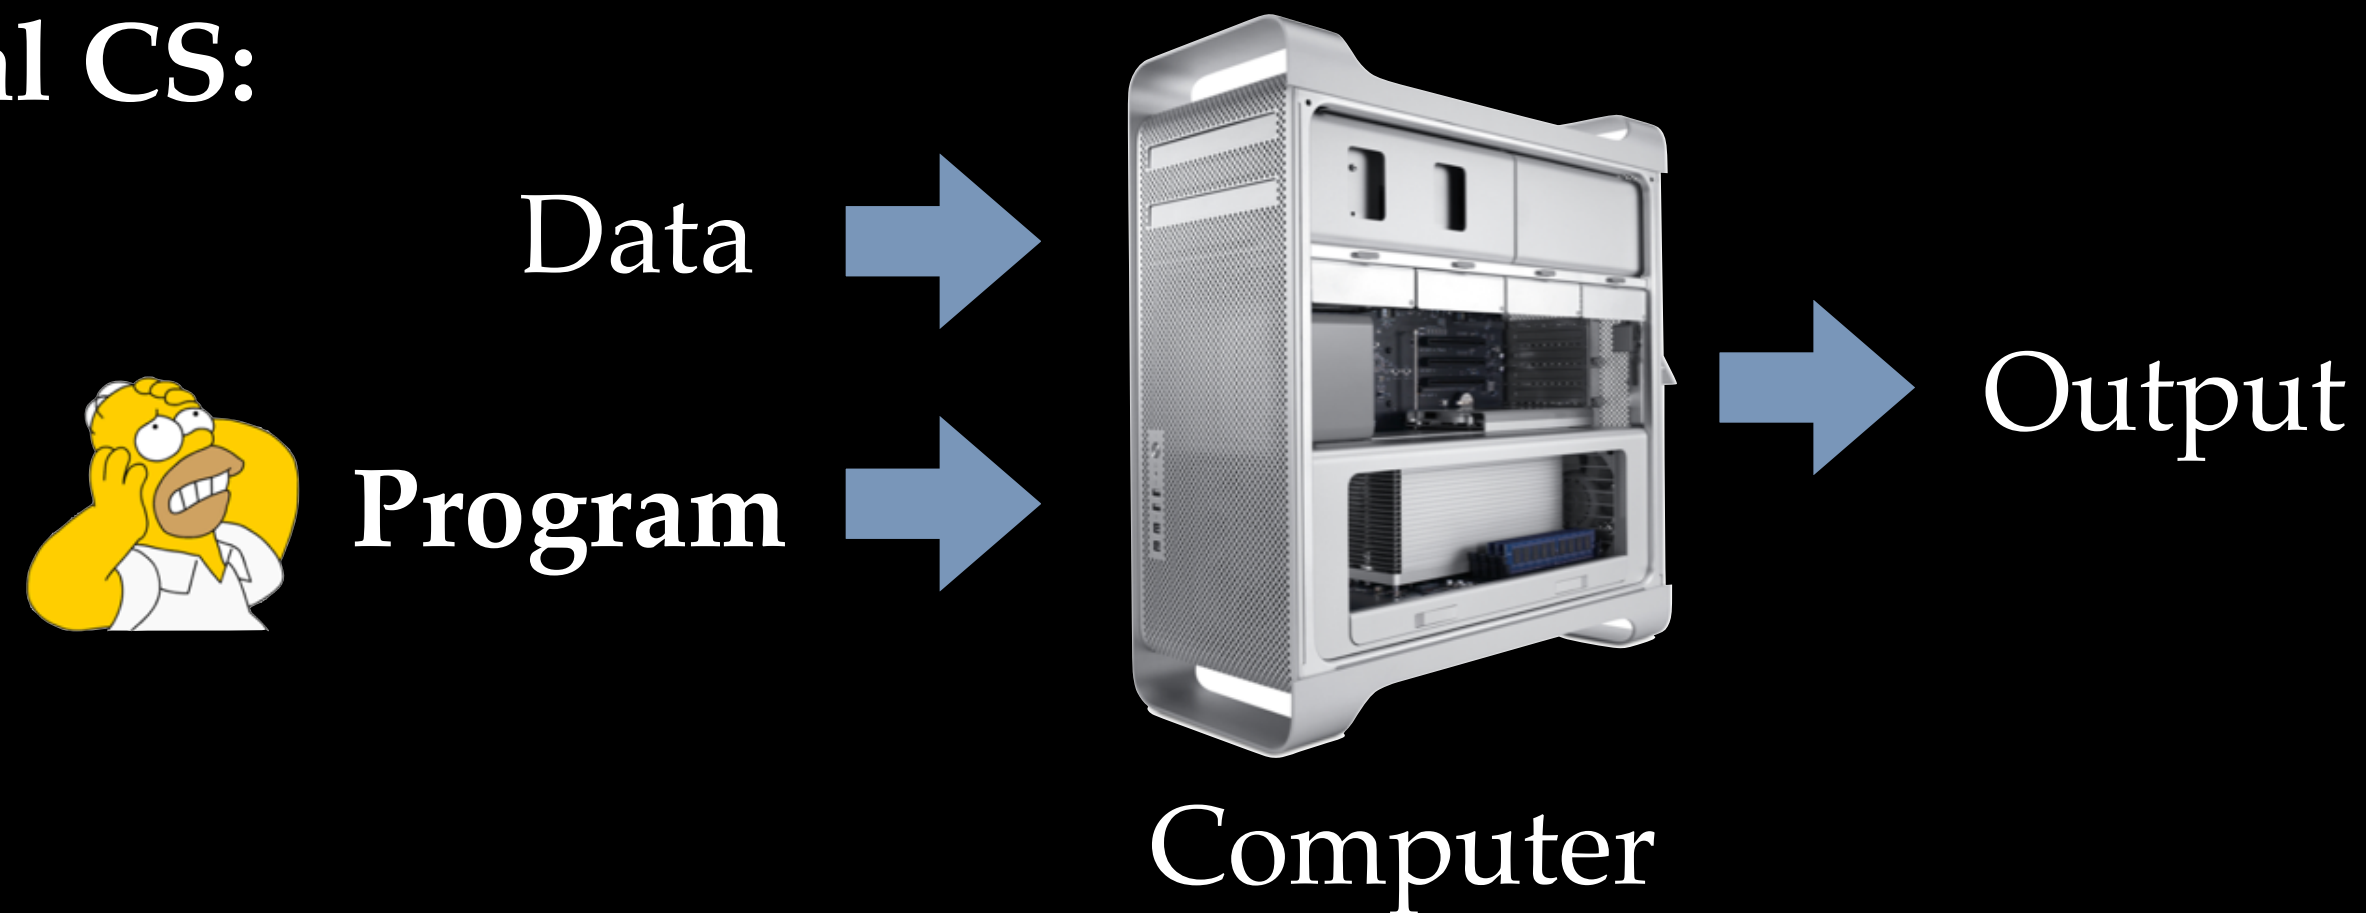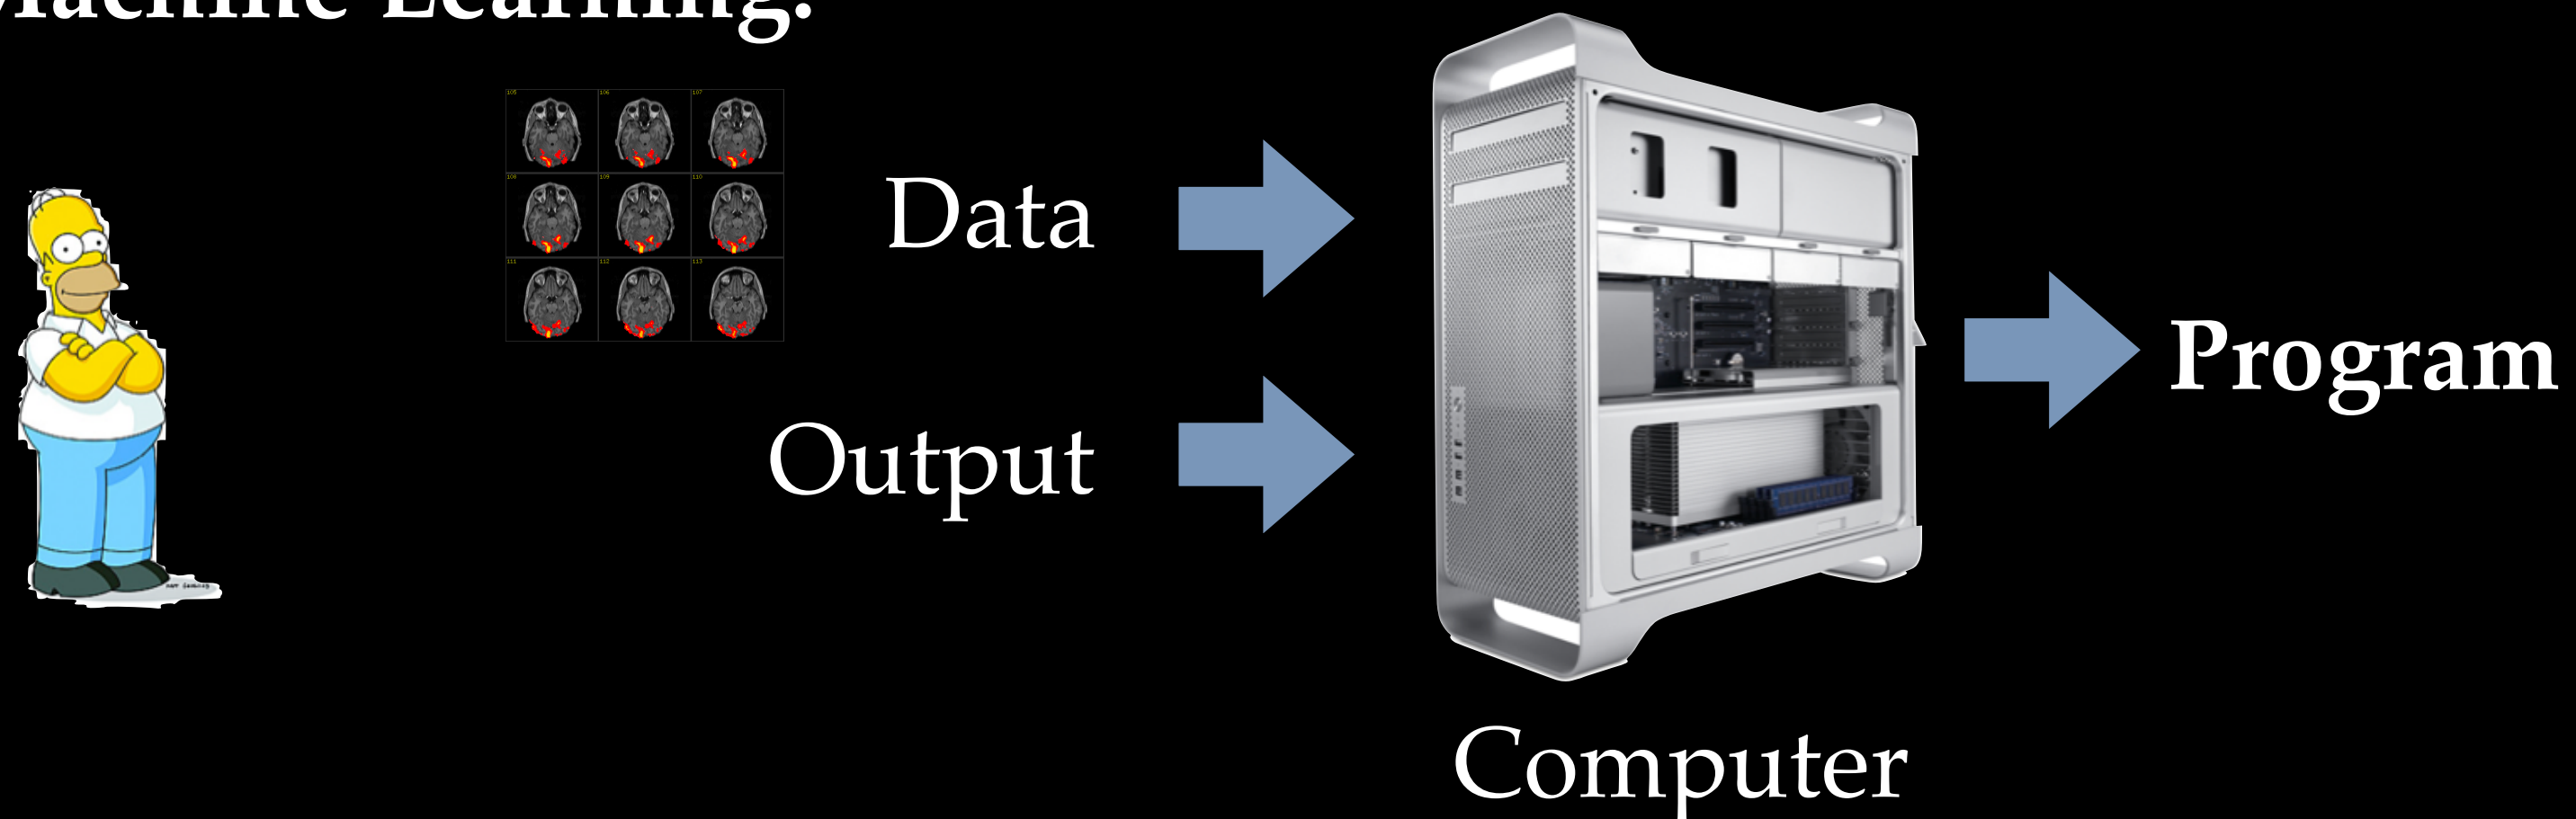
→ **Program**

# Machine Learning

# Machine Learning



Training:

Testing:

Data

Output

Computer

Program
(Model)

