Neural Networks and Deep Learning

Danna Gurari
University of Colorado Boulder
Fall 2022

https://home.cs.colorado.edu/~DrG/Courses/NeuralNetworksAndDeepLearning/AboutCourse.html

Today’s Topics

- Applications
- History of neural networks and deep learning
- How does a machine learn?
- Course logistics

Today’s Topics

- Applications
- History of neural networks and deep learning
- How does a machine learn?
- Course logistics

Key Motivation

Systems that support humans by either
improving upon existing human capabilities
or ***providing new capabilities***

Solutions – Spam Detection

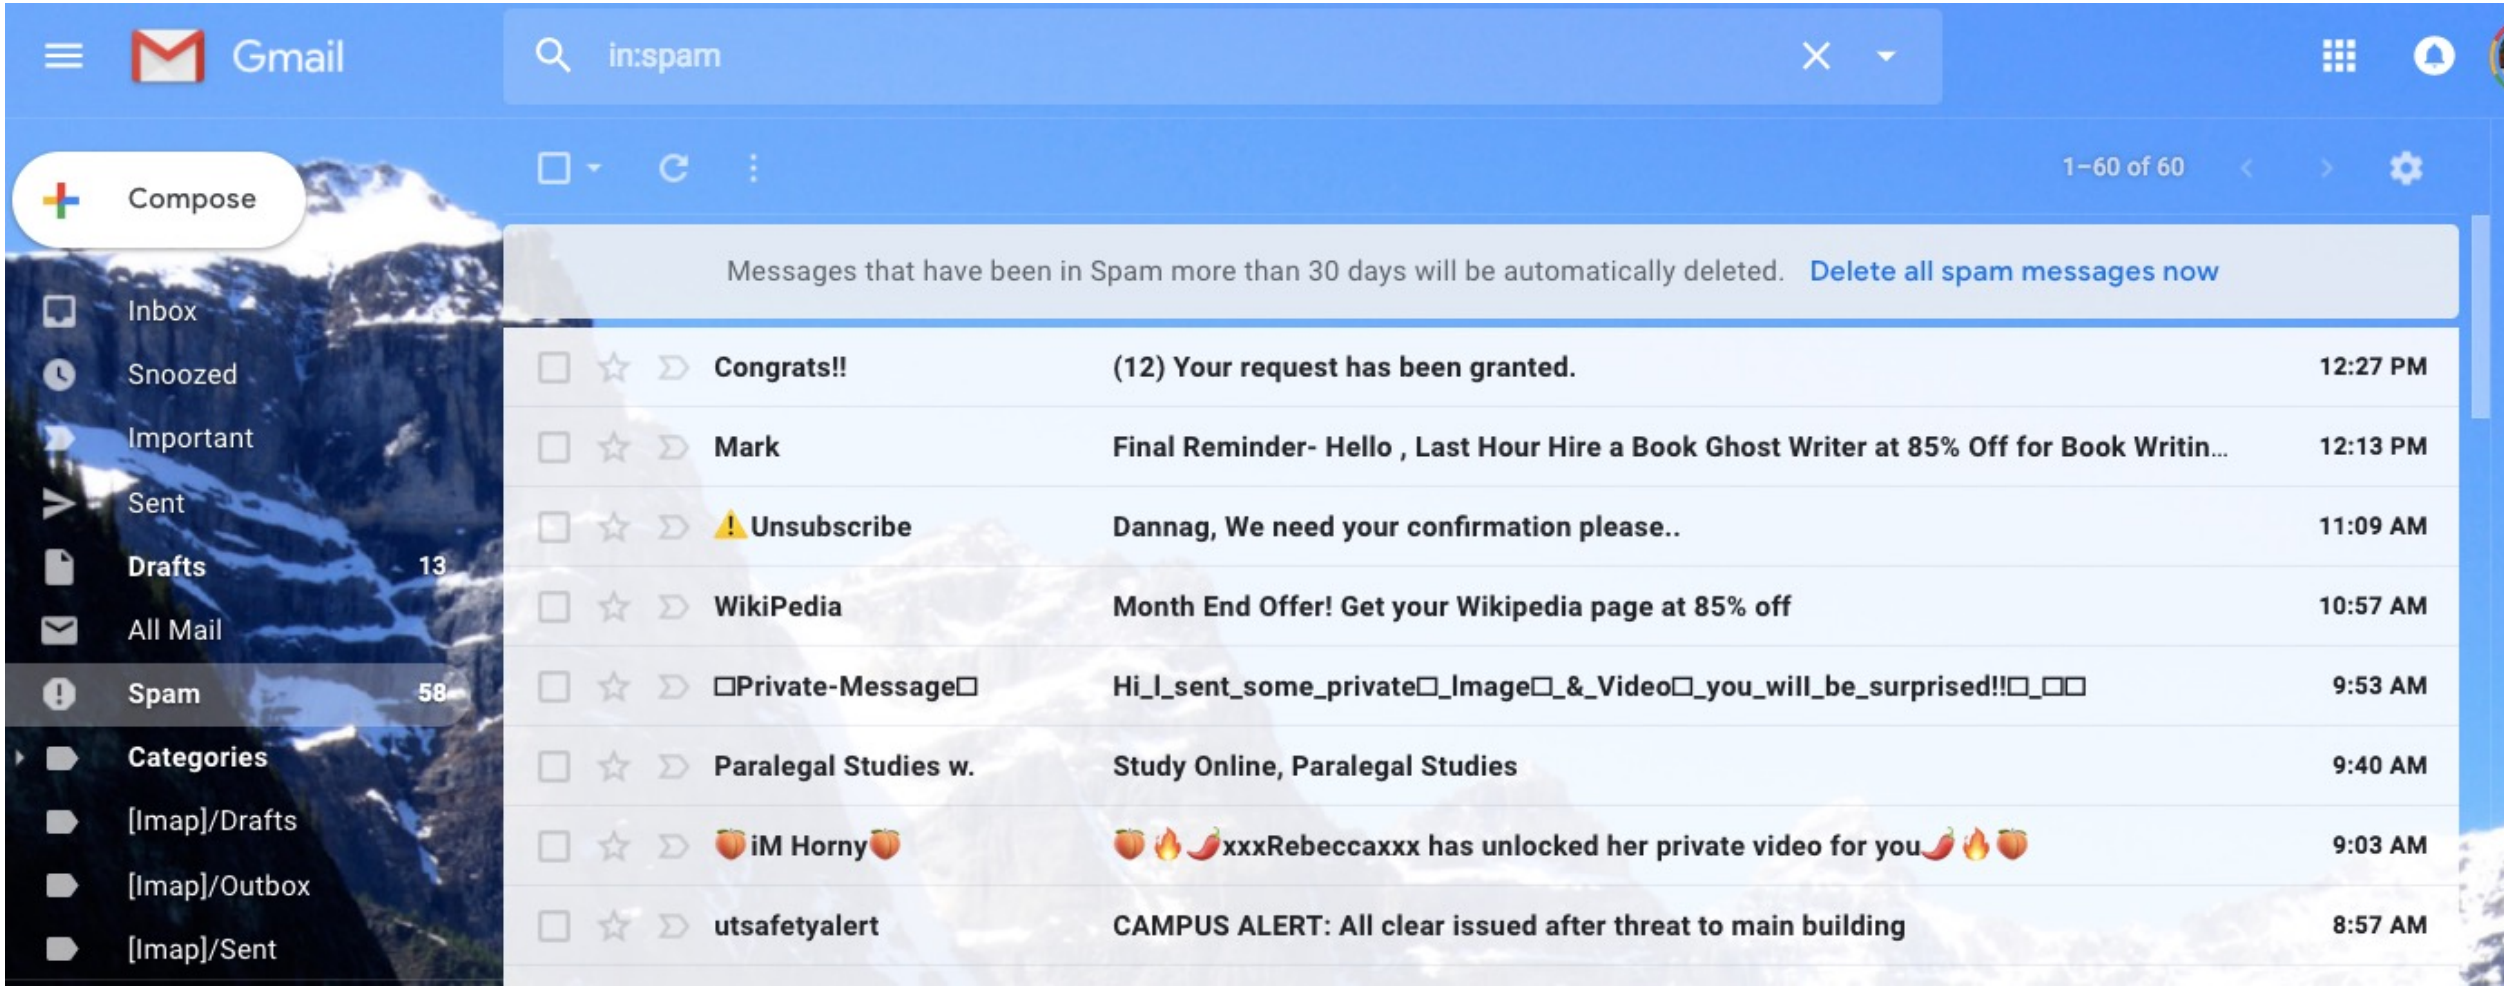

Solutions – Information Retrieval

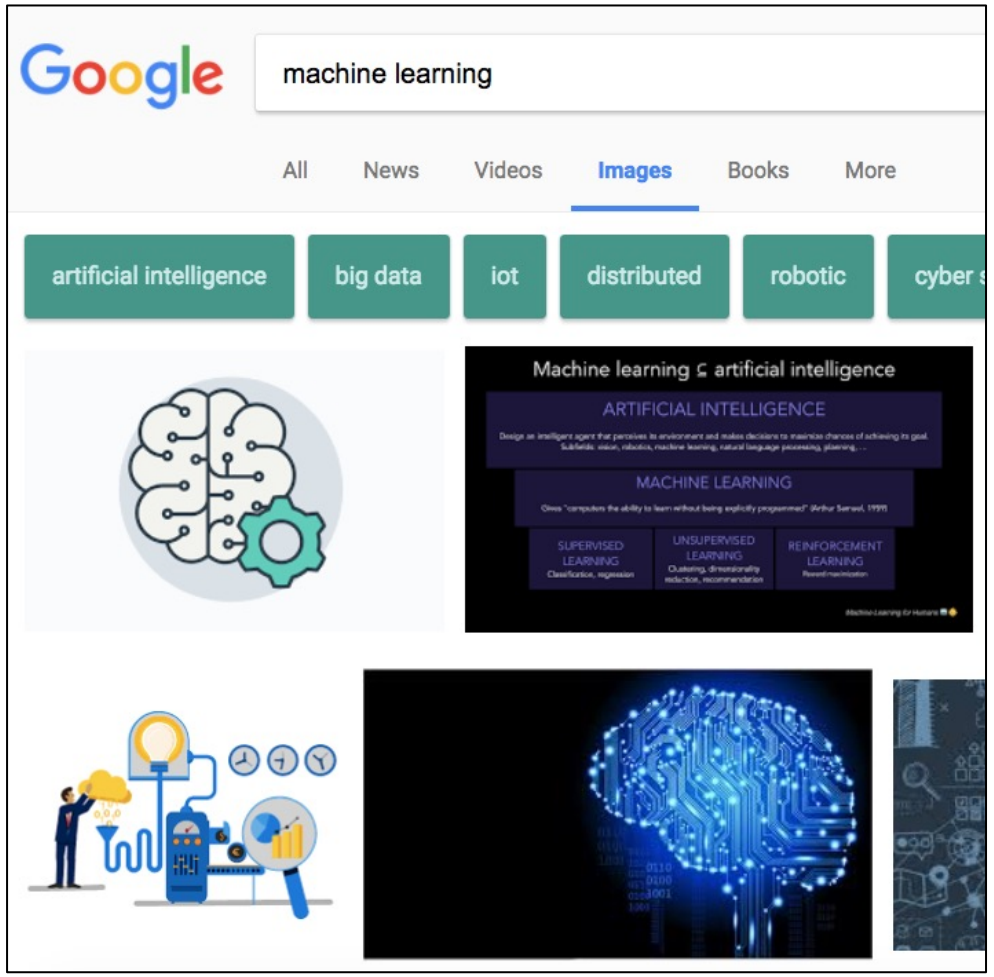

Solutions – Recognition

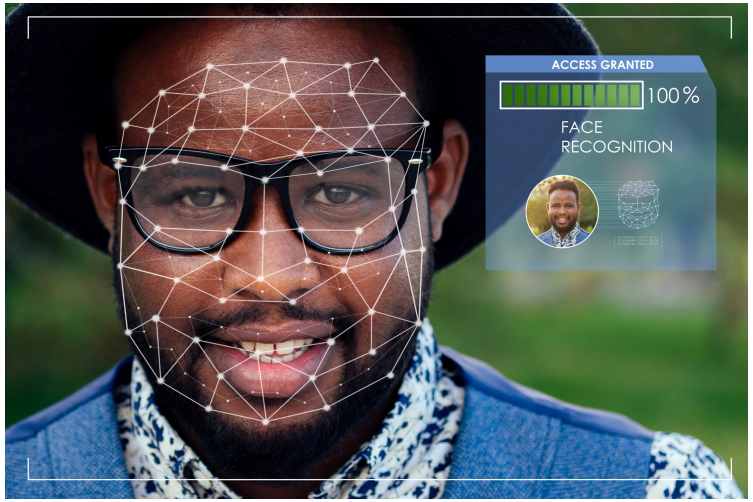

(Face)

(Speech)

(Fraud)

Solutions – Robotics

(Self-driving Vehicles)

(Medical Surgery)

(Manufacturing)

