

Introduction to Neural Networks and Deep Learning

Danna Gurari

University of Colorado Boulder

Spring 2025



<https://dannagurari.colorado.edu/course/neural-networks-and-deep-learning-spring-2025/>

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*

AI Companions; e.g.,

What can I help with?

Message ChatGPT



Welcome to **Gemini**, your personal AI assistant ×

[Google Terms](#) and the [Gemini Apps Privacy Notice](#) apply. Chats are reviewed and used to improve Google AI. [Learn about your choices](#). Gemini can make mistakes, so double-check it. Info [about your location](#) is also stored with your Gemini Apps activity.



Message Copilot



Speech-Based Text Messaging




Coding Assistants; e.g.,

Copilot Enterprise now available

The world's most widely adopted AI developer tool.

Get started with Copilot >



GITHUB COPILOT: CHAT

monalisa

Write unit tests for this function

GitHub Copilot

```
import unittest
import datetime

class TestParseExpenses(unittest.TestCase):
    def test_parse_expenses_with_valid_input(self):
        expenses_string = '2023-01-02 -34.01 USD\n2023-01-03 2.59 DKK\n2023-01-03 -2.72 EUR'
        expected_output = [(datetime.datetime(2023, 1, 2), -34.01, 'USD'), (datetime.datetime(2023, 1, 3), 2.59, 'DKK'), (datetime.datetime(2023, 1, 3), -2.72, 'EUR')]
        self.assertEqual(parse_expenses(expenses_string), expected_output)

    def test_parse_expenses_with_empty_input(self):
        expenses_string = ''
        expected_output = []
        self.assertEqual(parse_expenses(expenses_string), expected_output)

if __name__ == '__main__':
    unittest.main()
```

parse_expenses.py

```
import datetime

def parse_expenses (expenses_string):
    """Parse the list of expenses and return the list of triples (date, amount, currency)
    Ignore lines starting with #.
    Parse the date using datetime.
    Example expenses_string:
    2023-01-02 -34.01 USD
    2023-01-03 2.59 DKK
    2023-01-03 -2.72 EUR
    """
    expenses = []

    for line in expenses_string.splitlines():
        if line.startswith("#"):
            continue
        date, value, currency = line.split(" ")
        expenses.append((datetime.datetime.strptime(date, "%Y-%m-%d"),
                        float(value),
                        currency))

    return expenses

expenses_data = '''2023-01-02 -34.01 USD
2023-01-03 2.59 DKK
2023-01-03 -2.72 EUR'''
```

Writing Assistants; e.g.,

Try the **AI writing** tool with **unparalleled** story smarts.

Sudowrite is the non-judgemental AI writing partner you always wanted.

Try Sudowrite for free

 Sign up with Google

VANITY FAIR

THE NEW YORKER

The Atlantic

The Verge

The Washington Post

WIRED

Meeting Summaries; e.g., Zoom

The image displays a Zoom meeting window with four participants: Brock Davis, Jada Grimes, Rob, and Macy Holloway. A 'Start Summary' button is visible in the bottom right of the meeting controls. A chat window titled 'Meeting Summary with Zoom IQ Created' is open, showing a summary of a 'Marketing Meeting with Max'. The summary text reads: 'Julie met with Rob and Max to discuss the current state of the business and explore future plans to propel the company forward. Julie provided a comprehensive overview of the existing marketing strategy. With compelling data, she showcased the remarkable success the current strategy had achieved to boost brand awareness and drive impressive sales figures. Rob and Max were excited and appreciative of the outcomes that were presented. Rob asked about the current state of the business. Julie is confident that the innovative marketing strategies, commitment to data-driven decision-making, and a firm grasp on the competitive landscape, will put the company in a strong position for continued success. This summary is auto generated and may be inaccurate or misleading. Always check for accuracy. Have feedback?' A 'View More' button is at the bottom of the chat window. A 'FREE TRIAL Meeting Summary with Zoom IQ' notification is also present, stating: 'Meeting Summary with Zoom IQ automatically generates and shares a summary of your meeting.'

Meeting Summaries; e.g., Doctor's Visits

Clinical Documentation: Reimagined

Transform patient conversations into actionable clinical intelligence with DeepScribe Ambient AI.

Learn more

Request a demo

The screenshot displays the DeepScribe interface for a patient named Belinda Hills. The interface is divided into several sections:

- My Schedule:** A calendar view for July 30th (Thursday) showing appointments for Calvin Harvey (11:40pm), Cheryl Bartoletti (11:40pm), Belinda Hills (12:00pm), Phil West Sr. (02:00pm), Sherri Jakubowski (02:30pm), and Amanda Wolff (03:00pm).
- Patient Profile:** Belinda Hills, with a search bar for notes and navigation buttons for previous and next notes.
- Subjective:** A summary of the patient's presentation, including DOB, Existing Medications, Poor Appetite, and Hours of sleep. The main text describes the patient's lower back pain, its characteristics (dull, radiating down the leg), and the doctor's examination and recommendations.
- AI Notable:** A list of key points extracted from the transcript, such as "The pain? It's probably about lower left, right above my hip bone." and "It's a pretty dull pain for the most part, but if I move a certain way or a jolt out of bed, it's a pretty sharp pain."
- Transcript:** A list of timestamps and corresponding text from the patient-doctor conversation, including "Yeah, actually it is." and "Okay, and what would you describe your pain level on a scale from one to ten?"

Self-Driving Vehicles; e.g.,



Waymo Self-Driving Taxi

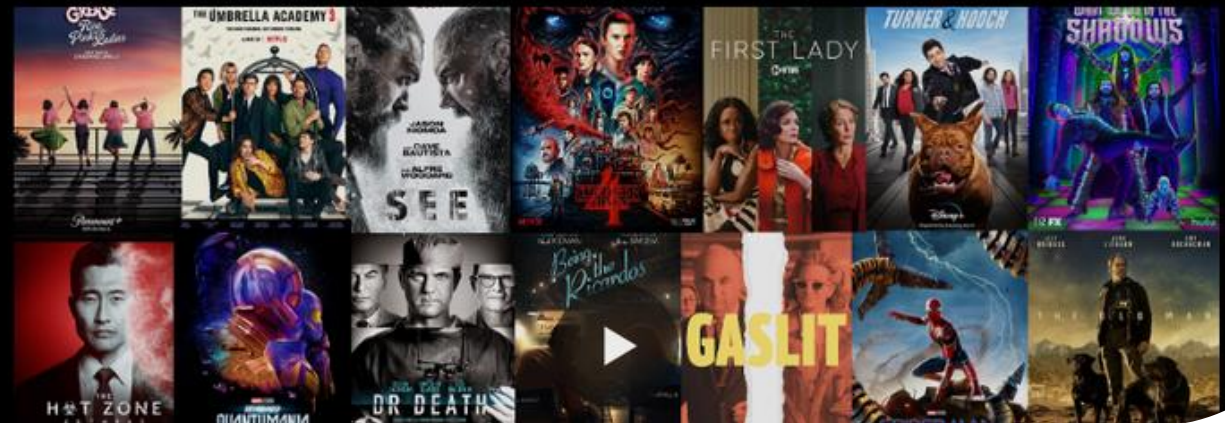
Appearance Editor; e.g., for Movies

V A N I T Y
AI

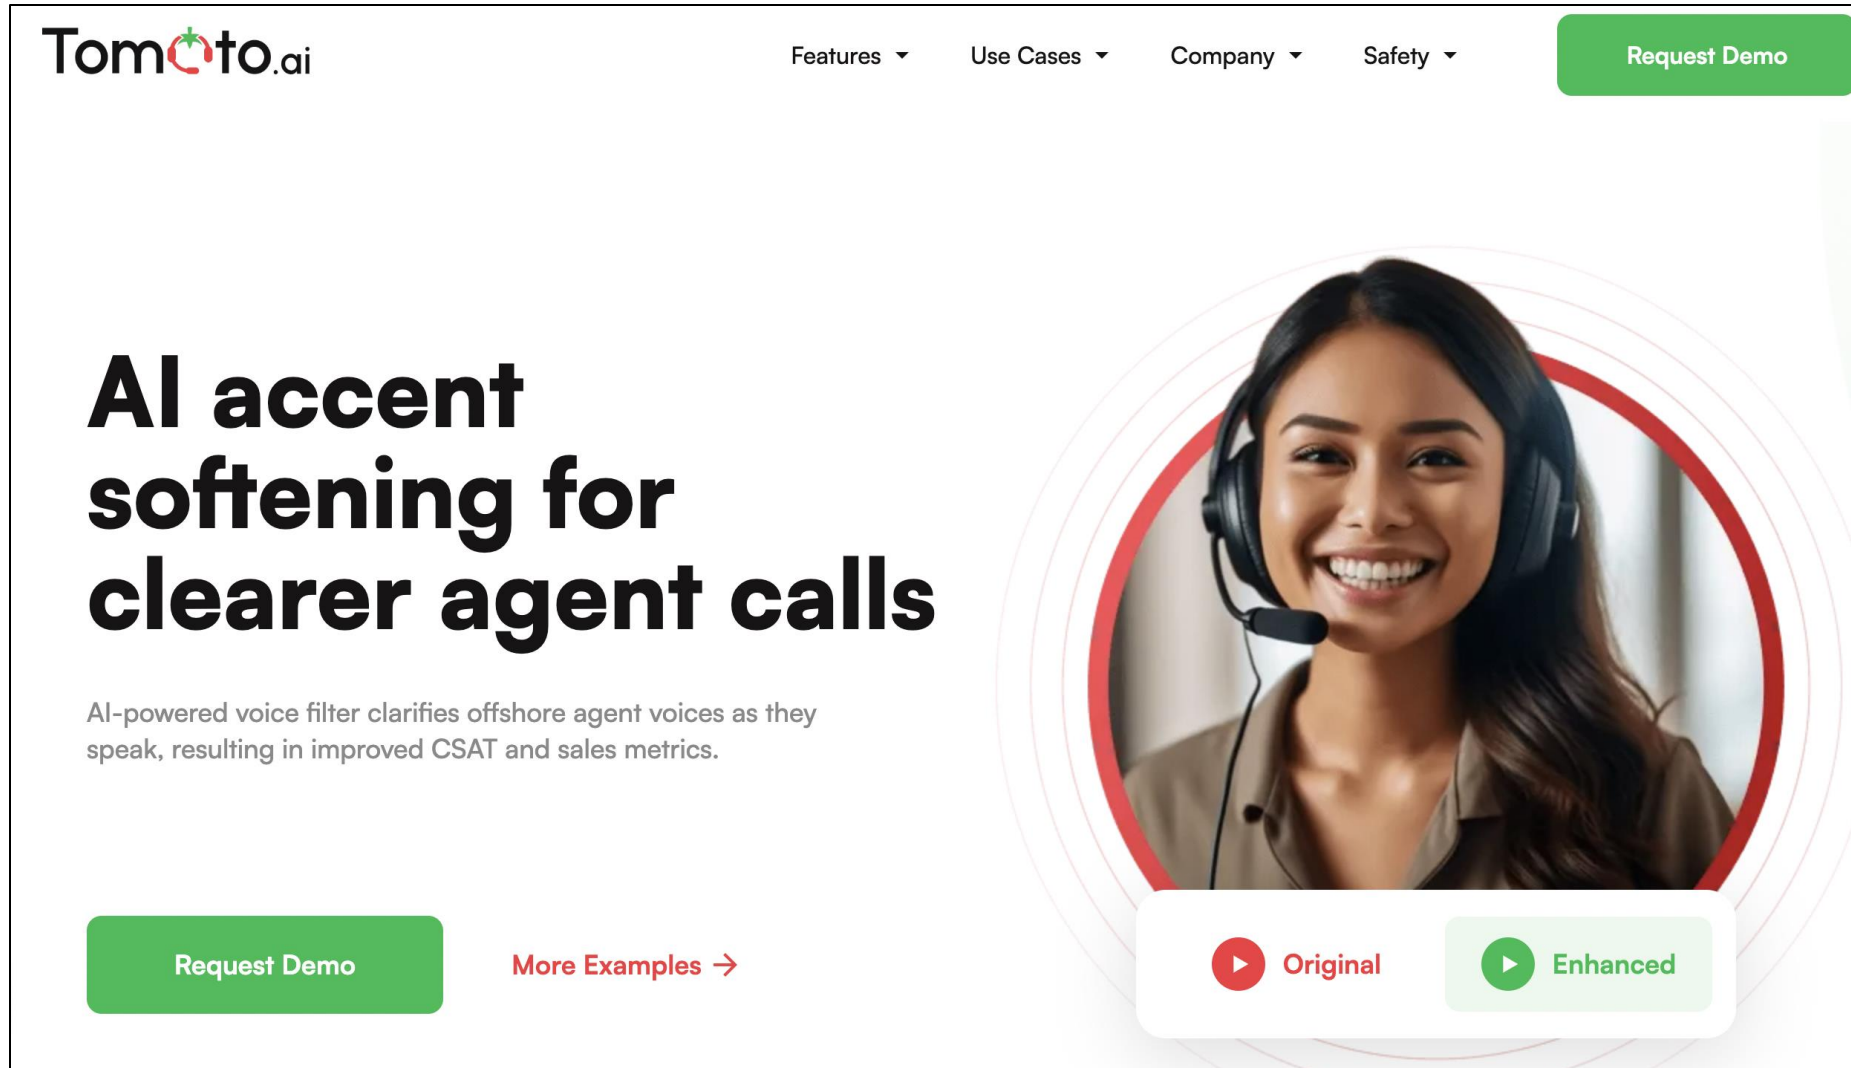
Book a Demo

The world's first end-to-end AI solution for Hollywood VFX.

Vanity AI is a production-ready solution that empowers VFX teams and Hollywood to deliver large volumes of high-end 2D aging, de-aging, cosmetic, wig, and prosthetic fixes. The technology is up to 300 times faster than traditional VFX pipelines, significantly more cost-effective, and has no



Accent Editor; e.g., for Customer Support



The screenshot displays the Tomato.ai website interface. At the top left is the logo "Tomato.ai" with a tomato icon. To the right are navigation links: "Features", "Use Cases", "Company", and "Safety", each with a dropdown arrow. A green "Request Demo" button is positioned in the top right corner. The main content area features a large heading: "AI accent softening for clearer agent calls". Below this heading is a sub-headline: "AI-powered voice filter clarifies offshore agent voices as they speak, resulting in improved CSAT and sales metrics." On the right side of the main content is a circular image of a smiling female customer support agent wearing a headset. Below the image are two buttons: "Original" with a red play icon and "Enhanced" with a green play icon. At the bottom left of the main content area, there is a green "Request Demo" button and a red "More Examples ->" link.