Data

Output

Computer
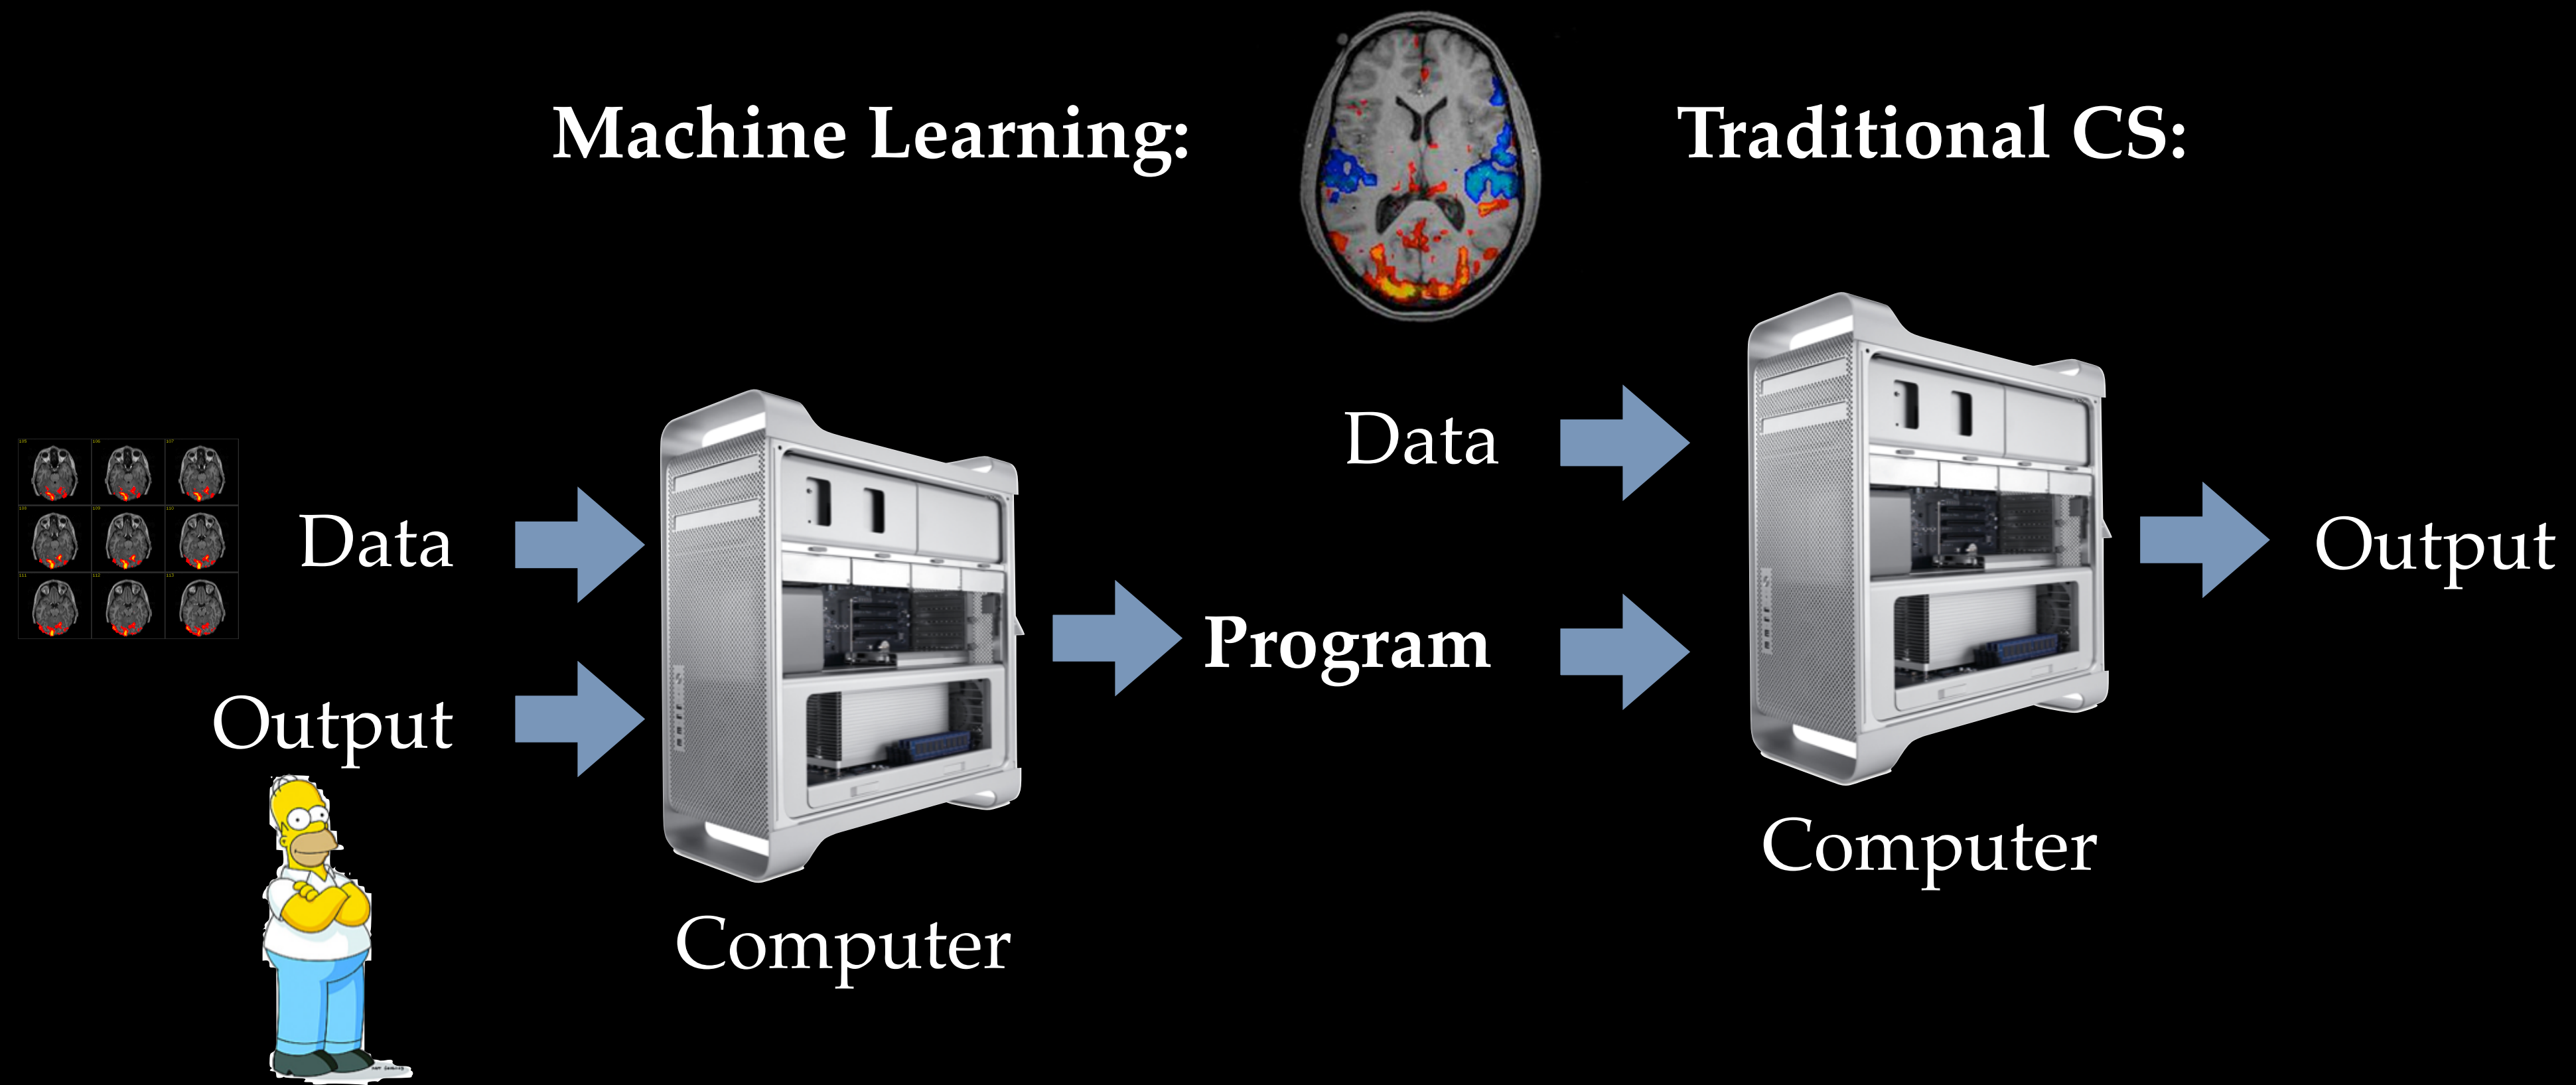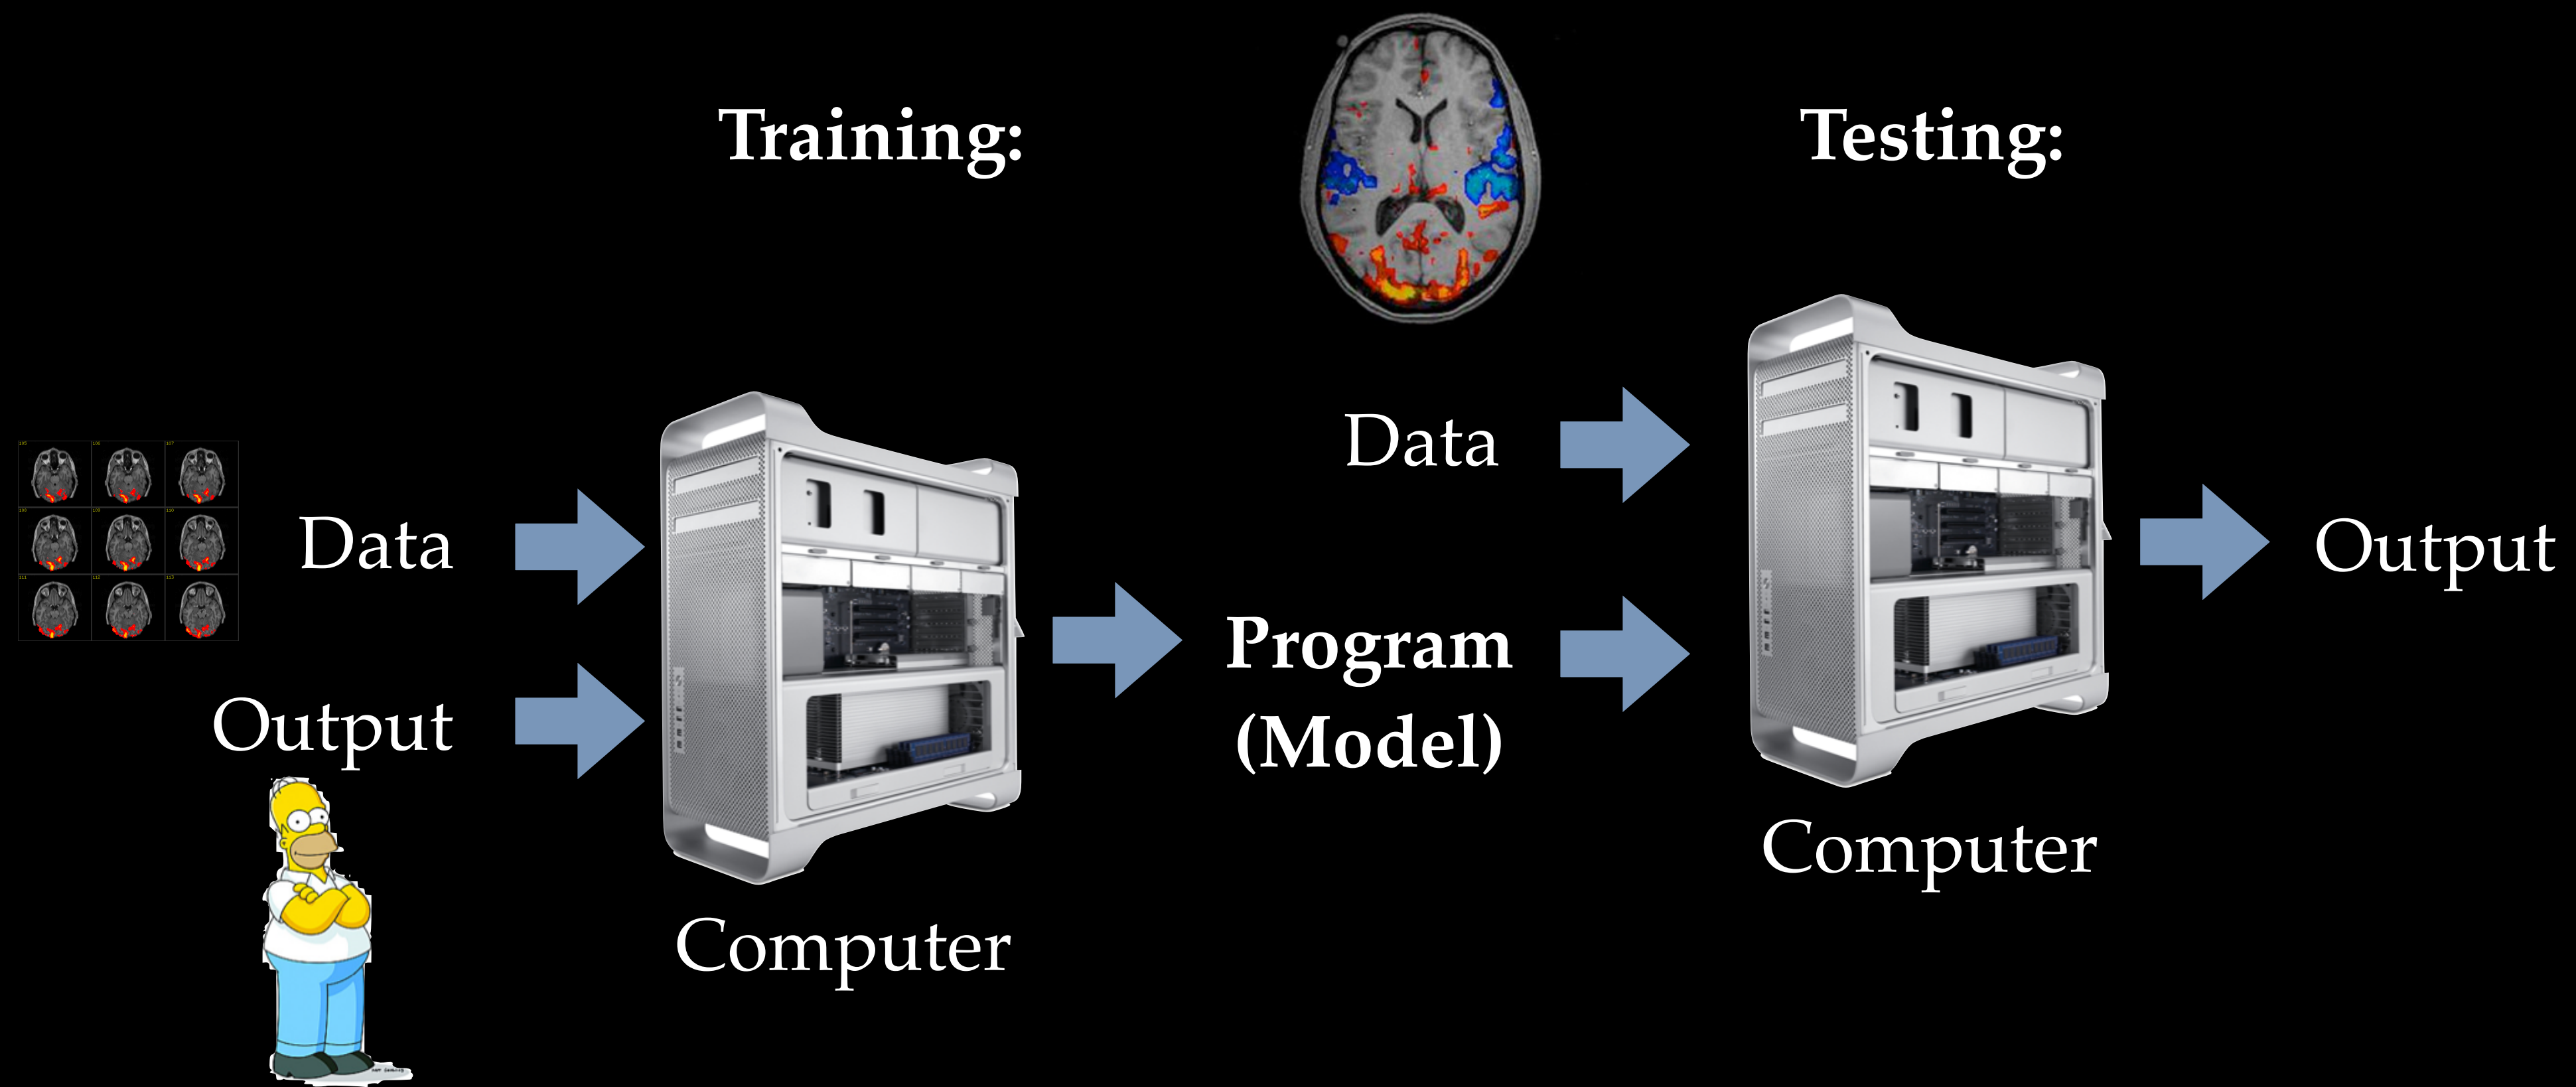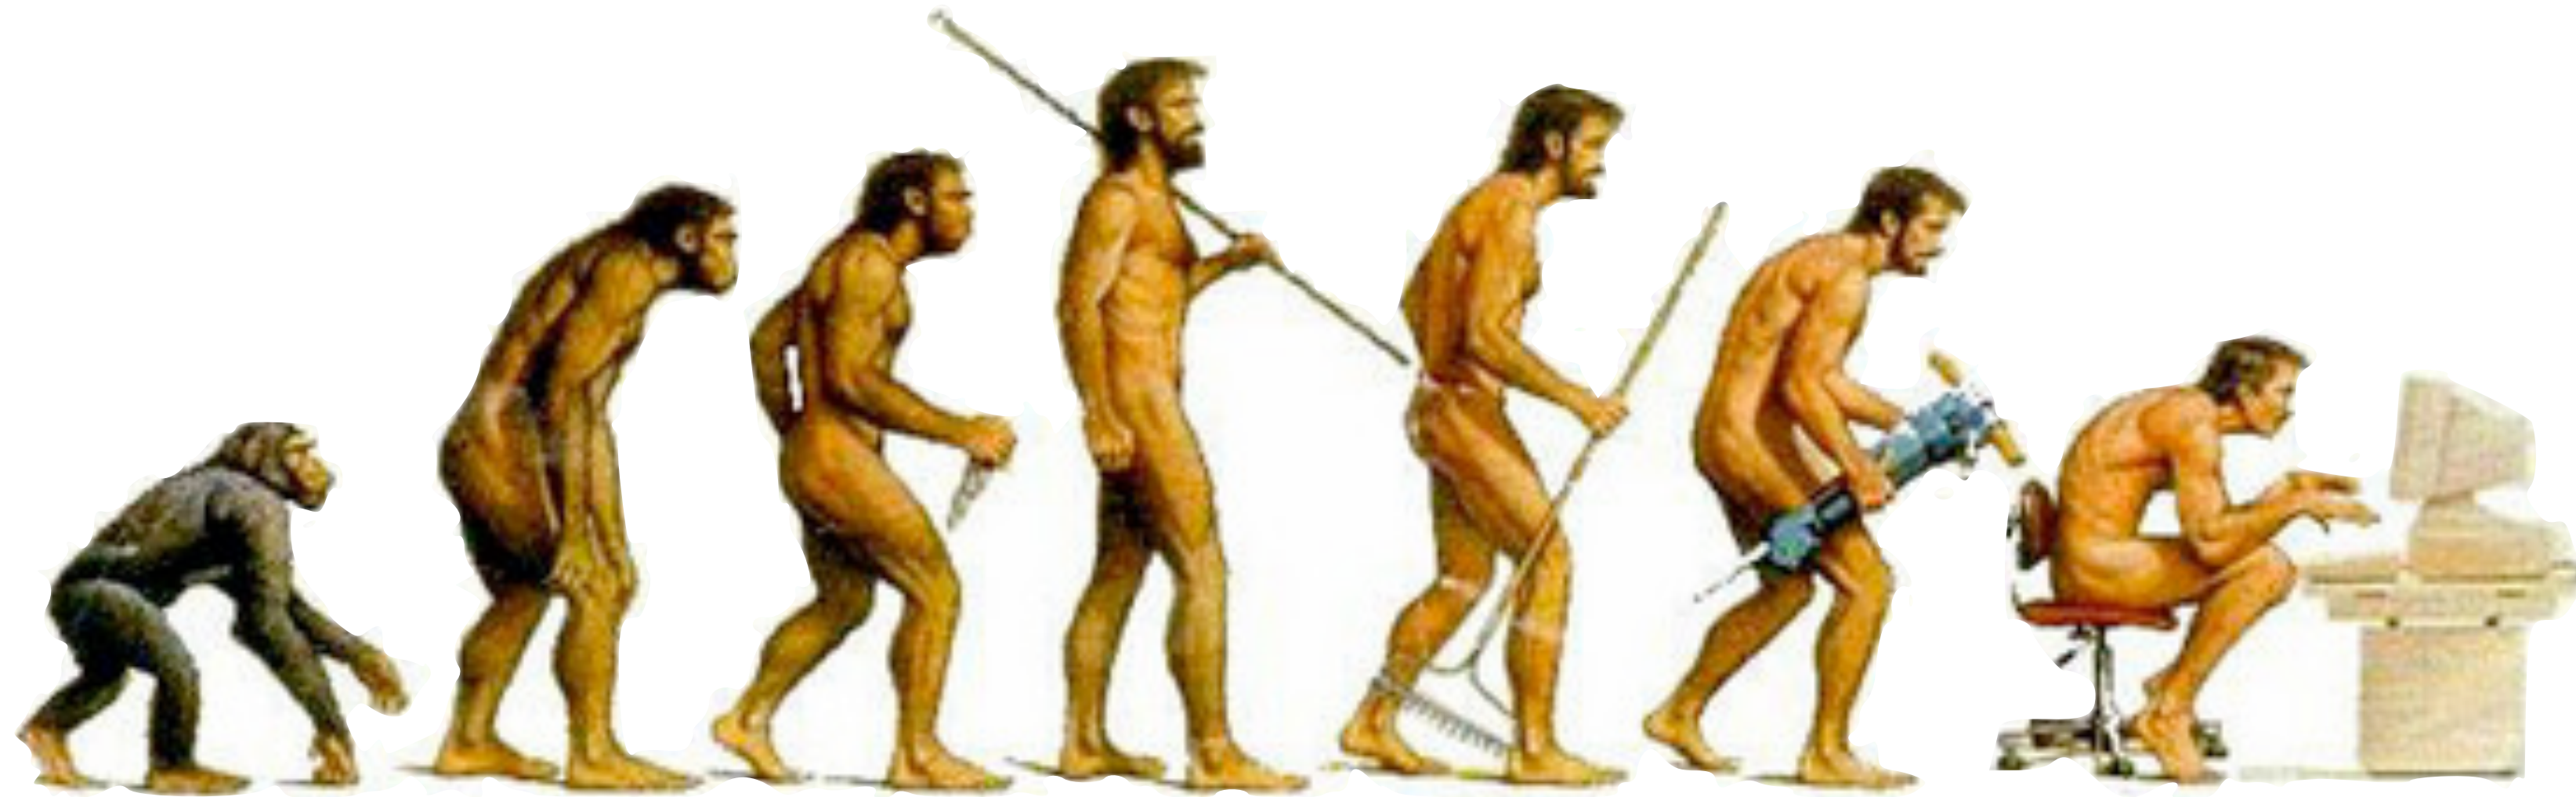
A (very brief) History of ML