Solutions – Recommendation Systems

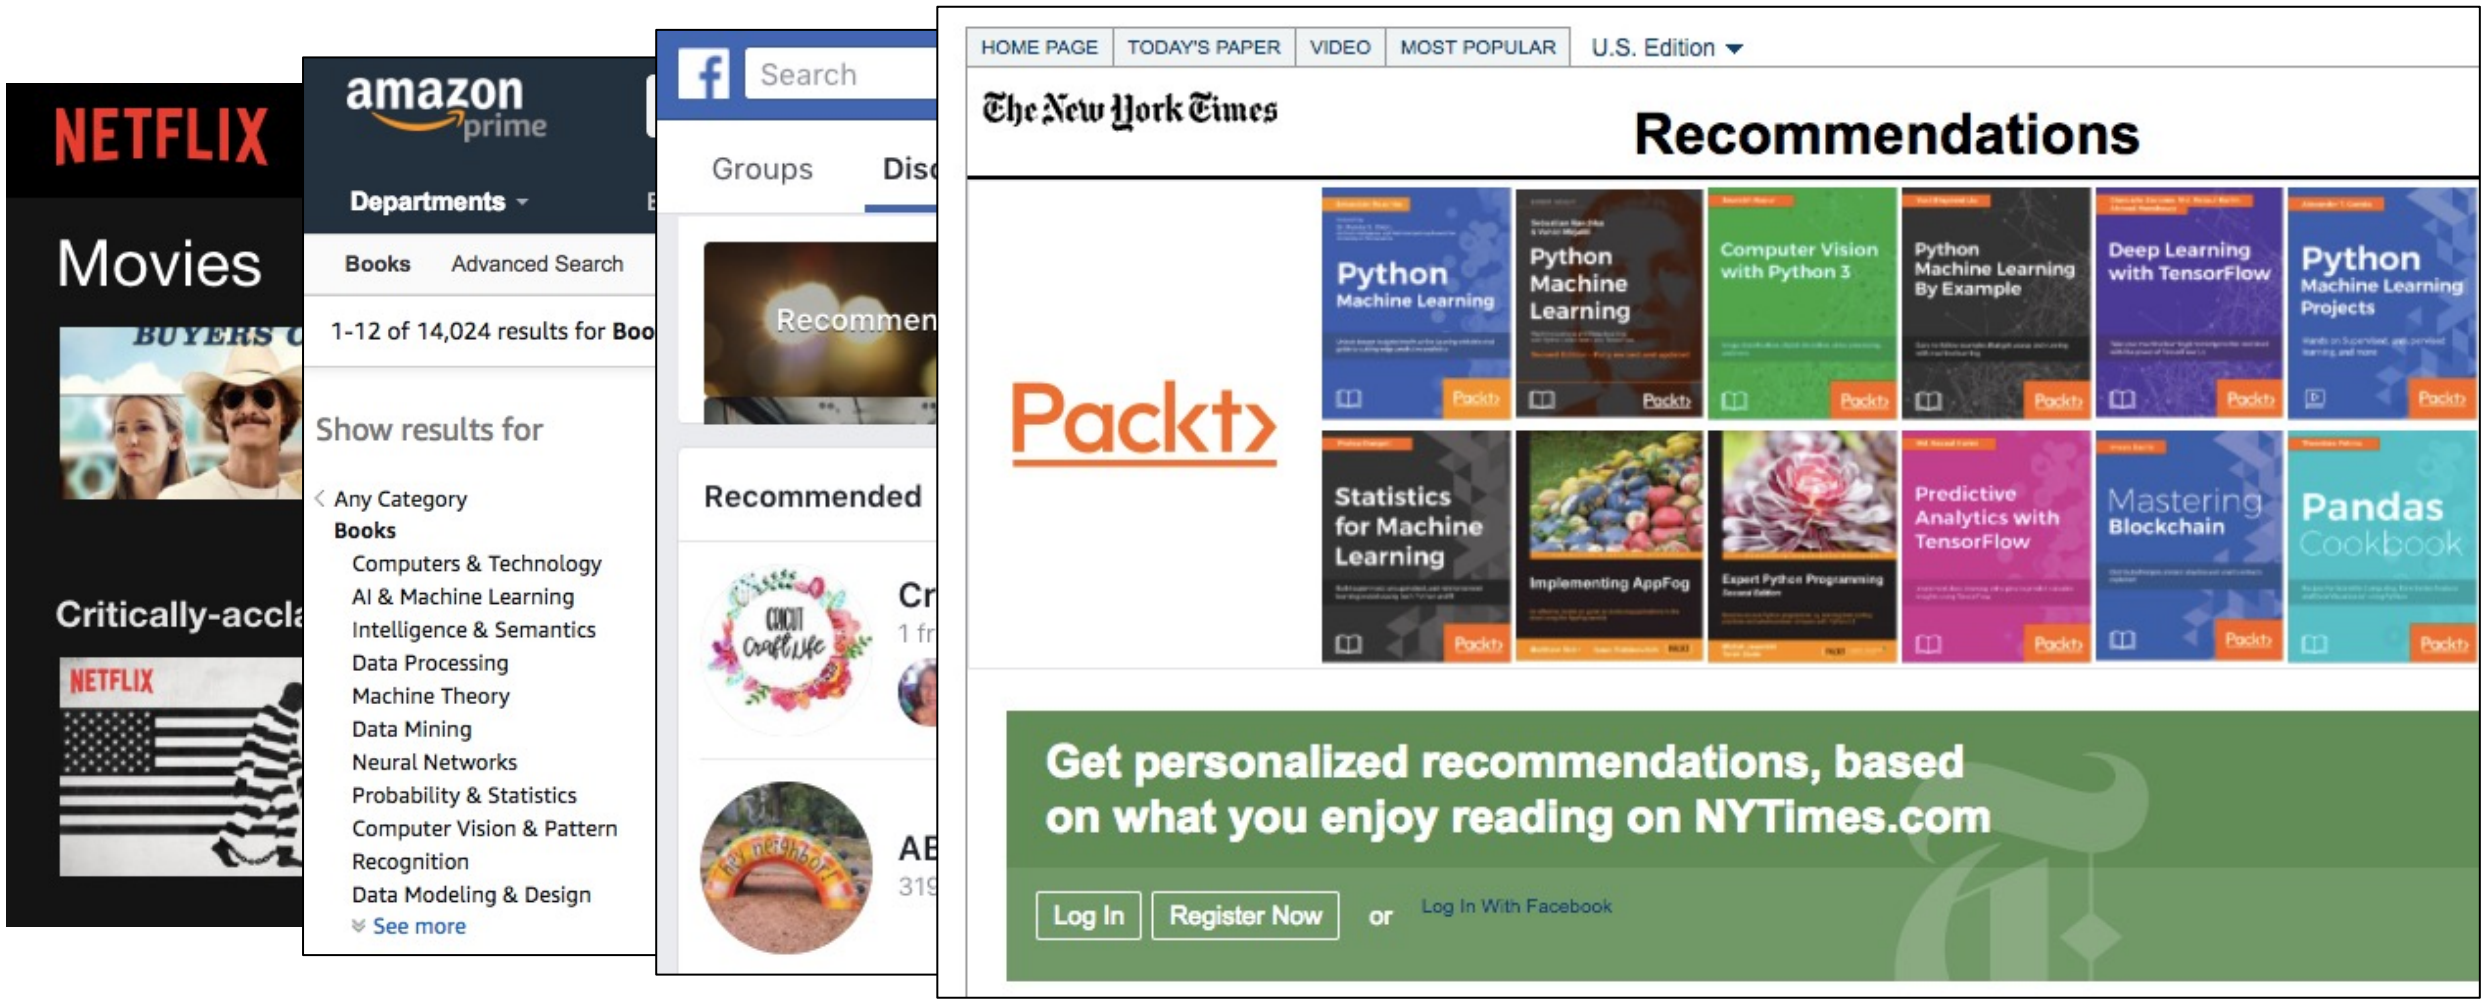

Solutions – Advertising

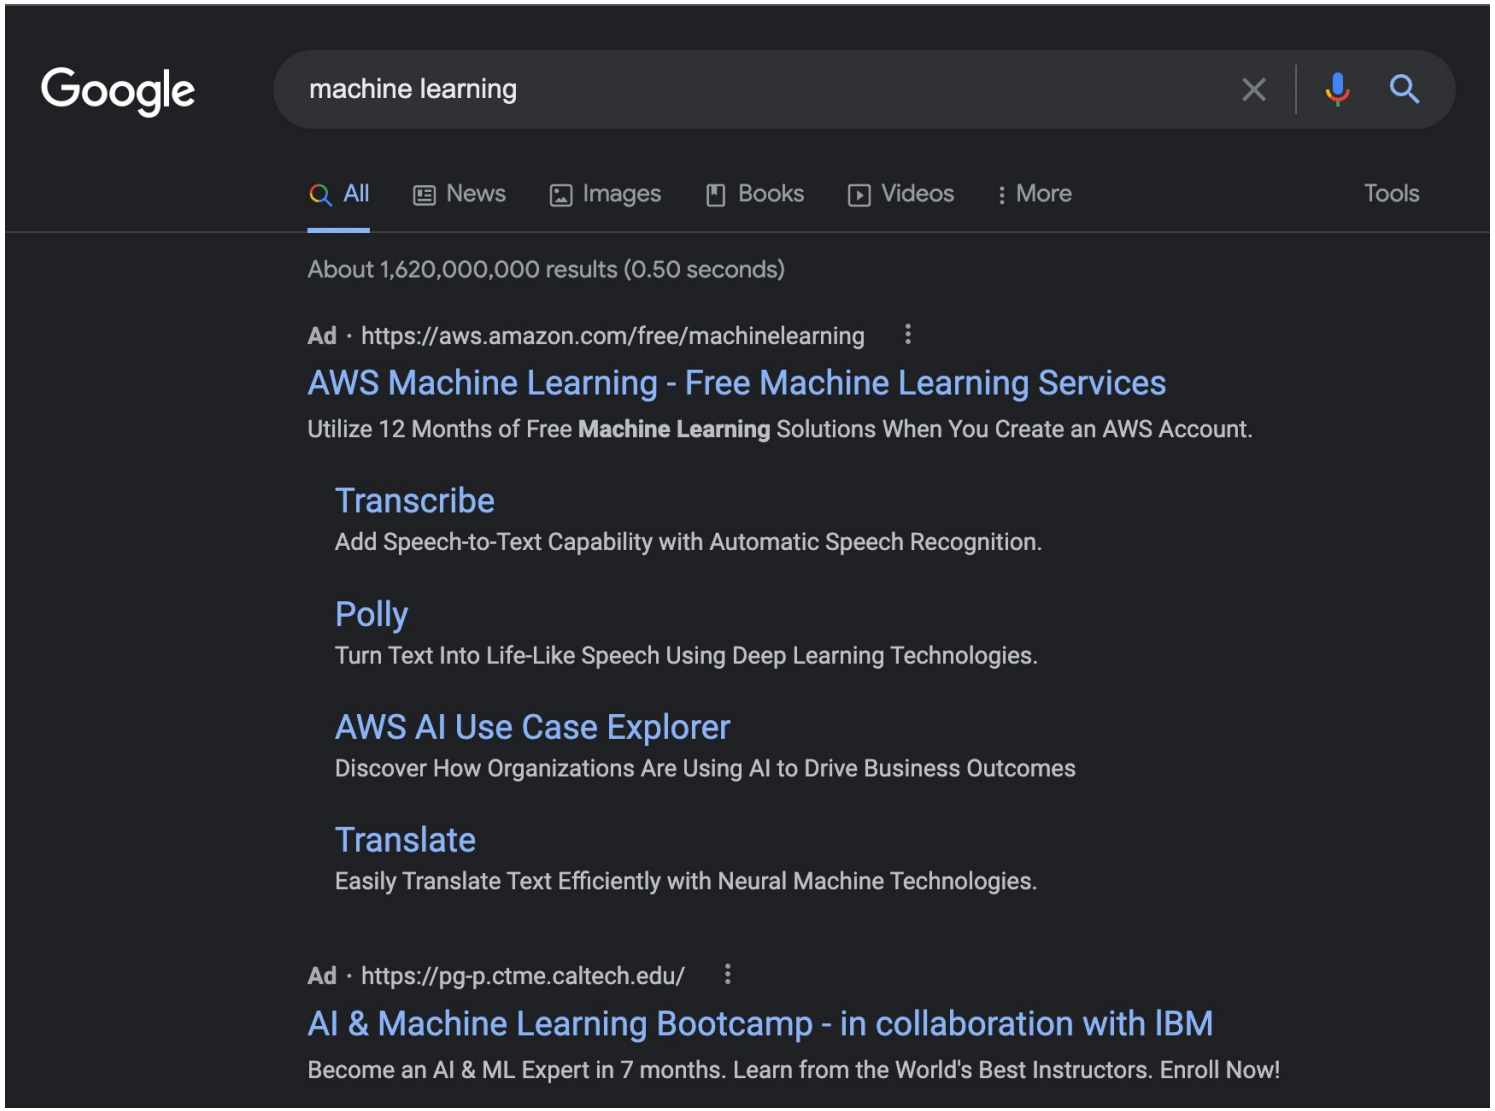

Solutions – Home Virtual Assistants

e.g., Amazon’s Echo with Alexa

e.g., Google Home

Today’s Topics

- Applications
- History of neural networks and deep learning
- How does a machine learn?
- Course logistics

Origins: Computers

1945

First programmable machine

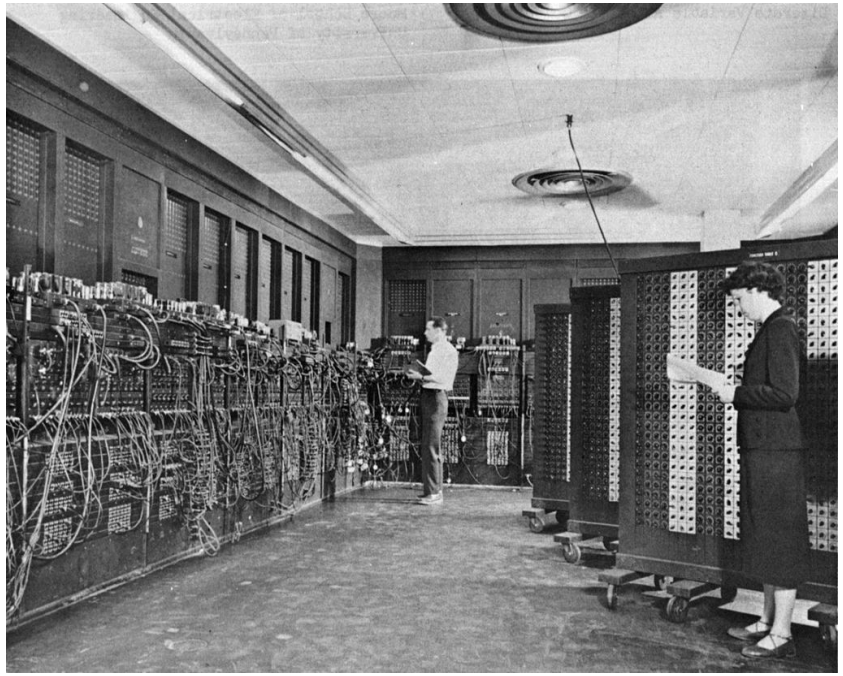

ENIAC (Electronic Numerical Integrator and Computer) created during World War II
(could compute 5,000 additions in one second)

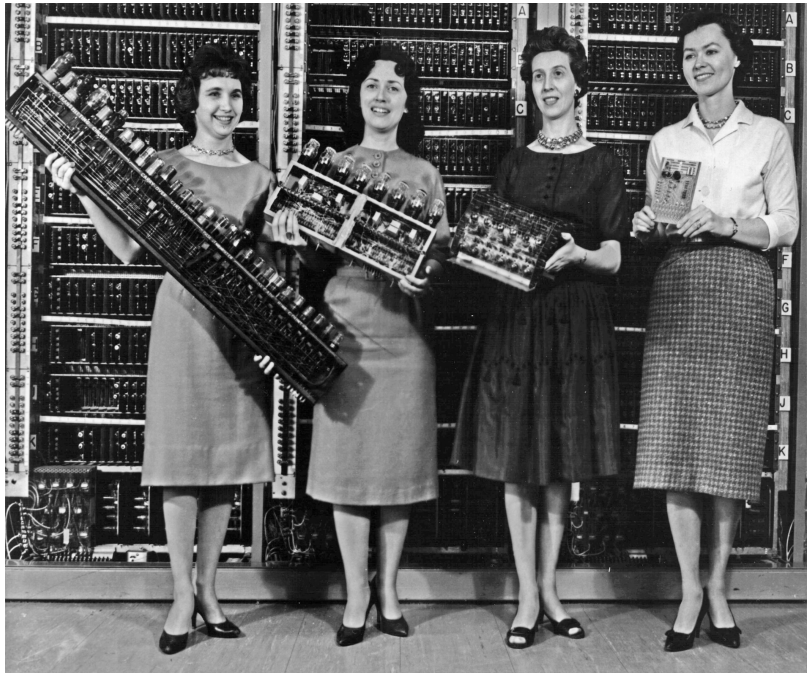

First programmers

Origins: Conceptual Framework

1945 1950

First programmable machine

Turing test

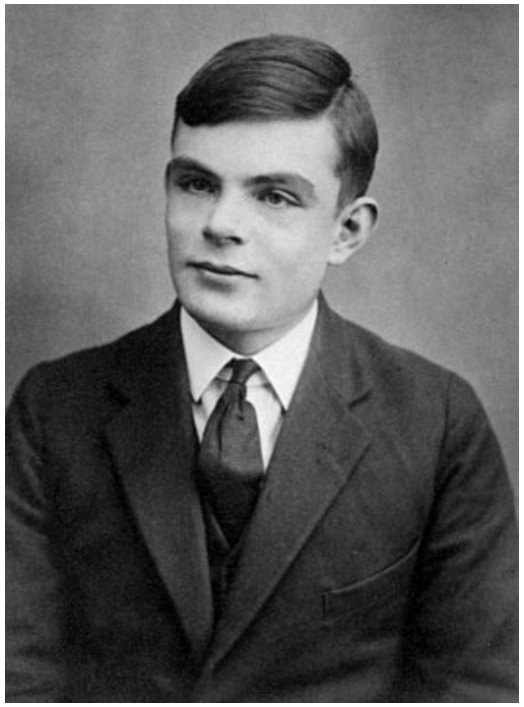

Alan Turing
(1912-1954)

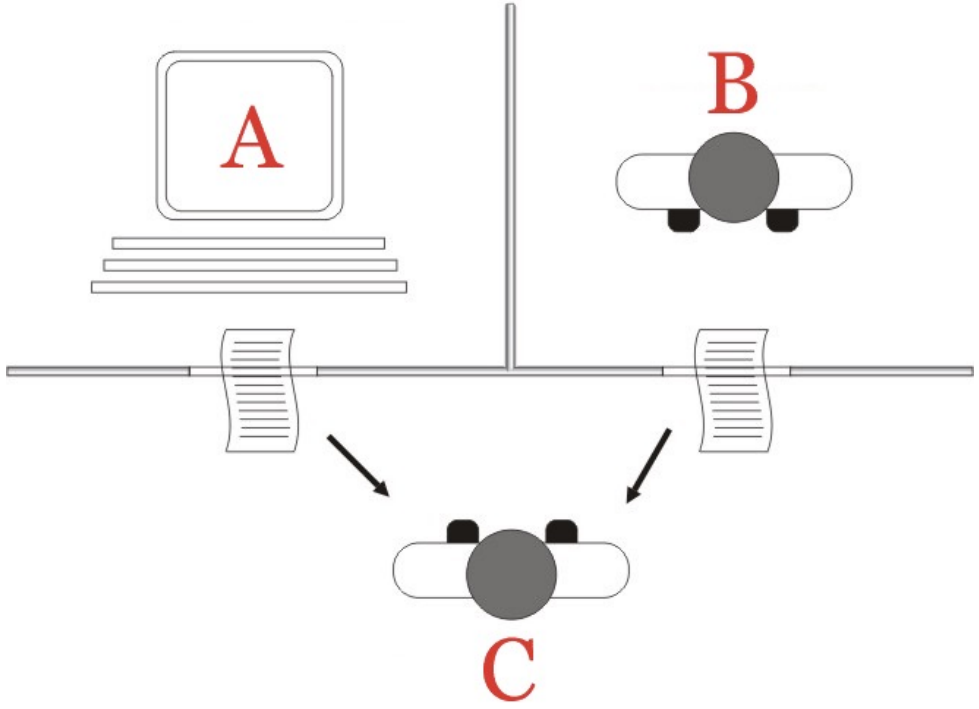

Turing Test: can ”C” decide whether text responses come from a machine or human

Origins: Conceptual Framework

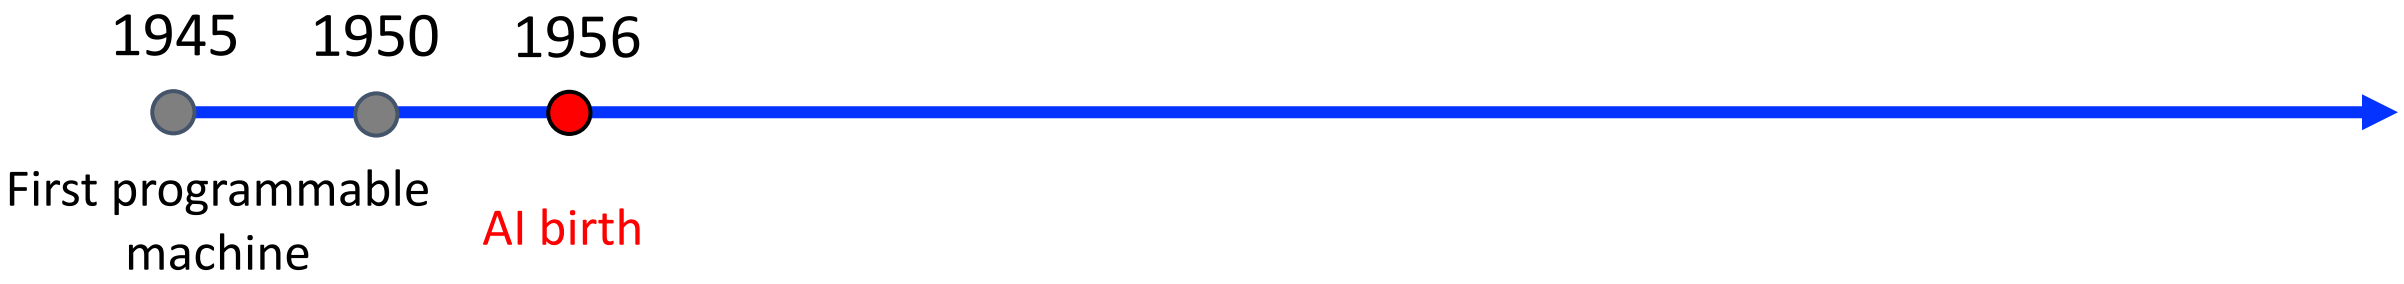

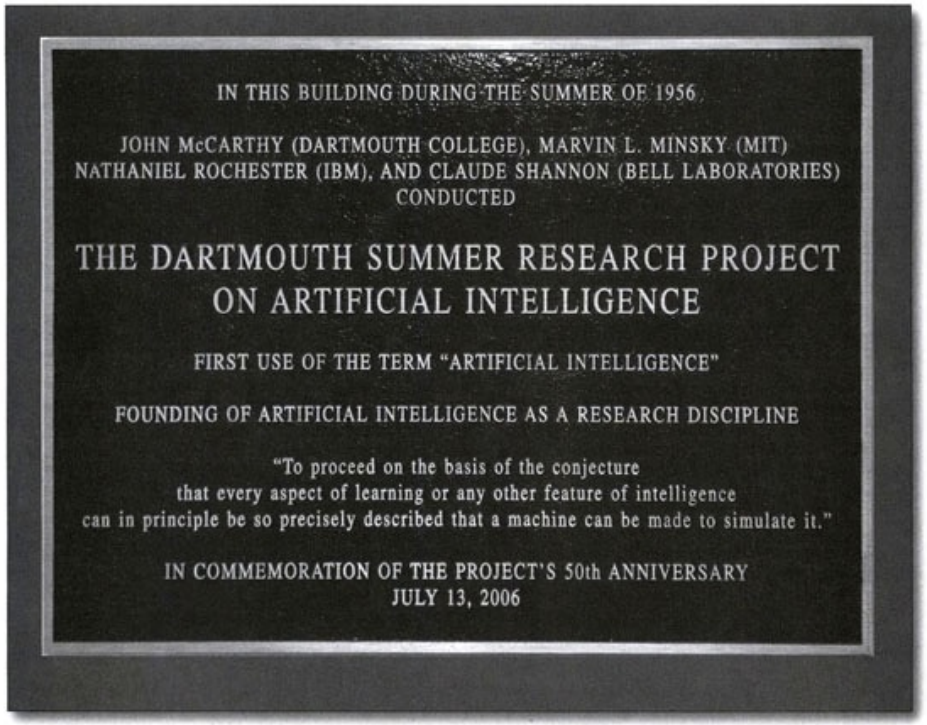

"Artificial intelligence" established as a field at a workshop

Origins: Conceptual Framework

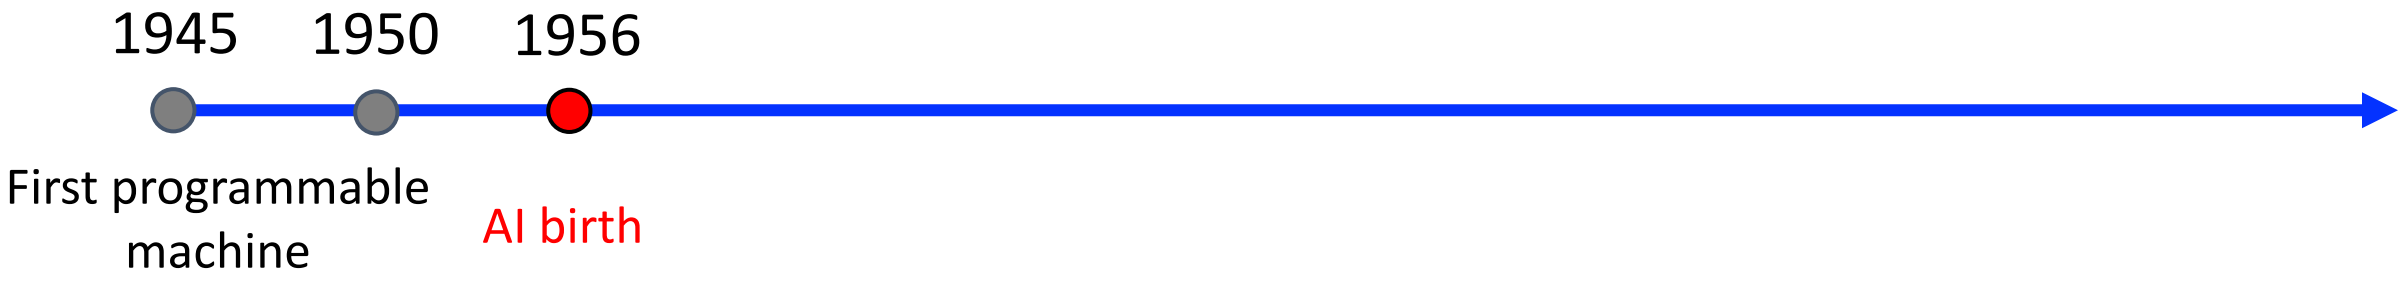

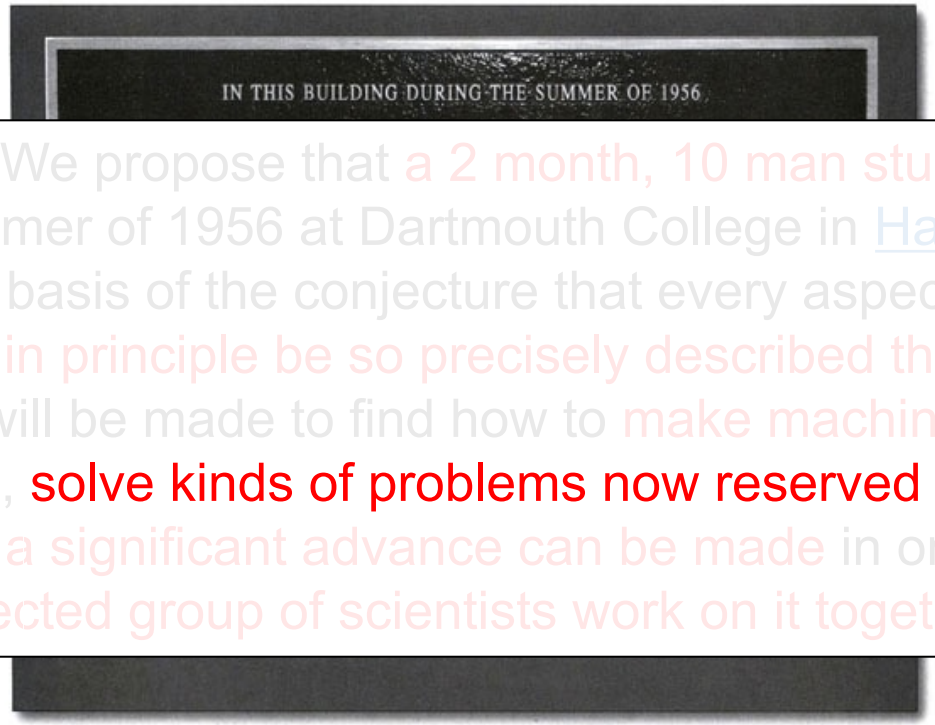

Workshop Proposal: "… We propose that a 2 month, 10 man study of artificial intelligence be carried out during the summer of 1956 at Dartmouth College in Hanover, New Hampshire. The study is to proceed on the basis of the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it. An attempt will be made to find how to make machines use language, form abstractions and concepts, solve kinds of problems now reserved for humans, and improve themselves. We think that a significant advance can be made in one or more of these problems if a carefully selected group of scientists work on it together for a summer…"