Tomato.ai

Features ▾ Use Cases ▾ Company ▾ Safety ▾

Request Demo

AI accent softening for clearer agent calls

AI-powered voice filter clarifies offshore agent voices as they speak, resulting in improved CSAT and sales metrics.

Request Demo

More Examples →

Original Enhanced

Music Composer; Demo with Suno



Video Generation; Demo with DeepAI

AI Video Generator

This AI video generator converts text to video, or image to video. Use it to convert images to animations. Or, Create videos from scratch from text!

Prompt

Image

Inspirational video about a course on neural networks and deep learning



And Many More Applications...

*Also available at TAAFT.com

THERE'S AN AI FOR THAT^{®*}

26,910 AIs for 🤖 15,911 tasks and 4,958 jobs.


Sponsor: CleeAI (Sales Agent)

Find AIs using AI



#1 AI aggregator. Updated daily. Used by 40M+ humans.

New

 **VOTE NOW!:** The Best AI Tool of 2024

Home

For You

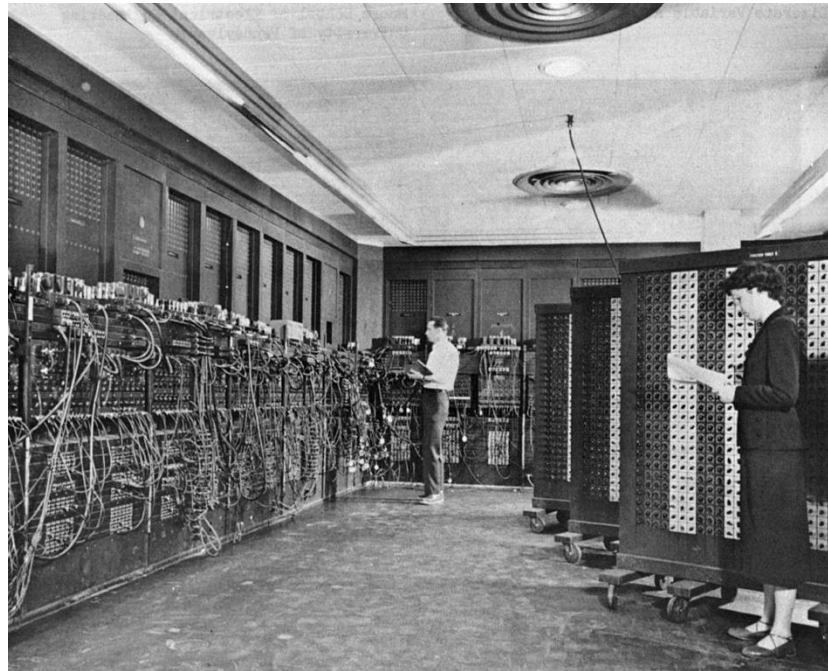
Today's Topics

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

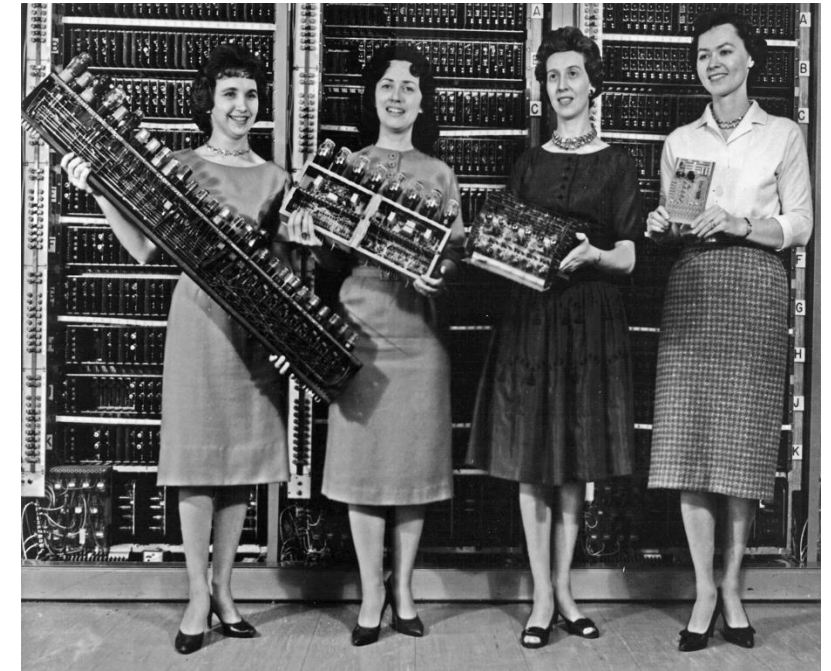
Origins: Computers

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

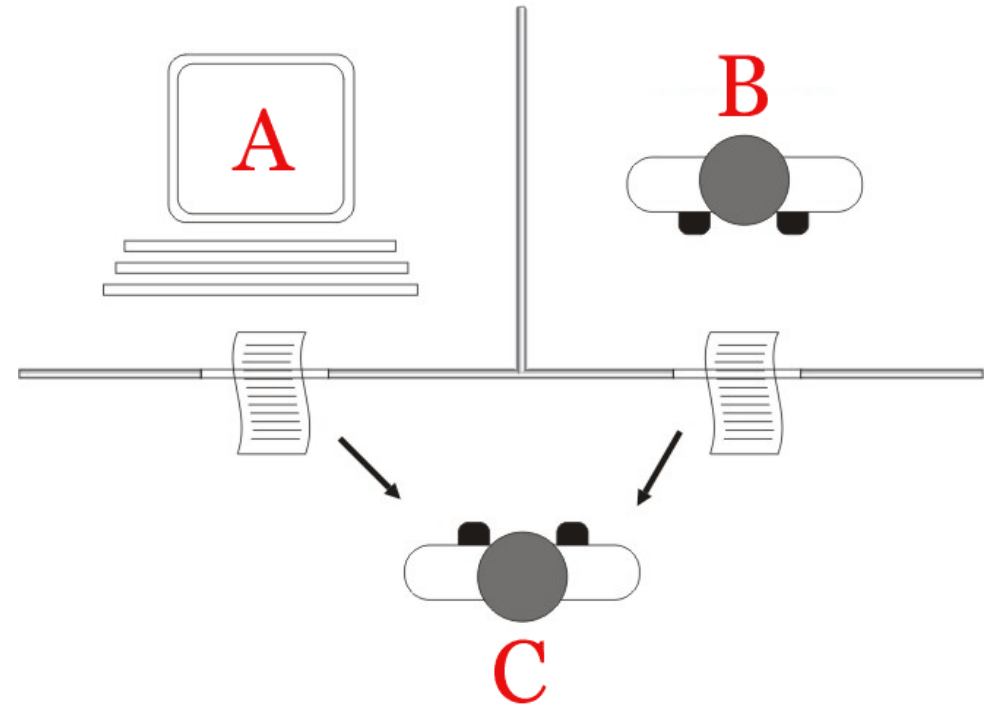


First programmable
machine

Turing test



Alan Turing
(1912-1954)



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

Origins: Conceptual Framework

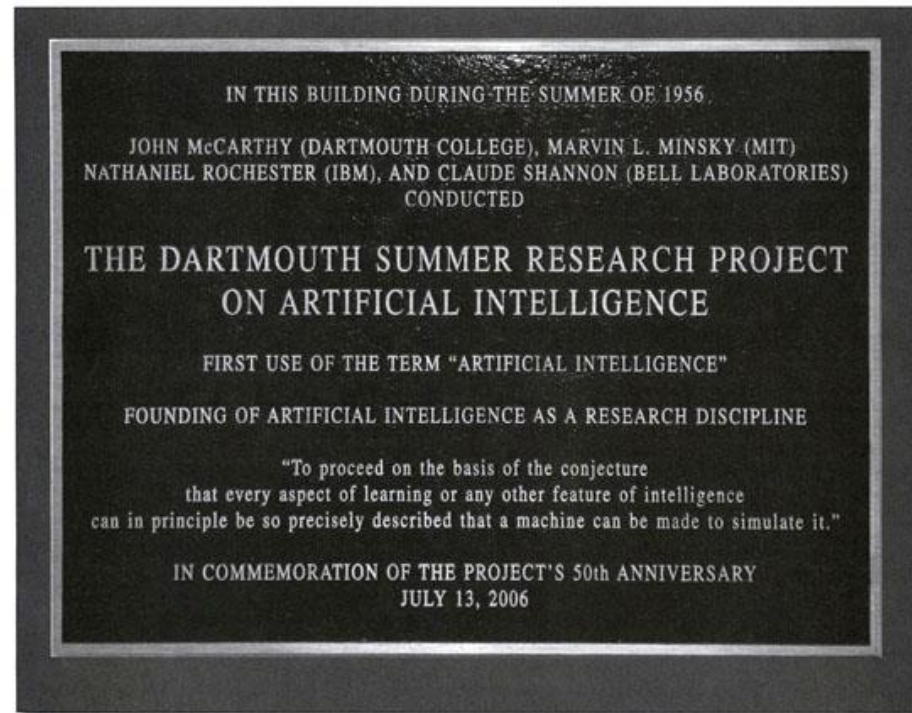
1945 1950 1956



First programmable
machine

AI birth

Turing test



“Artificial intelligence” established as a field at a workshop

Origins: Conceptual Framework

1945 1950 1956



First programmable
machine

AI birth

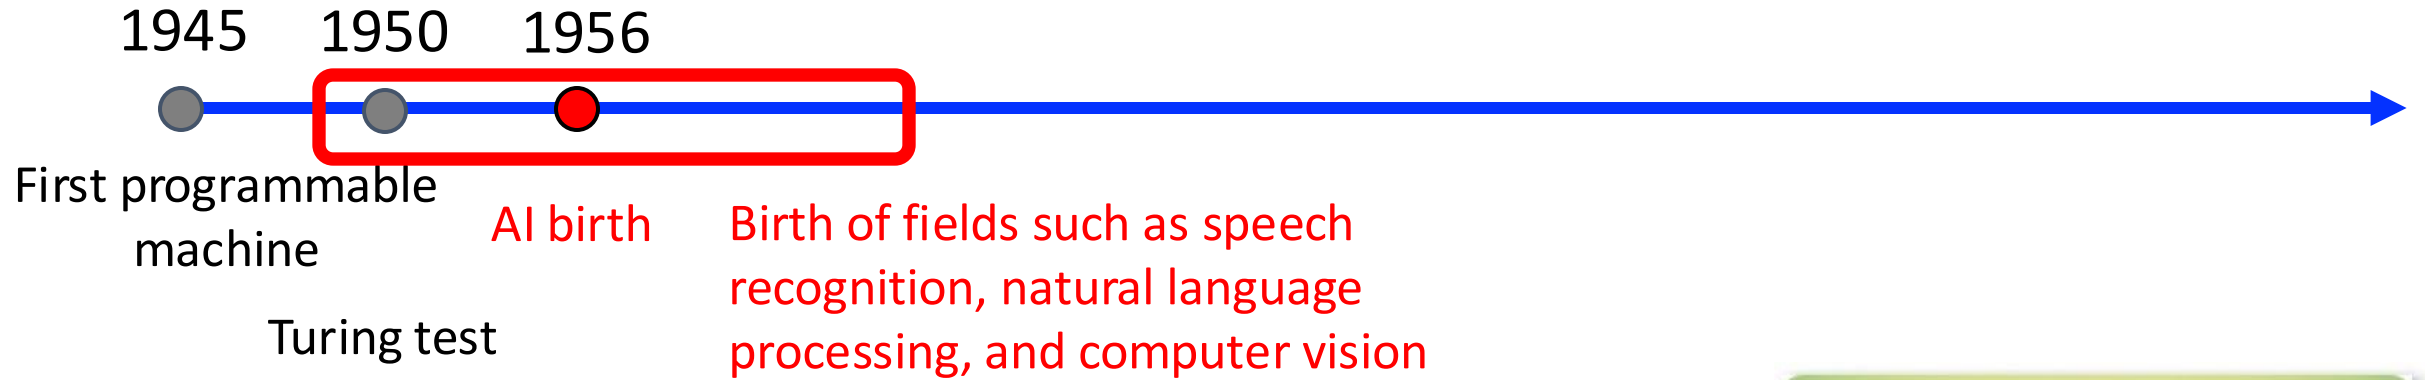
Turing test

IN THIS BUILDING DURING THE SUMMER OF 1956

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

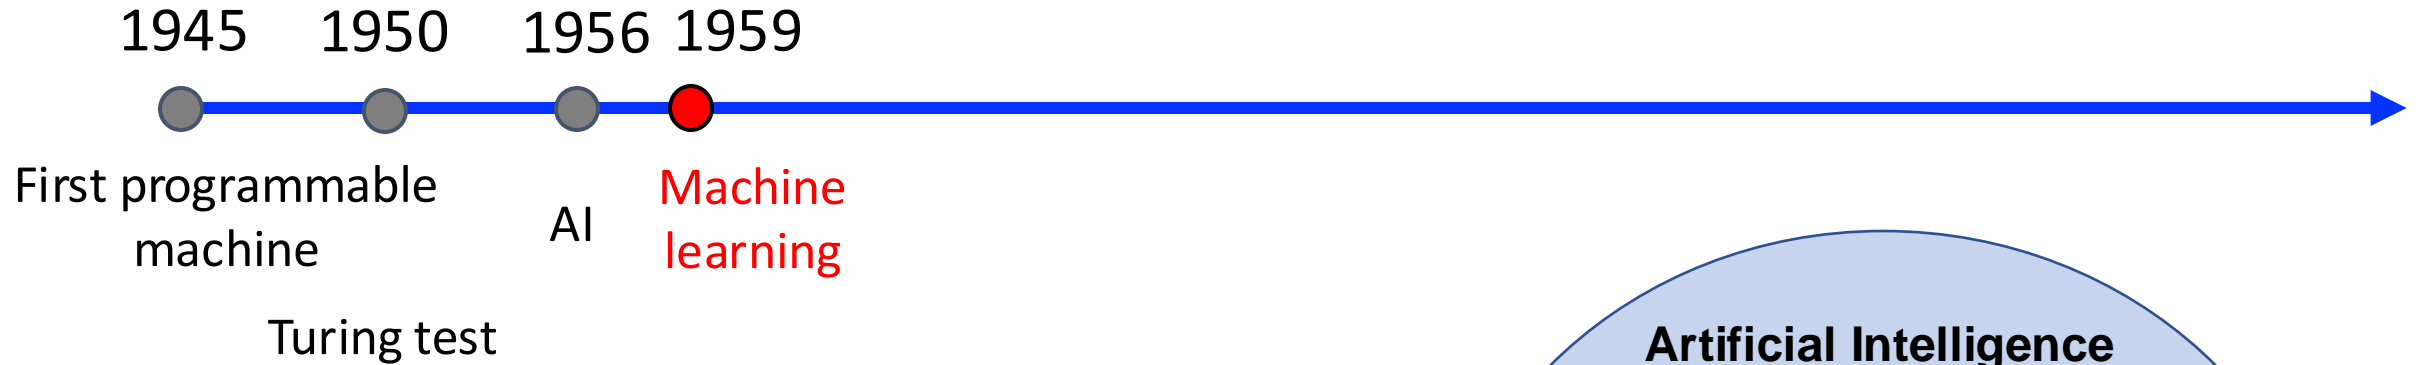
Origins: Conceptual Framework



What human intelligence might computers imitate?