# The Turing Test, 1950

**Alan Turing**

A machine is intelligent if its answers are indistinguishable from a human's
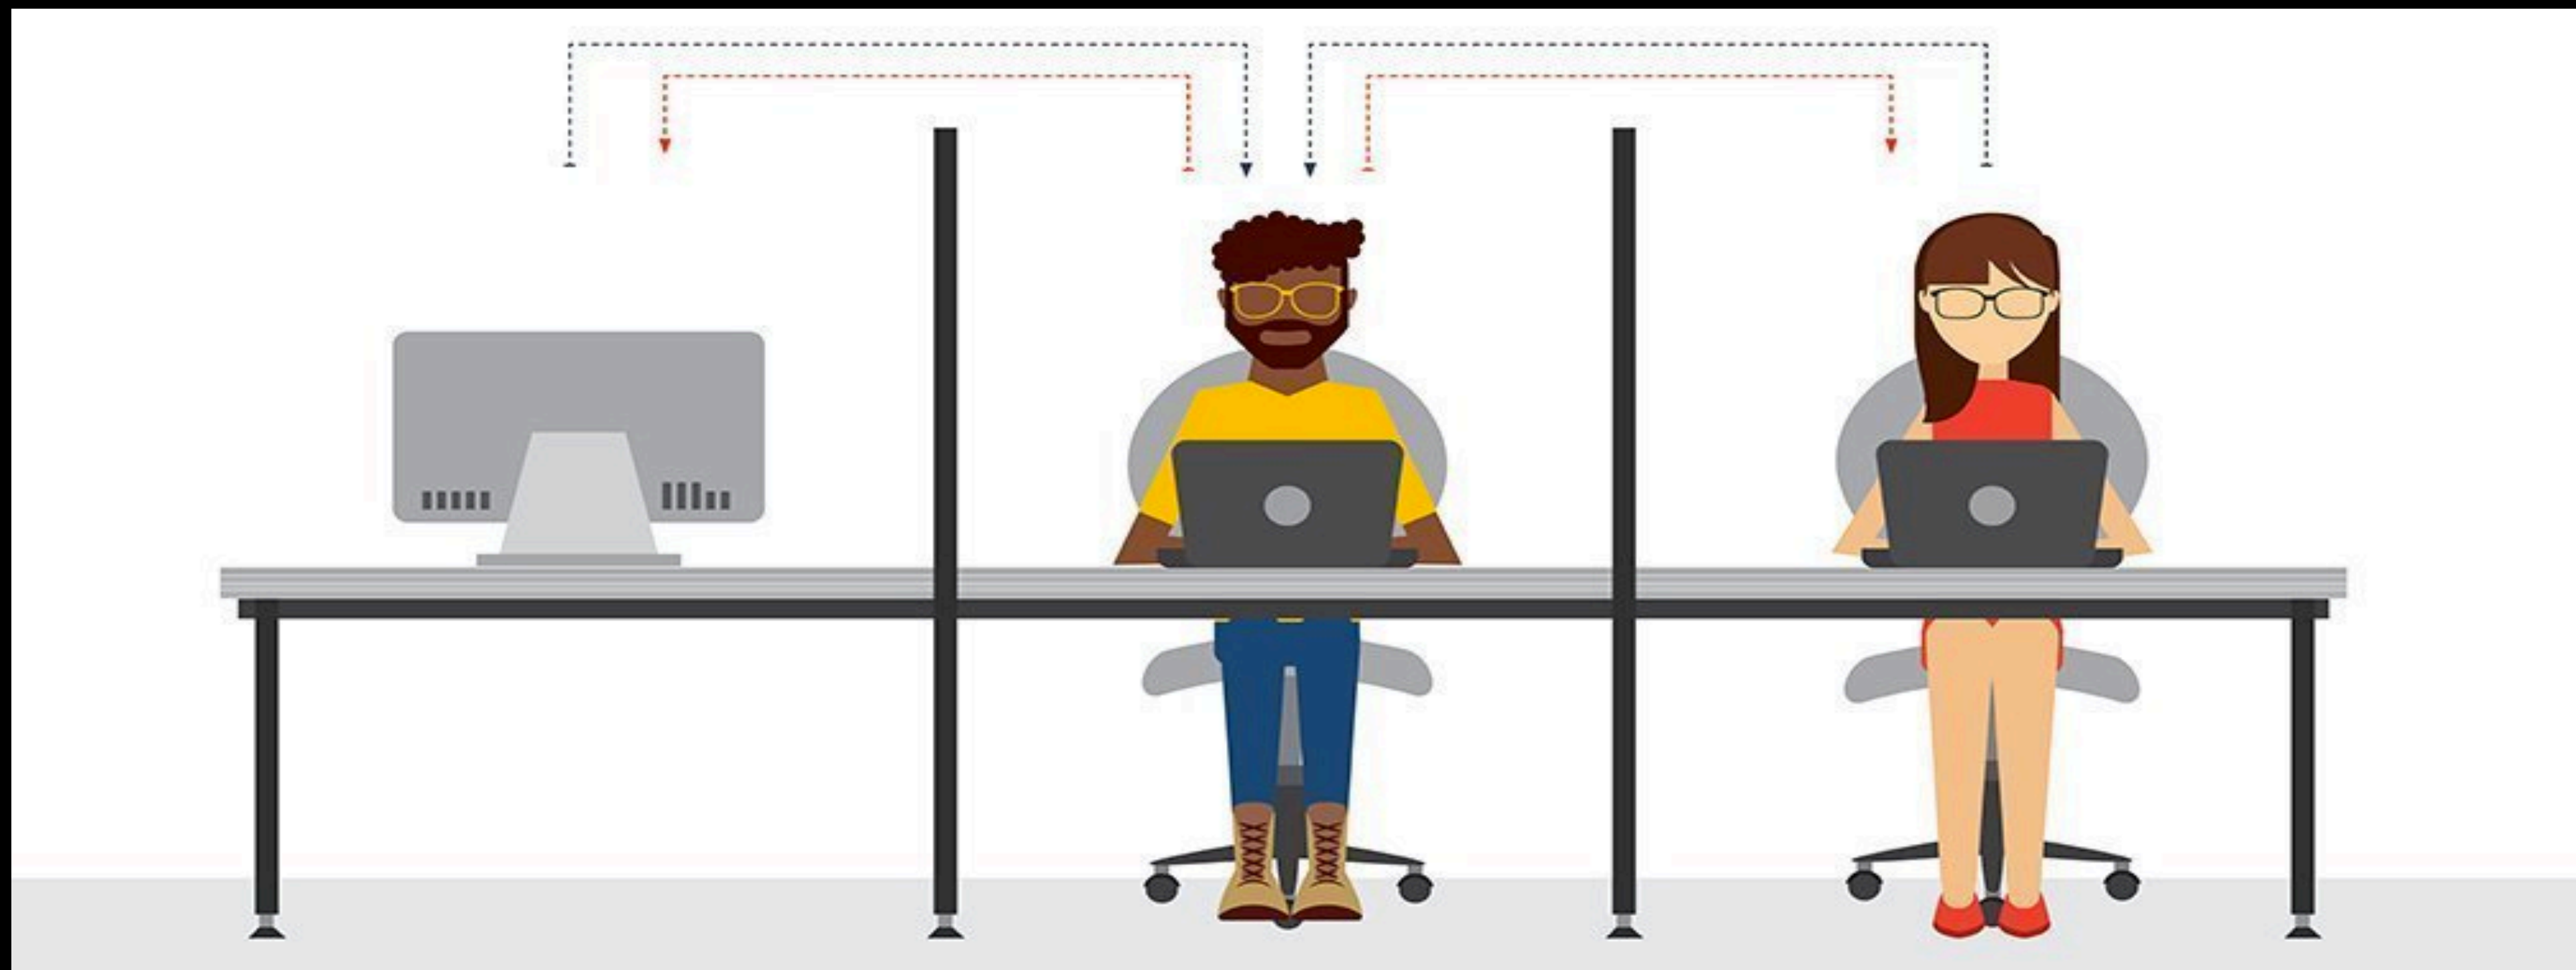
# Checkers Program, 1952

Arthur Samuel

Created a Checkers-playing program that got better overtime.

Also introduced the term "**Machine Learning**."
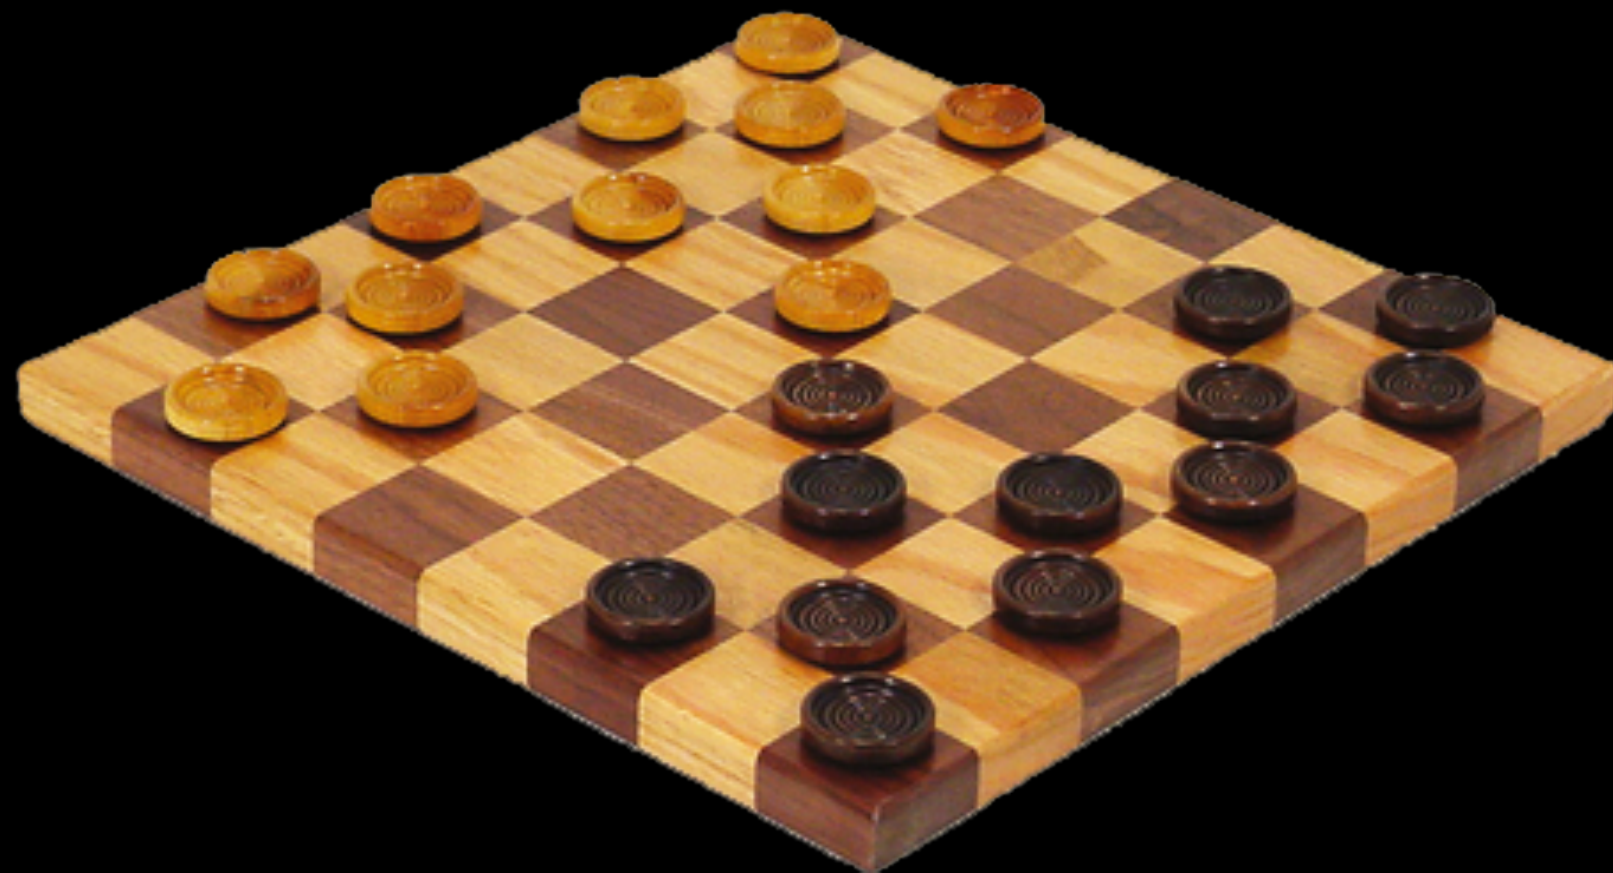
The term **Artificial Intelligence** originated in 1956 at a Workshop at Dartmouth

# Perceptron, 1957



Frank Rosenblatt
@ Cornell!