"Artificial intelligence" established as a field at a workshop

Origins: Conceptual Framework

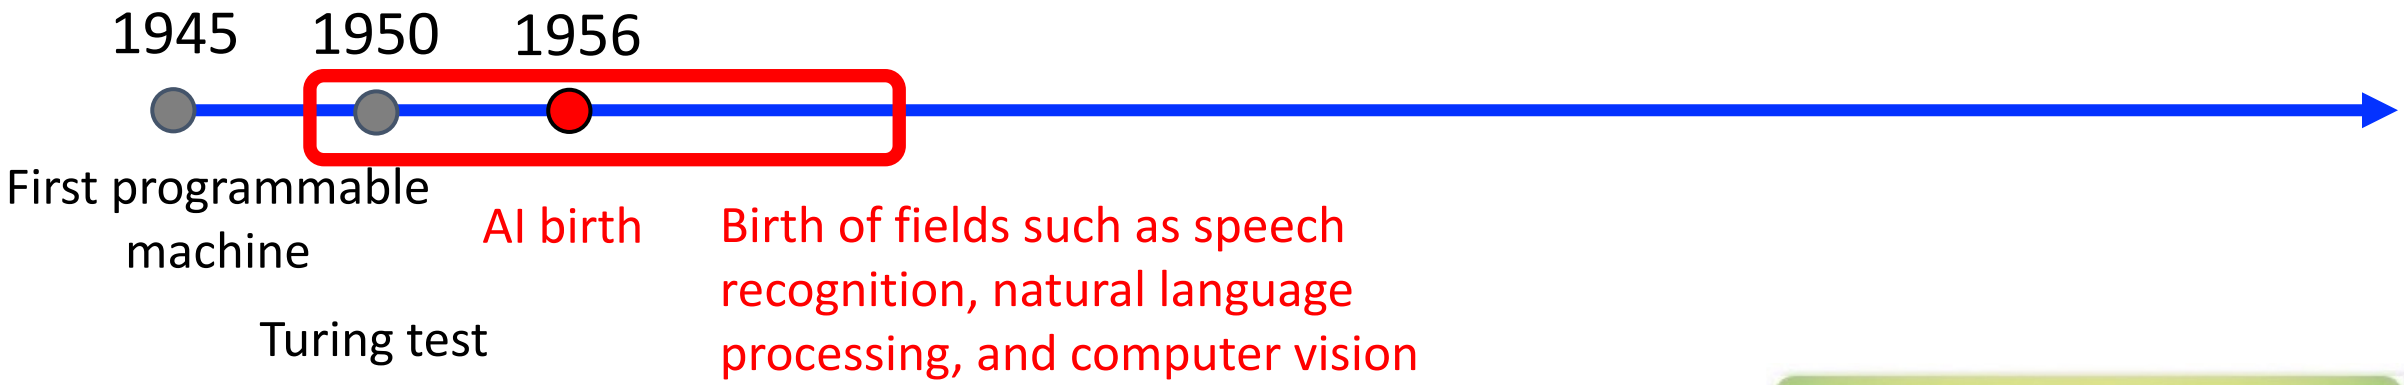

What human intelligence might computers imitate?

Origins: Conceptual Framework

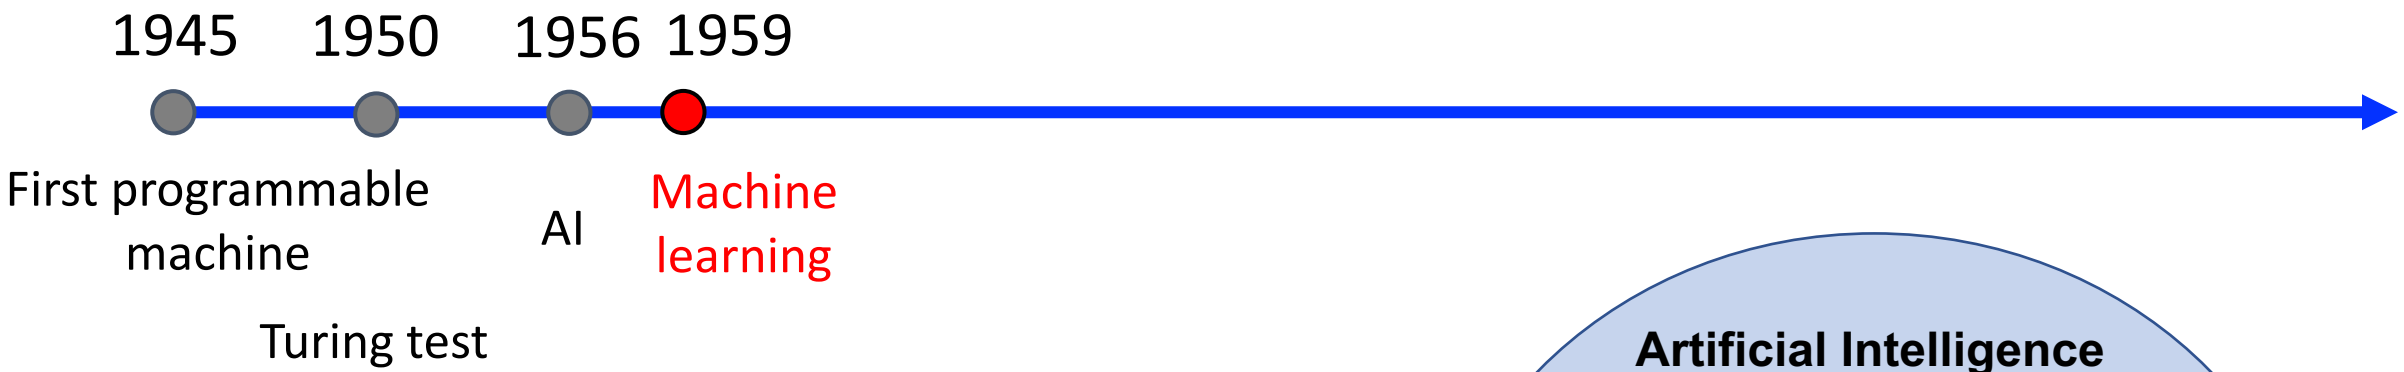

AI researcher Arthur Samuel coins the term "machine learning" as:

"Field of study that gives computers the ability to learn without being explicitly programmed."

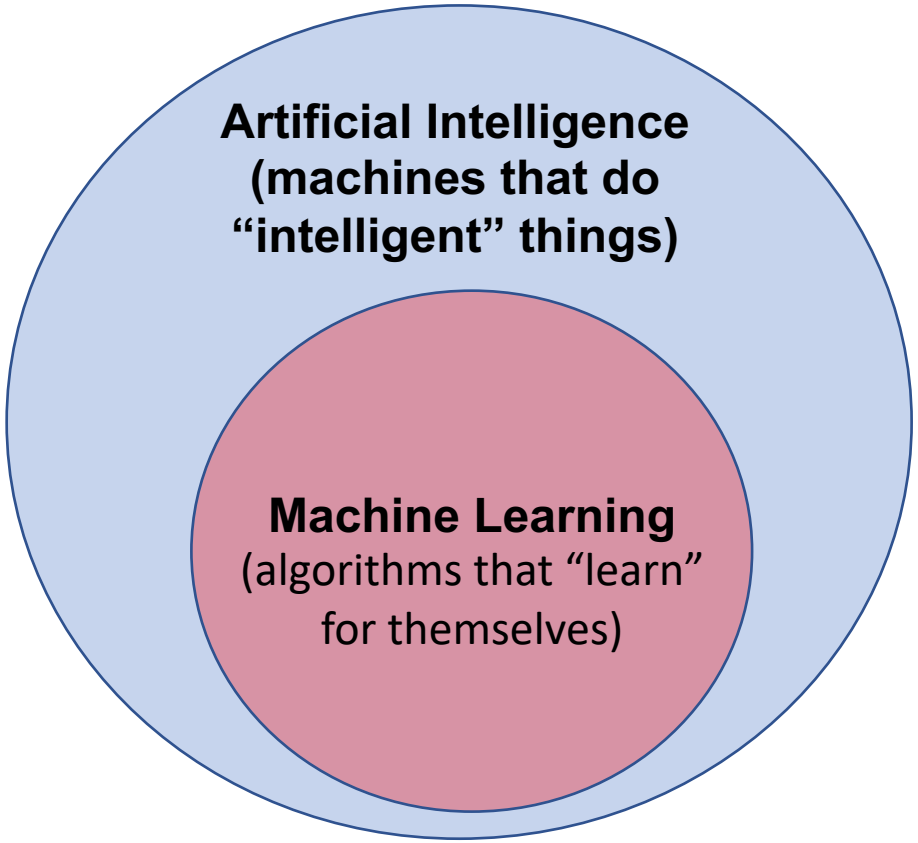

Motivation for Machines that “Learn”

- Process for hand-crafted rules:

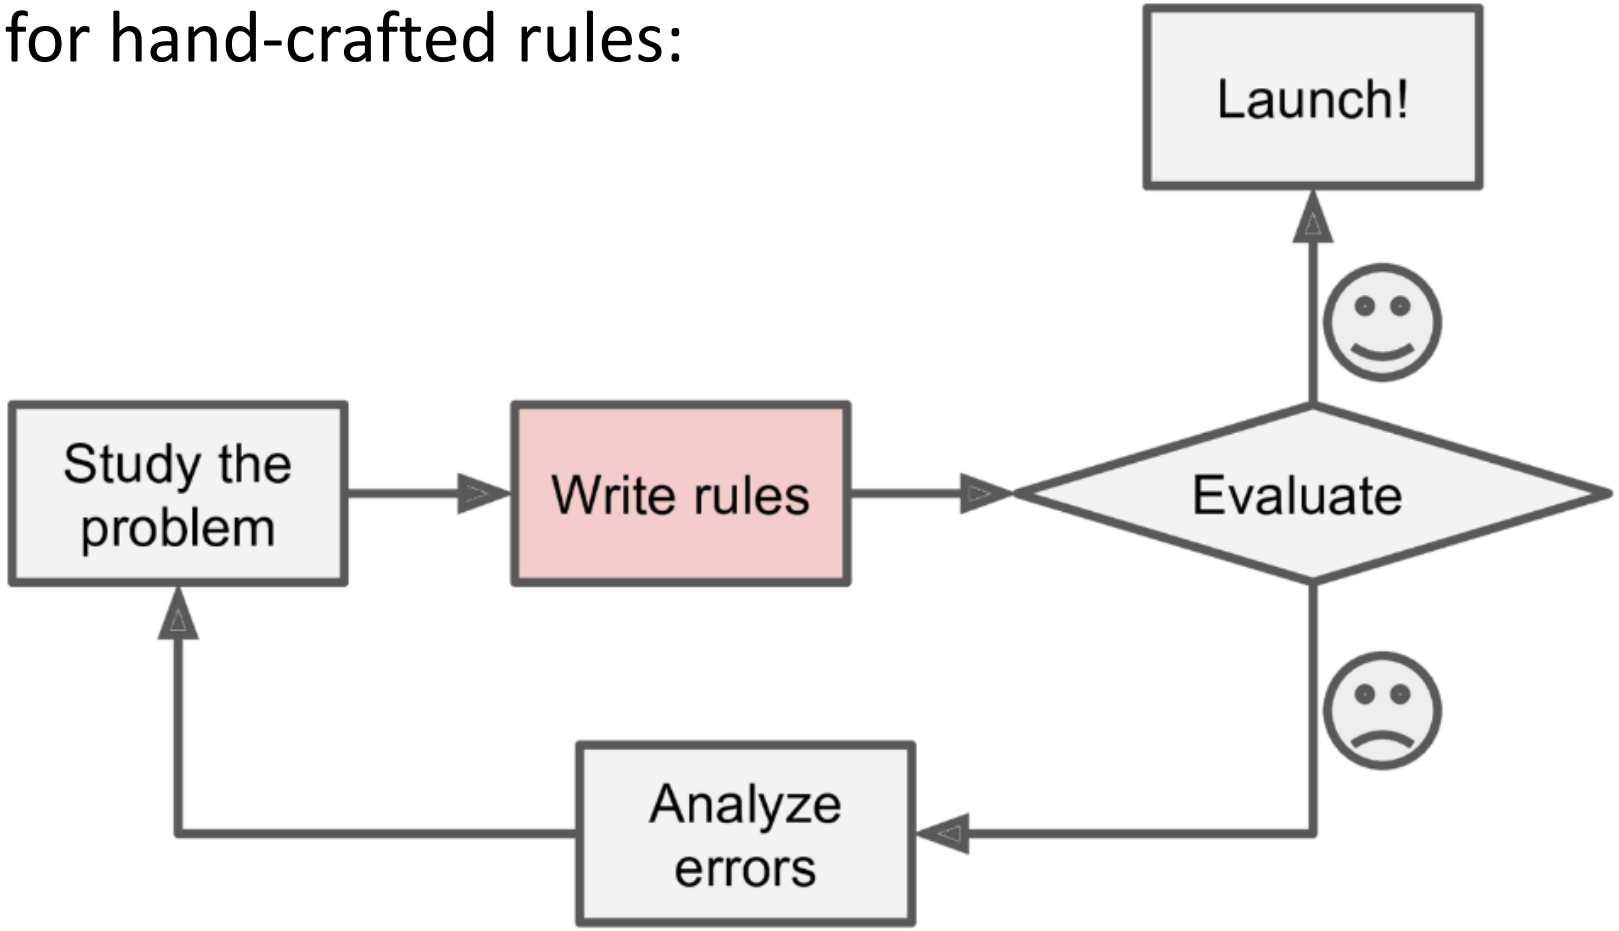

https://www.oreilly.com/library/view/hands-on-machine-learning/9781491962282/ch01.html

Motivation for Machines that “Learn”: Class Task

e.g., What rules would you use to answer: “Is a person in the image?”

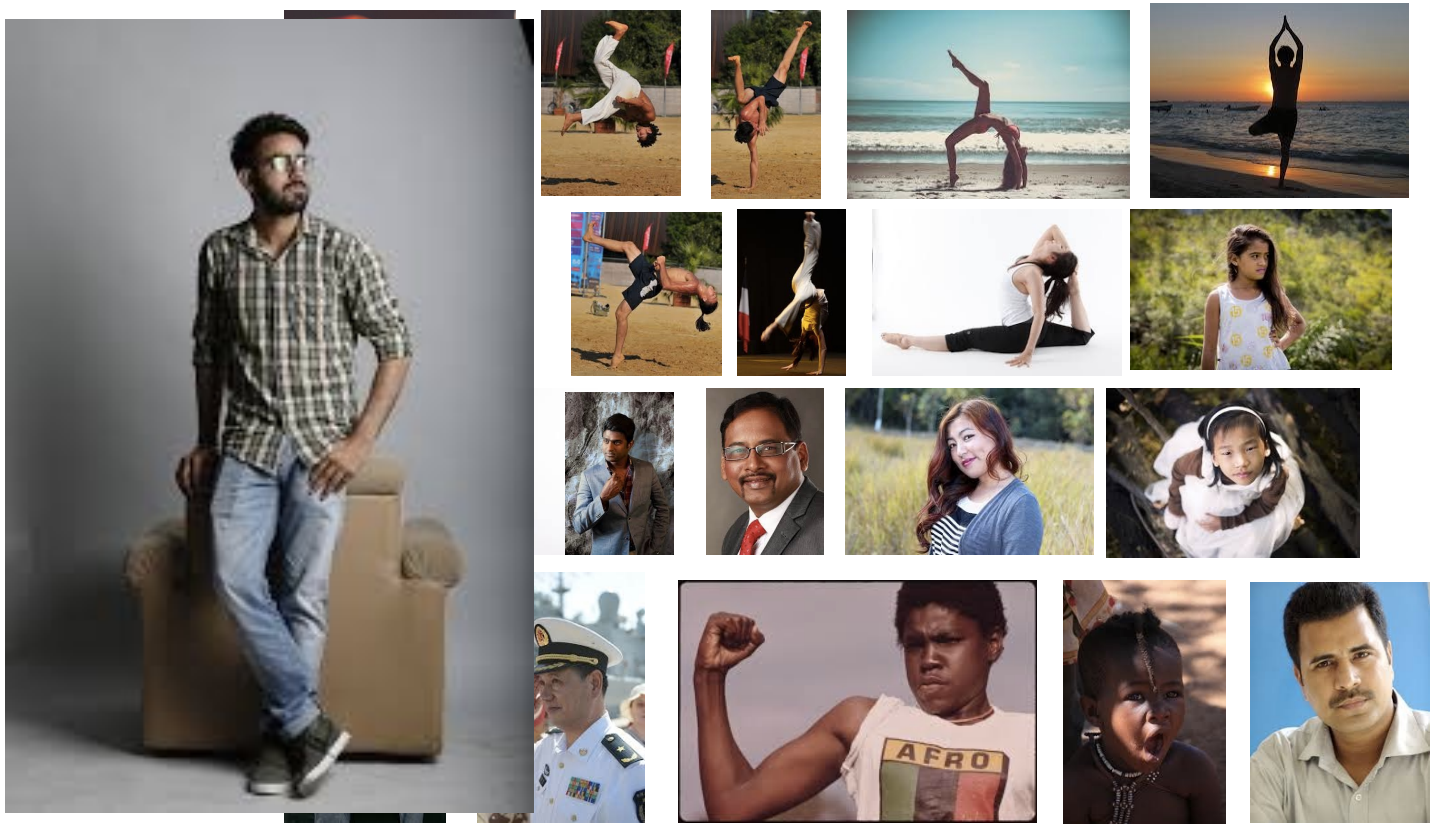

Motivation for Machines that “Learn”

e.g., are these lines parallel?