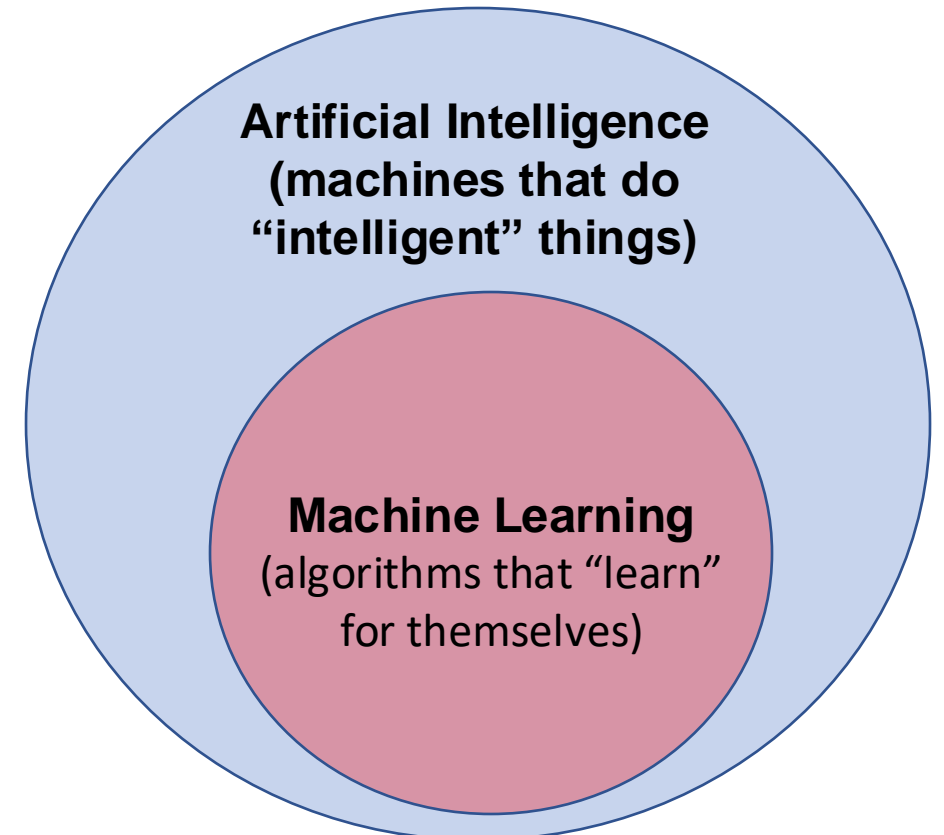


Origins: Conceptual Framework



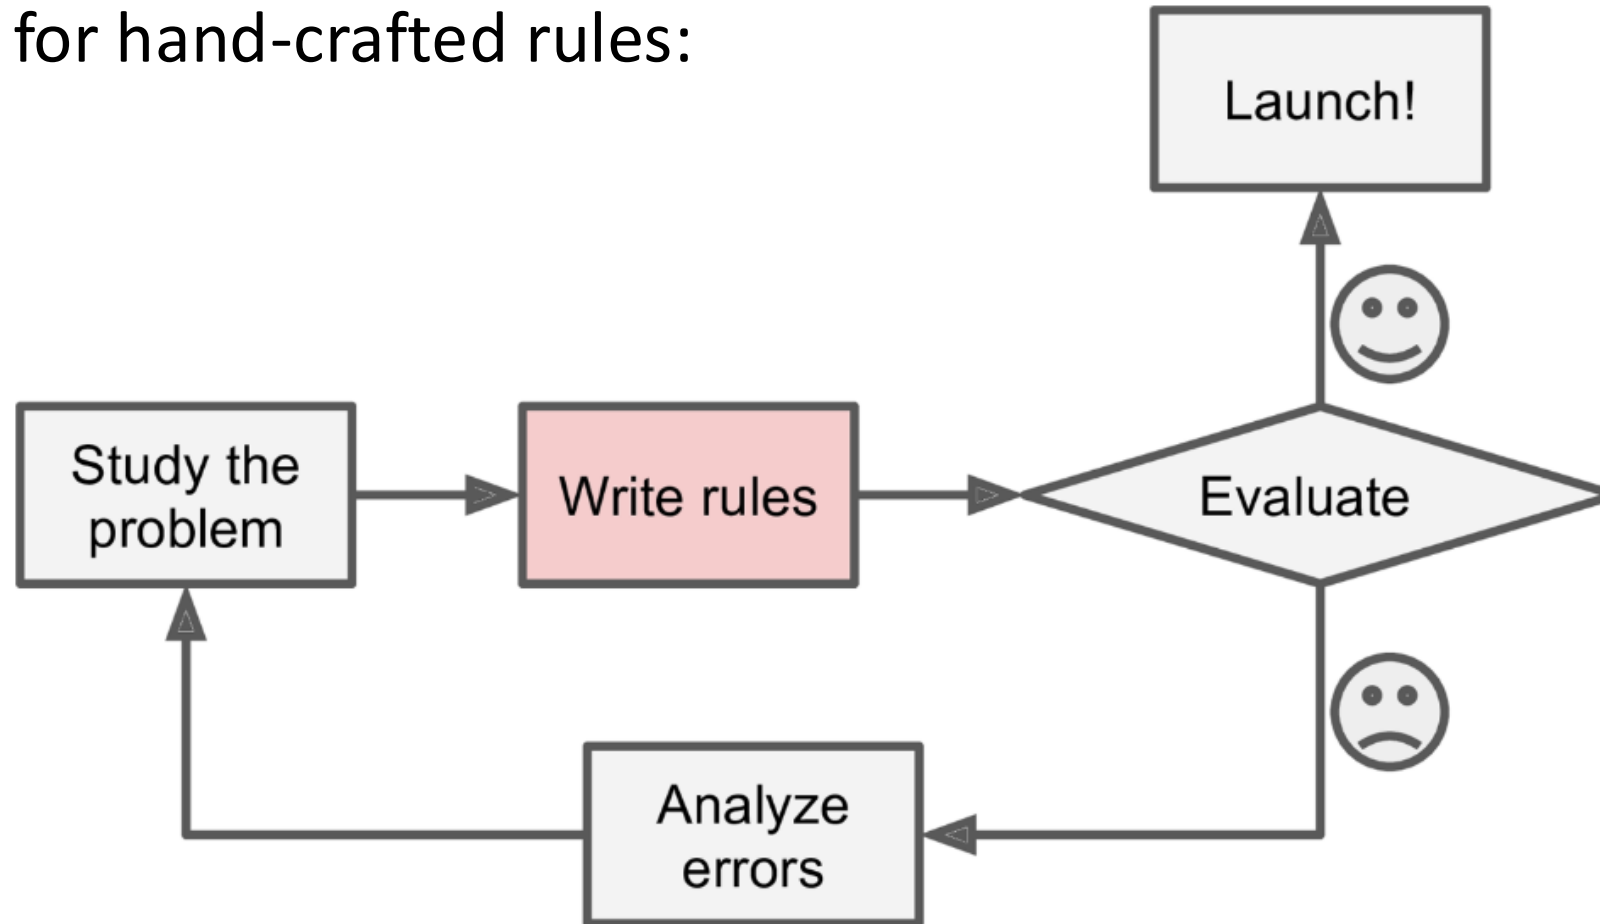
AI researcher Arthur Samuel coins the term “machine learning” as:

“Field of study that gives computers **the ability to learn without being explicitly programmed.**”



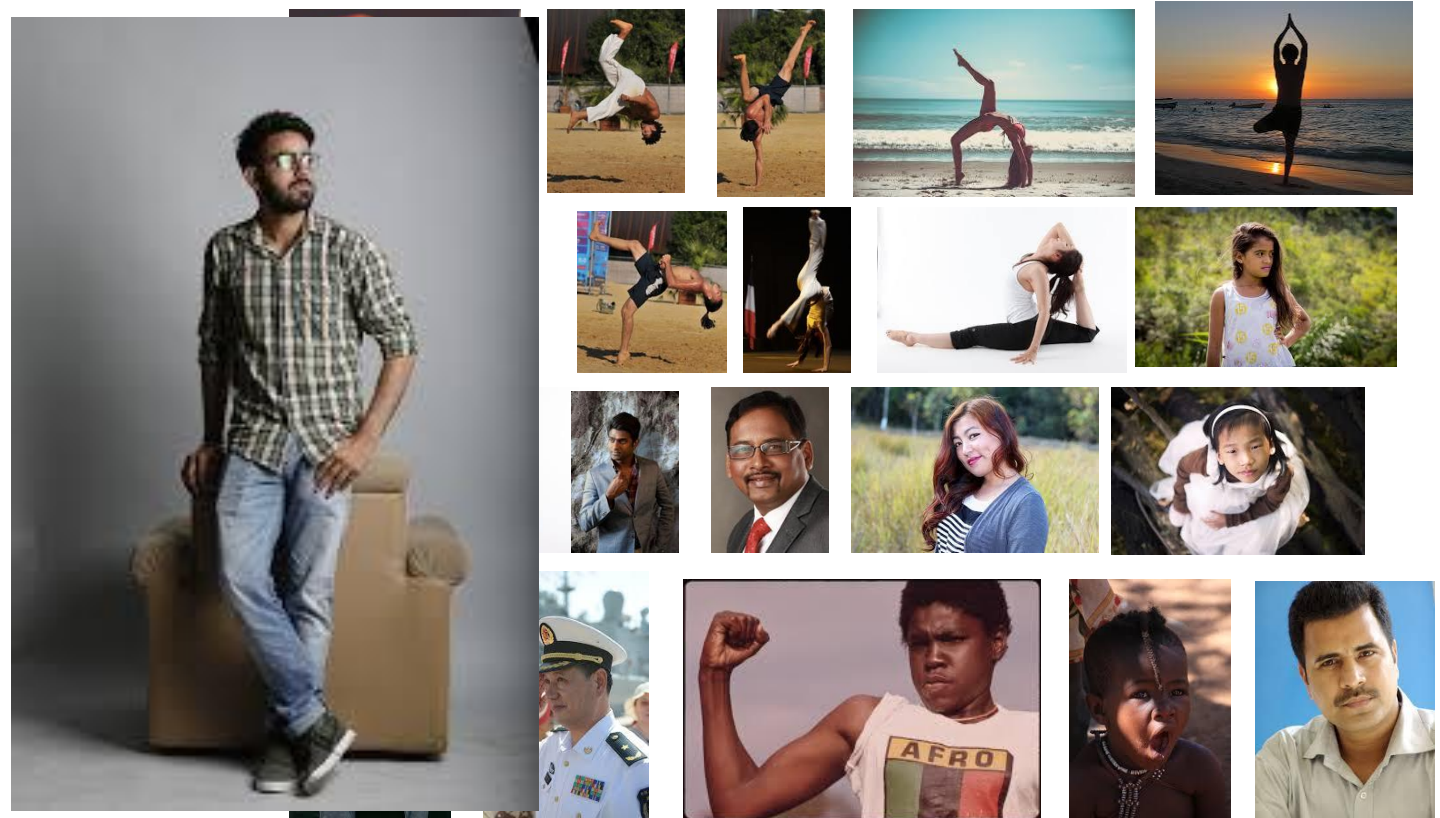
Motivation for Machines that “Learn”

- Process for hand-crafted rules:



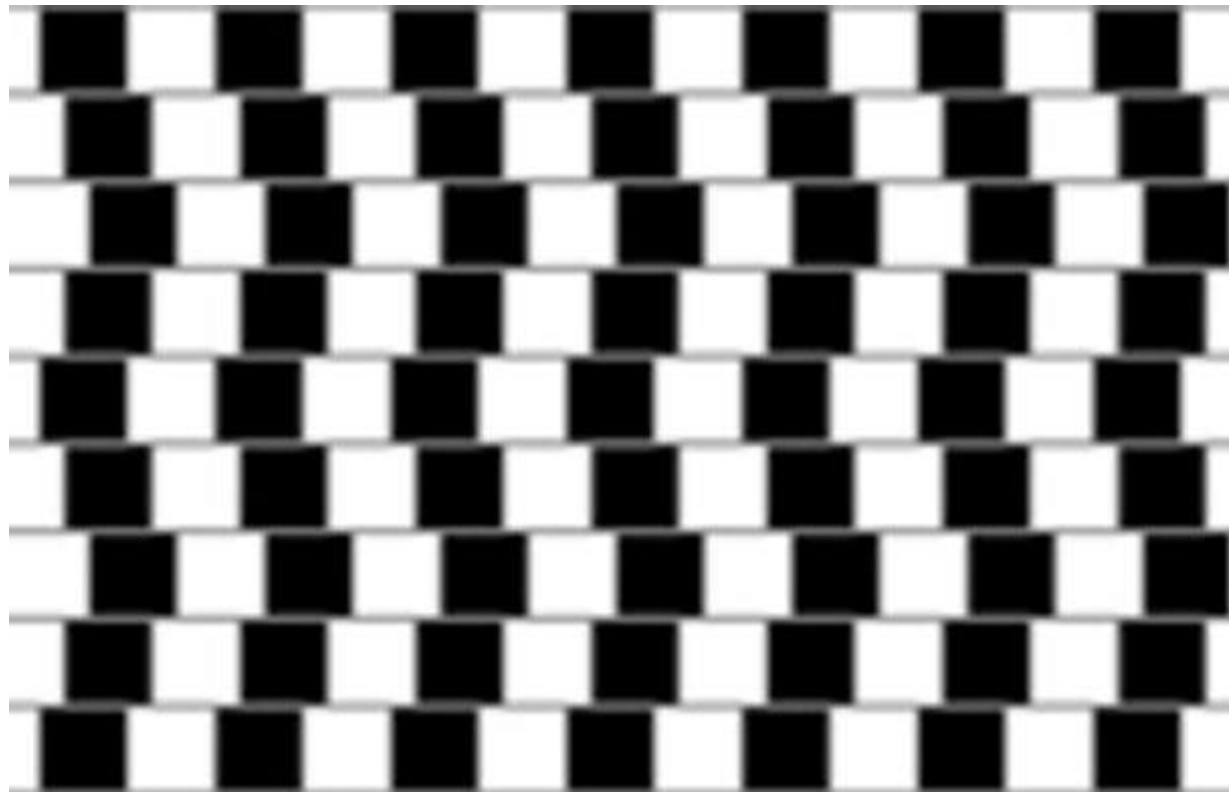
Motivation for Machines that “Learn”