Predecessor of deep networks.

Separating two classes of objects using a linear threshold classifier.

# Perceptron, 1957

## Frank Rosenblatt
### @ Cornell!

*New Navy Device Learns by Doing*
   - The New York Times (July 8, 1958)

"Later perceptrons will be able to recognize people and call out their names and instantly translate speech in one language to speech or writing in another language, it was predicted."

---

**NEW NAVY DEVICE LEARNS BY DOING**

Psychologist Shows Embryo of Computer Designed to Read and Grow Wiser

WASHINGTON, July 7 (UPI) —The Navy revealed the embryo of an electronic computer today that it expects will be able to walk, talk, see, write, reproduce itself and be conscious of its existence.

The embryo—the Weather Bureau's $2,000,000 "704" computer—learned to differentiate between right and left after fifty attempts in the Navy's demonstration for newsmen.

The service said it would use this principle to build the first of its Perceptron thinking machines that will be able to read and write. It is expected to be finished in about a year at a cost of $100,000.

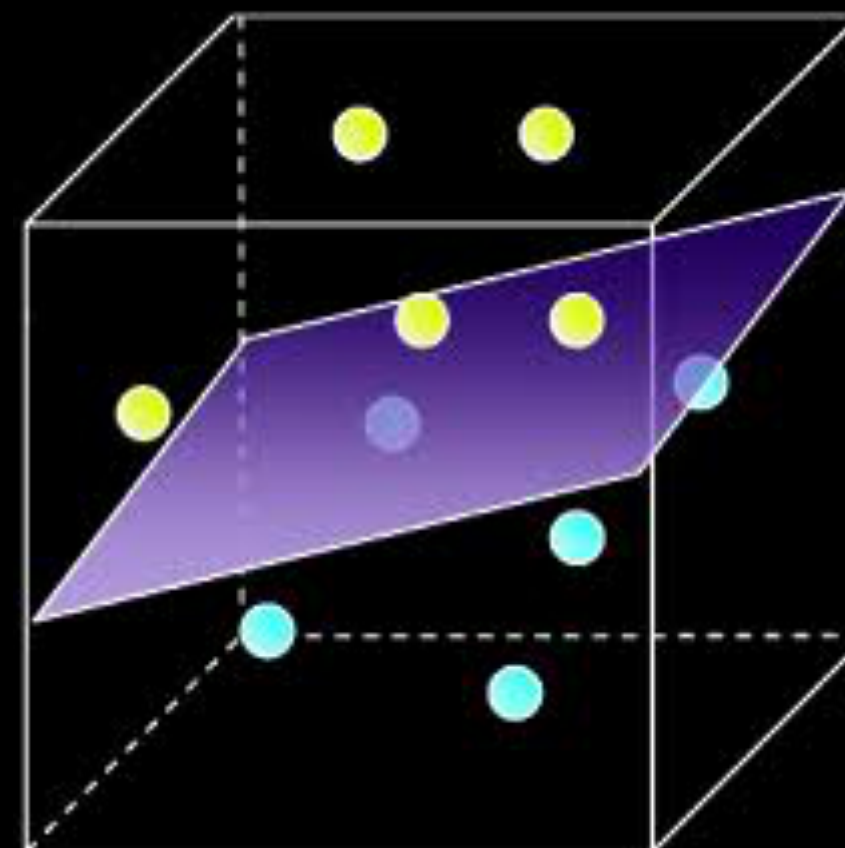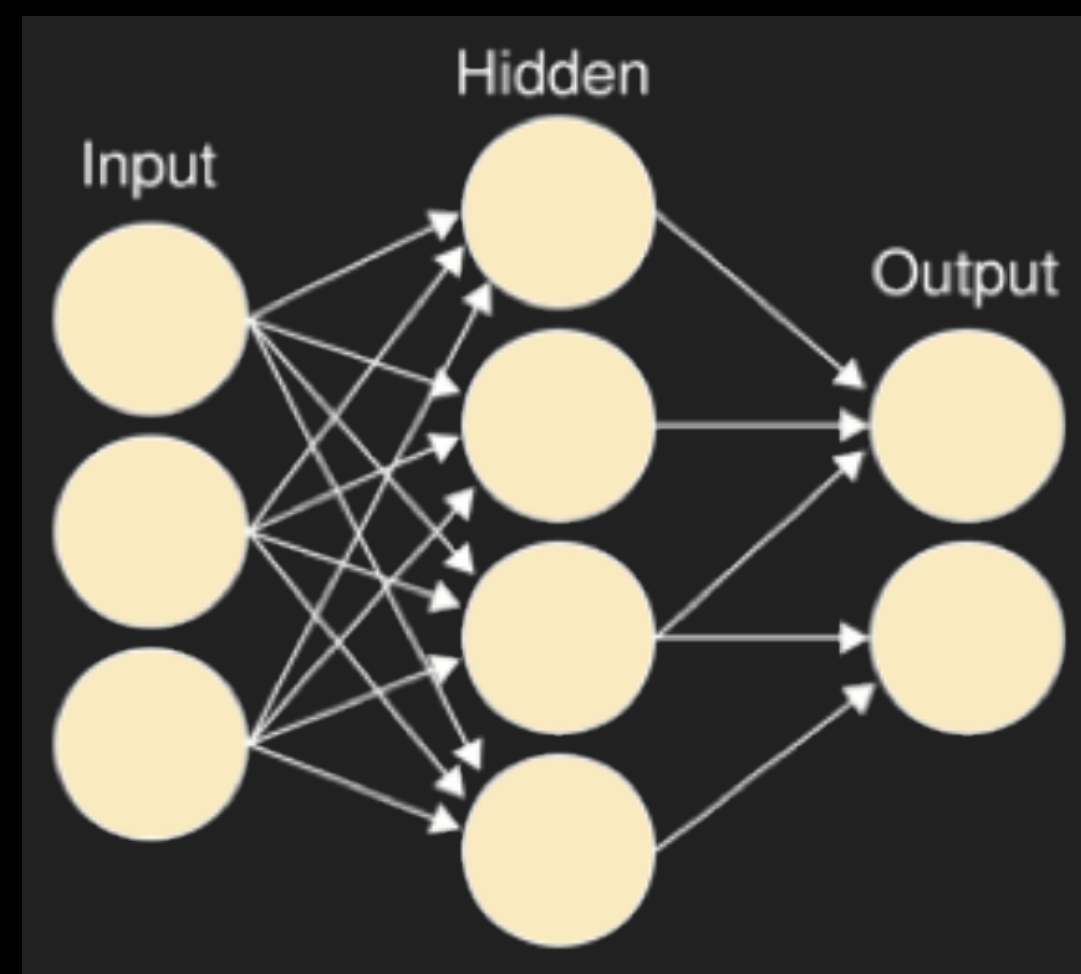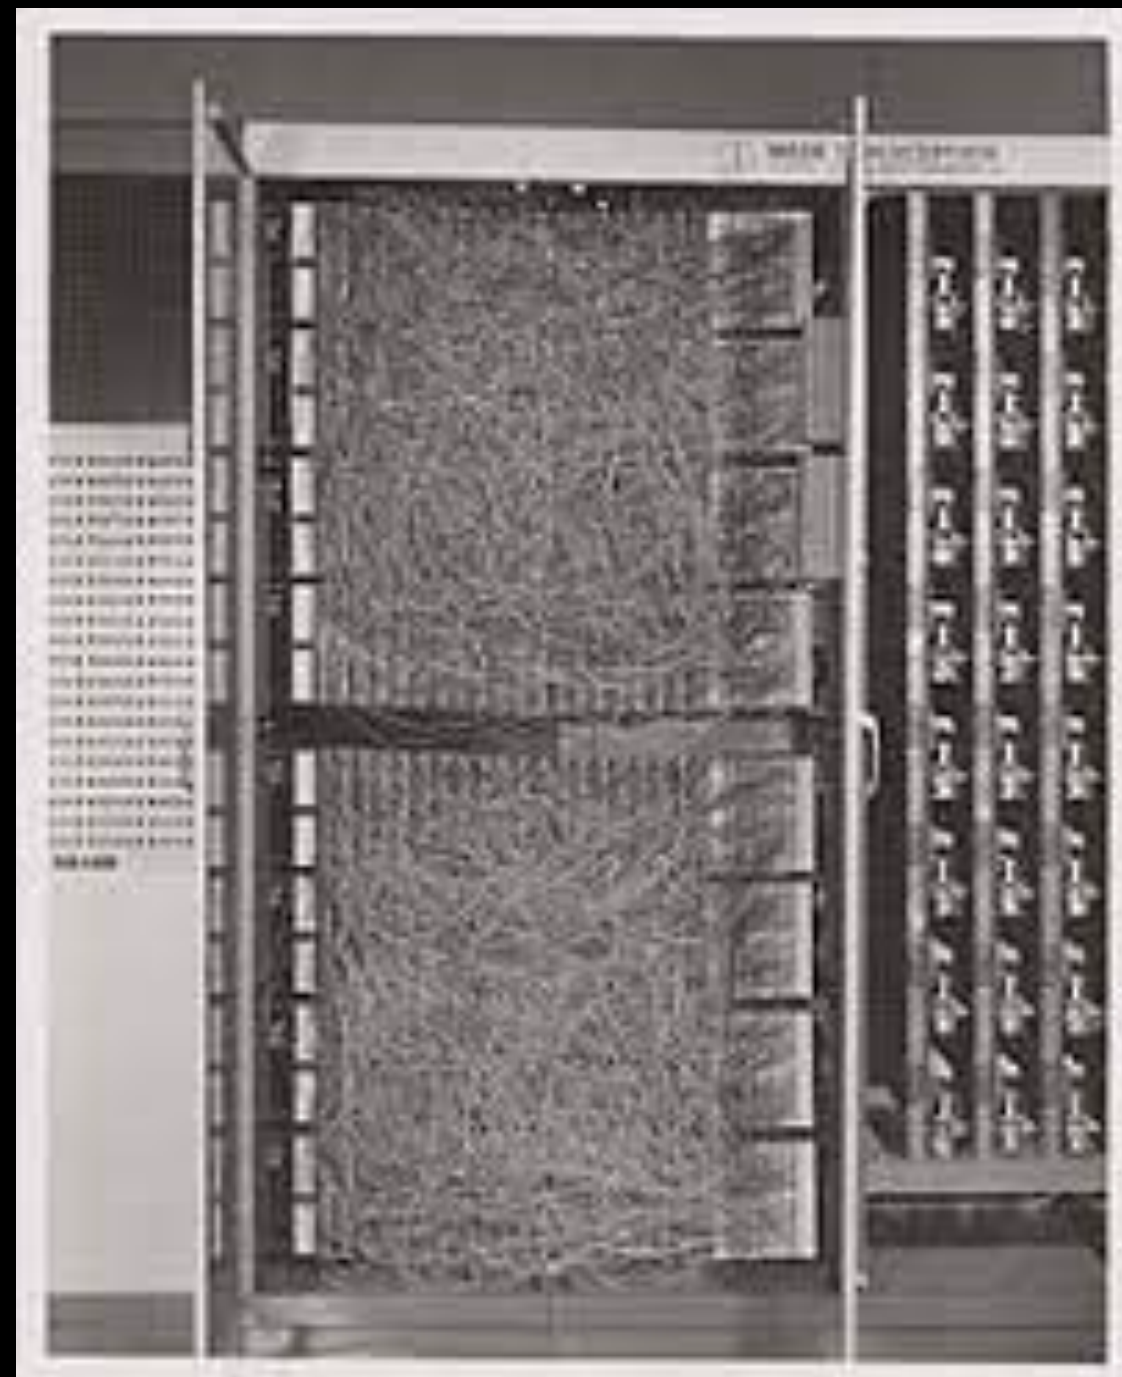Dr. Frank Rosenblatt, designer of the Perceptron, conducted the demonstration. He said the machine would be the first device to think as the human brain. As do human beings, Perceptron will make mistakes at first, but will grow wiser as it gains experience, he said.

Dr. Rosenblatt, a research psychologist at the Cornell Aeronautical Laboratory, Buffalo, said Perceptrons might be fired to the planets as mechanical space explorers.

**Without Human Controls**

The Navy said the perceptron would be the first non-living mechanism "capable of receiving, recognizing and identifying its surroundings without any human training or control."

The "brain" is designed to remember images and information it has perceived itself. Ordinary computers remember only what is fed into them on punch cards or magnetic tape.

Later Perceptrons will be able to recognize people and call out their names and instantly translate speech in one language to speech or writing in another language, it was predicted.

Mr. Rosenblatt said in principle it would be possible to build brains that could reproduce themselves on an assembly line and which would be conscious of their existence.

In today's demonstration, the "704" was fed two cards, one with squares marked on the left side and the other with squares on the right side.

**Learns by Doing**

In the first fifty trials, the machine made no distinction between them. It then started registering a "Q" for the left squares and "O" for the right squares.

Dr. Rosenblatt said he could explain why the machine learned only in highly technical terms. But he said the computer had undergone a "self-induced change in the wiring diagram."

The first Perceptron will have about 1,000 electronic "association cells" receiving electrical impulses from an eye-like scanning device with 400 photo-cells. The human brain has 10,000,000,000 responsive cells, including 100,000,000 connections with the eyes.

The New York Times

# Neural Network Boom, 1960-1979



Frank Rosenblatt
@ Cornell!