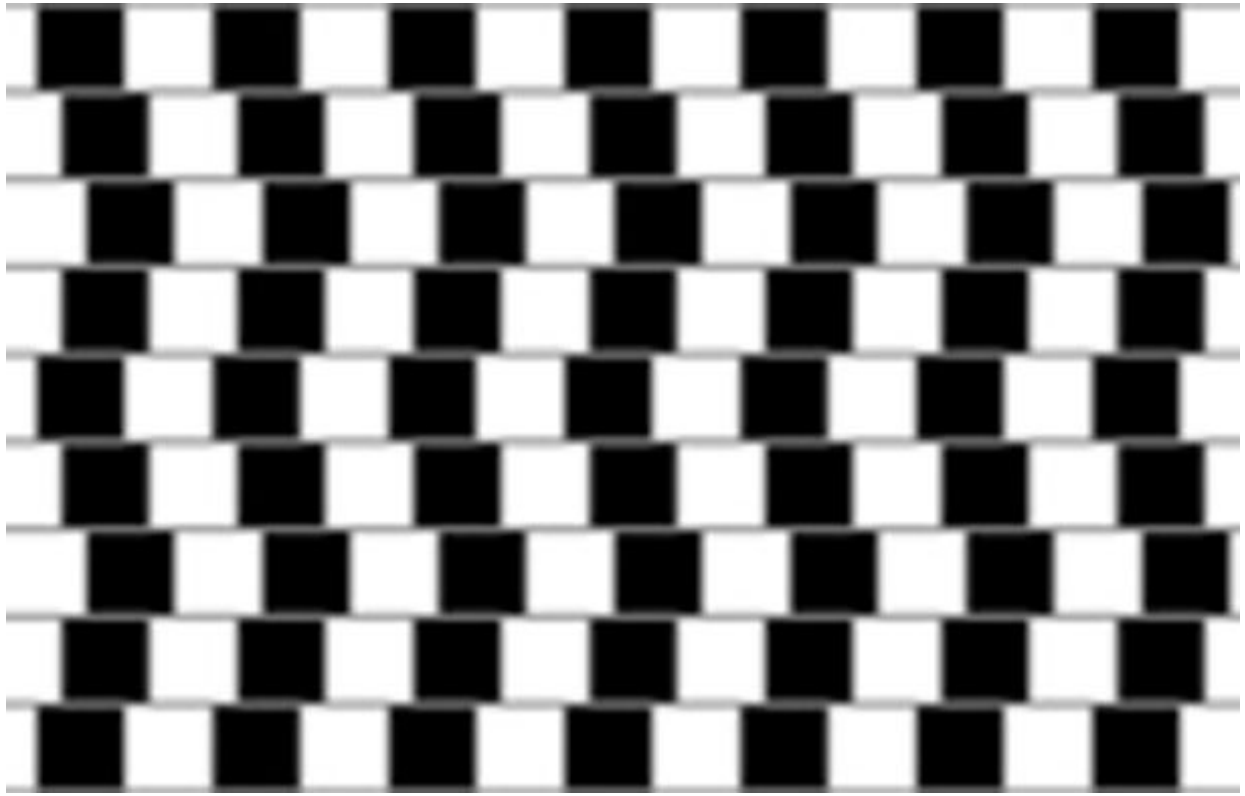

Motivation for Machines that “Learn”

e.g., are these lines parallel?

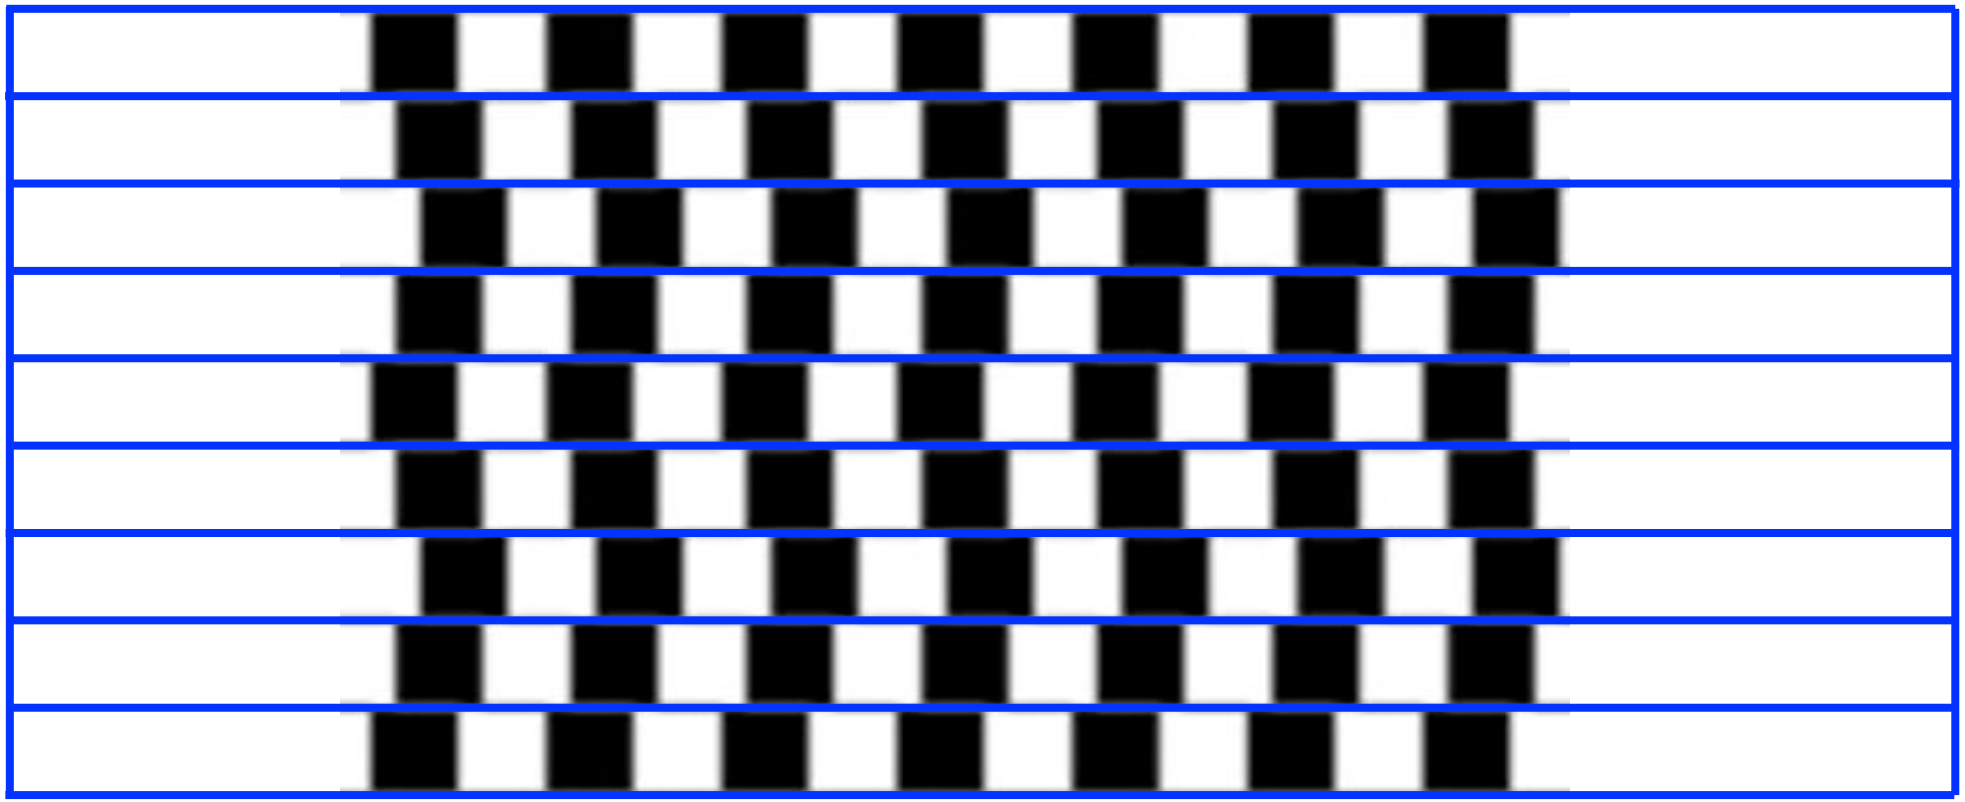

Motivation for Machines that “Learn”

1. It is hard to hand-craft a complete set of rules

2. We, as humans, may not devise the best rules for a machine since our brains (unconsciously) pre-process the data we sense

Origins: Neural Networks with Deep Learning

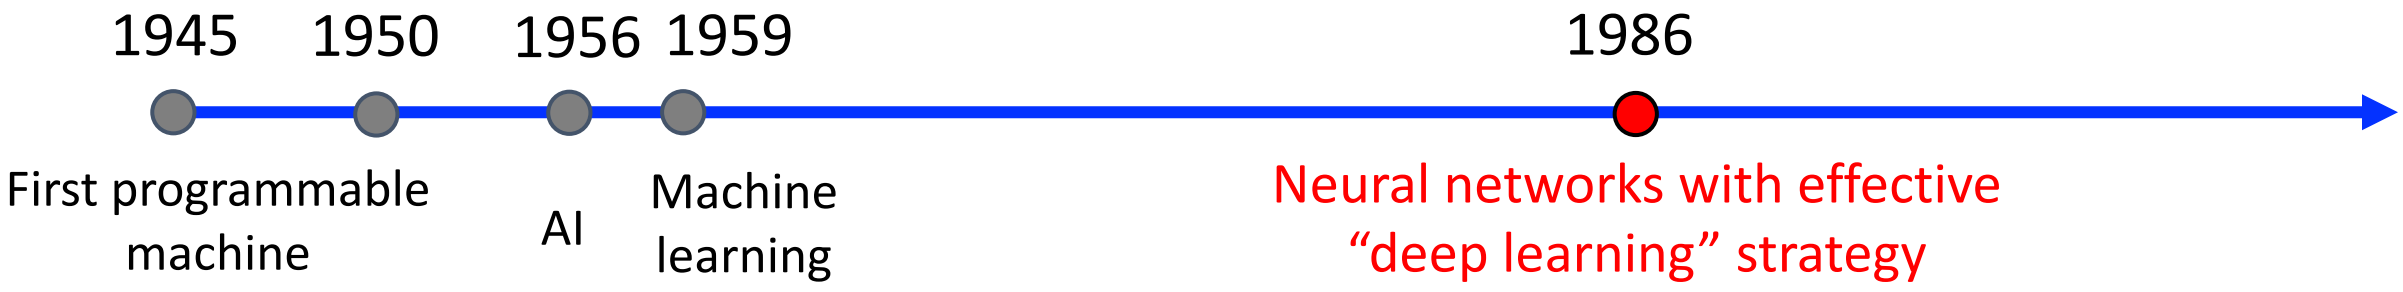

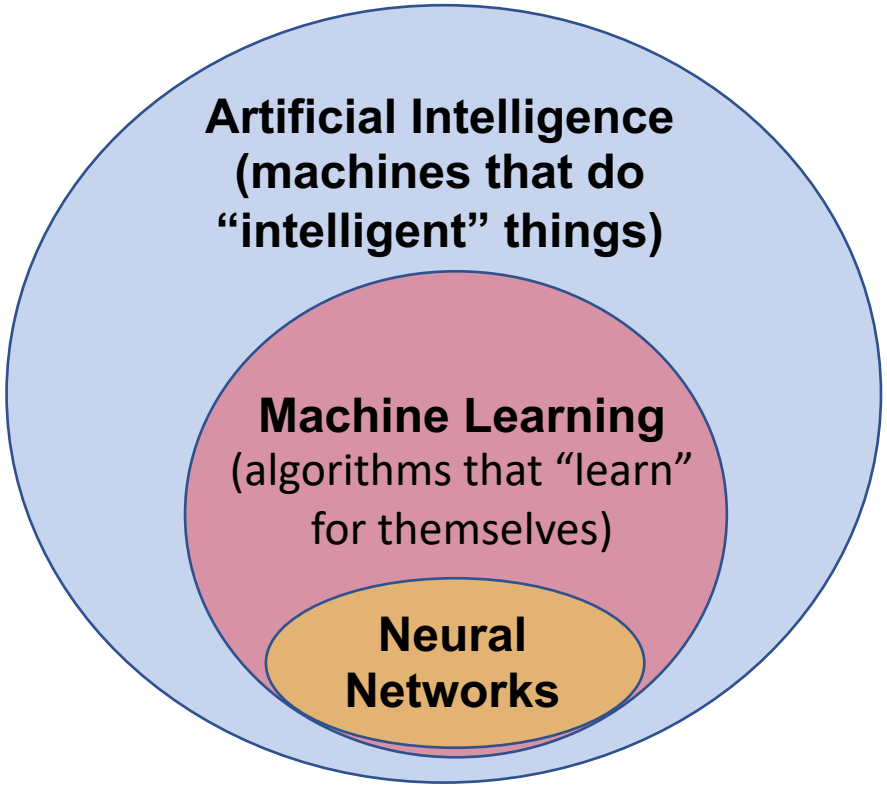

Motivation for Neural Networks (NNs) Over Other Machine Learning (ML) Approaches

INPUT

e.g., Is a person present?

OUTPUT

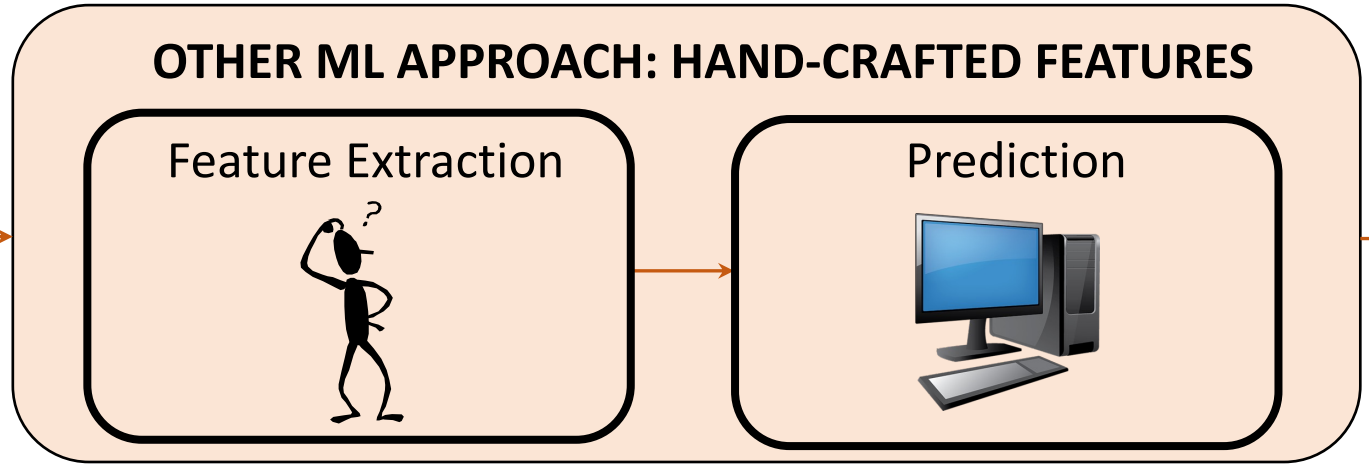

What features would help predict yes/no?

e.g., corners, lines, and model of expected body parts as connected shapes

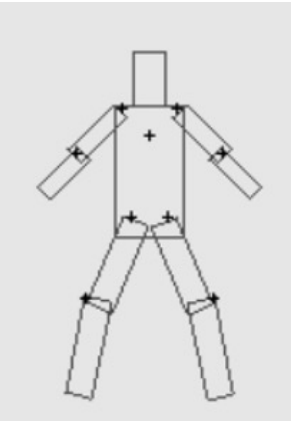

e.g., Pedro F Felzenszwalb and Daniel P Huttenlocher, IJCV 2004

Motivation for Neural Networks (NNs) Over Other Machine Learning (ML) Approaches

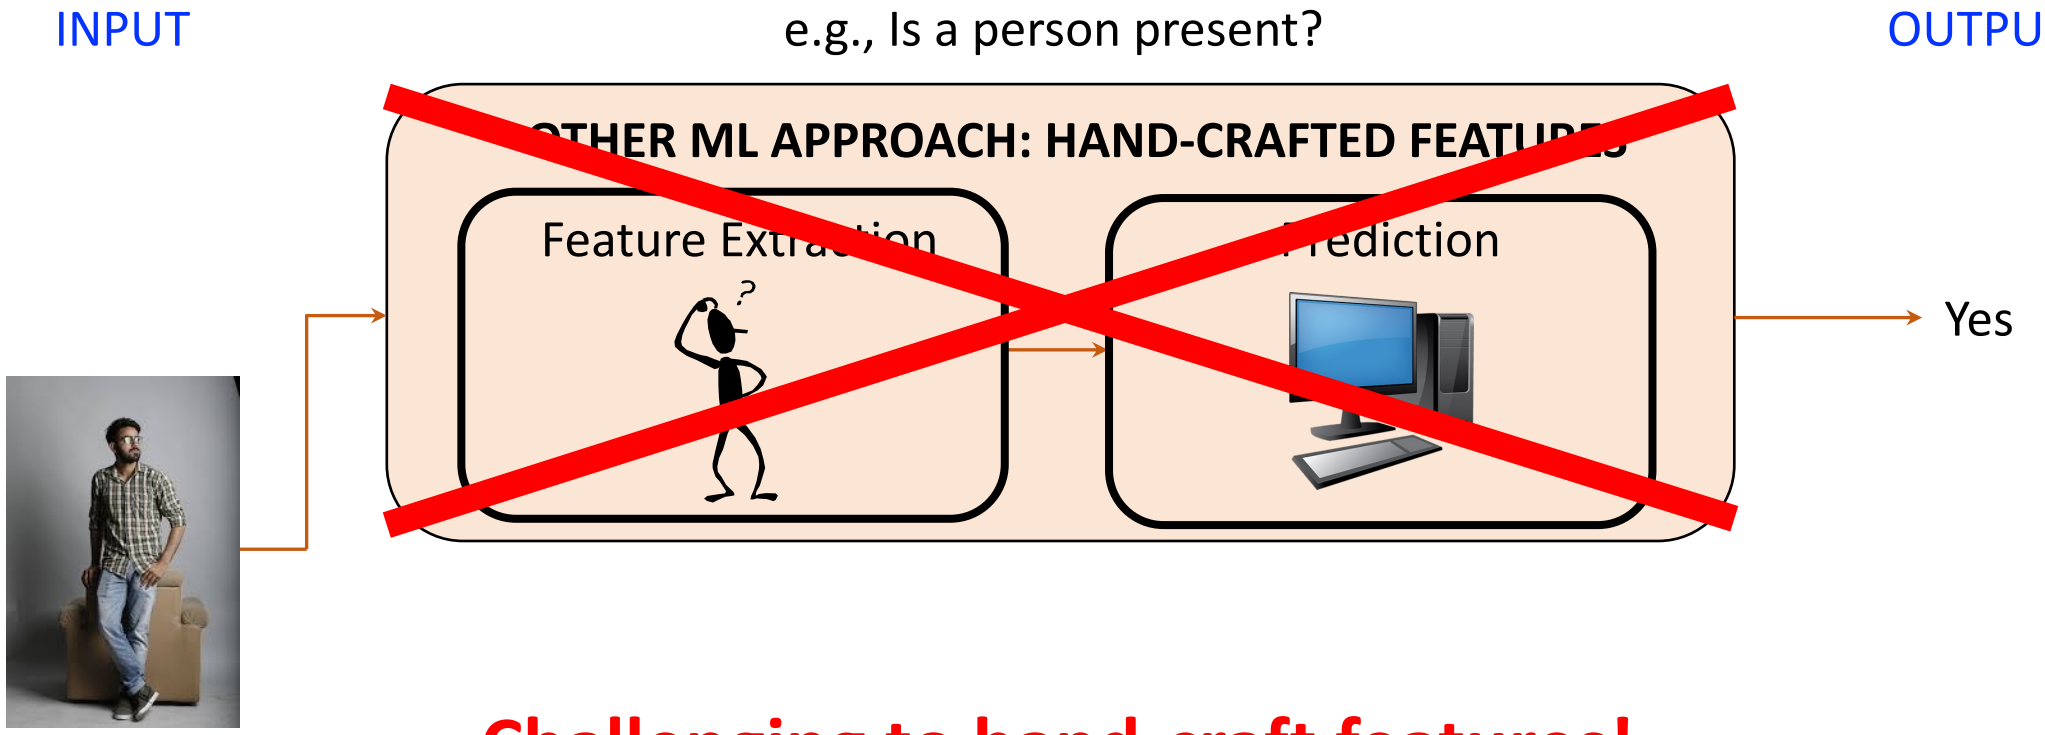

Challenging to hand-craft features!