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



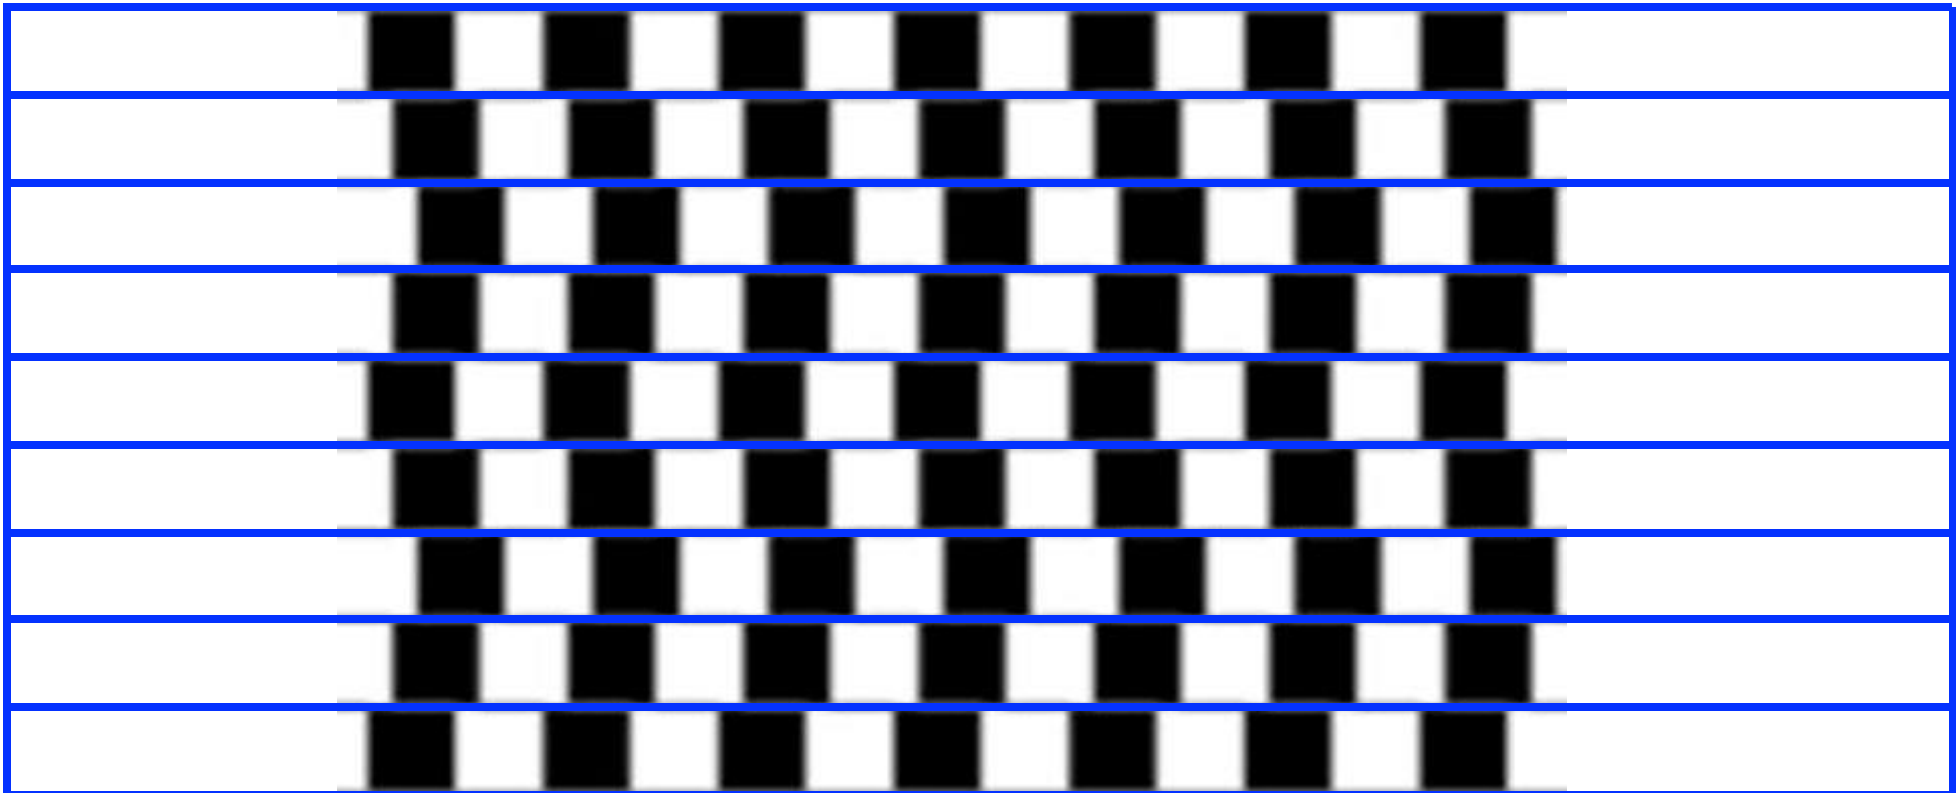
Motivation for Machines that “Learn”

e.g., are the horizontal lines parallel?



Motivation for Machines that “Learn”

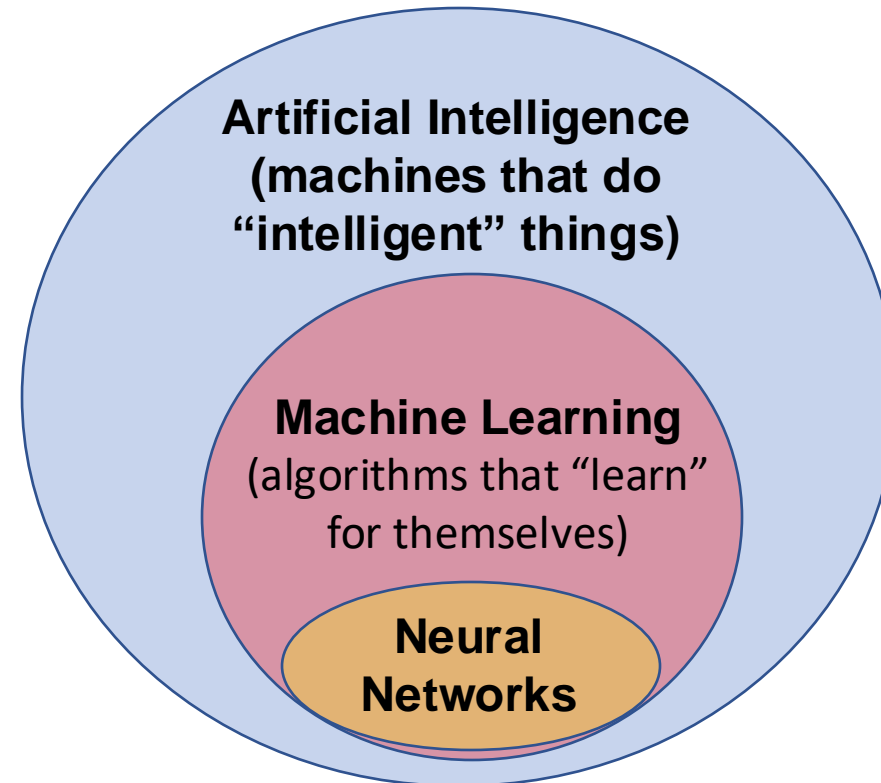
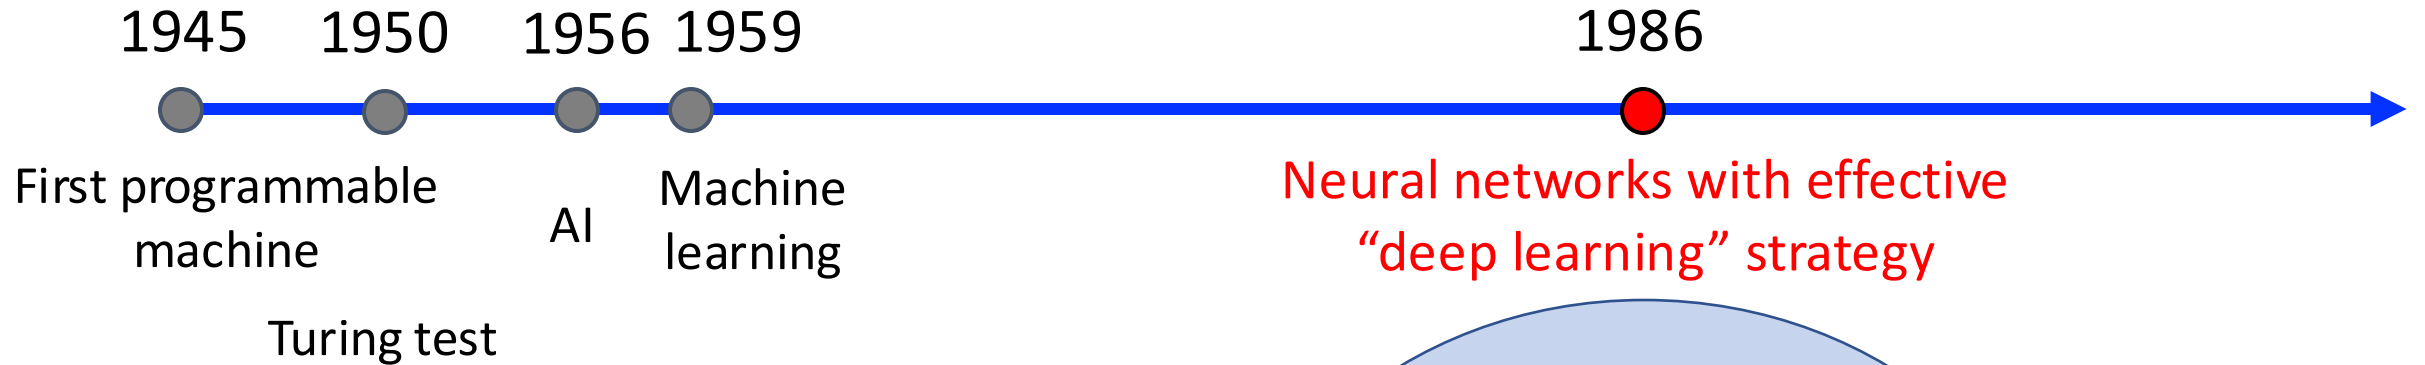
e.g., are the horizontal lines parallel?



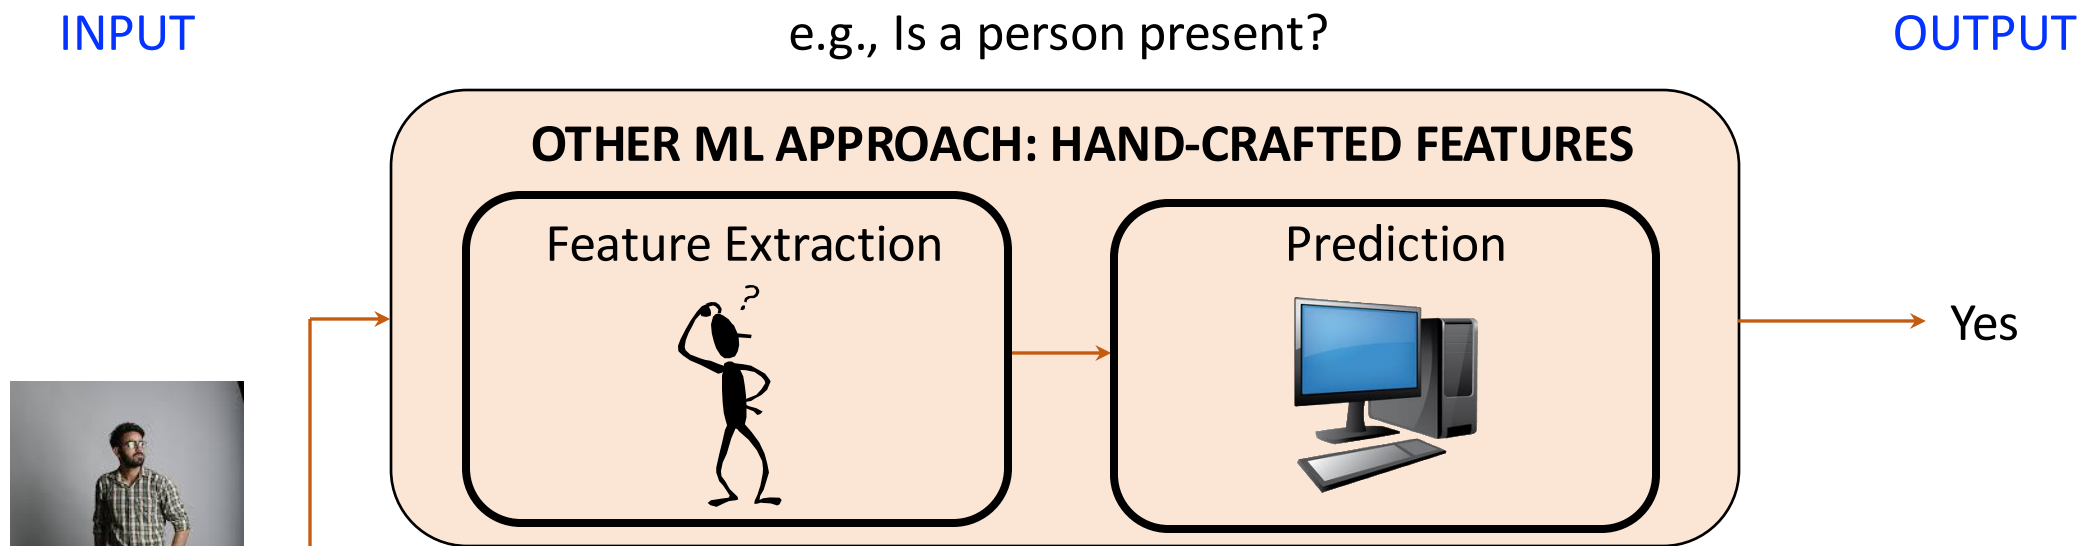
Motivation for Machines that “Learn”