RETINA OF
S UNITS

A UNITS

R → OUTPUT SIGNAL
+1 OR -1

Fig. 2. Network organization of a typical elementary perceptron

SEPTEMBER 1964          Rosenblatt—Analytic Techniques for the Study of Neu
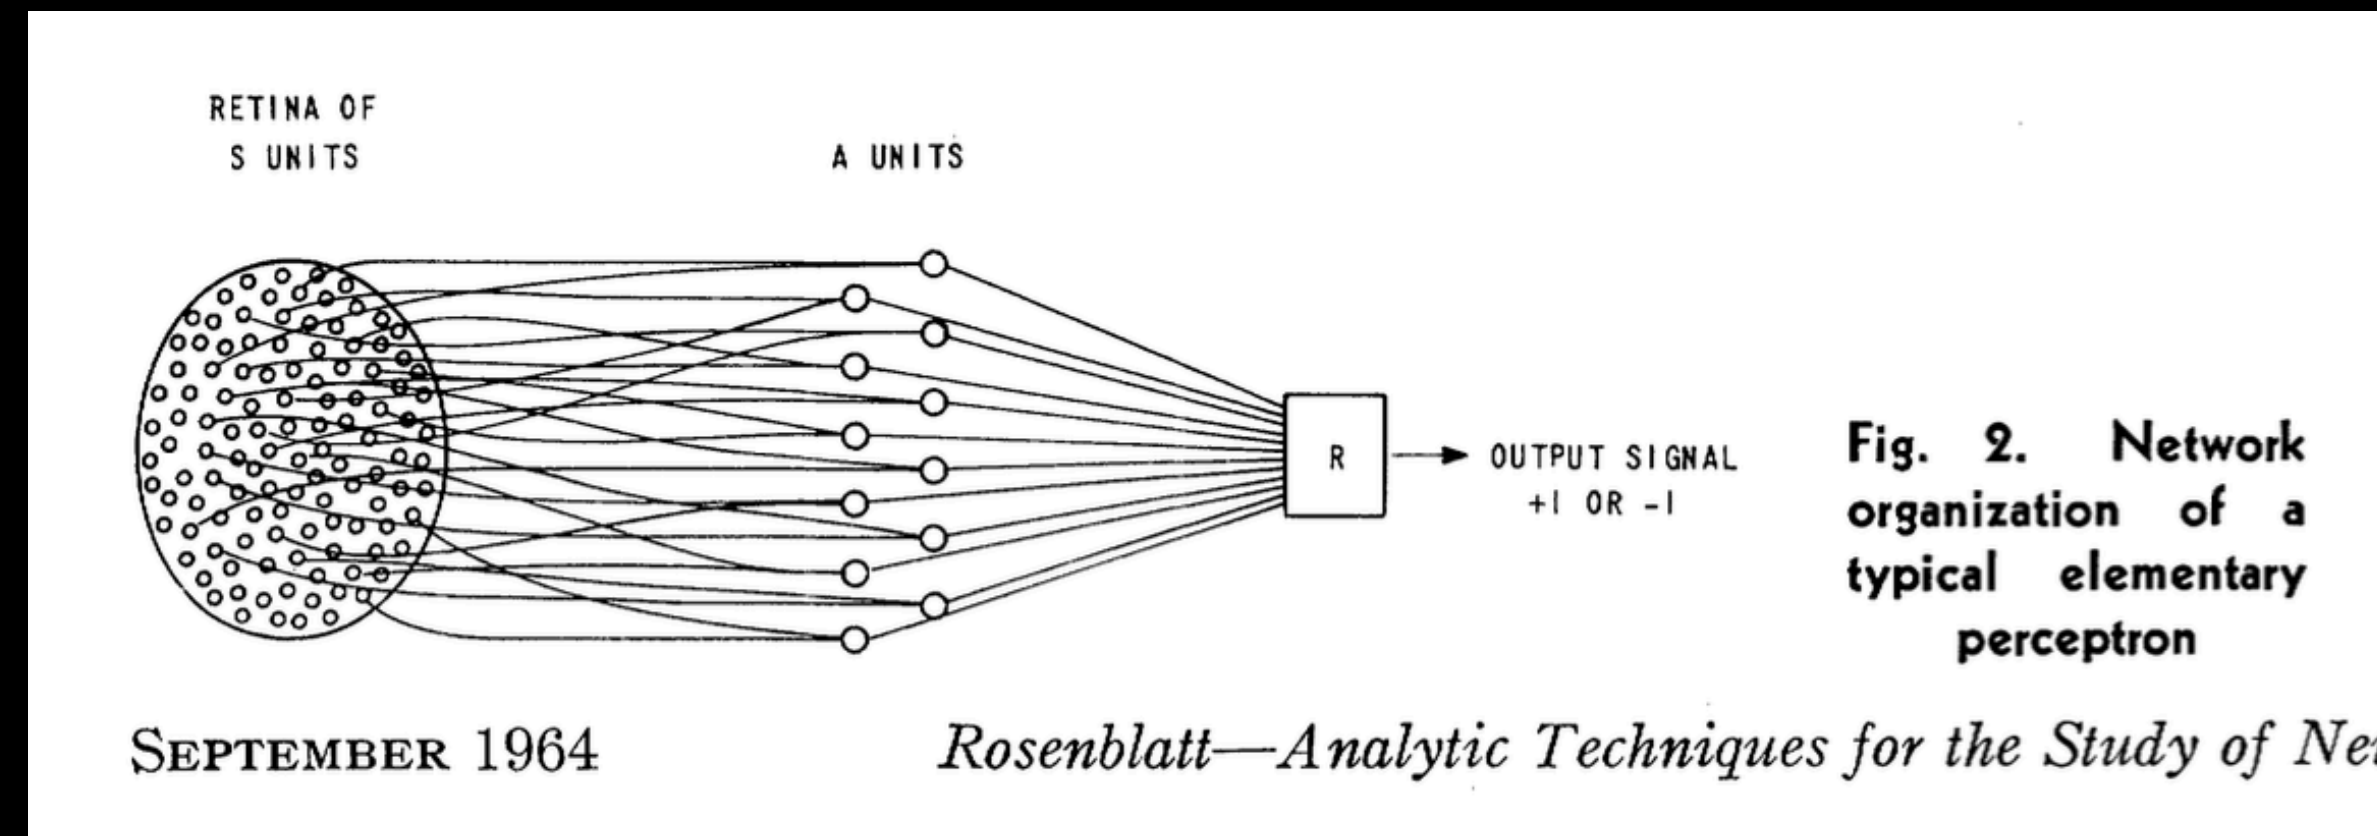
- 1962 Rosenblatt invents Multi-Layer Perceptron (MLP) (fixed hidden layer)
- 1965 Ivakhnenko and Lapa introduce first Feed Forward Neural Net (FFNN)
- 1967 First FFN trained with SGD [Amari 1967]
- 1970 Modern back-propagation is introduced by Seppo Linnainmaa

# First AI Winter (1974-1980)

- (1969) Minsky & Papert "killed" AI

  - Perceptron cannot learn XOR function

- Burst huge expectation bubble

- Speech understanding / translation fails

- UK and US stop funding AI research

- Neural Networks and AI become "bad words"

- Rise of Rule based Systems
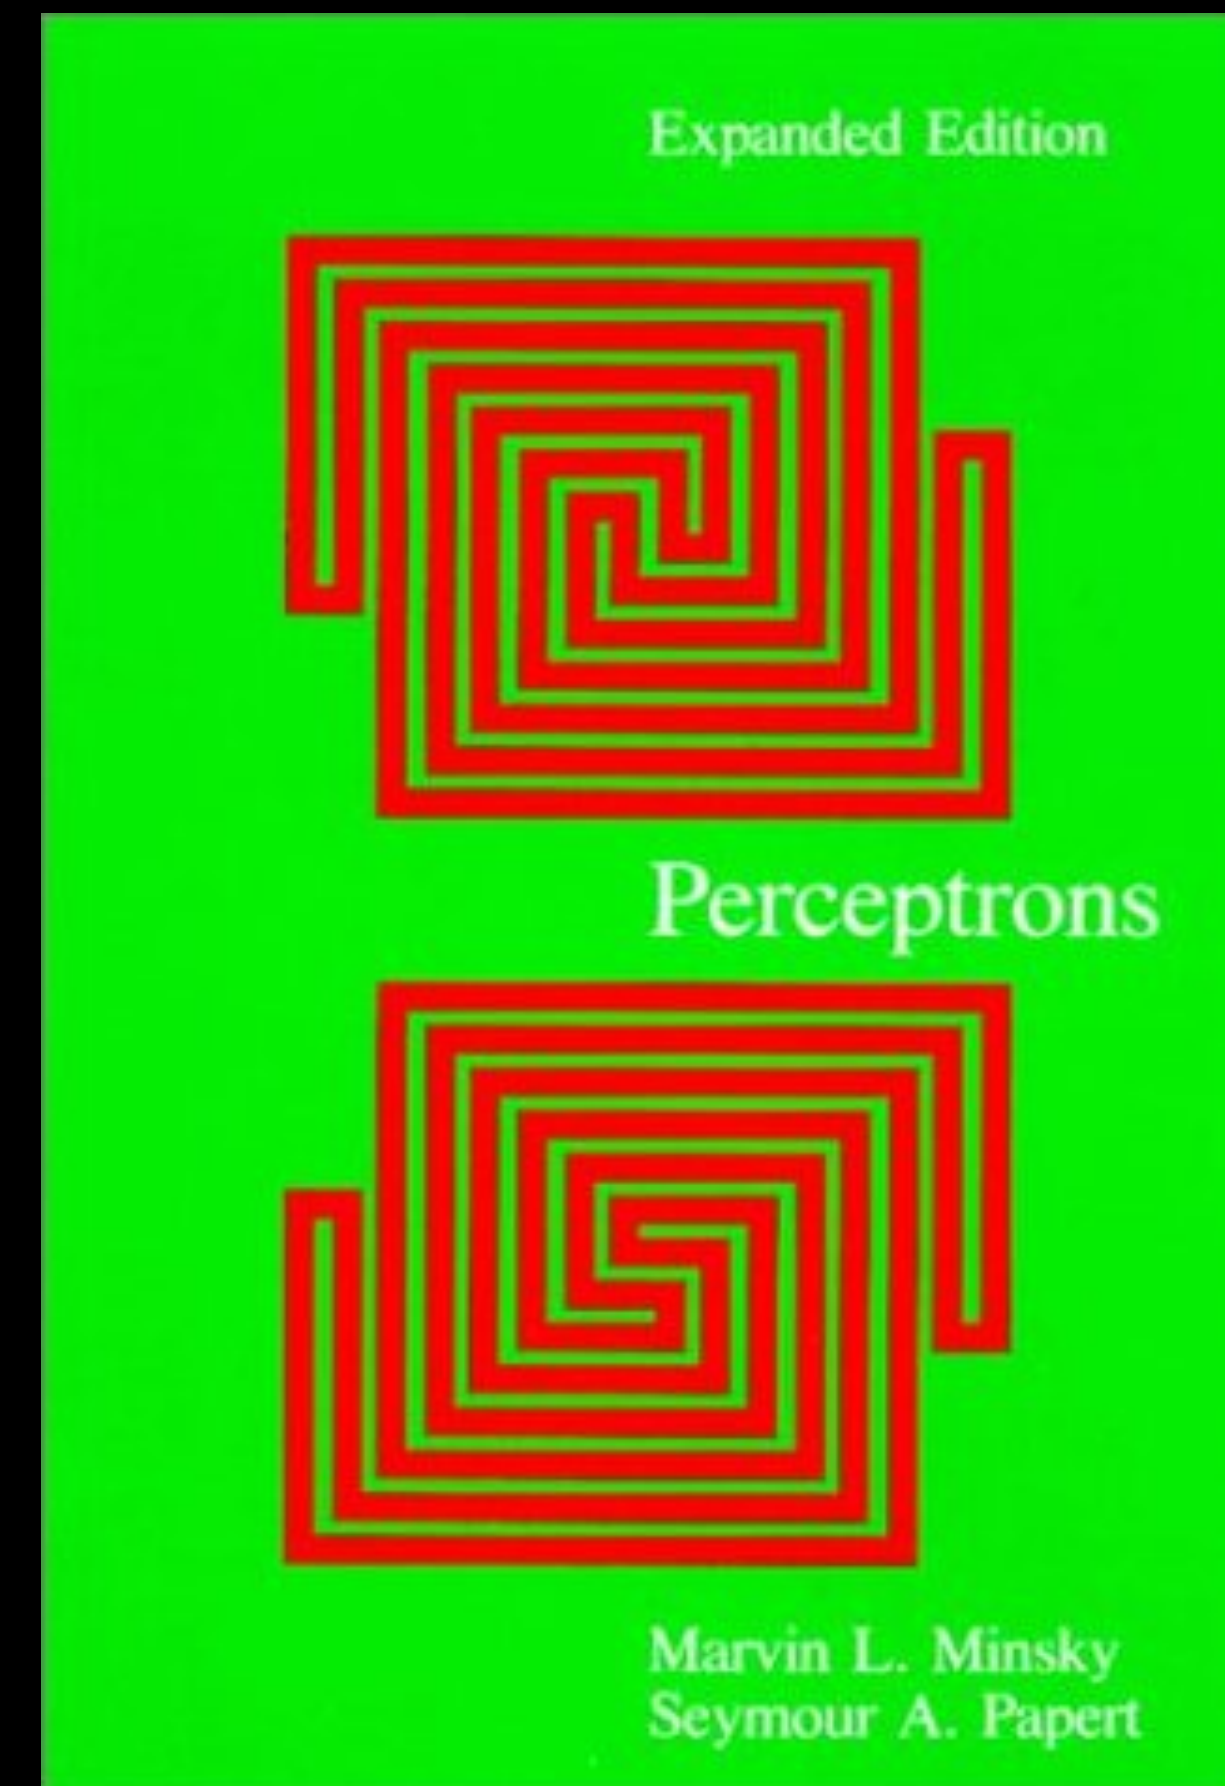
# A.I. Boom - then Winter  [1980-1993]

- AI BOOM: 1980-1987

  - Search Algorithms (No Learning)

  - Expert Systems

- AI winter: 1987-1993

  - Expectations too overhyped

  - Exponential algorithms misunderstood

  - A.I. Bubble bursts

# Rebirth of A.I. as Machine Learning

AI

- Machine Learning:

  - Originally: Mostly a name game to get funding.

- Differences in approach:

  - ML: Bottom up, AI: Top down

  - ML: More practical smaller goals

  - Based on **Statistics and Optimization, not Logic and Search**

ML
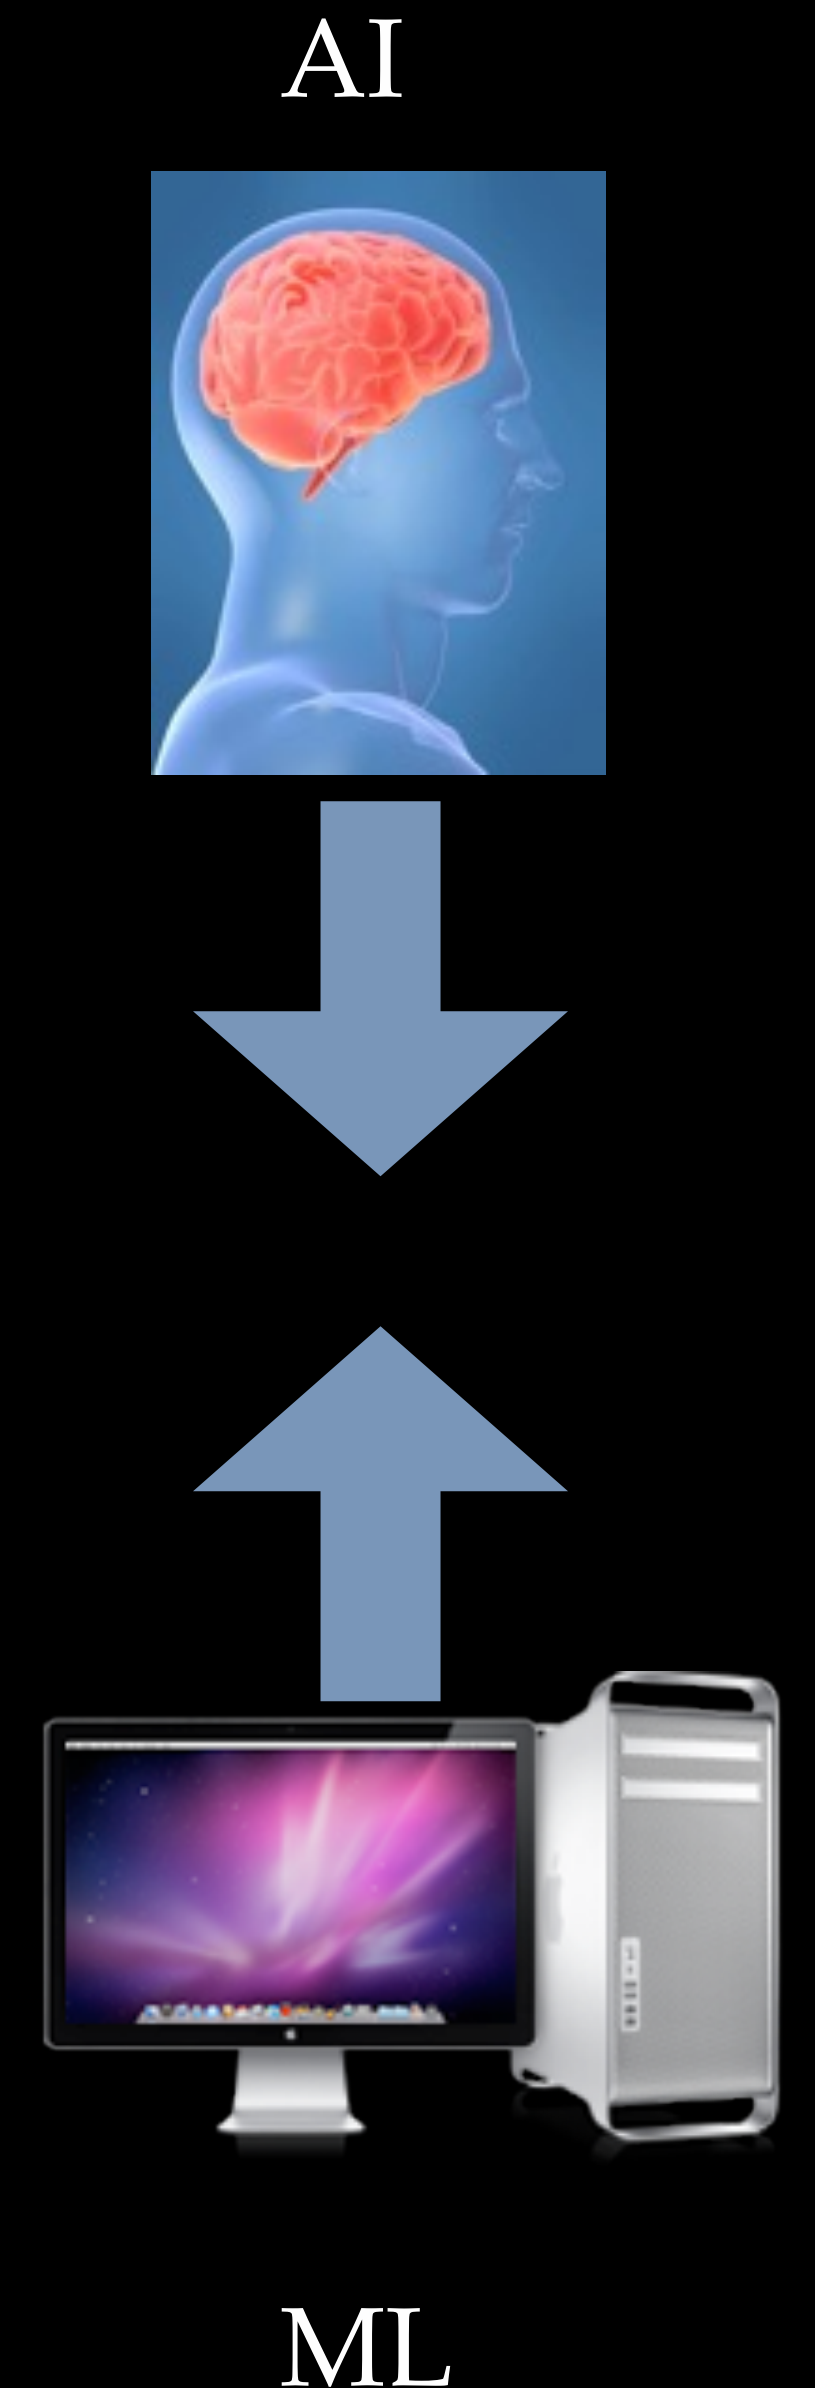
# TD-Gammon (1994)

- Gerry Tesauro (IBM) teaches a neural network to play Backgammon. The net plays 100K+ games **against itself** and beats world champion [Neurocomputation 1994]

- Algorithm teaches **itself** how to play so well!!!