Motivation for Neural Networks (NNs) Over Other Machine Learning (ML) Approaches

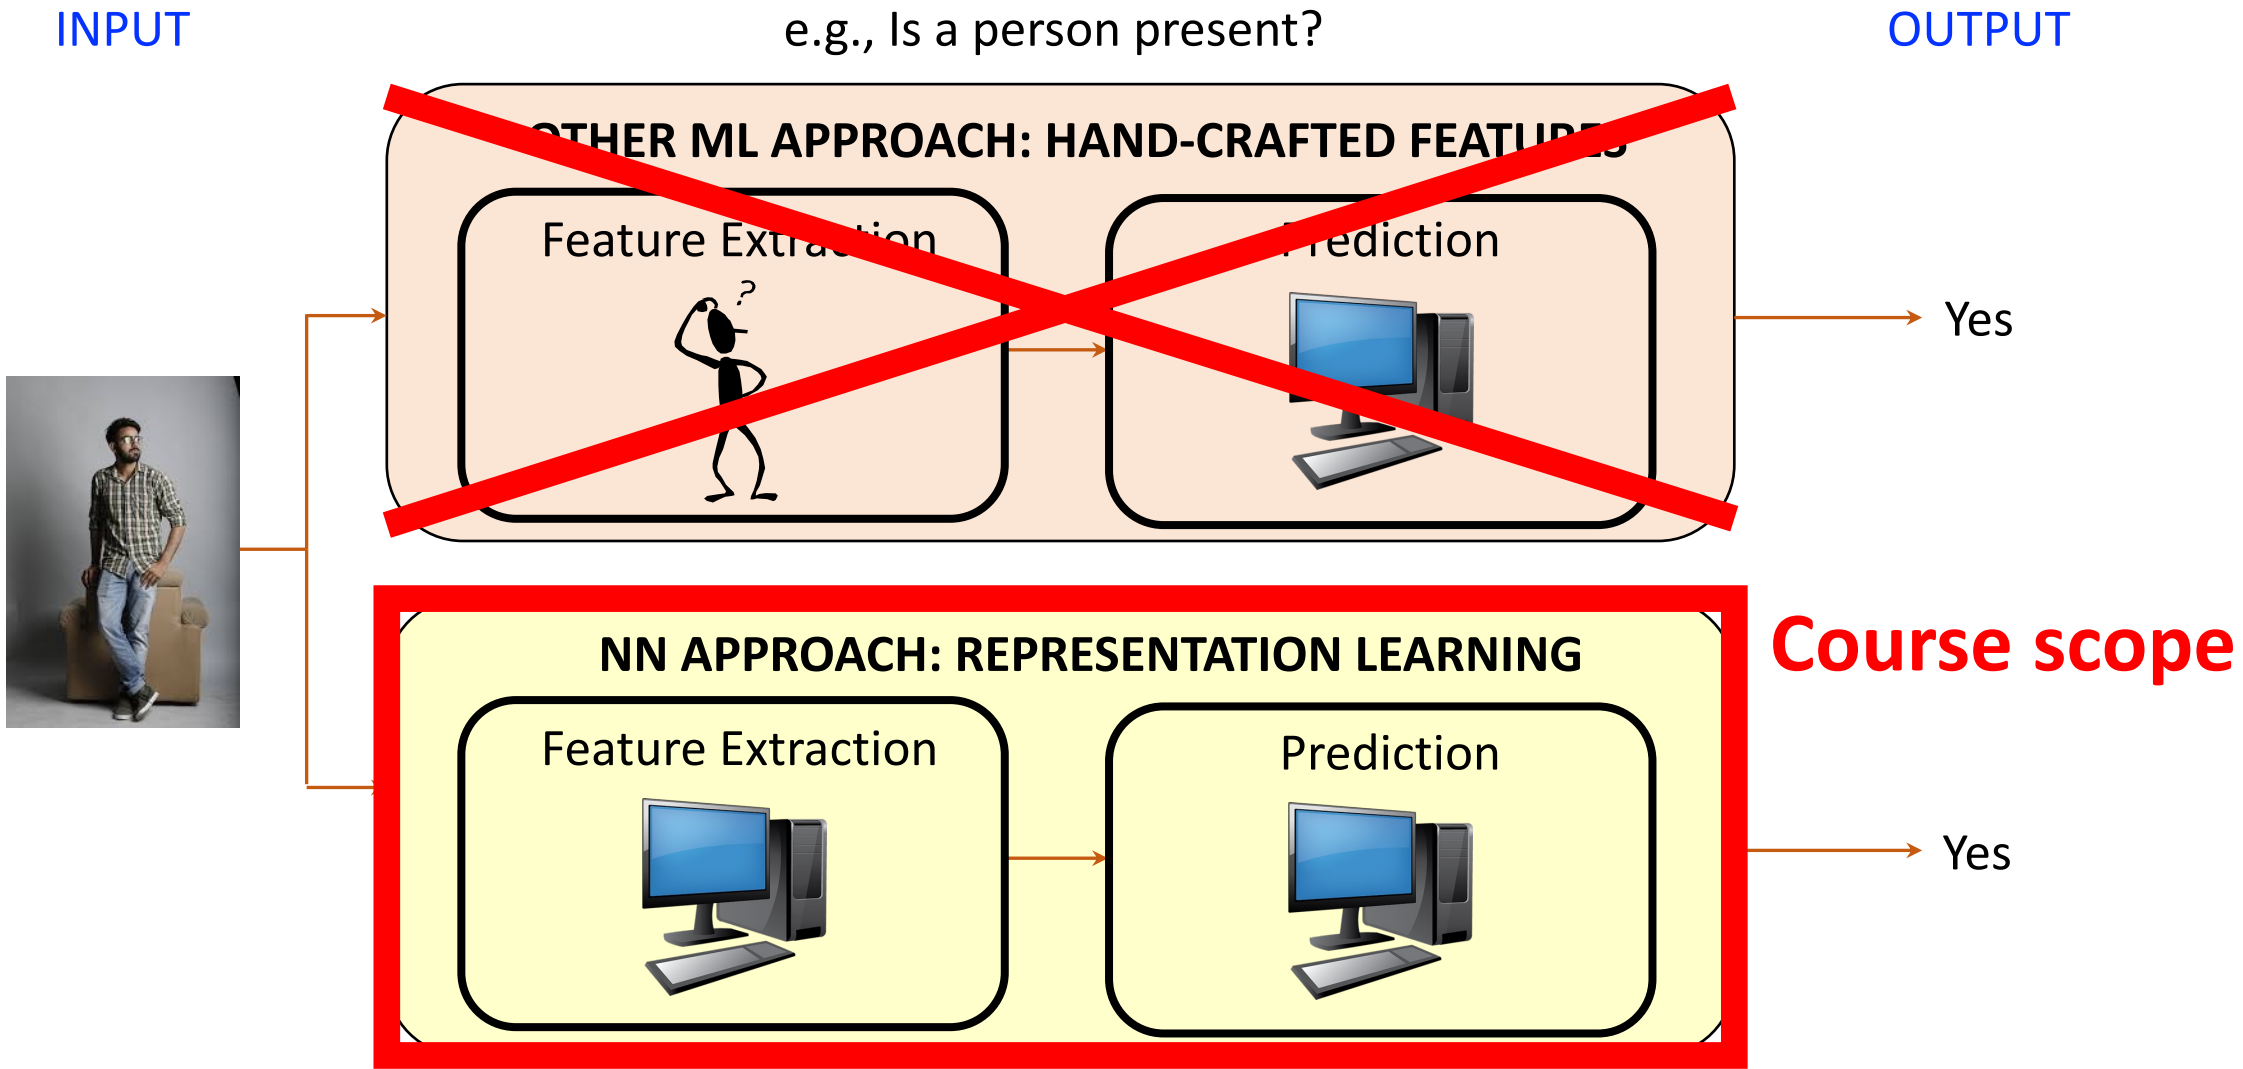

Origins: Rises/Falls of Neural Network Popularity

1945 1950 1956 1959 1986

First programmable machine

Turing test

AI

Machine learning

Neural networks with effective "deep learning" strategy

Neural networks are not new and have been called many names:

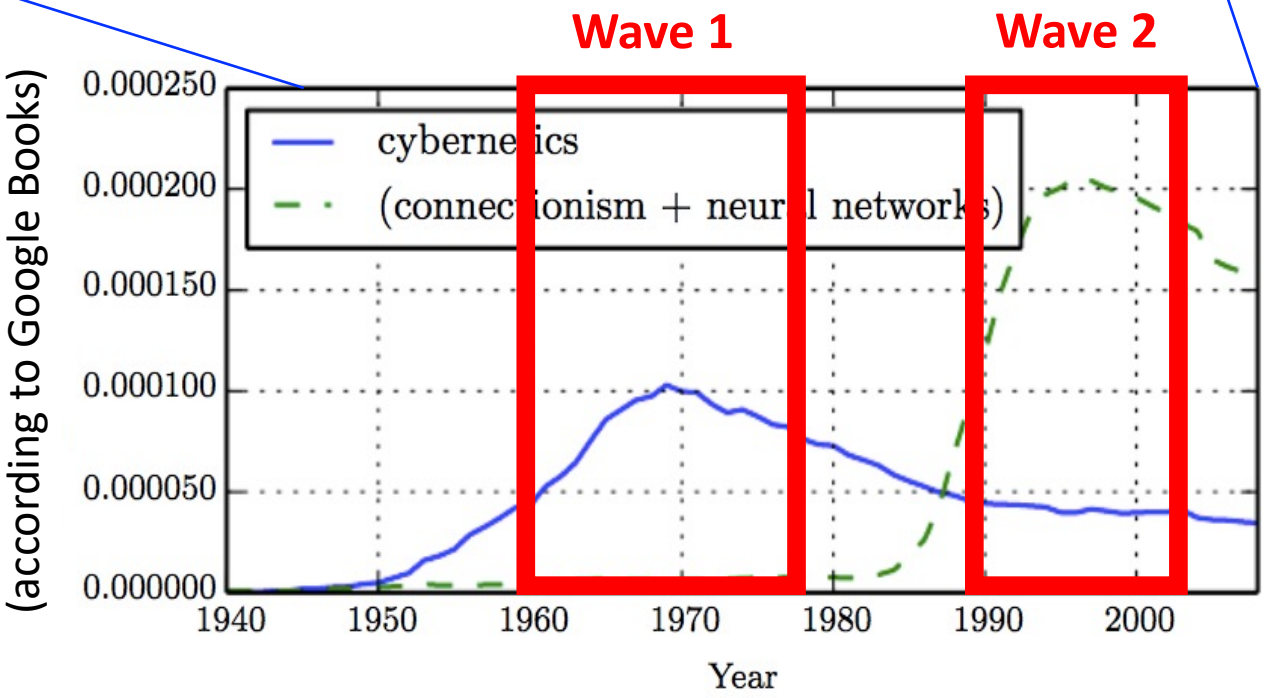

Ian Goodfellow, Yoshua Bengio, and Aaron Courville; Deep Learning, 2016

Origins: Rises/Falls of Neural Network Popularity

1945 First programmable machine

1950 Turing test

1956 AI

1959 Machine learning

1986 Neural networks with effective "deep learning" strategy

2012 Wave 3: rise of "deep learning"

(Google Trends)

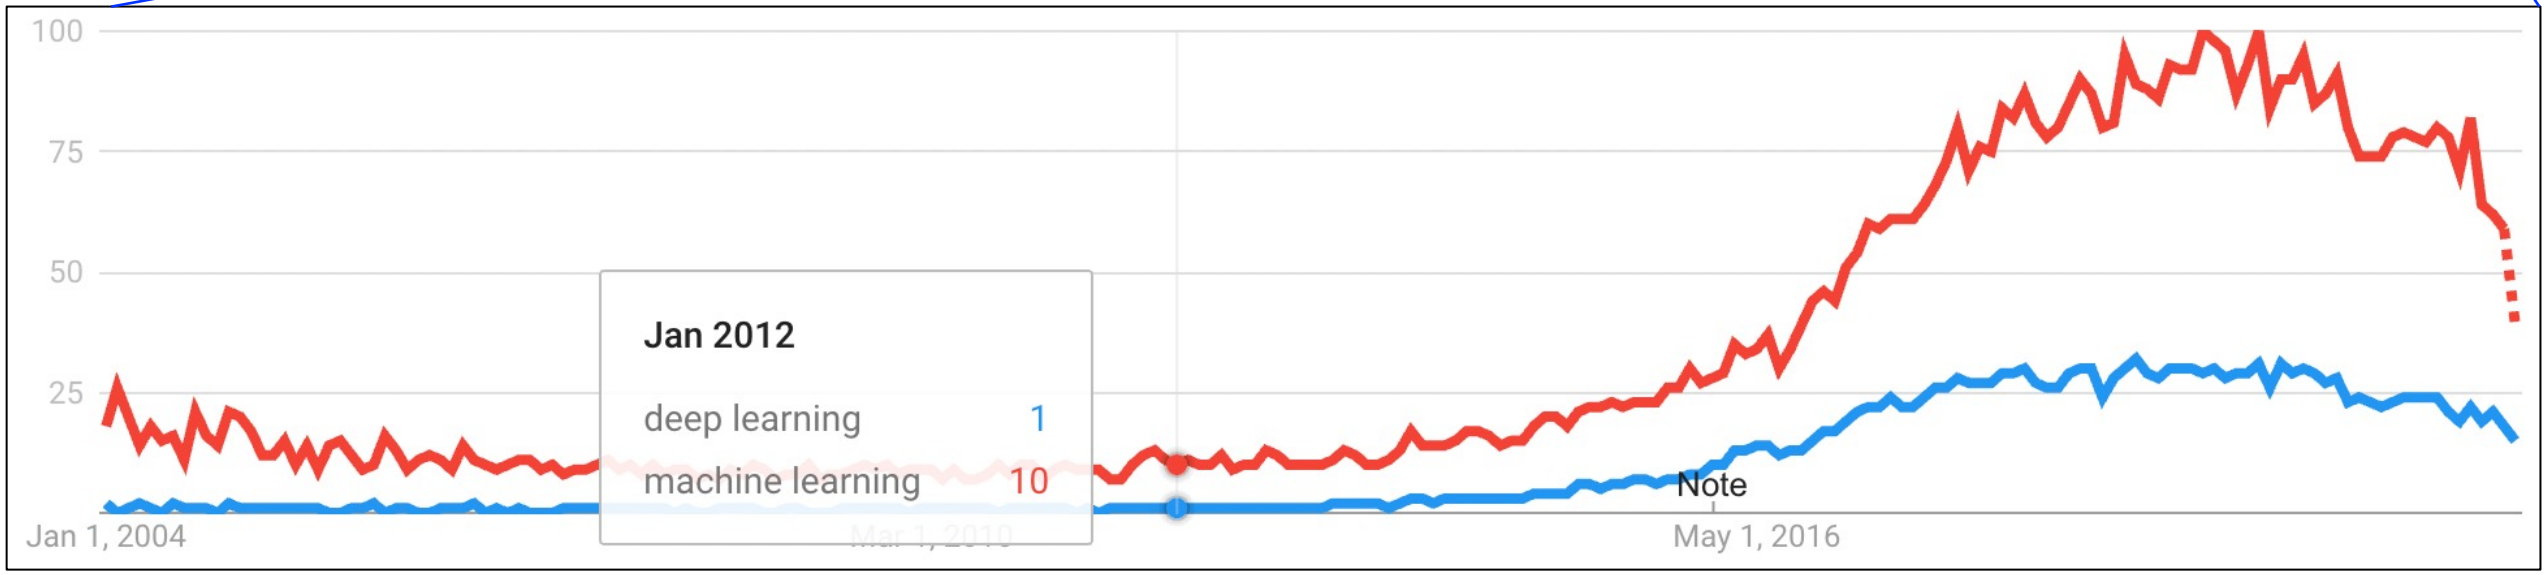

Machine learning popularity has paralleled rise of deep learning popularity

Today’s Topics

- Applications
- History of neural networks and deep learning
- How does a machine learn?
- Course logistics

General Idea

An **algorithm** learns from **data** patterns that will be used to make a prediction

General Idea

An **algorithm** learns from data patterns that will be used to make a prediction

Typical Algorithm Design

- Unsupervised Learning
  - Identify patterns by observing *unstructured* data

- Supervised Learning
  - Identify patterns by studying *structured* data with labels of expected outputs

Types of “Unsupervised” Learning Tasks

Clustering

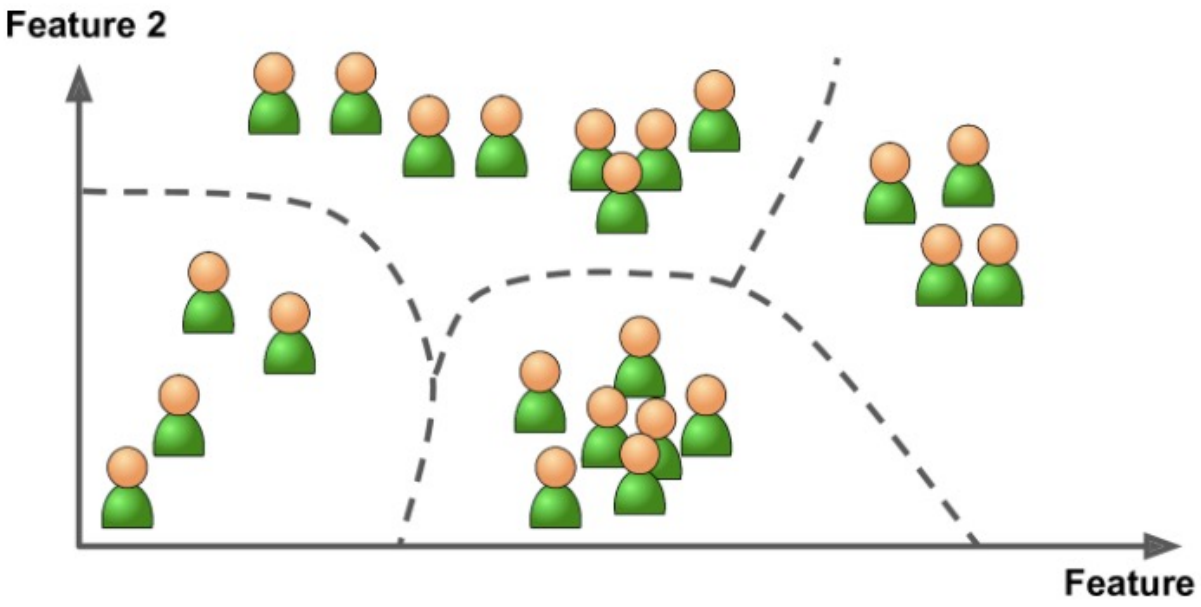

Anomaly Detection

Hands-on Machine Learning with Scikit-Learn & TensorFlow, Aurelien Geron

Types of “Supervised” Learning Tasks

Regression
(predict **continuous** value)

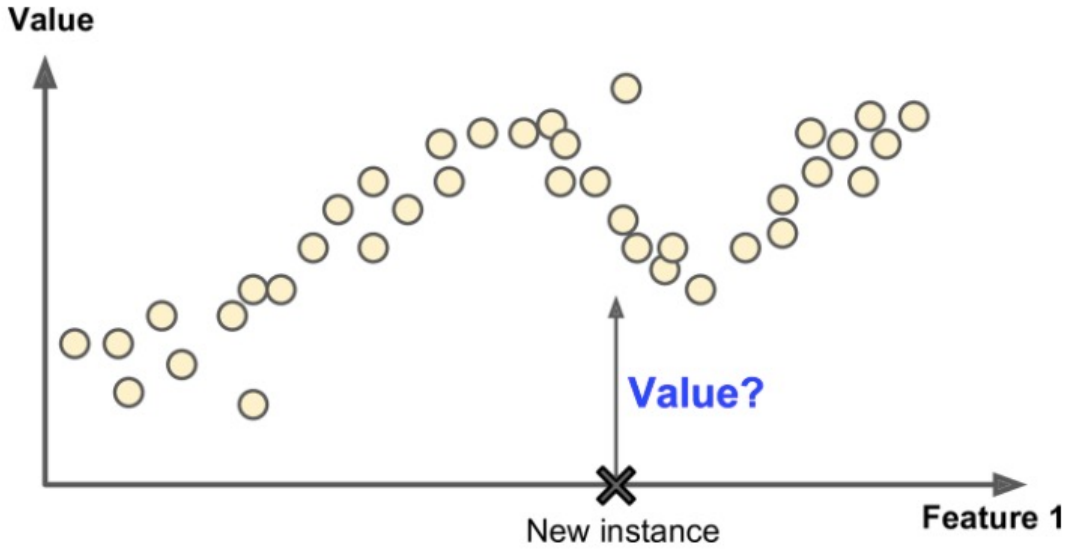

Classification
(predict **discrete** value)