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

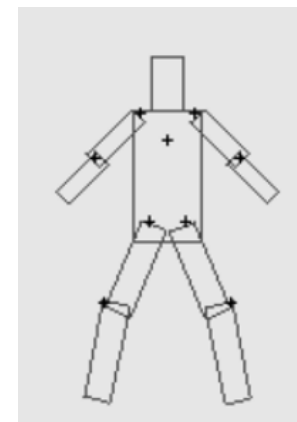


Motivation for Neural Networks (NNs) Over Other Machine Learning (ML) Approaches

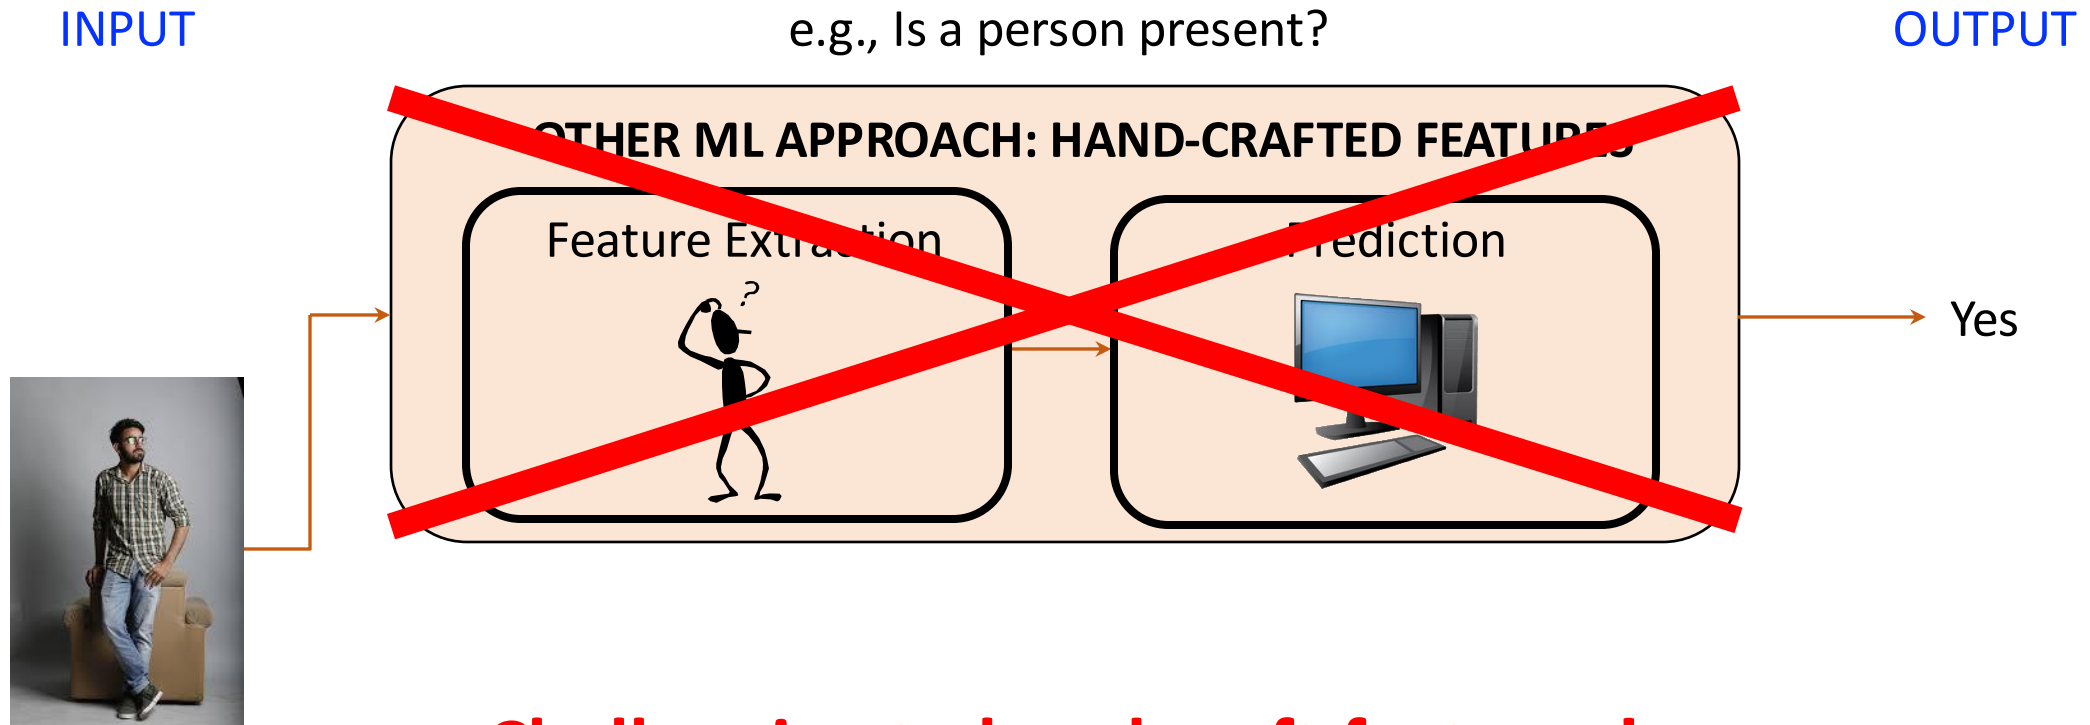


Features that may help predict yes/no:

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



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



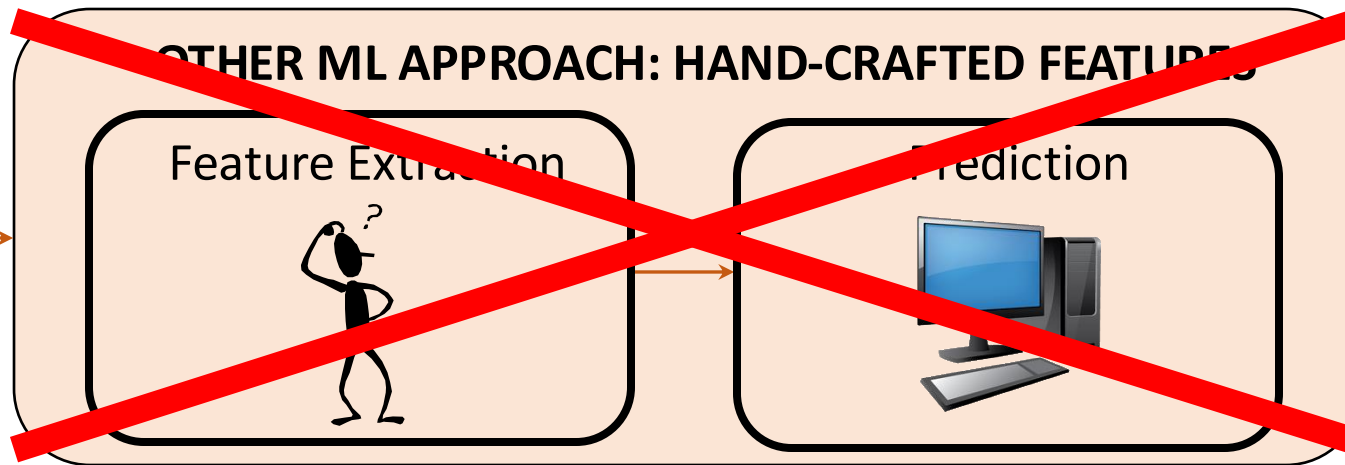
Challenging to hand-craft features!

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

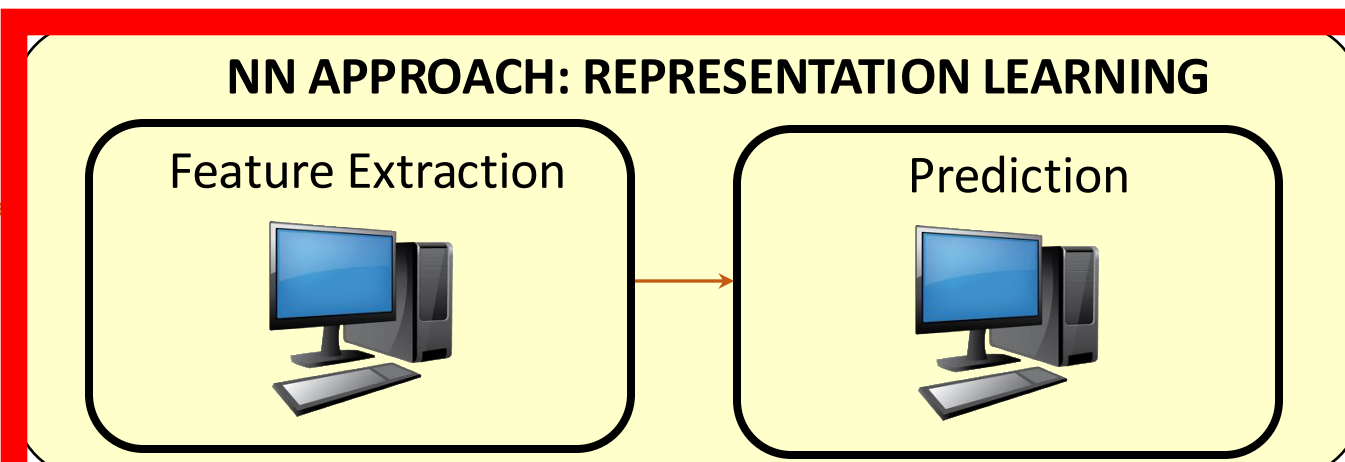
INPUT

e.g., Is a person present?