- Algorithm found new techniques that people had erroneously ruled out.

# Deep Blue (1997)

- IBM's Deep Blue wins against Kasparov in chess. Crucial winning move is made due to Machine Learning (G. Tesauro).

    - (Mostly a more classical AI system)

# Expanding the reach, 2000s



Learning to rank

- Powering search engines: Yahoo, Google, Bing, …

Topic Modeling:

- Detecting and organizing documents by subject matter.

- Making sense of the unstructured data on the web.

Online economy:

- Ad placement and pricing.

- Product recommendation.

**Machine learning became profitable!
Euphoria! Excitement!**

# 2000-2010 The Convex Years

- Success of **Support Vector Machines** (SVM) and **Convex** Optimization

- **Kernelization** enables non-linear classifiers that are still convex

- Allows **rigorous proofs** about convergence, learning guarantees

- Finally **solid theory** and **empirical success**

**-** (Non-convex algorithms, e.g. Neural Networks fell out of favor and were frowned upon.)
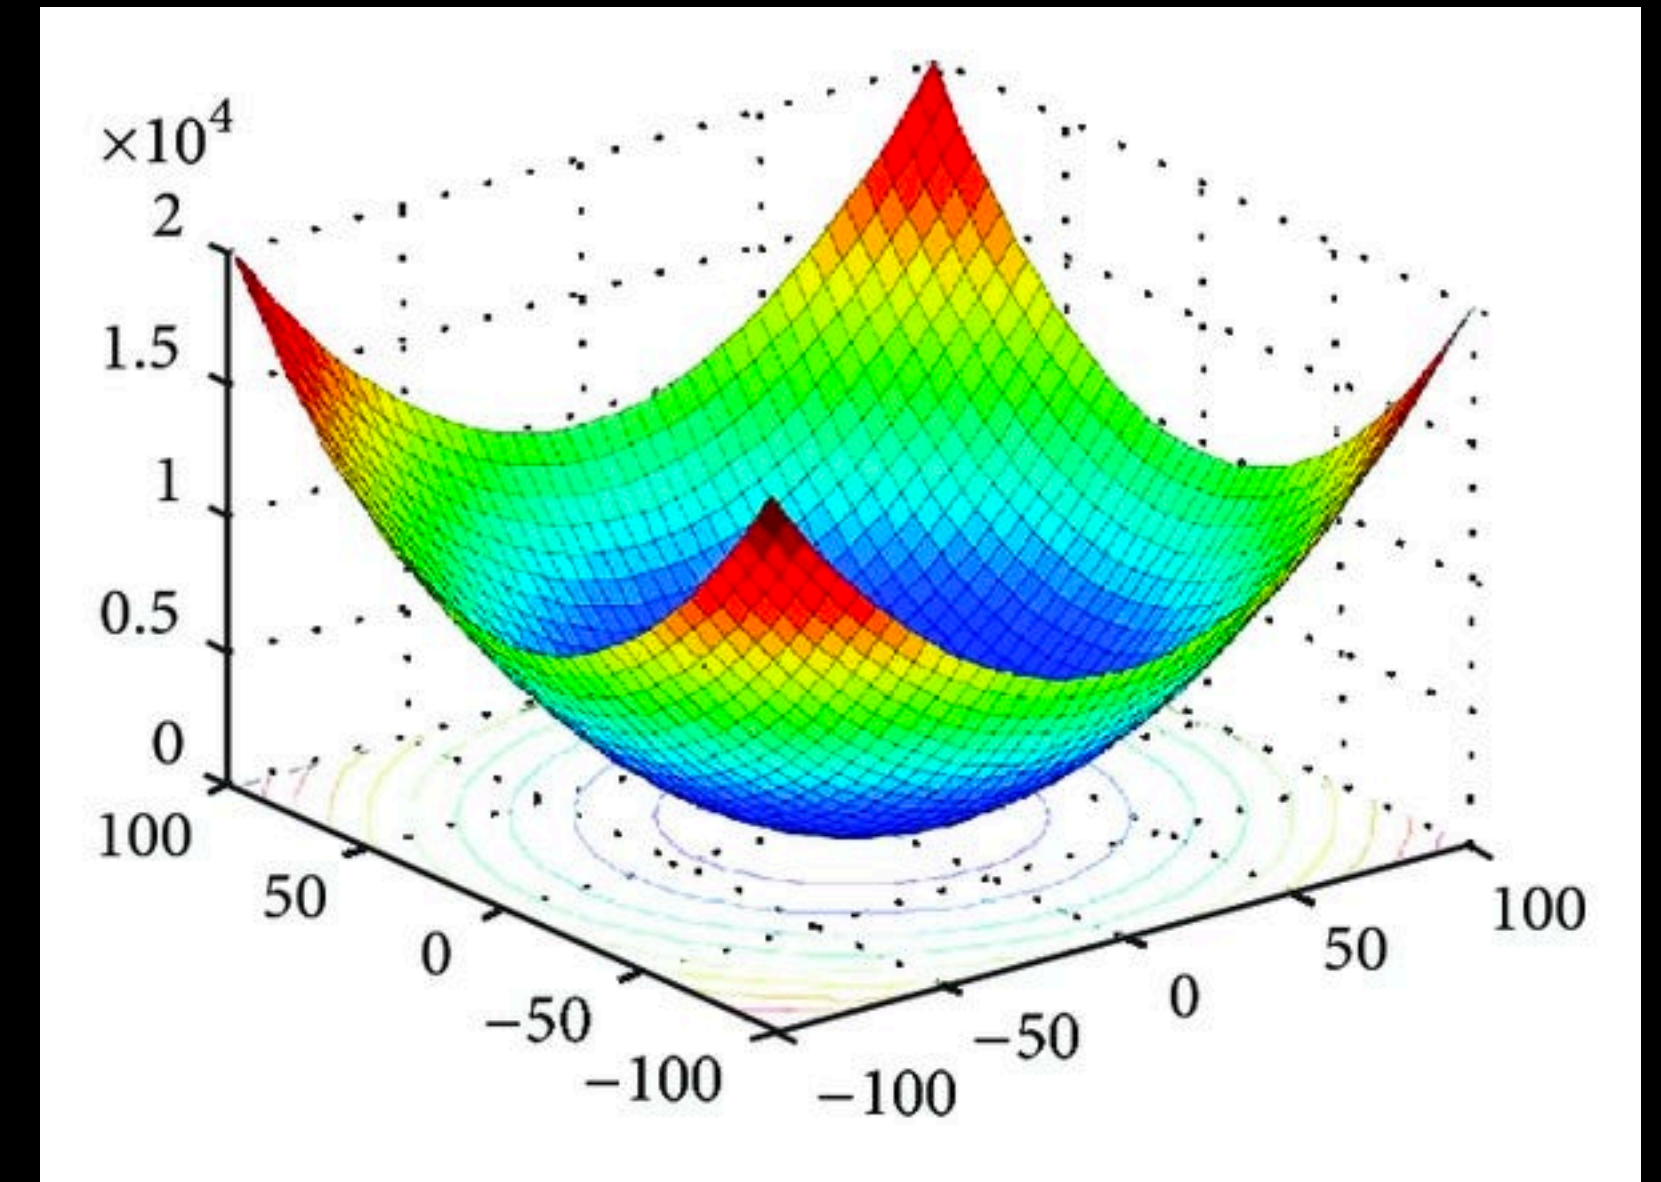
# Return of Neural Networks, 2010s

Neural networks (aka Multi-Layer Perceptrons) return and excel at image recognition, speech recognition, …

The 2018 Turing award was given to Yoshua Bengio, Geoff Hinton, and Yann LeCun.
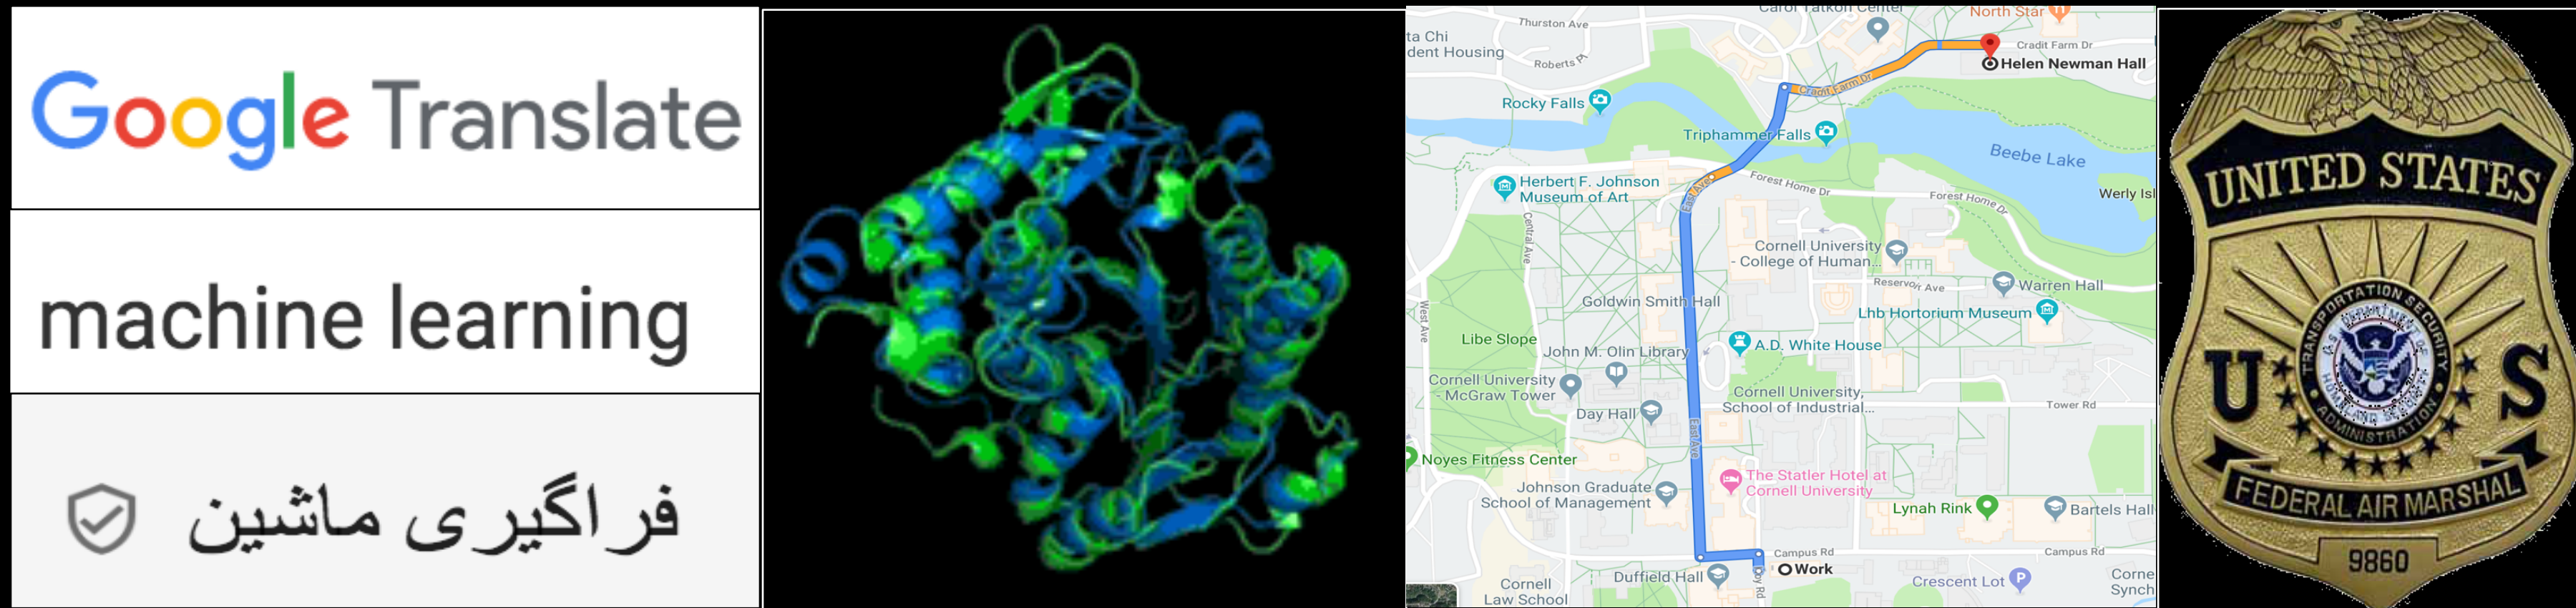
# 2016 Alpha Go

- 1920 CPUs and 280 GPUs

- Deep Mind's Alpha Go wins against Lee Sedol 5:1

- Big shock

  - China invests heavily in AI research

- Beginning of "AI arms race"