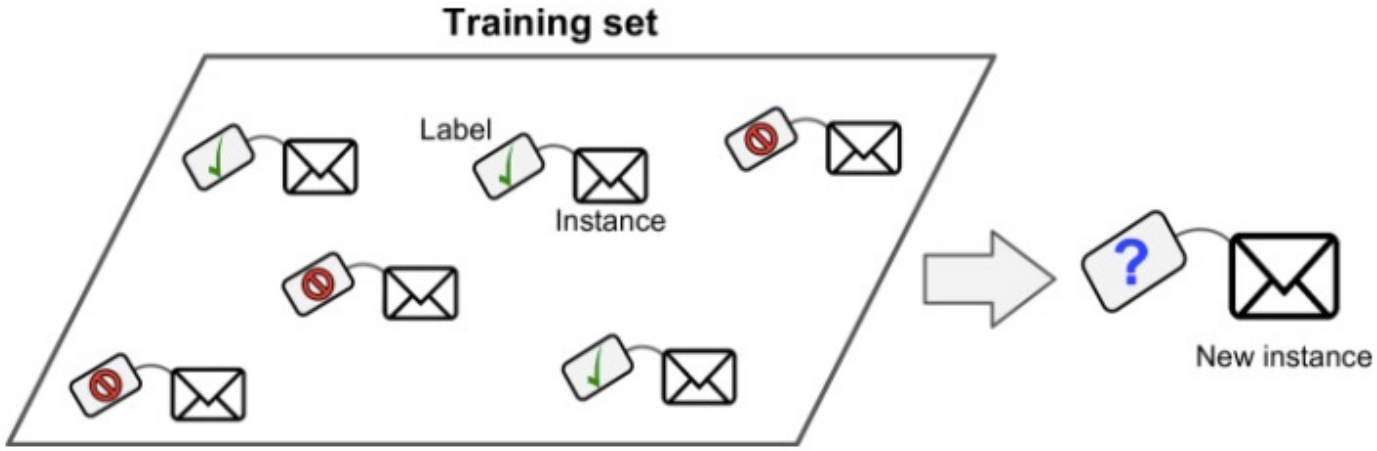

Hands-on Machine Learning with Scikit-Learn & TensorFlow, Aurelien Geron

Typical Supervised Learning Algorithm Design

- Model-based classification approach
  - e.g., create model to separate x from o

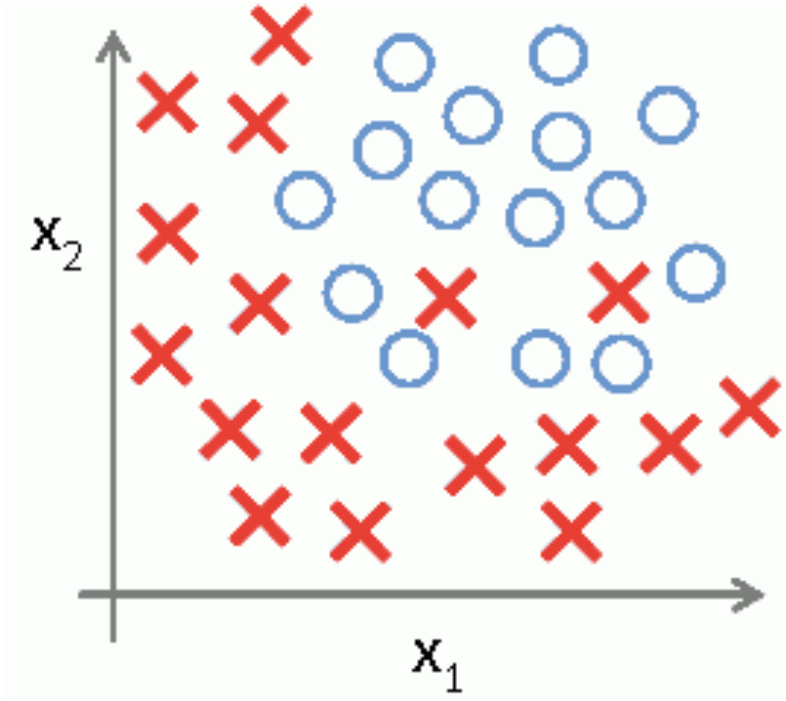

Class volunteer:
1) Draw a straight line (linear equation)
2) Draw a parabola (quadratic equation)
3) Draw any curve

Models with increasing representational capacity

Figure source: https://medium.com/greyatom/what-is-underfitting-and-overfitting-in-machine-learning-and-how-to-deal-with-it-6803a989c76

Algorithm Scope for Course: Last 65 Years

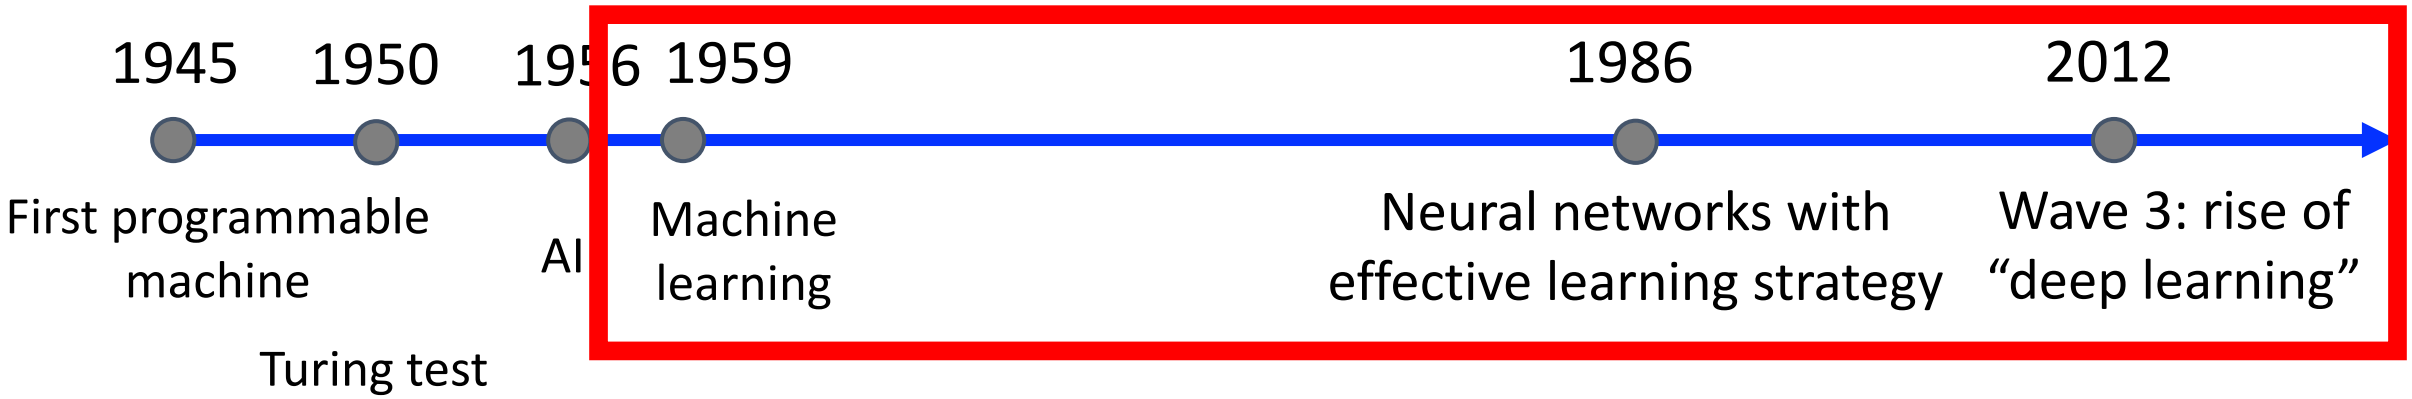

Analogous to understanding how houses we live-in work:

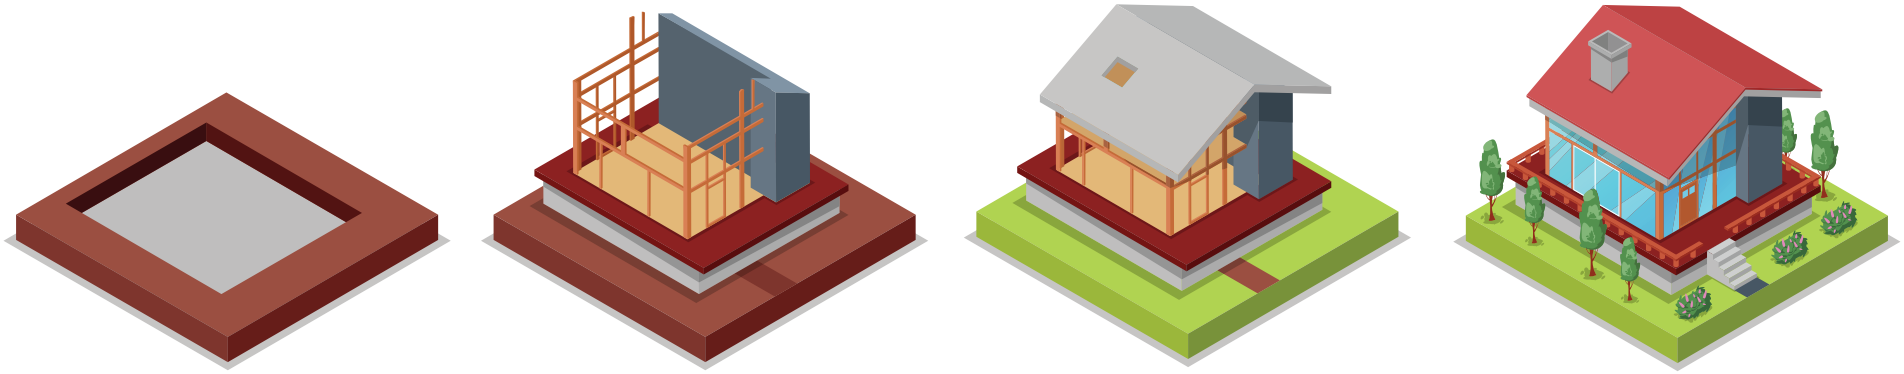

We study older algorithms because modern deep learning algorithms rely on techniques developed over the past 65 years.

Algorithm Scope for Course: Last 65 Years

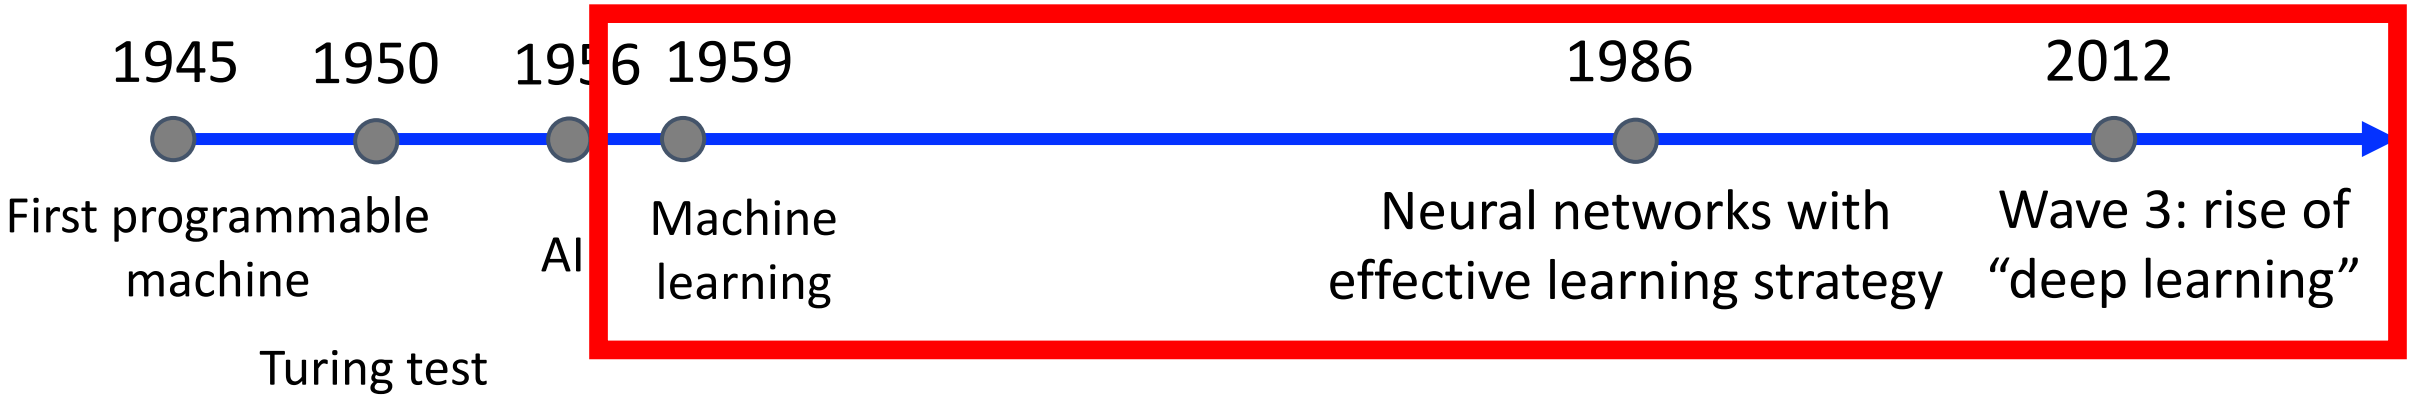

Week	Topic(s)
1	Introduction, Artificial neurons
2	Feedforward neural networks (NN), NN training
3	NN training
4	Convolutional neural networks (CNN), Introduction to CV
5	Training CNN algorithms
6	Regularization, Pretrained CNN features, Fine-tuning
7	Object detection, Semantic segmentation, Recurrent neural networks
8	Introduction to natural language proecssing, Neural word embeddings, Attention

Week	Topic(s)
9	Transformers
10	Multimodal NN (vision + language)
11	Multimodal NN, Self-supervised learning, GANs
12	Few/zero-shot learning, Responsible/ethical learning
13	Deep learning in industry (guest speakers)
14	Model compression, Efficient learning
15	NNs for speech processing & reinforcement learning
16	NN for information retrieval & course summary

General Idea

An algorithm learns from **data** patterns that will be used to make a prediction

Data Types: What a Machine Learns From?

- Audio
  - Input?

e.g.,

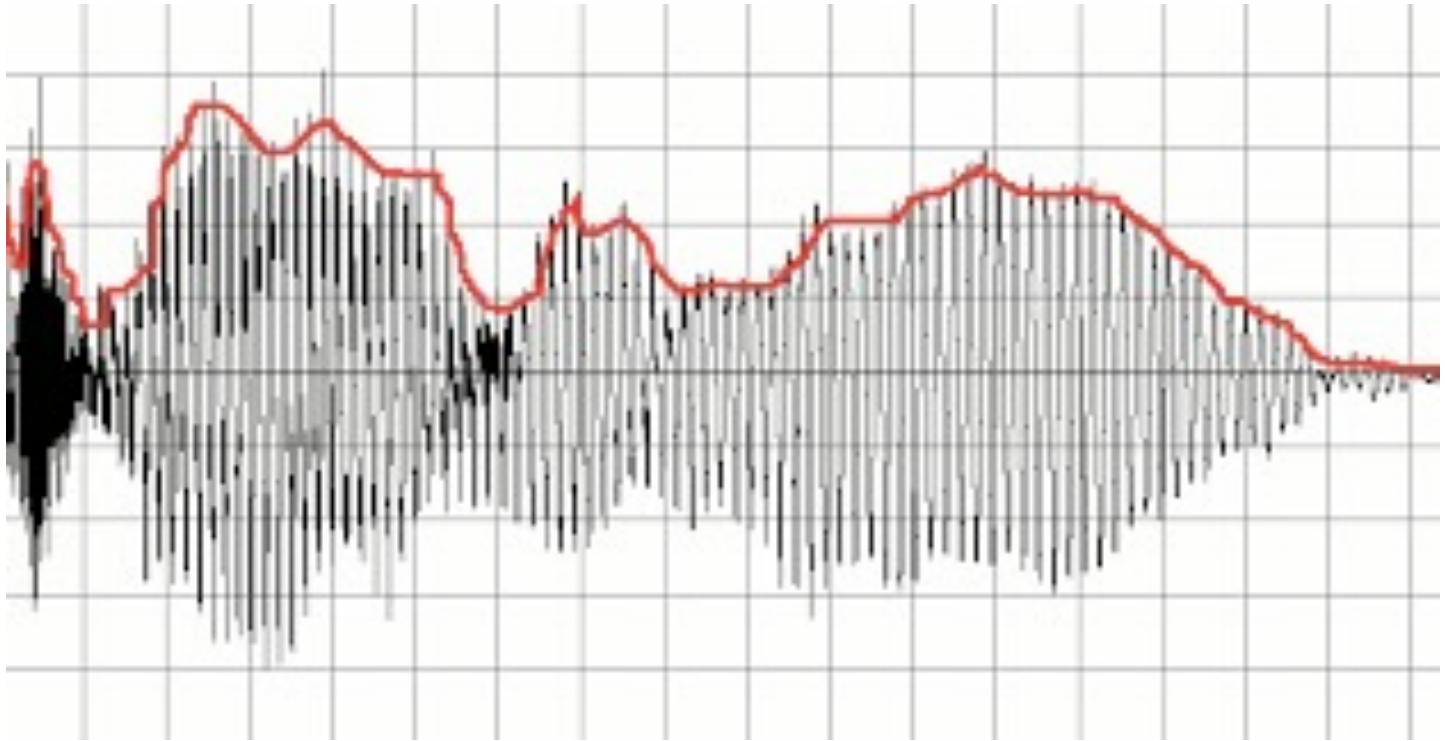

Data Types: What a Machine Learns From?

- Audio
  - Input?
- Images
  - Input?

157	153	174	168	150	152	129	151	172	161	155	156
155	182	163	74	75	62	33	17	110	210	180	154
180	180	50	14	34	6	10	33	48	106	159	181
206	109	5	24	131	111	120	204	166	15	56	180
194	68	137	251	237	239	239	228	227	87	71	201
172	105	207	233	233	214	220	239	228	98	74	206
188	88	179	209	185	215	211	158	139	75	20	169
189	97	165	84	10	168	134	11	31	62	22	148
199	168	191	193	158	227	178	143	182	106	36	190
205	174	155	252	236	231	149	178	228	43	95	234
190	216	116	149	236	187	86	150	79	38	218	241
190	224	147	108	227	210	127	102	36	101	255	224
190	214	173	66	103	143	96	50	2	109	249	215
187	196	235	75	1	81	47	0	6	217	255	211
183	202	237	145	0	0	12	108	200	138	243	236
195	206	123	207	177	121	123	200	175	13	96	218

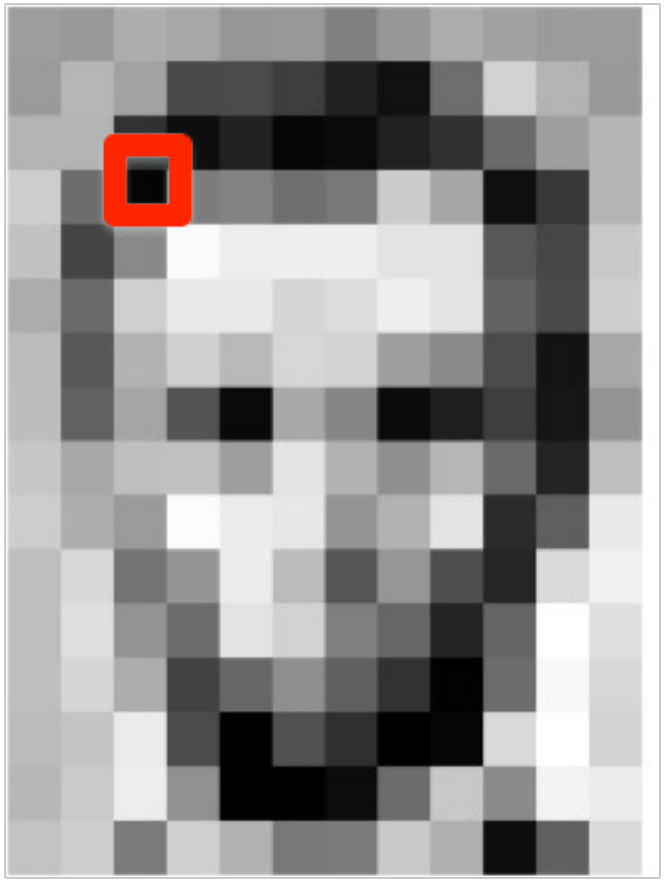

Data Types: What a Machine Learns From?