OUTPUT



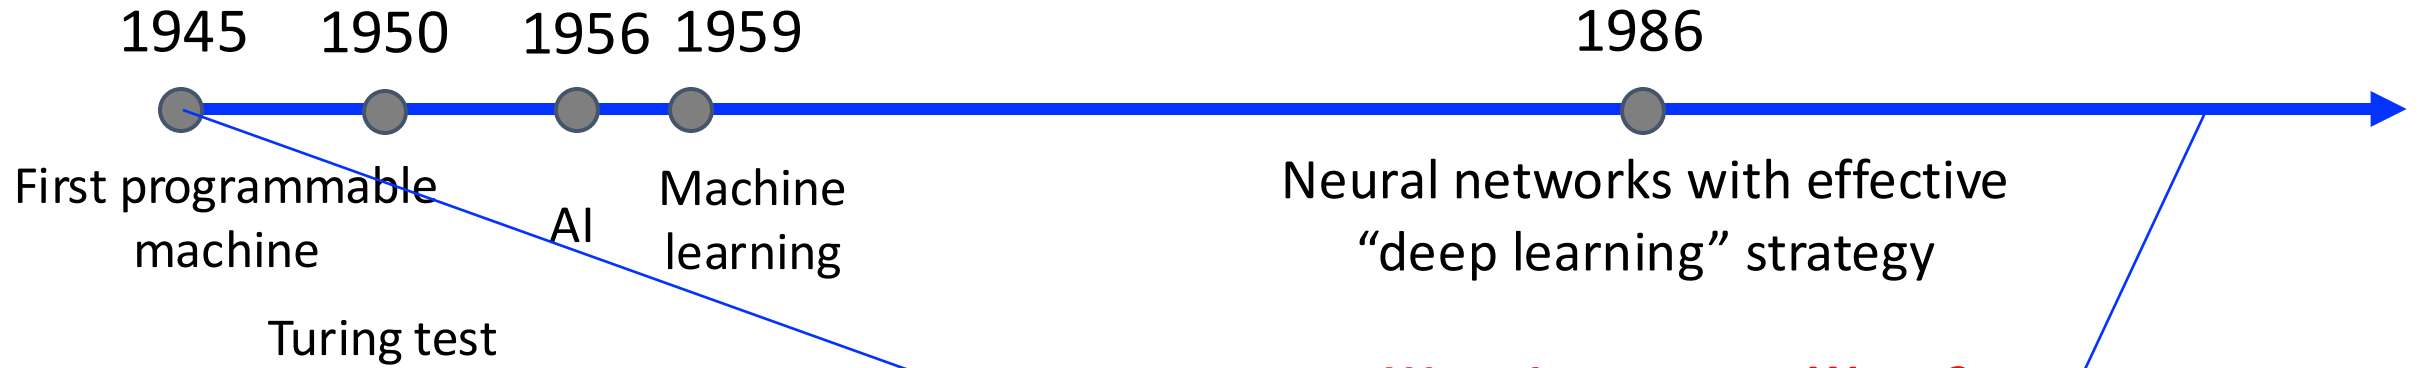
Yes



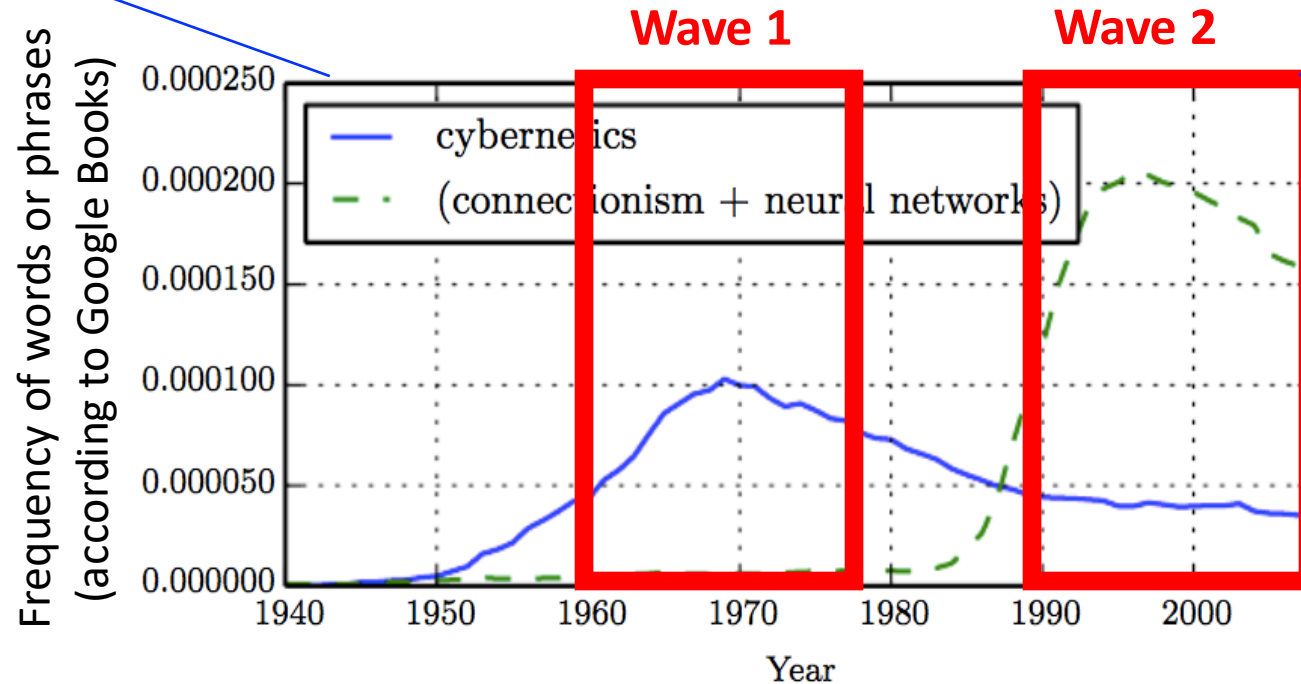
Course scope

Yes

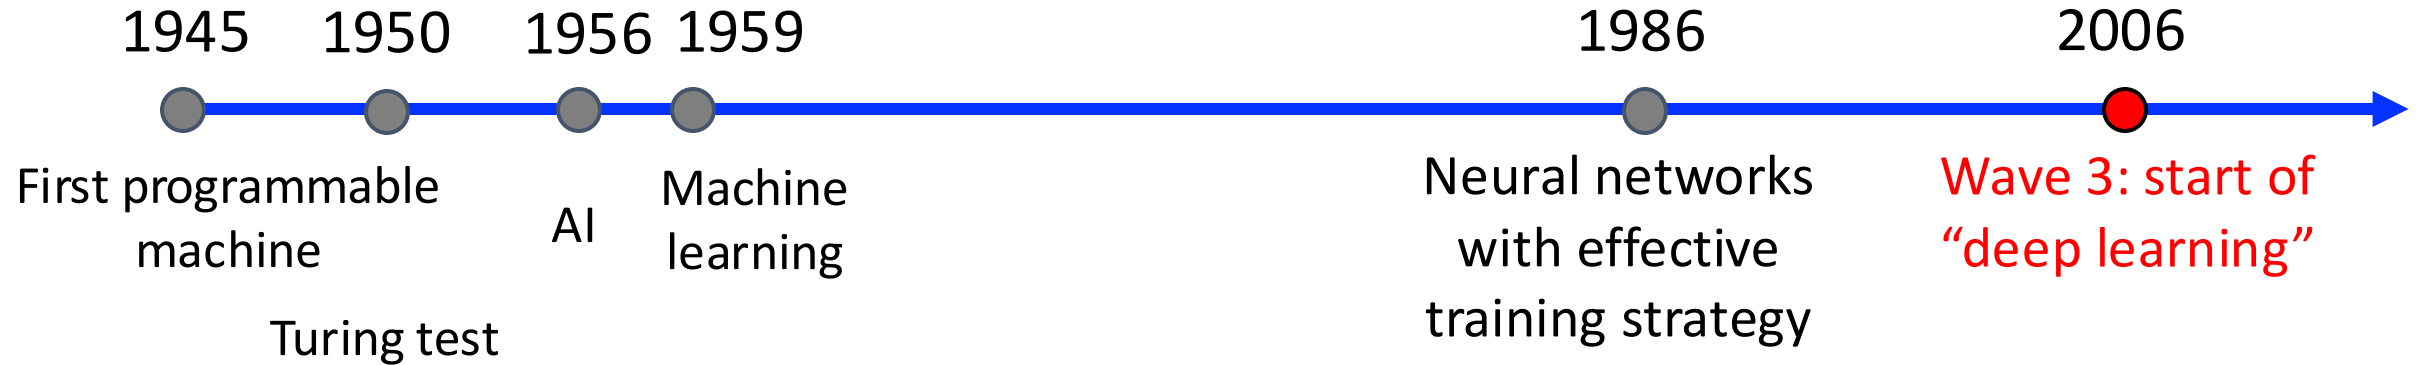
Origins: Rises/Falls of Neural Network Popularity



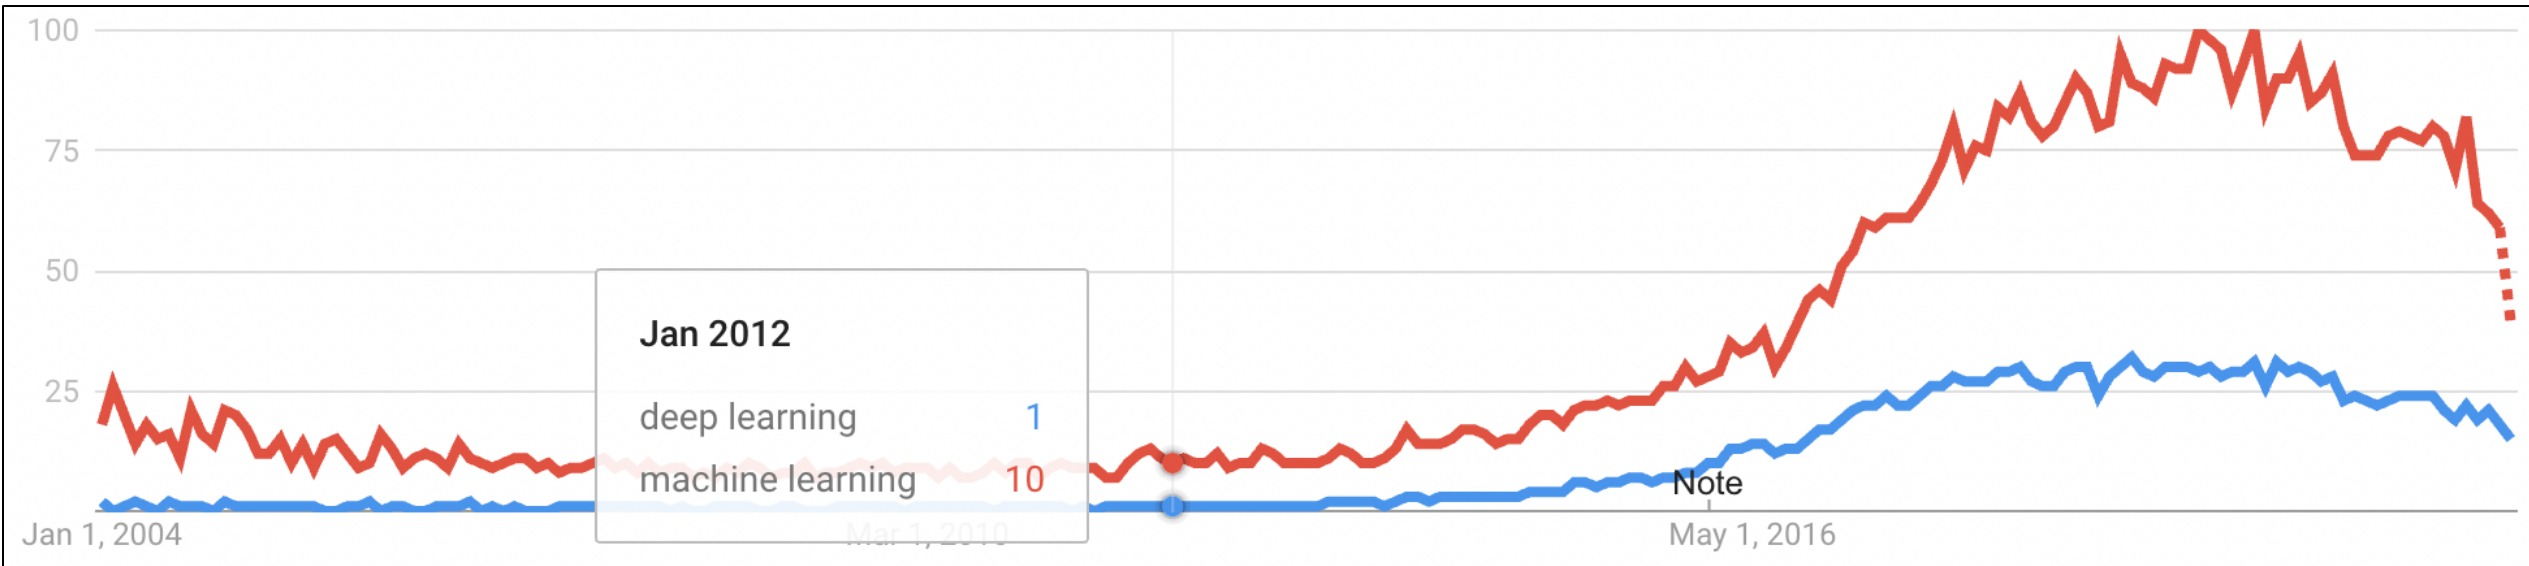
Neural networks have been called many names:



Origins: Rises/Falls of Neural Network Popularity



(Google Trends)



Machine learning popularity has paralleled rise of deep learning popularity

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”

Today's Topics

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

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

Neural Networks: Key Ingredients for Success

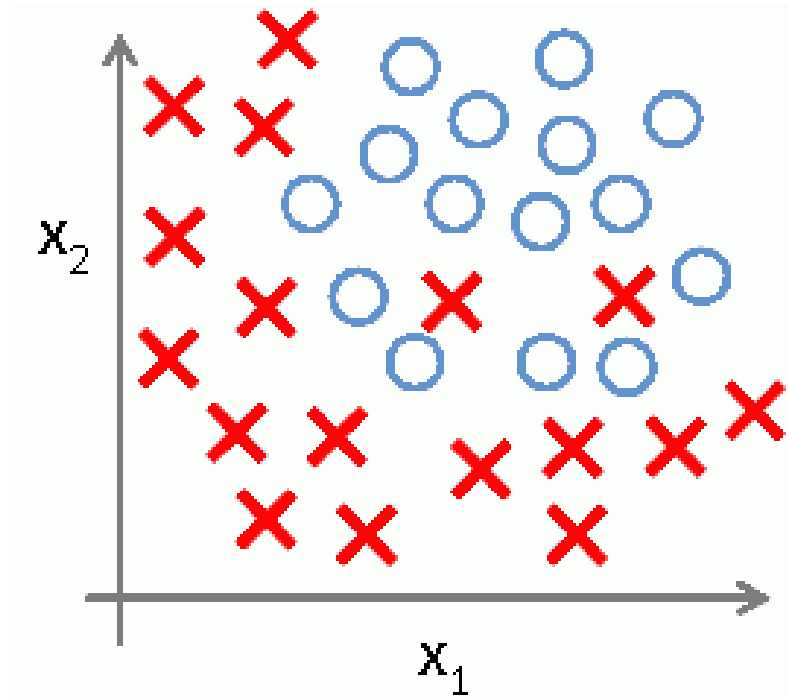
An **algorithm** learns from data on a processor the patterns that will be used to make a prediction



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

Algorithm Design: Model-Based

- e.g., create model to separate x from o

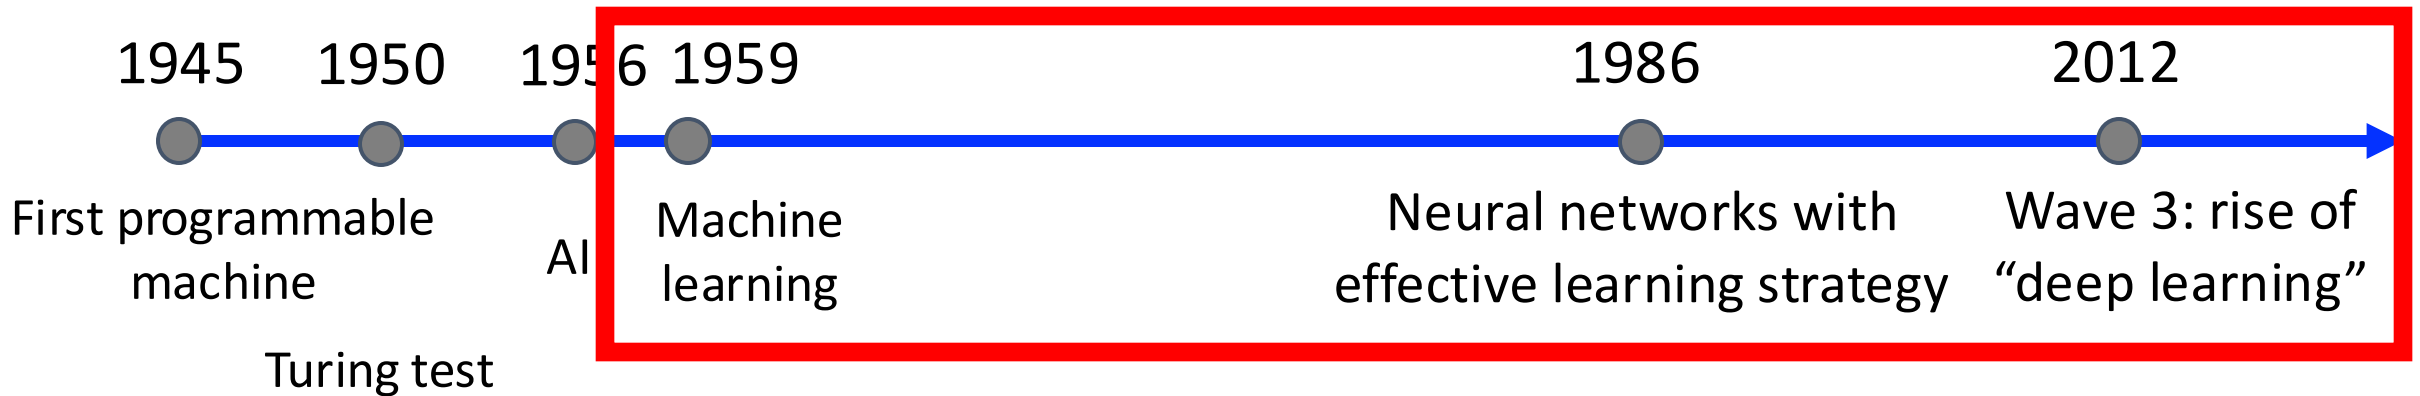


Class volunteer:

- 1) Draw a **straight line (linear equation)**
- 2) Draw a **parabola (quadratic equation)**
- 3) Draw any curve

Models with increasing
representational capacity

Algorithm Scope for Course: Last 65 Years



We will cover the history of deep learning because modern algorithms use techniques developed over the past 65 years.