# Now: Surrounded by Machine Learning
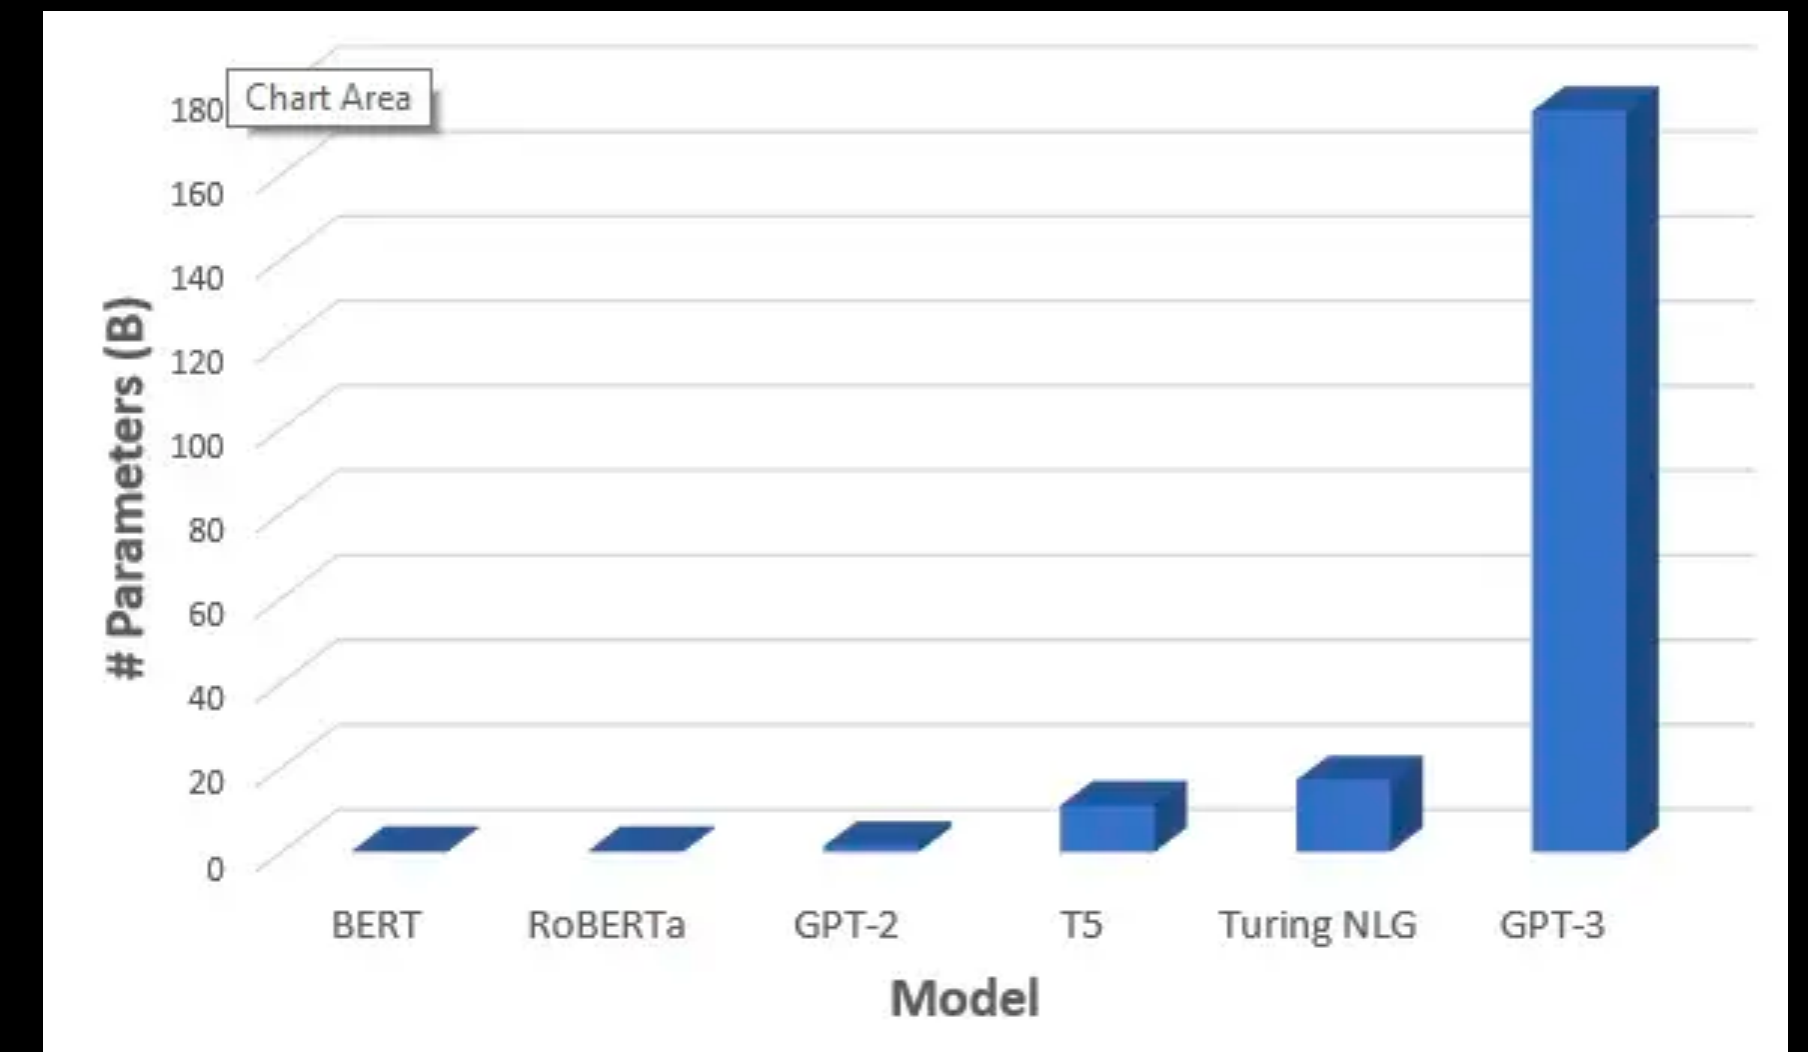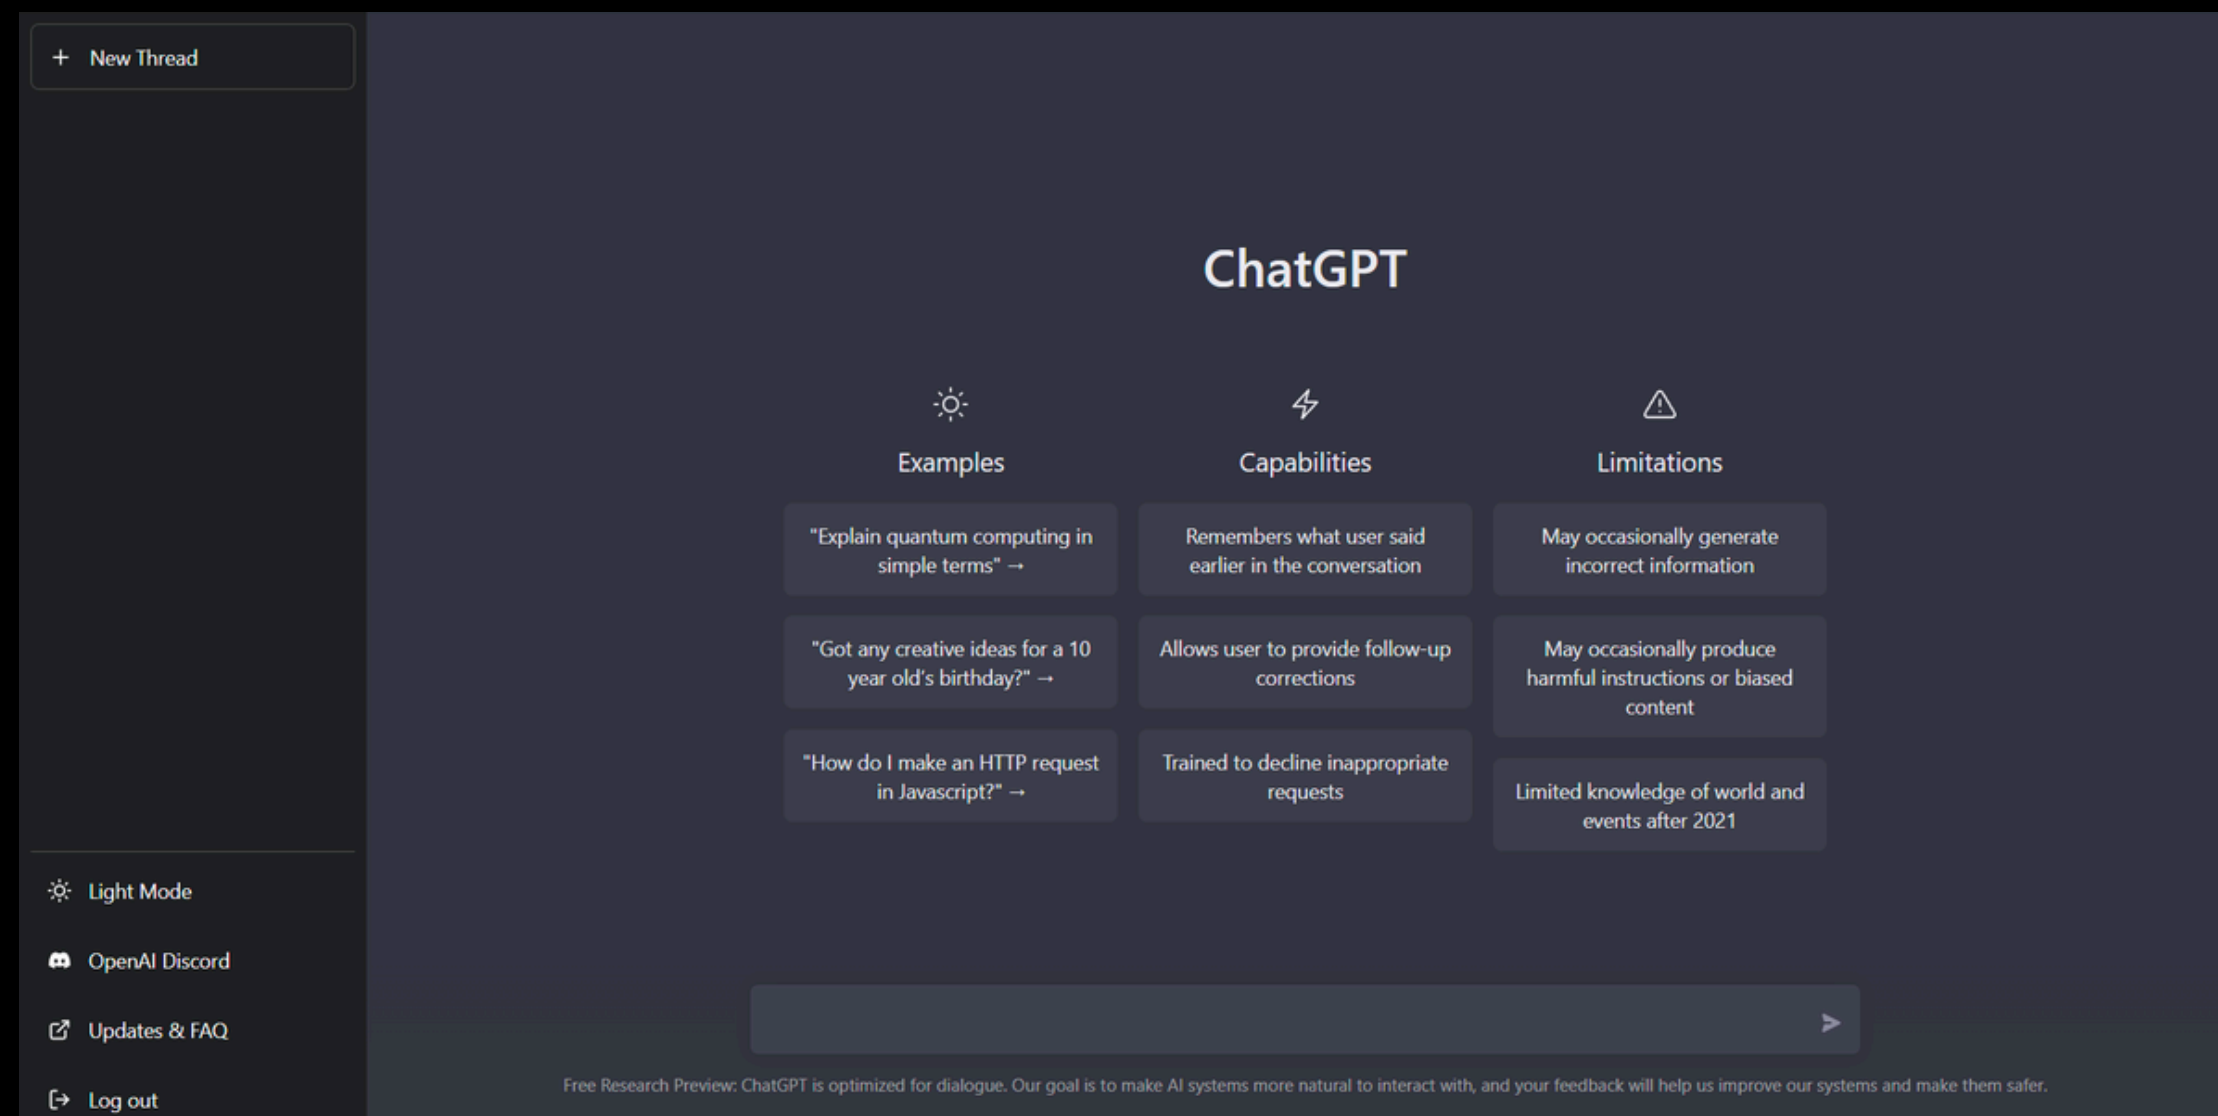
**Euphoria mixed with fear and anxiousness**

# 2021 Generative Models become Mainstream

- 178 Billion Parameters

- Cost: $4.6 Million (est.)

- 355 GPU years of training (est.)

- Trained on 45 TB of text

**T**    write a biblical verse in the style of the king james bible explaining how to remove a peanut butter sandwich from a VCR

And it came to pass that a man was troubled by a peanut butter sandwich, for it had been placed within his VCR, and he knew not how to remove it.

And he cried out to the Lord, saying, "Oh, Lord, how can I remove this sandwich from my VCR, for it is stuck fast and will not budge?"

And the Lord spoke unto him, saying, "Fear not, my child, for I shall guide thy hand and show thee the way. Take thy butter knife, and carefully insert it between the sandwich and the VCR, and gently pry them apart. And with patience and perseverance, the sandwich shall be removed, and thy VCR shall be saved."

And the man did as the Lord commanded, and lo and behold, the sandwich was removed from the VCR, and the man was saved.

And the Lord said, "Verily I say unto thee, seek not to put thy peanut butter sandwiches in thy VCR, for it is not a suitable place for such things. Rather, keep thy sandwiches in thy refrigerator or on thy plate, where they belong."

And the man heeded the Lord's words, and from that day forth, he kept his sandwiches in their proper place, and was saved from trouble and woe. Amen.

# DALLE Text2Images - ML is creative



Teddy bears
mixing sparkling chemicals as mad scientists
in a steampunk style
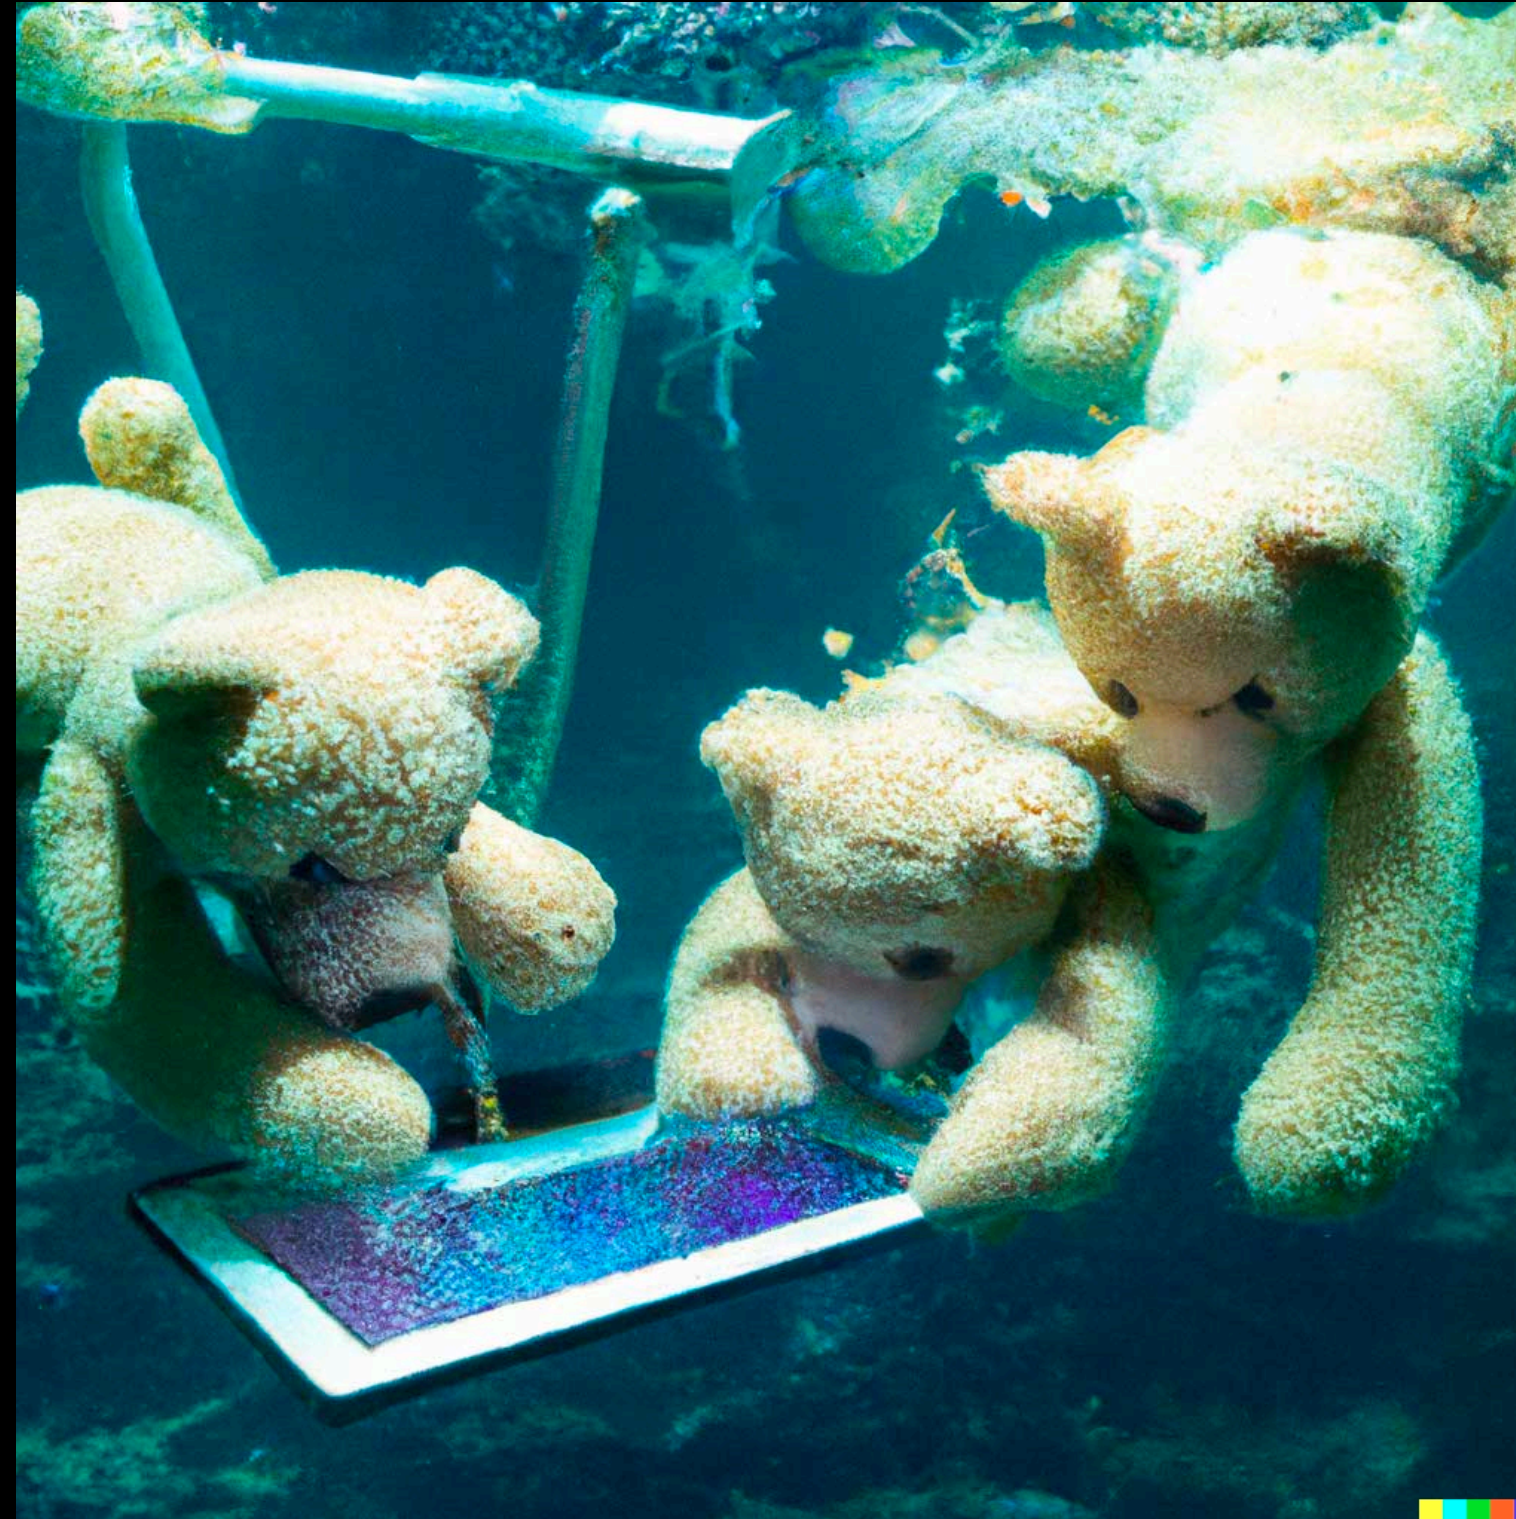
Teddy bears
working on new AI research
underwater with 1990s technology