- Audio
  - Input?
- Images
  - Input?
- Video
  - Input?

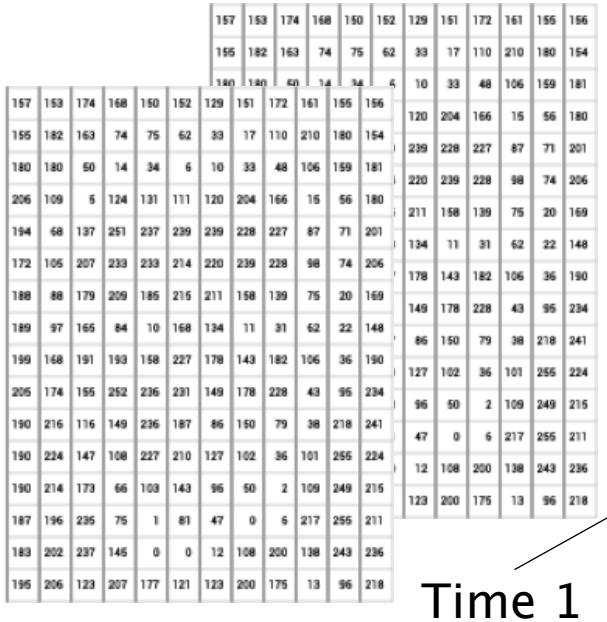

Time 1

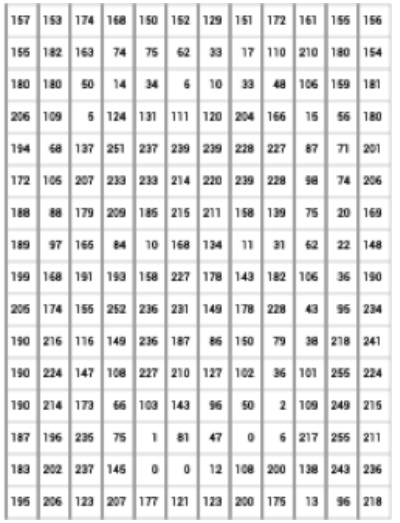

1 hour

Analogous to:

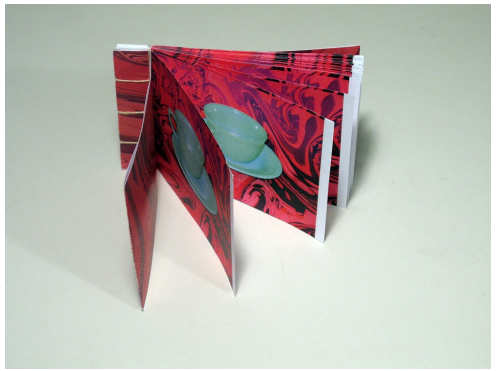

Data Types: What a Machine Learns From?

- Audio
  - Input?
- Images
  - Input?
- Video
  - Input?
- Text
  - Input?

e.g.,

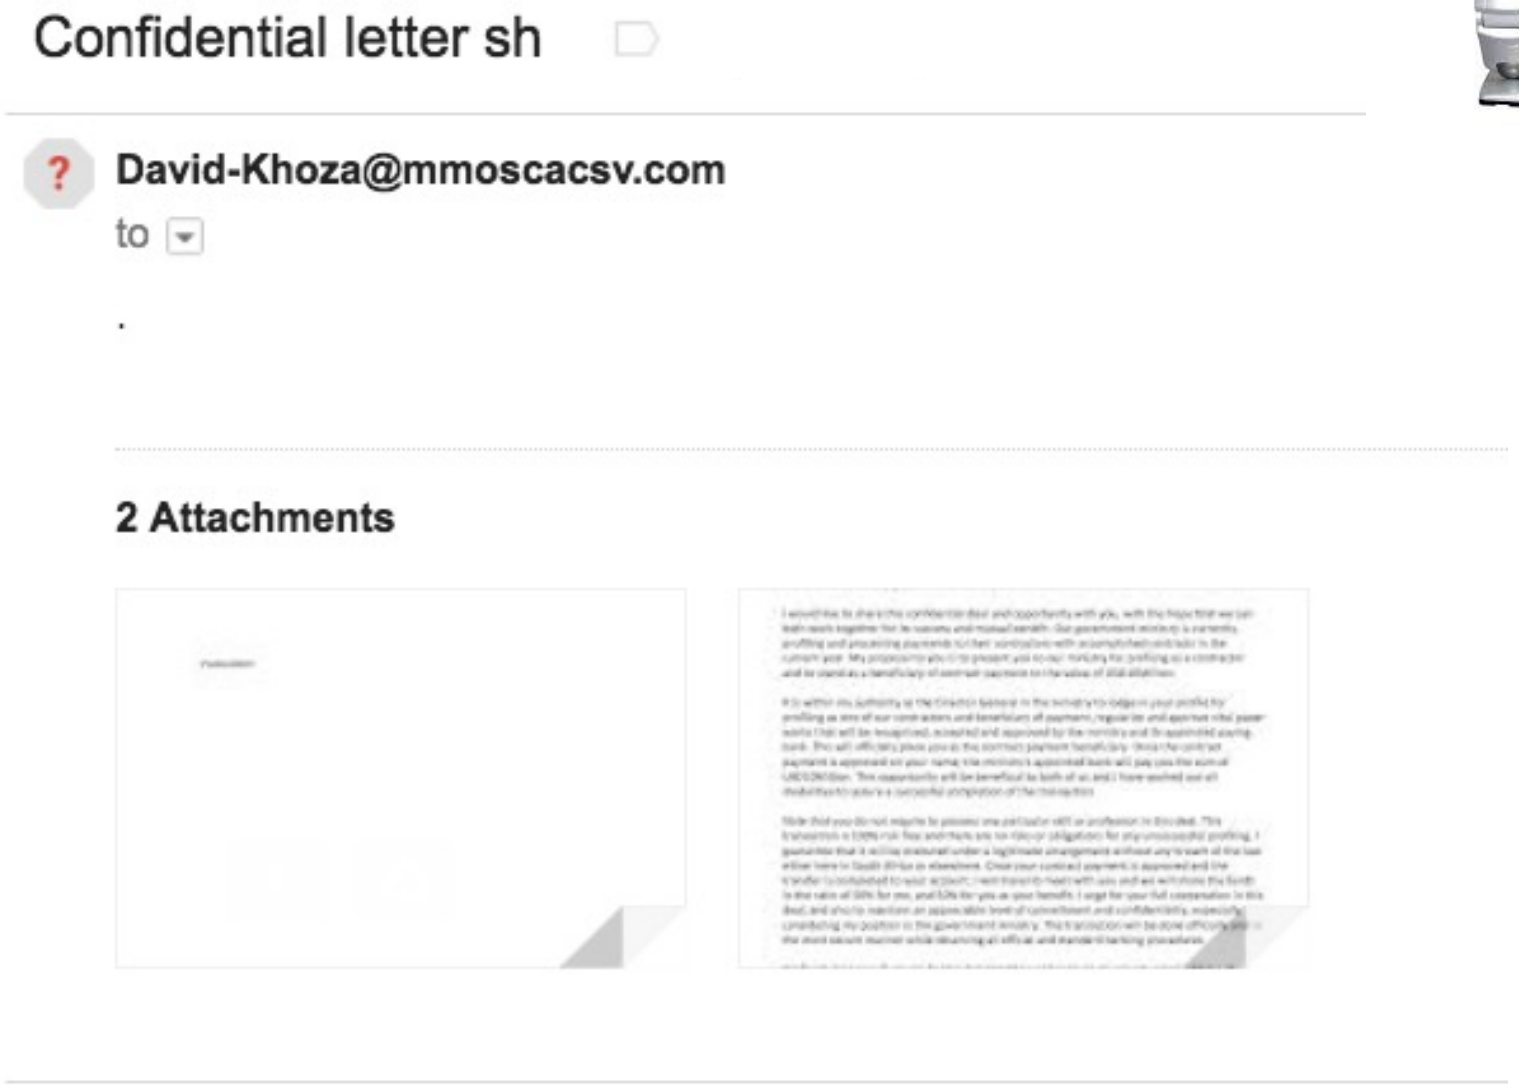

Data Types: What a Machine Learns From?

- Audio
  - Input?
- Images
  - Input?
- Video
  - Input?
- Text
  - Input?
- Multi-modal
  - Input? - combination of the above

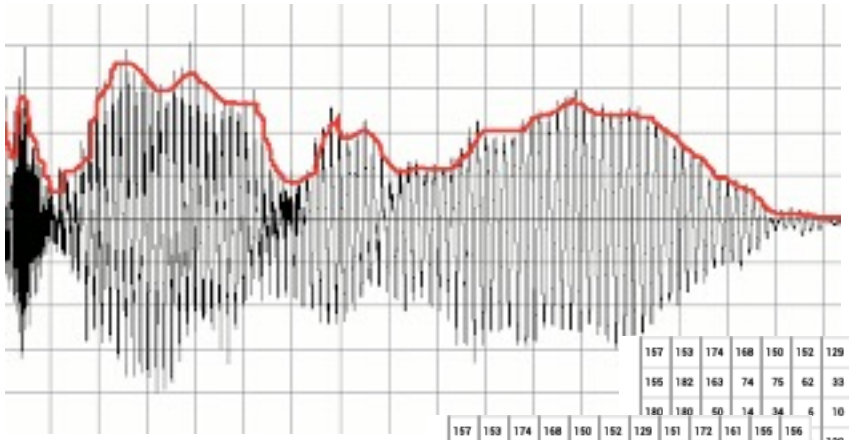

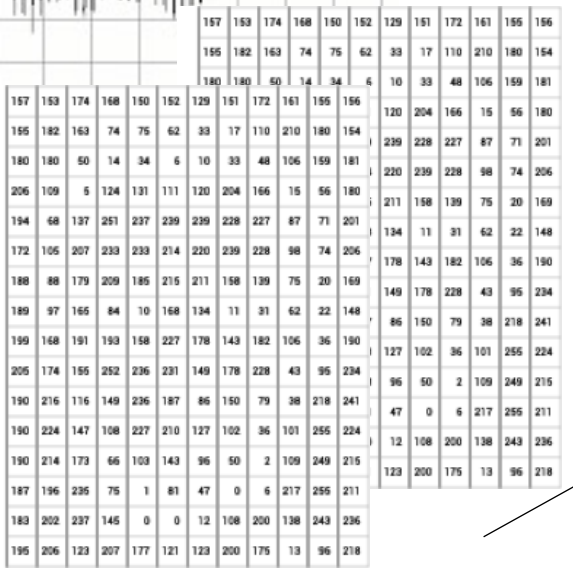

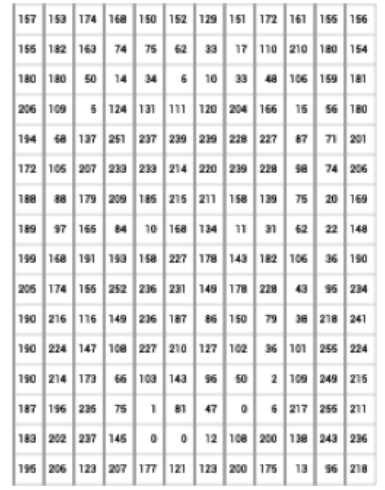

Data Types: Many Public Datasets Available

- Dataset creation is beyond the scope of this class
- We will benefit from other people's efforts:
  - Google Dataset Search
  - Amazon's AWS datasets
  - Kaggle datasets
  - Wikipedia's list
  - UC Irvine Machine Learning Repository
  - Quora.com
  - Reddit
  - Dataportals.org
  - Opendatamonitor.eu
  - Quandl.com

General Idea

An **algorithm** learns from **data** patterns that will be used to make a prediction

Why Are Neural Networks and Deep Learning So Popular? – Its Success in Practice!

It's success was realized with the relatively recent onset of:

1. **Big data**: originally, often from the Internet
2. **Better hardware**: faster hardware and more storage enabled practically fast "deep learning"

Neural Networks: Key Ingredients for Success

An **algorithm** learns from **data**
on a **processor** the patterns that
will be used to make a prediction

Analogous to a Love Story of Partnering Up and Road Tripping Somewhere

Key Challenge 1: How Long Does Learning Take?

An **algorithm** learns from **data**
on a **processor** the patterns that
will be used to make a prediction

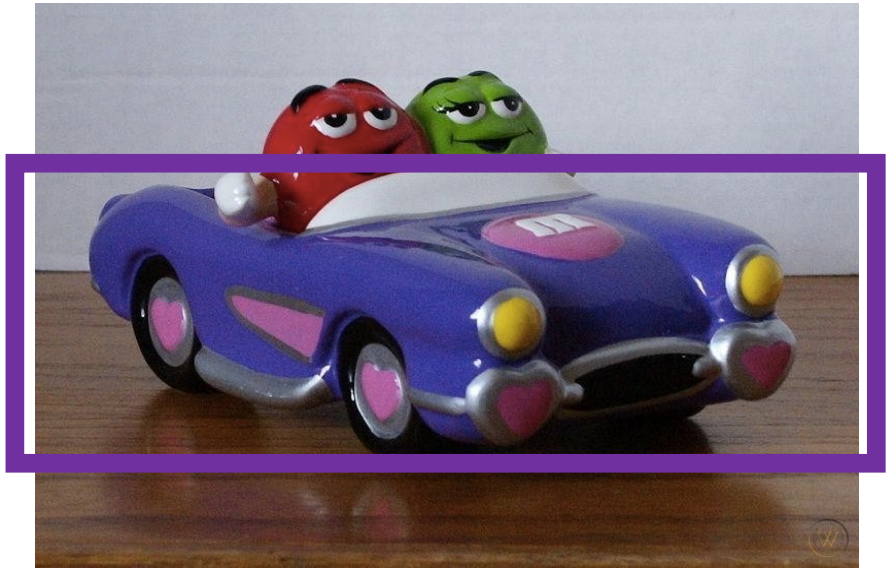

Analogous to a Love Story of Partnering Up and Road Tripping Somewhere

Key Challenge 1: How Long Does Learning Take?

e.g., Train Algorithms Using
GPUs (think Porsche) Instead of CPUs (think Golf Cart)

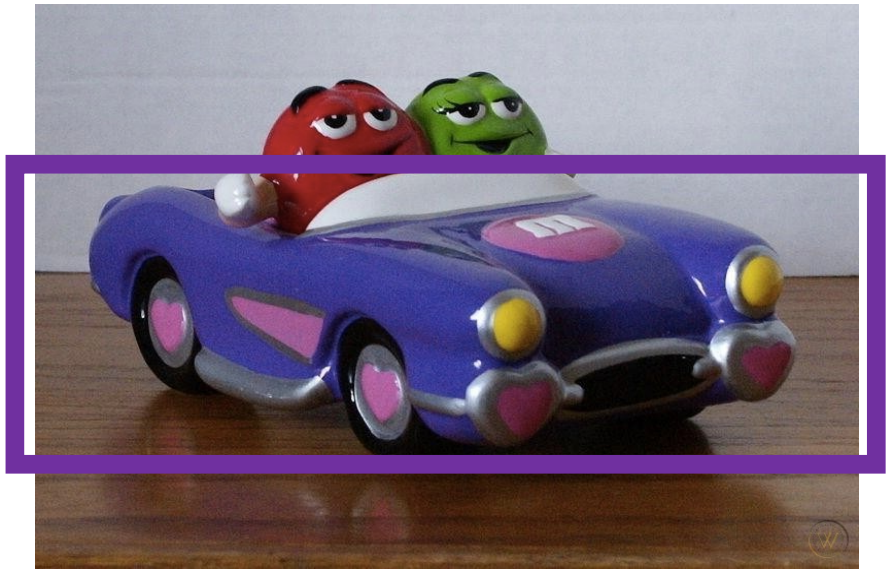

Key Challenge 2: Where Will You Go?

https://www.theverge.com/2015/7/1/8880363/google-apologizes-photos-app-tags-two-black-people-gorillas

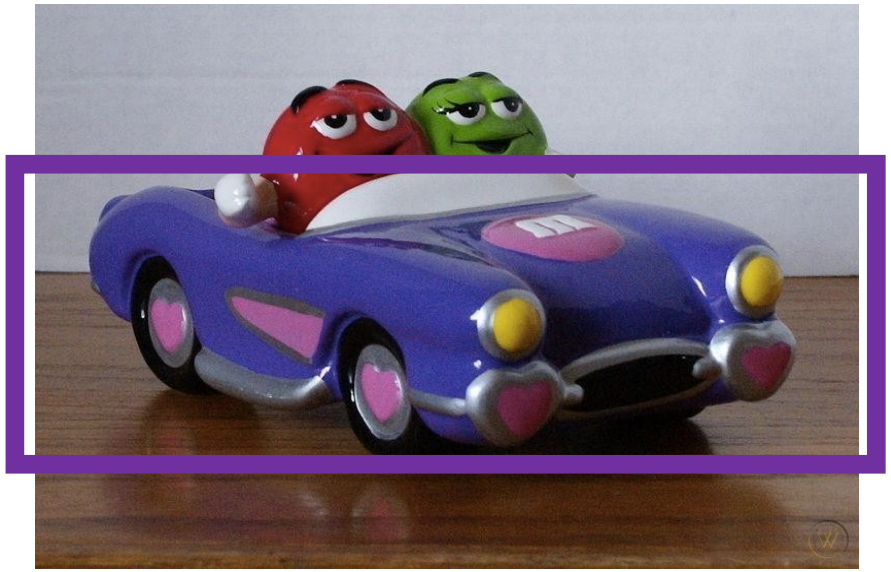

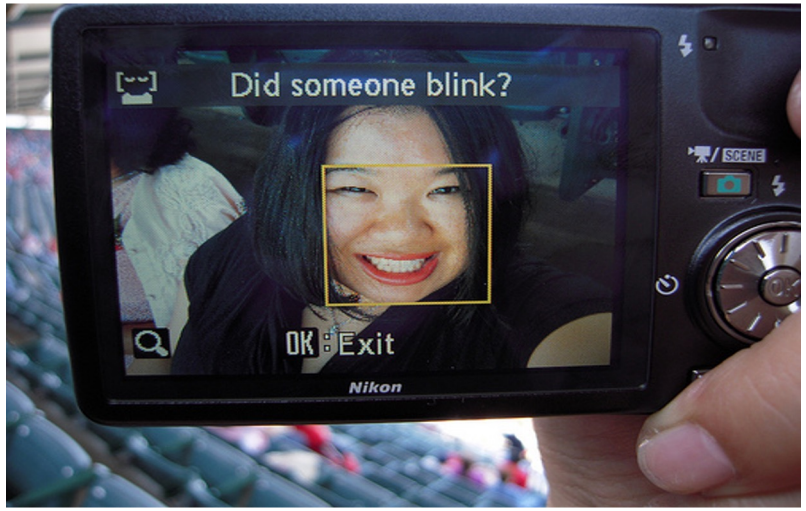

Two kids bought their mom a camera for Mother's Day… when they took portrait pictures of each other, a message flashed across the screen asking, "Did someone blink?"

http://content.time.com/time/business/article/0,8599,1954643,00.html

Today’s Topics

- Applications
- History of neural networks and deep learning
- How does a machine learn?
- Course logistics

Introductions

Instructor: Danna Gurari; aka, Dr. G

(preferred pronouns: she/her)

Teaching assistant: Samreen Anjum

(preferred pronouns: she/her)

Teaching assistant: Lucas Hayne

(preferred pronouns: he/him)

Introductions

NameCoach: share your name pronunciation in Canvas

To record your name:

1. Find NameCoach in Canvas courses page
2. Click on record button to start
3. Review recording by clicking on play button