Neural Networks: Key Ingredients for Success

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



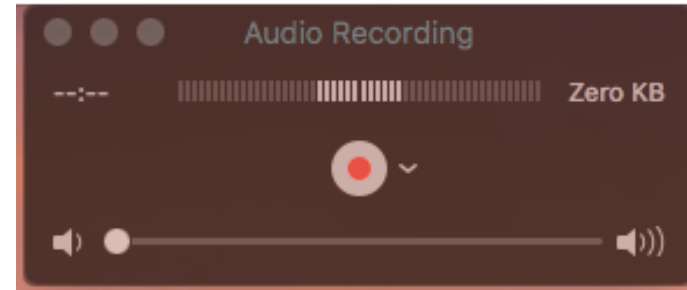
Analogous to a Love Story of Partnering Up and Road Tripping Somewhere

Data Types: What a Machine Learns From?

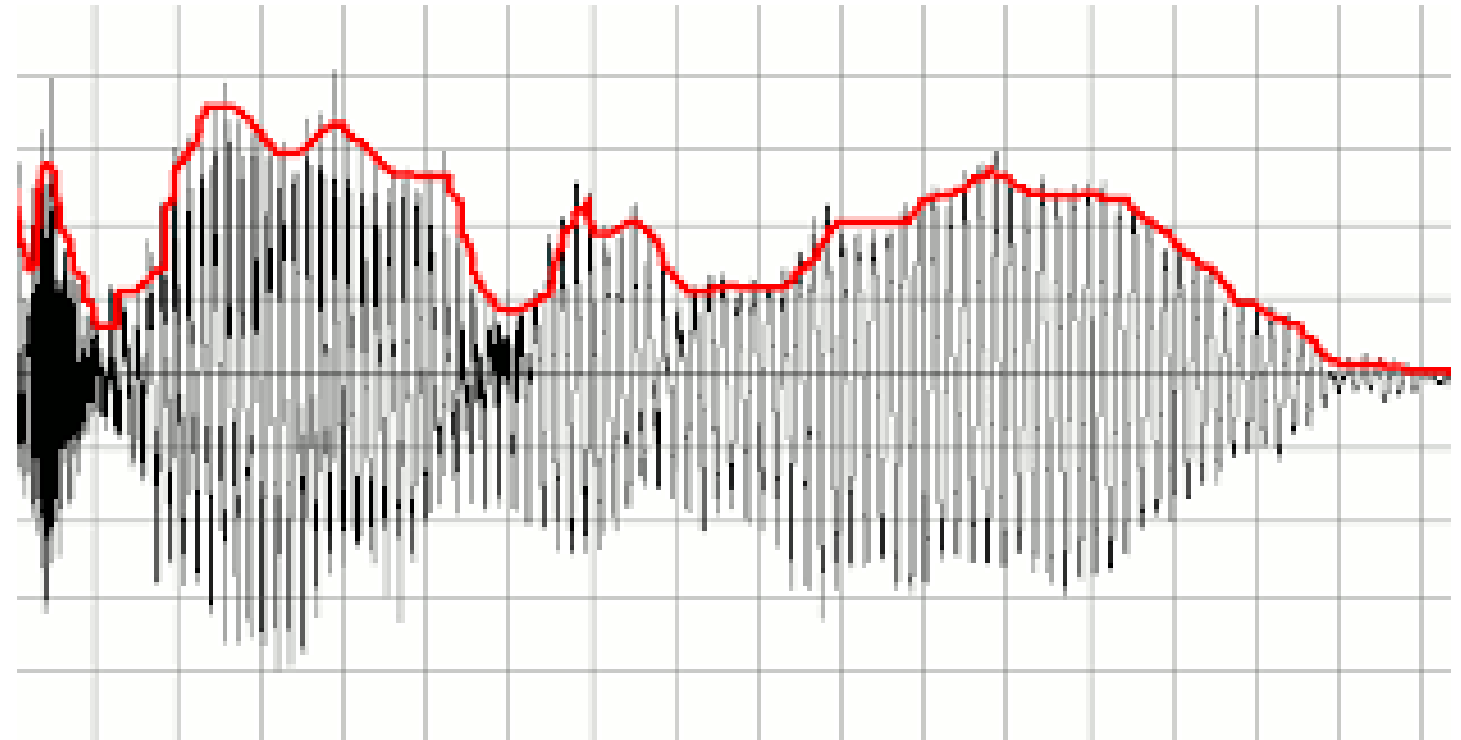


- Audio

- Input?



e.g.,

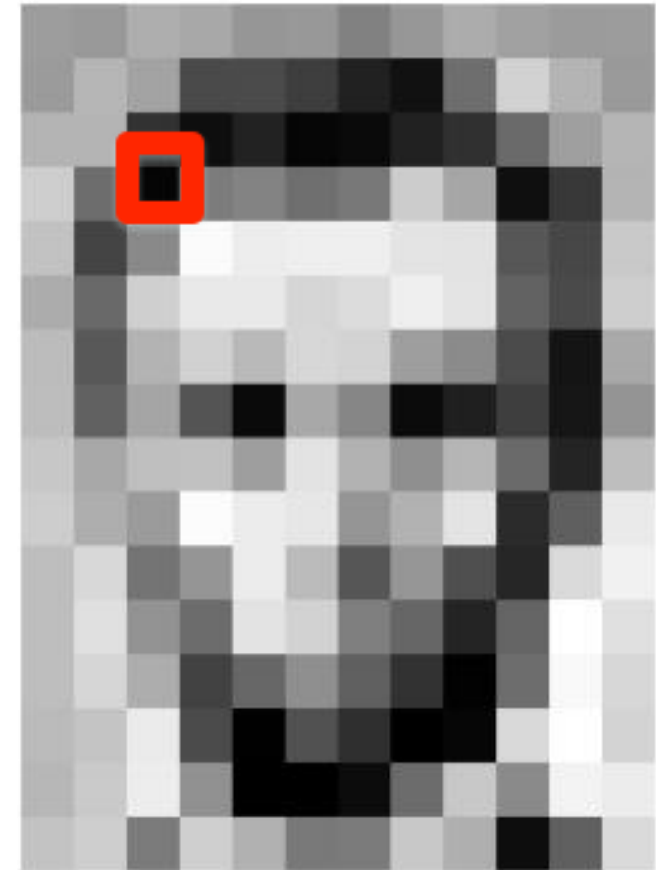


Data Types: What a Machine Learns From?



- 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	105	5	14	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?

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	124	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

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	124	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

Time 1

1 hour

Analogous to:







Data Types: What a Machine Learns From?

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

e.g.,

Confidential letter sh 

 David-Khoza@mmoscacs.v.com
to 

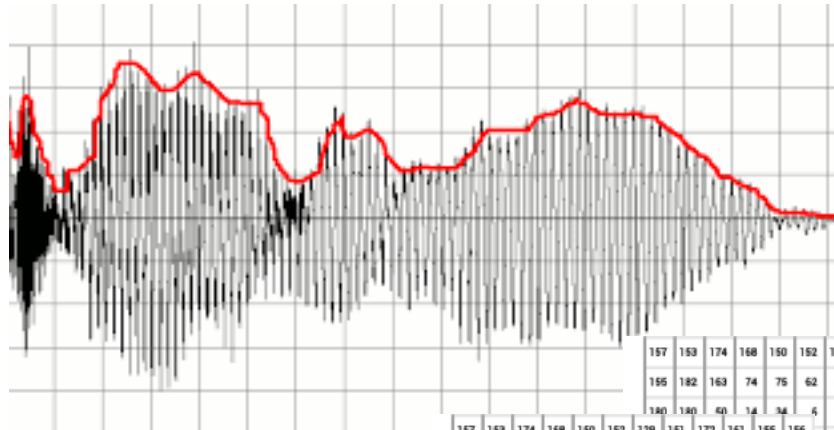
2 Attachments





Data Types: What a Machine Learns From?

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



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	124	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	168	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

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	124	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	168	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

Confidential letter sh

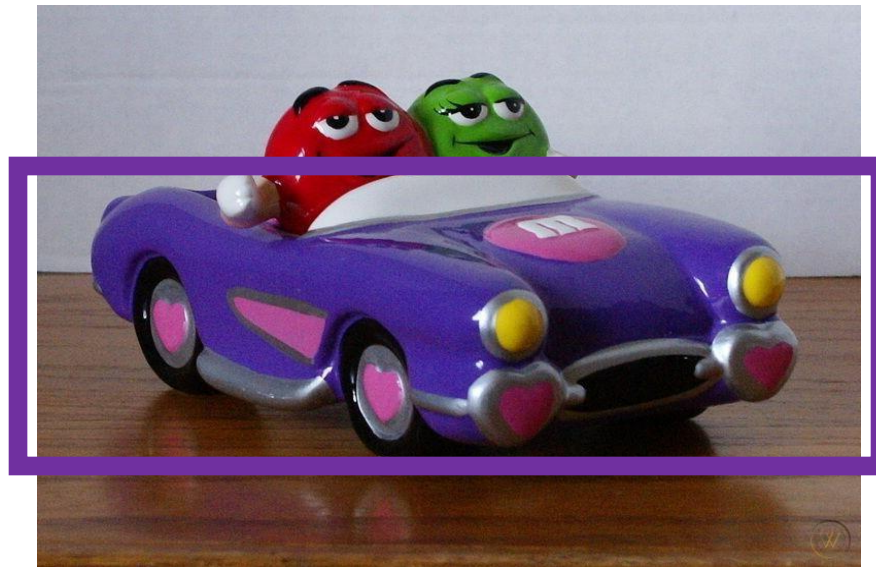
? David-Khoza@moscacs.com
to

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

Neural Networks: Key Ingredients for Success

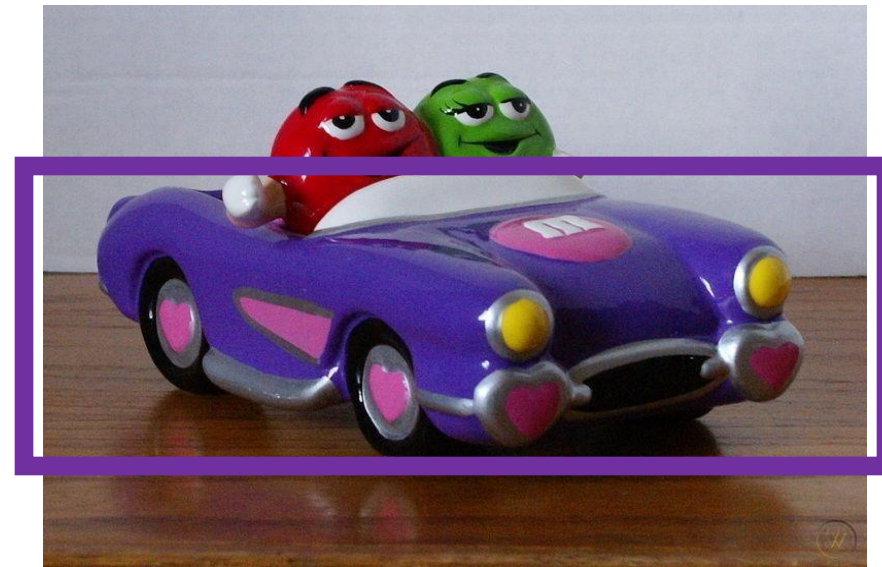
An algorithm learns from data
on a **processor** the patterns that
will be used to make a prediction



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

Neural Networks: Key Ingredients for Success

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



Summary

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

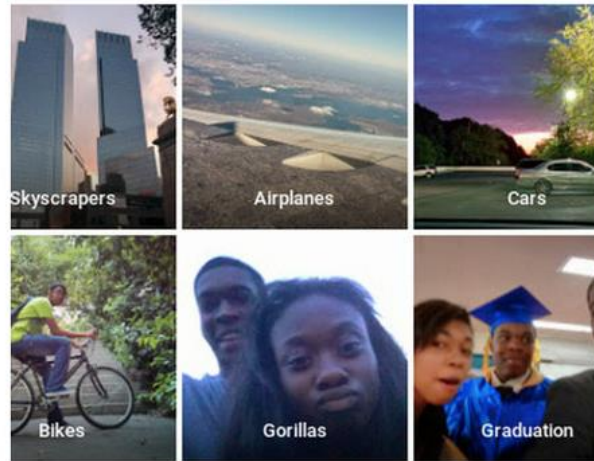


Analogous to a Love Story of Partnering Up and Road Tripping Somewhere

Key Challenge: Where Will You Go?



Google Photos, y'all fucked up. My friend's not a gorilla.



Using Twitter to call out Google's algorithmic bias

<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**

Course Objectives

- Understand the key concepts for designing deep learning models:
 1. Characterize the key architectures used in designing neural networks
 2. Recognize techniques used to train and evaluate deep learning models
 3. Identify strengths and weaknesses of different neural network architectures

Course Objectives

- Apply deep learning models to perform various AI tasks:
 1. Develop deep learning models from scratch
 2. Experiment with established deep learning libraries
 3. Evaluate deep learning algorithms for tasks in various application domains, including for analyzing text and images

Course Objectives

- Conduct and communicate about a novel project:
 1. Propose a novel project idea
 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 Resources

- Website with Syllabus:
 - <https://dannagurari.colorado.edu/course/neural-networks-and-deep-learning-spring-2025/>
- Lecture Slides: hyperlinked from course website
- Lecture Recordings: available on Canvas immediately after each lecture

Q&A: “How does this course differ from other courses?”