An astronaut
playing basketball with cats in space
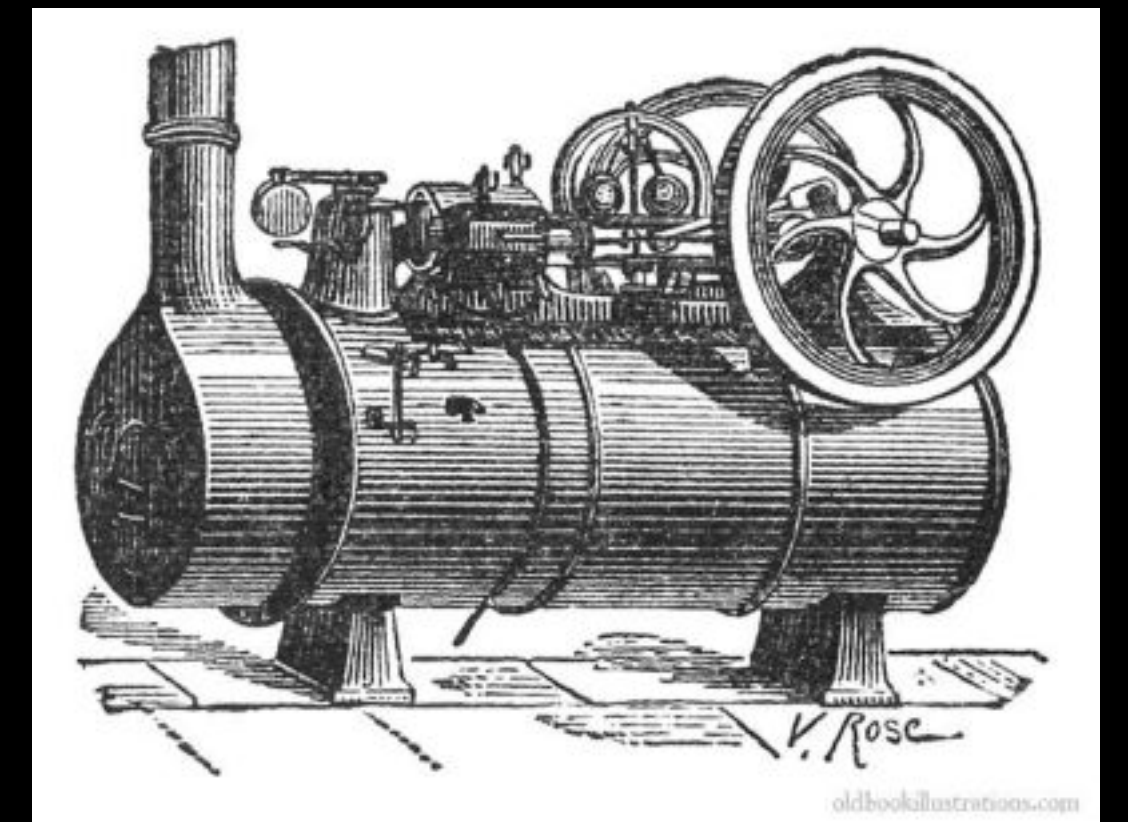in a watercolor style

https://openai.com/dall-e-2/

**DANGER**

# ML in Society

Change of Problem Formulation:

- 1980-2005: Make ML accurate enough for Society

- 2005-now:  Prevent ML misuse in Society

- New Technology that can **transform** society can also **overwhelm** it

- E.g. Steam engine late 1800s (lead to sweat shops, exploitation)

- E.g. Cars: out-of-control traffic accidents until 1970s

# Data privacy / misuse

Learning models leak training data (Fredrickson et al. '15)

Learning algorithms detect sexual orientation better than people

(Wang & Kosinski'17)

Leaked data

Real image

Formal definitions of data privacy:

- K- anonymity (Sweeney)

- Differential Privacy (Dwork, McSherry, Nissim, Smith).

Latanya Sweeney     Cynthia Dwork     Frank McSherry     Kobbi Nissim     Adam Smith
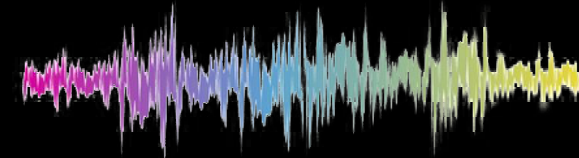
# Robust and Secure ML

Image Recognition
Misreading traffic signs
(Eykholt et al)

Speech recognition
Hide commands in noise (Carlini & Wagner)

Poisoning Attacks
Tay (chat bot) became inflammatory in 16 hr.



How to create robust and secure machine learning algorithms?

# Learning and the Society

- Bad dynamics, perpetuating and worsening stereotypes and biases.

- Who carries the burden of bad prediction?

- How to design good dynamics?

**The Best Algorithms Struggle to Recognize Black Faces Equally**

Google's algorithm shows prestigious job ads to men, but not to women. Here's why that should worry you.

Gender and racial bias found in Amazon's facial recognition technology (again)

Do Google's 'unprofessional hair' results show it is racist?

**How Amazon Accidentally Invented a Sexist Hiring Algorithm**

A company experiment to use artificial intelligence in hiring inadvertently favored male candidates.

unprofessional hairstyles for work
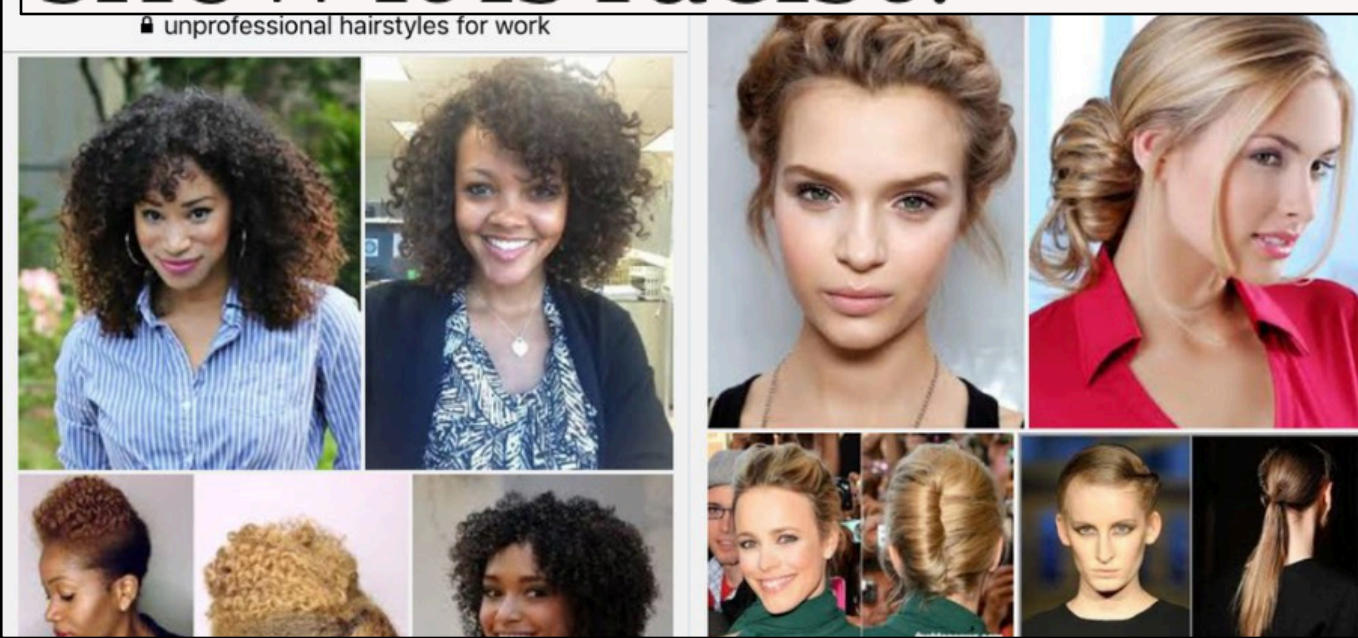
When an Algorithm Helps Send You to Prison

By Ellora Thadaney Israni

# AI and Online advertising

- Advertisers pay AI companies to induce change in people's behaviors

- Learn from clicks what people want

- People don't always want what they click on

  - Promotion of fake / sensational news

  - Leads to fast clicks/ prolonged engagement (more advertising time)

  - Causes (social) anxiety, fear, undesired behavior, elevates misinformation

A man can do what he wants, but not want what he wants.

Arthur Schopenhauer

*"Der Mensch kann tun was er will; er kann aber nicht wollen was er will."*

BrainyQuote®

AI
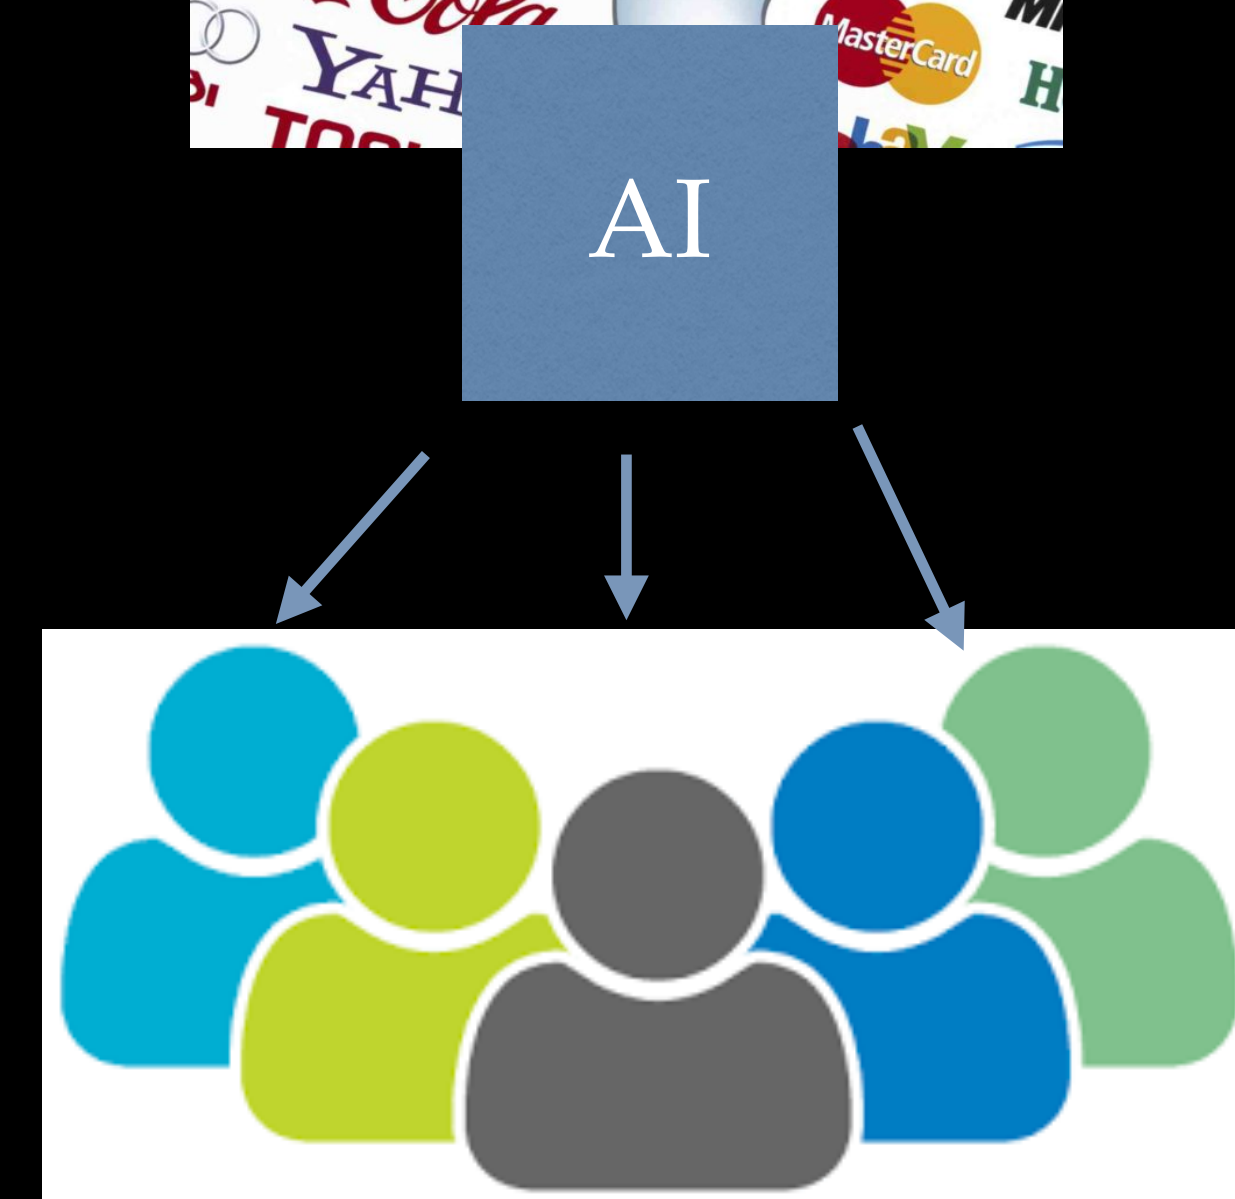
Now let's get crack'n

# To-do action items:

- Pass Placement Exam!

- Find study group

- Autograder Setup:

  - Find Teammate

  - Start Project -1

  - Start Project 0