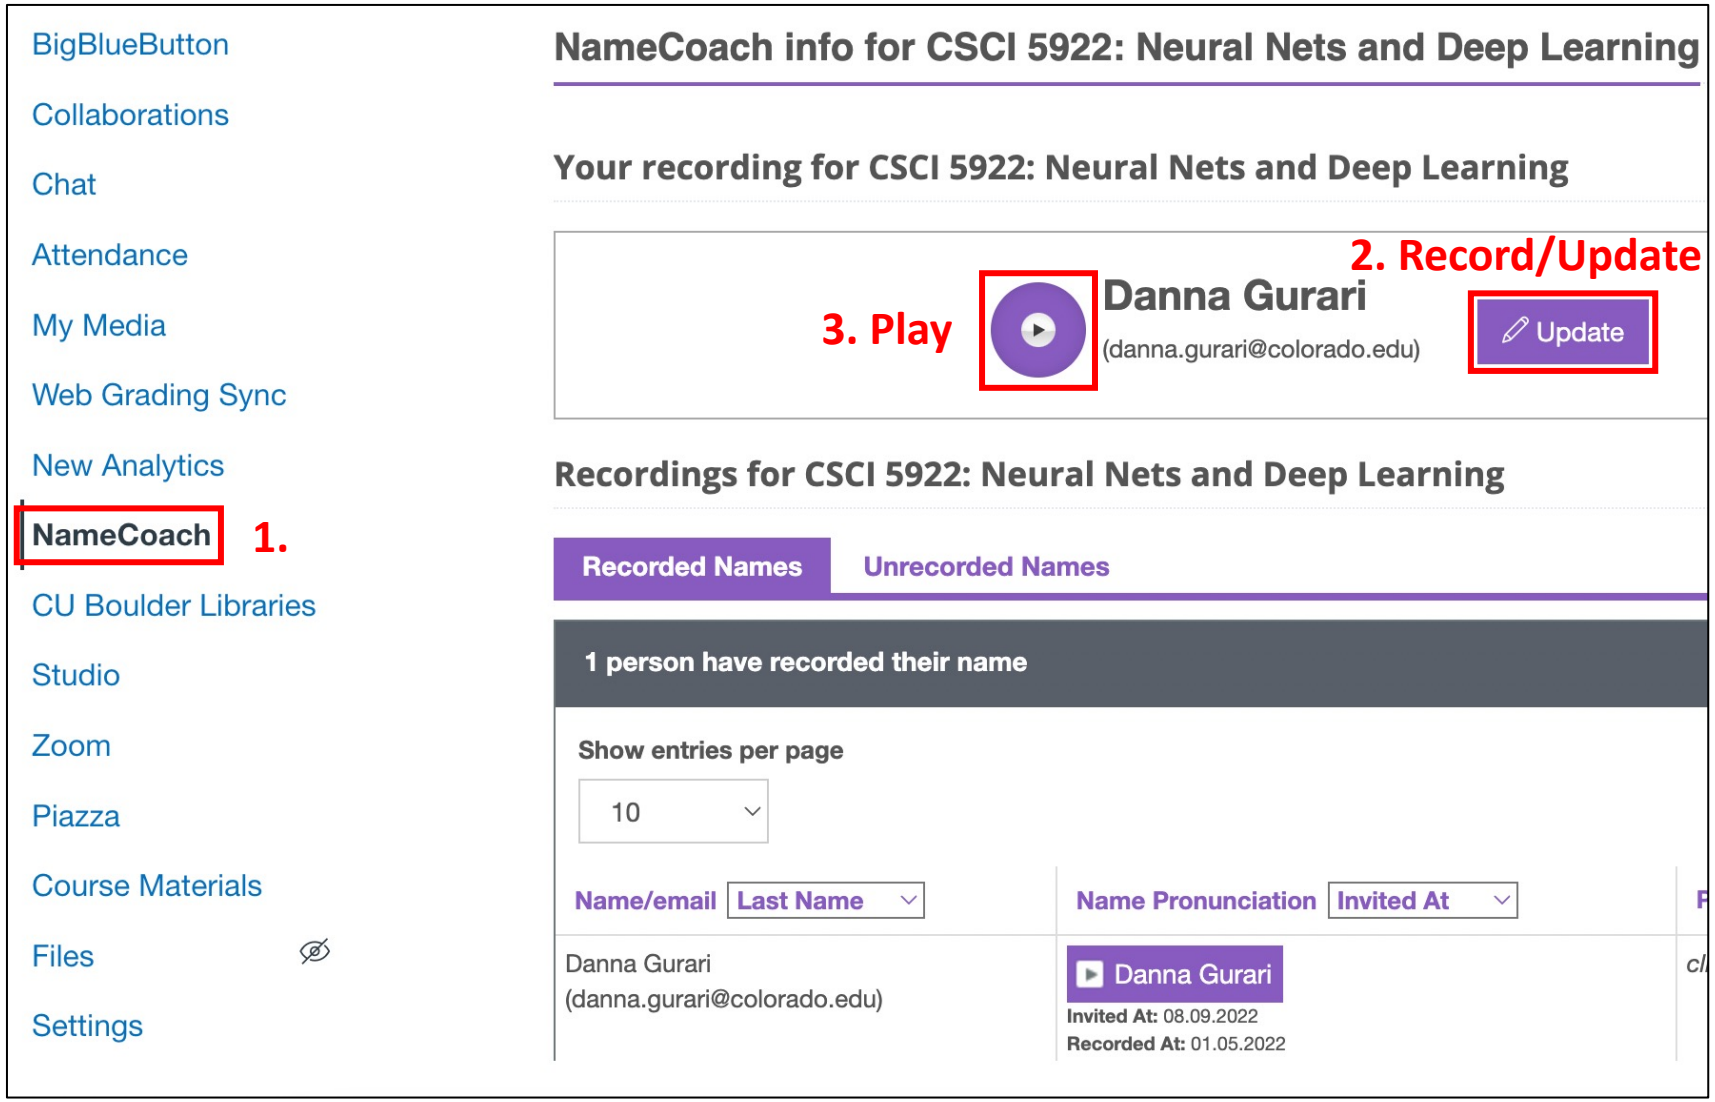

Course Objectives

- Understand the key concepts for designing deep learning models:
  1. Characterize the process to train and test deep learning algorithms
  2. Identify the challenges for designing modern deep learning algorithms that can harness today's 'big' data
  3. Recognize strengths and weaknesses of different deep learning algorithms

- Apply deep learning models to perform various AI tasks:
  1. Experiment with deep learning libraries, including scikit-learn and Keras
  2. Evaluate deep learning algorithms for tasks in various application domains, including for analyzing text and images

Course Objectives

- Conduct and communicate a novel project:
  1. Propose a novel project idea (this will be an iterative process)
  2. Design and execute experiments to support the proposed idea
  3. Create a presentation about the project
  4. Write a report about the project

Course Overview

- Website
  - https://home.cs.colorado.edu/~DrG/Courses/NeuralNetworksAndDeepLearning/AboutCourse.html

- Syllabus on website

- Grading :

	% of Final Class Grade
Problem Sets	25%
Lab Assignments	35%
Final Project	40%

Q&A: “What are the assignments?”

- 3 problem sets (first assignment due next week)
- 4 lab assignments
- Final project
  - Proposal
  - Outline
  - Presentation
  - Final report
- Grading policies:
  - Late policy: Penalized 1% per hour for up to 2 hours and no credit afterwards
  - Regrade requests: Must be submitted within 2 weeks of receiving the grade to the TA, Samreen Anjum. After the allotted time frame, regrade requests will not be considered.
- Grading timeline
  - Due to many students (100+) and 2 graders, expect to receive grades ~1 week after submission

Q&A: “Do I have the appropriate pre-requisites/background?”

- You are expected to have programming competency as well as experience with probability/statistics and linear algebra.

Q&A: “What are required textbooks?”

(available online for free)

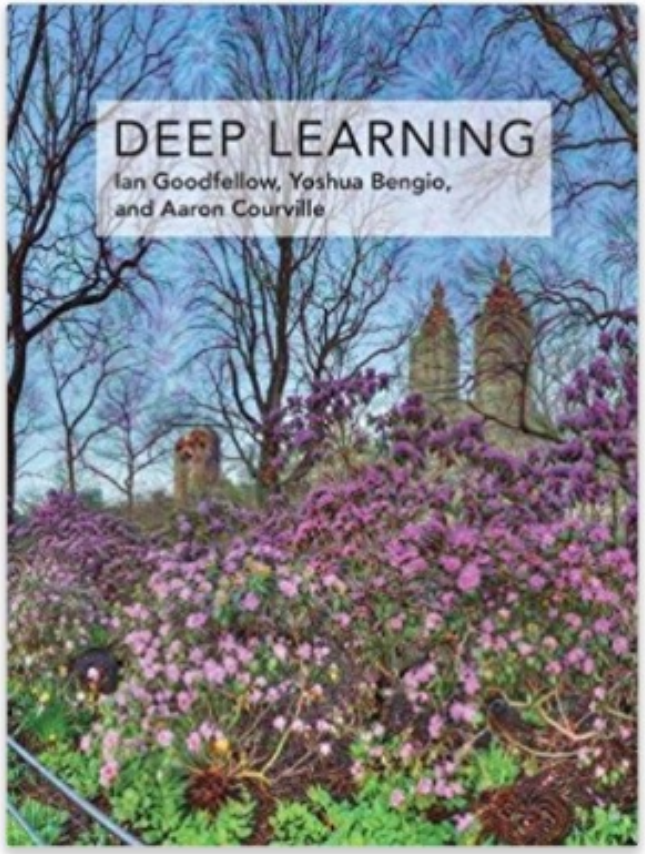

(available online for free when connected to CU Boulder’s network or VPN)

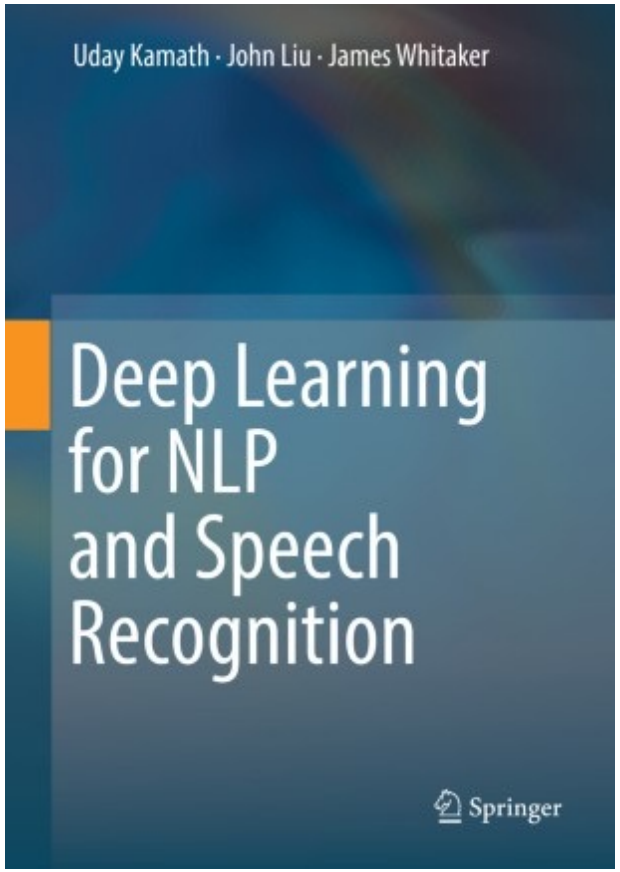

Q&A: “How Do I Contact You for Questions?”

- **Questions for Instructor**: I will stay after each class lecture to answer questions.

- **Piazza:** We encourage you to first post any questions to Piazza. This can benefit other students, who may have similar questions, to see the answers posted on Piazza and it also allows peers to answer your questions. You can post questions to Piazza from Canvas.

- **Office Hours**: The TAs will host office hours every day from Monday to Friday. To attend, please enter your name on the office hours spreadsheet shared in Canvas and Piazza and then join the video meeting link shared on Canvas.

- **Appointments**: Email the TA to make an appointment or solicit an answer. Please note that at least 24 hours notice will typically be needed before the TA will be able to meet. The TA will involve the instructor for any items she is unable to address.

- **Regrade requests**: All requests must be emailed to the TA, Samreen Anjum, within 2 weeks of receiving the grade to be considered.

My Experience Related to Deep Learning

2007-2010

Software developer and project manager helping to record and analyze visible and infrared video

Source: Boulder Imaging

2010-2015

PhD student designing computer vision methods to segment and track cells in images and videos

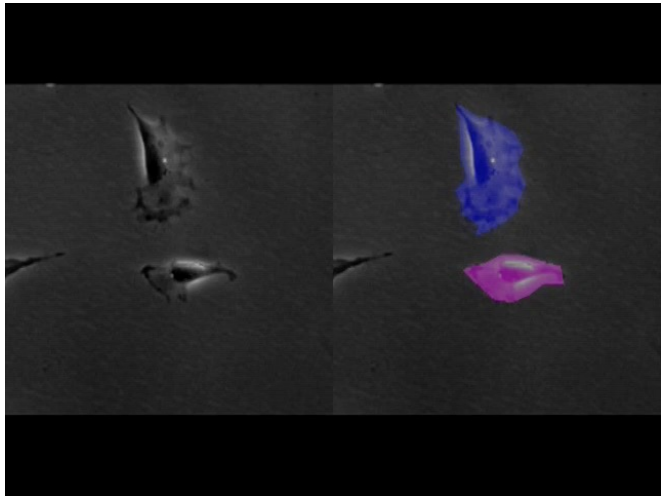

2015-2017

Postdoctoral fellow conducting research projects related to vision and language problems

Is it edible or poisonous?

2017-Present

Assistant professor overseeing research projects related to many vision and language problems

e.g., image classification, object detection, semantic segmentation, object tracking, image captioning, visual question answering, style transfer, image inpainting, and image search

My Experience Related to Deep Learning

2007-2010	2010-2015	2015-2017	2017-Present
Software developer and project manager helping to record and analyze visible and infrared video	PhD student designing computer vision methods to segment and track cells in images and videos	Postdoctoral fellow conducting research projects related to vision and language problems	Assistant professor overseeing research projects related to many vision and language problems

Source: Boulder Imaging

Is it edible or poisonous?

e.g., image classification, object detection, semantic

style transfer, image inpainting, and image search

20 publications involving deep learning

What is My “Why” for Teaching You…

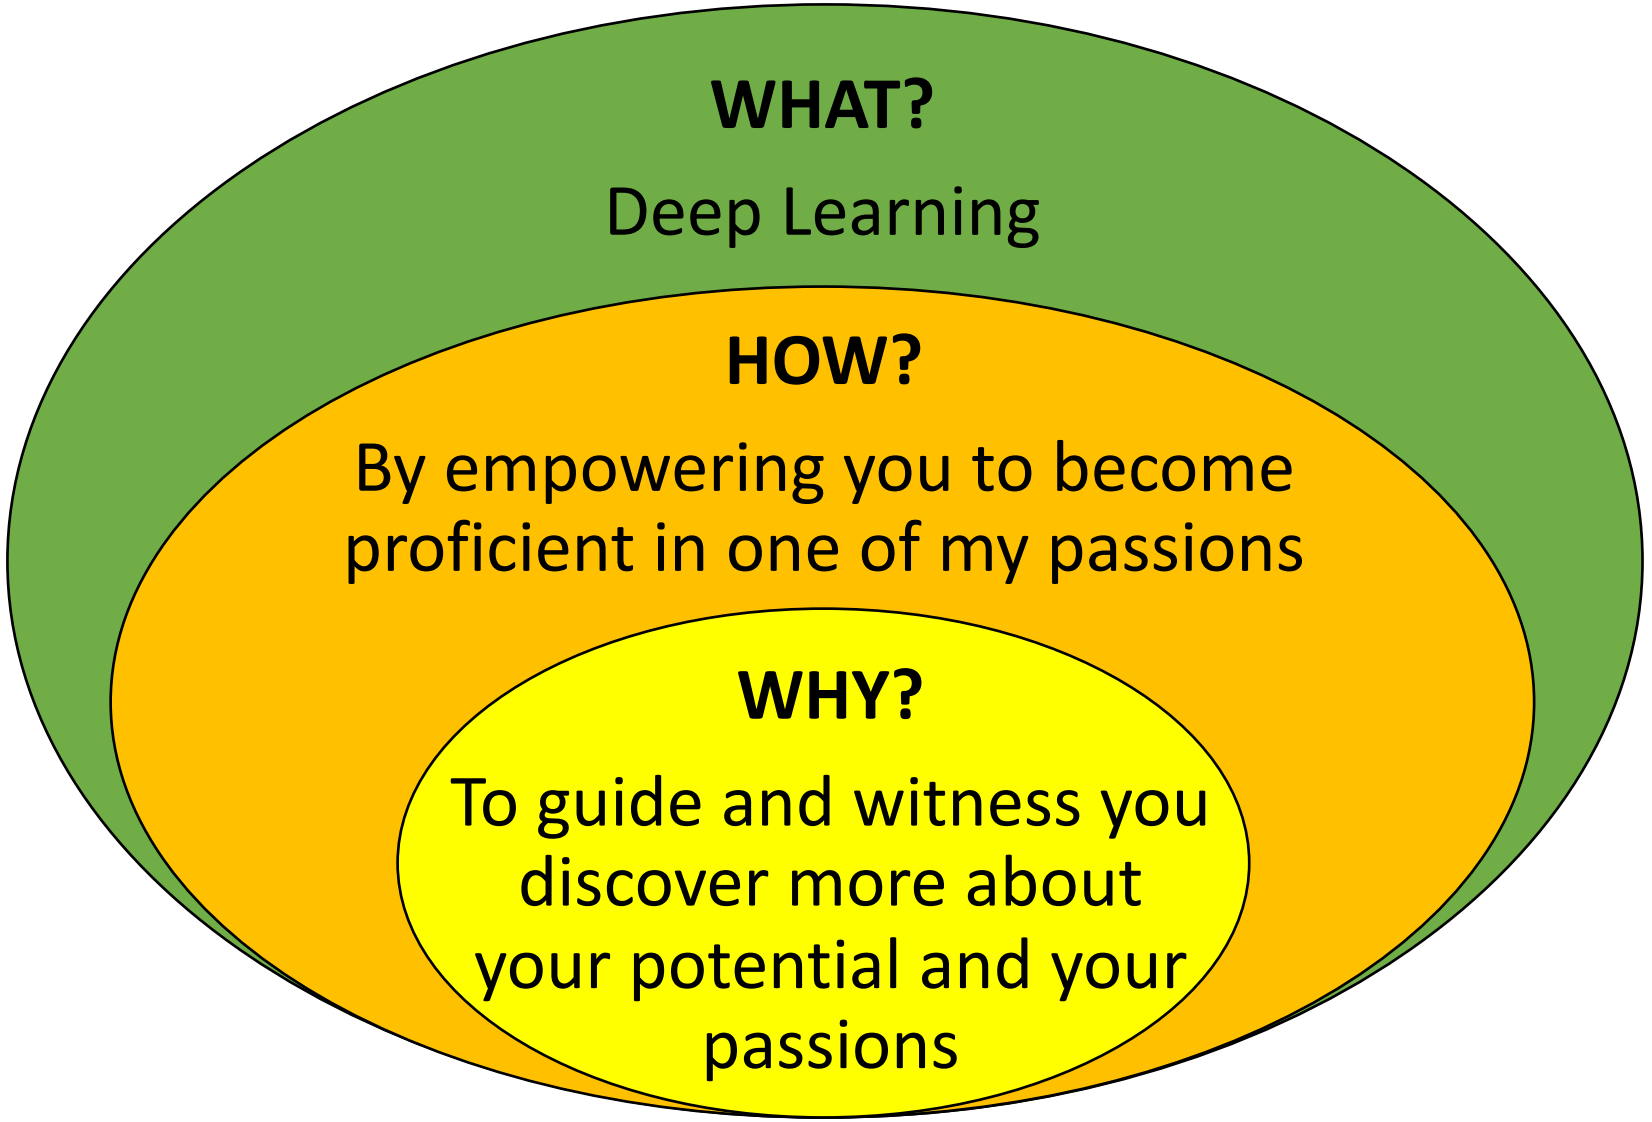

Today’s Topics

- Applications
- History of neural networks and deep learning
- How does a machine learn?
- Course logistics


The End