- Other courses also cover neural networks and deep learning, including about:
 - Computer vision
 - Natural language processing
- Unique benefit of this course:
 - Focus is on the fundamentals, analogous to learning how to build a car instead of only on how to drive a car

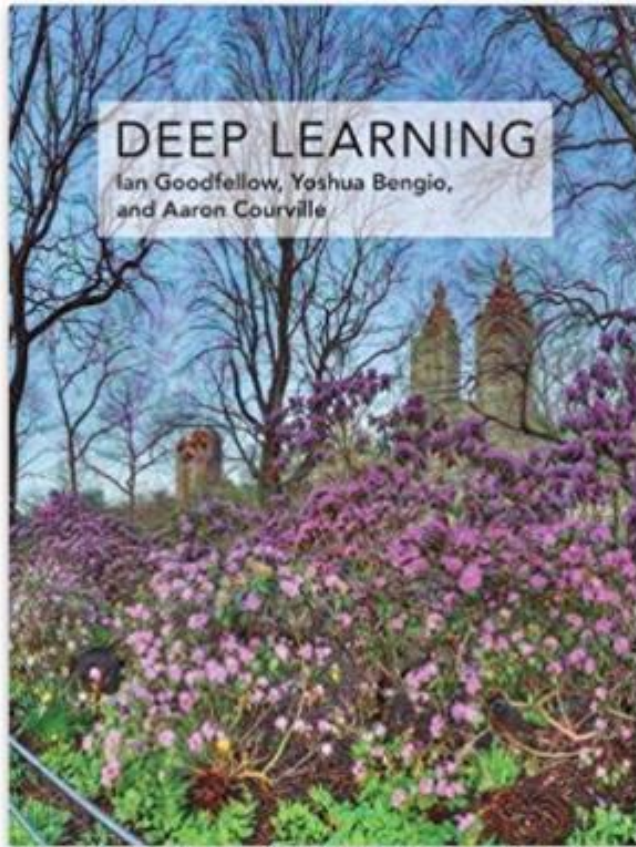


Q&A: “Do I have the appropriate pre-requisites/background?”

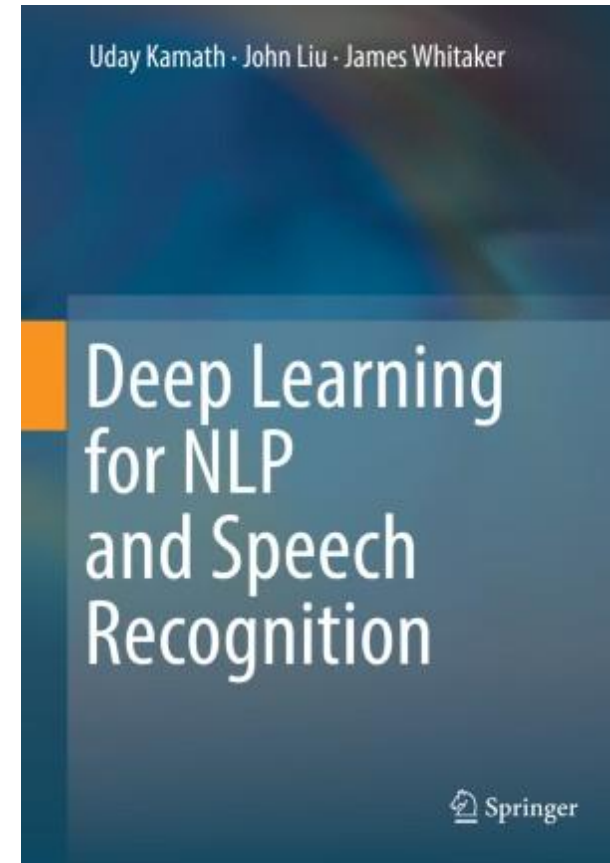
- Programming competency
- Experience with probability/statistics and linear algebra

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: “What are the assignments?”

- 4 problem sets (first assignment due in 2 weeks)
- 3 lab assignments
- Final project
 - Outline
 - Presentation
 - Peer evaluation
 - Final report
- Grading policies:
 - Late policy: Penalized 1% per hour for up to 2 hours and no credit afterwards
 - Regrade requests: must submit to TA within 2 weeks of receiving the grade
- Grading timeline
 - Due to the many students (150+), grades will be released ~1 week after submission

Q&A: “How is my final grade determined?”

	% of Final Class Grade
Problem Sets	35%
Lab Assignments	35%
Final Project	30%

Introductions

Instructor: Danna Gurari; aka, Dr. G; preferred pronouns: she/her:



Teaching assistants (will hold office hours)

- Everley Tseng; preferred pronouns: she/her:
(programming tutorials, Q&A on Piazza)



- Nick Cooper (preferred pronouns: he/him):
(assignment review sessions, regrade requests)



- Neelima Prasad (preferred pronouns: she/her):
(lead grader)



Grader: Mohammad Qazim Bhat

Q&A: “How Do I Get Answers to my 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 approach can benefit other students who have similar questions, as they can view the answers on Piazza. It also provides an opportunity for peers to answer your questions.
- **Office Hours:** The TAs will host office hours every day of the week at times posted on the course website and Canvas. Instructions for how to attend are provided on the home page of Canvas.
- **Appointments:** For matters that cannot be addressed in office hours, first email the TAs Everley Tseng and Nick Cooper to make an appointment or solicit an answer. Please note that at least 24 hours notice will typically be needed before a TA will be able to meet. The TAs will involve the instructor for any items they are unable to address.
- **Regrade requests:** All requests must be emailed to the TA, Nick Cooper, within 2 weeks of receiving the grade to be considered.

Introductions

NameCoach: share your name pronunciation

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

BigBlueButton

Collaborations

Chat

Attendance

My Media

Web Grading Sync

New Analytics

NameCoach 1.

CU Boulder Libraries

Studio

Zoom

Piazza

Course Materials

Files

Settings

NameCoach info for CSCI 5922: Neural Nets and Deep Learning

Your recording for CSCI 5922: Neural Nets and Deep Learning

3. Play **2. Record/Update** **Danna Gurari**
(danna.gurari@colorado.edu)

Recordings for CSCI 5922: Neural Nets and Deep Learning

Recorded Names **Unrecorded Names**

1 person have recorded their name

Show entries per page
10

Name/email	Last Name	Name Pronunciation	Invited At	
Danna Gurari (danna.gurari@colorado.edu)		Danna Gurari	Invited At: 08.09.2022 Recorded At: 01.05.2022	cl

My Experience

2007-2010

2010-2015

2015-2017

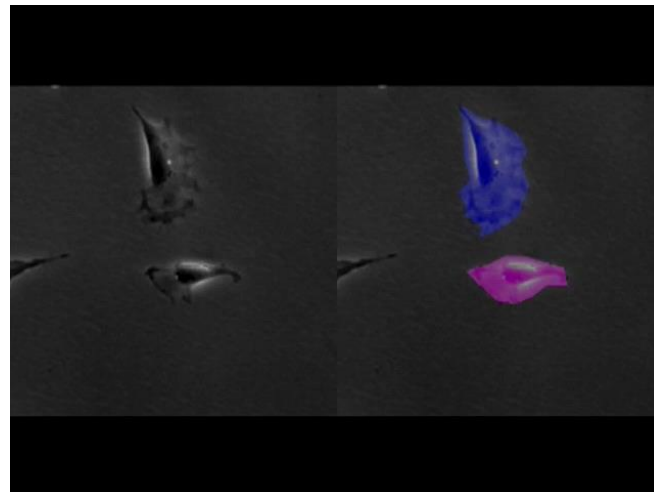
2017-Present

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



Source: Boulder Imaging

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



Is it edible or poisonous?

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

2007-2010

2010-2015

2015-2017

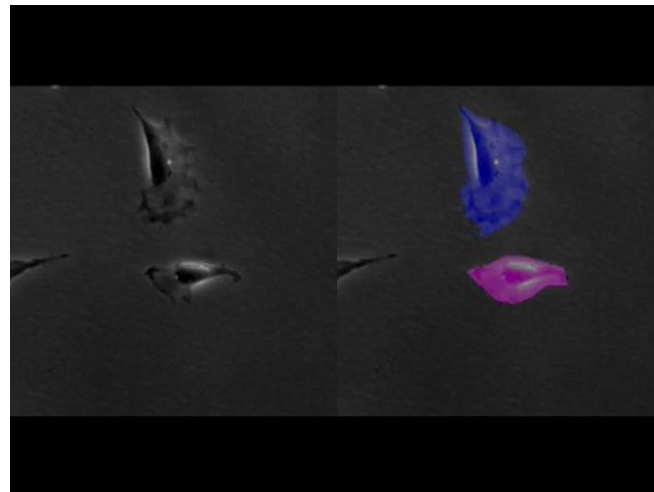
2017-Present

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



Source: Boulder Imaging

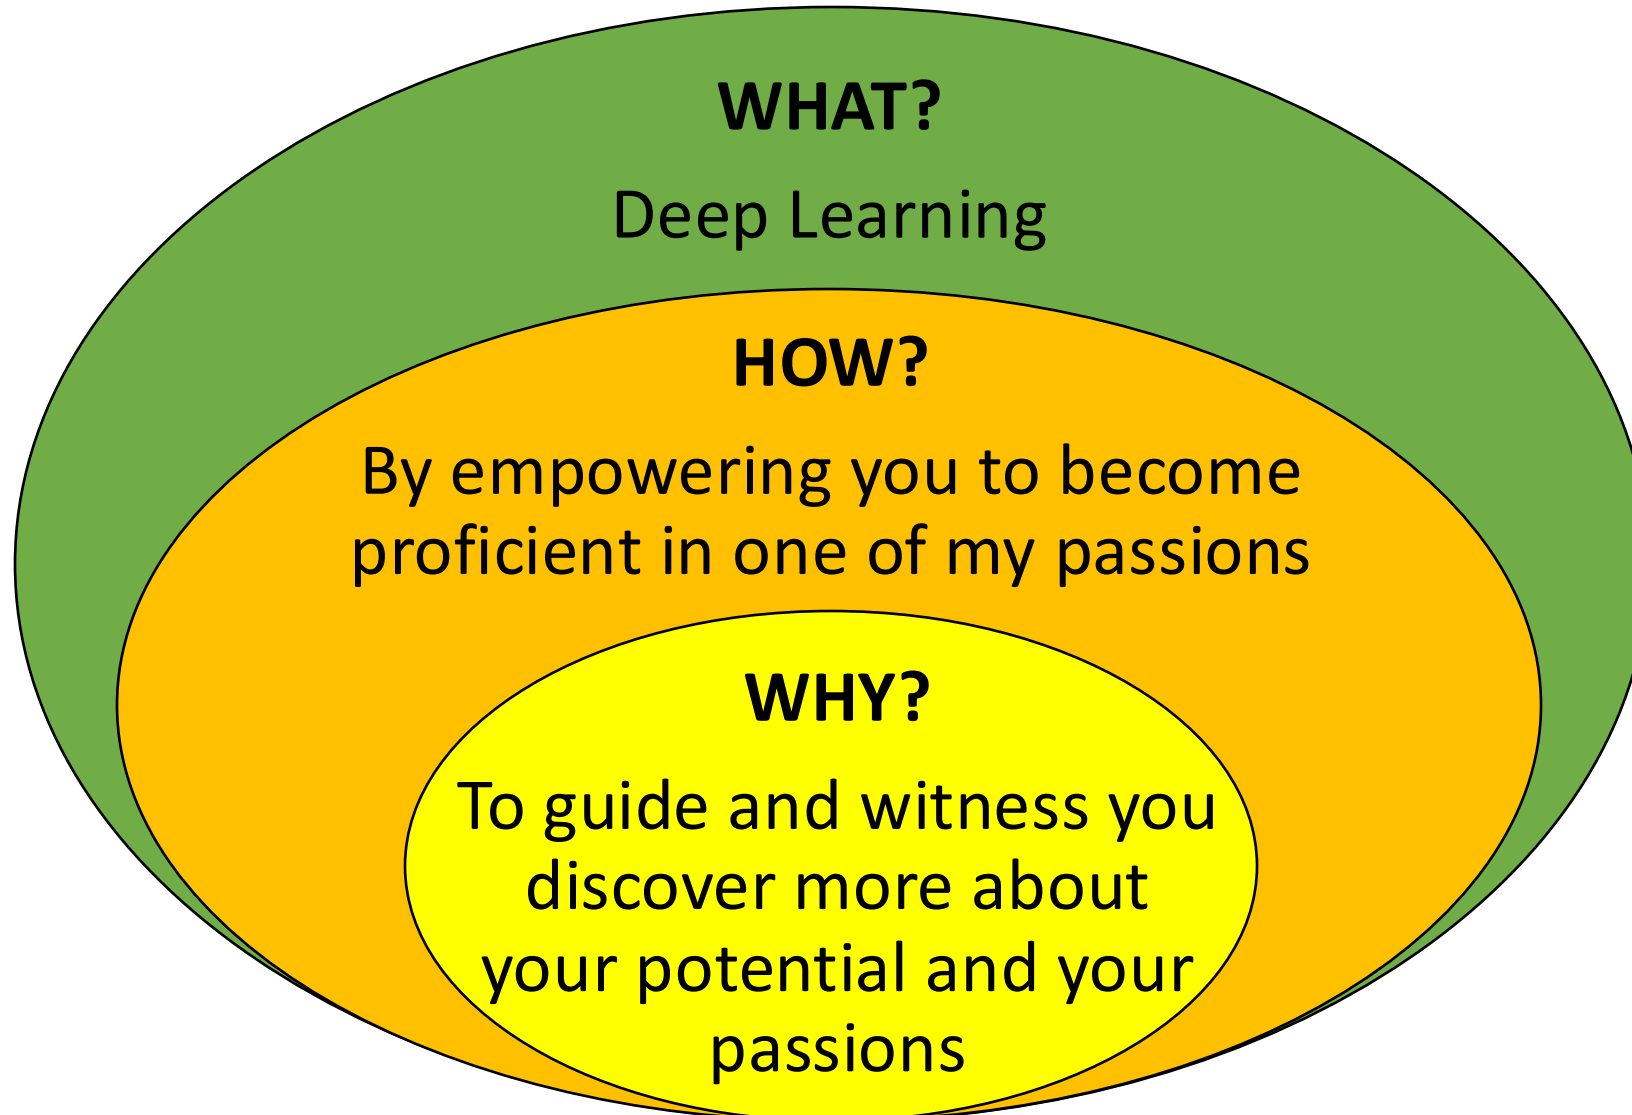
PhD student designing computer vision methods to segment and track cells in images and videos



34 publications involving deep learning

poisonous?

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