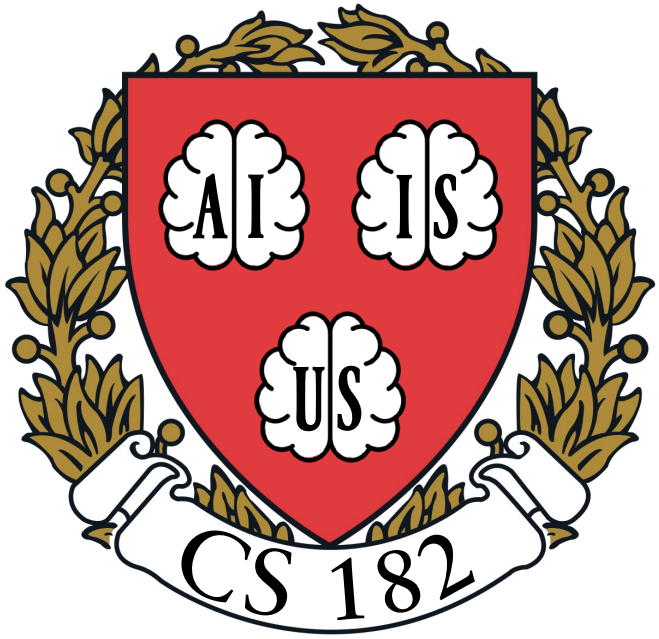

Fall 2021 | Lecture 1

Introduction

Ariel Procaccia | Harvard University

ORIGINS OF AI

1956 Darthmouth workshop participants, from left: Oliver Selfridge, Nathaniel Rochester, Ray Solomonoff, Marvin Minsky, Trenchard More, John McCarthy, Claude Shannon

A Proposal for the

DARTMOUTH SUMMER RESEARCH PROJECT ON ARTIFICIAL INTELLIGENCE

June 17 - Aug. 16

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.

The following are some aspects of the artificial intelligence problem:

1) Automatic Computers

If a machine can do a job, then an automatic calculator can be programmed to simulate the machine. The speeds and memory capacities of present computers may be insufficient to simulate many of the higher functions of the human brain, but the major obstacle is not lack of machine capacity, but our inability to write programs taking full advantage of what we have.

2) How Can a Computer be Programmed to Use a Language

It may be speculated that a large part of human thought consists of manipulating words according to rules of reasoning

WHAT IS AI?

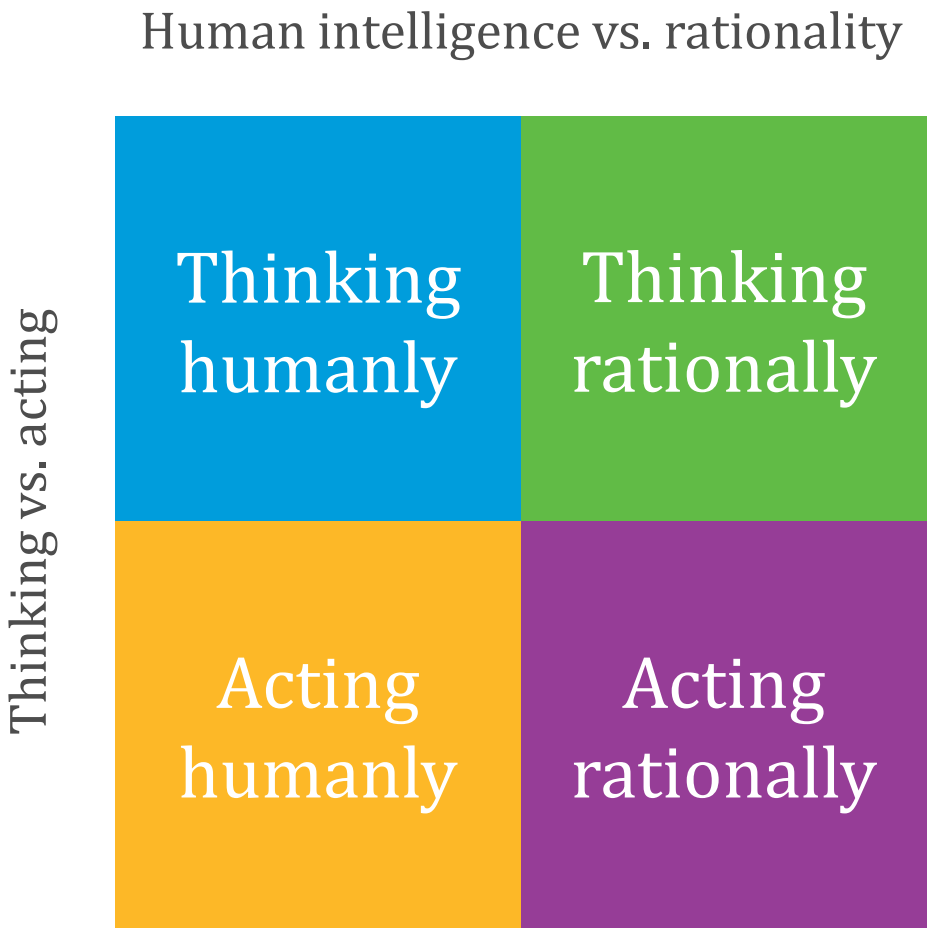
Human intelligence vs. rationality
Thinking vs. acting
Thinking humanly
Thinking rationally
Acting humanly
Acting rationally

ACTING HUMANLY

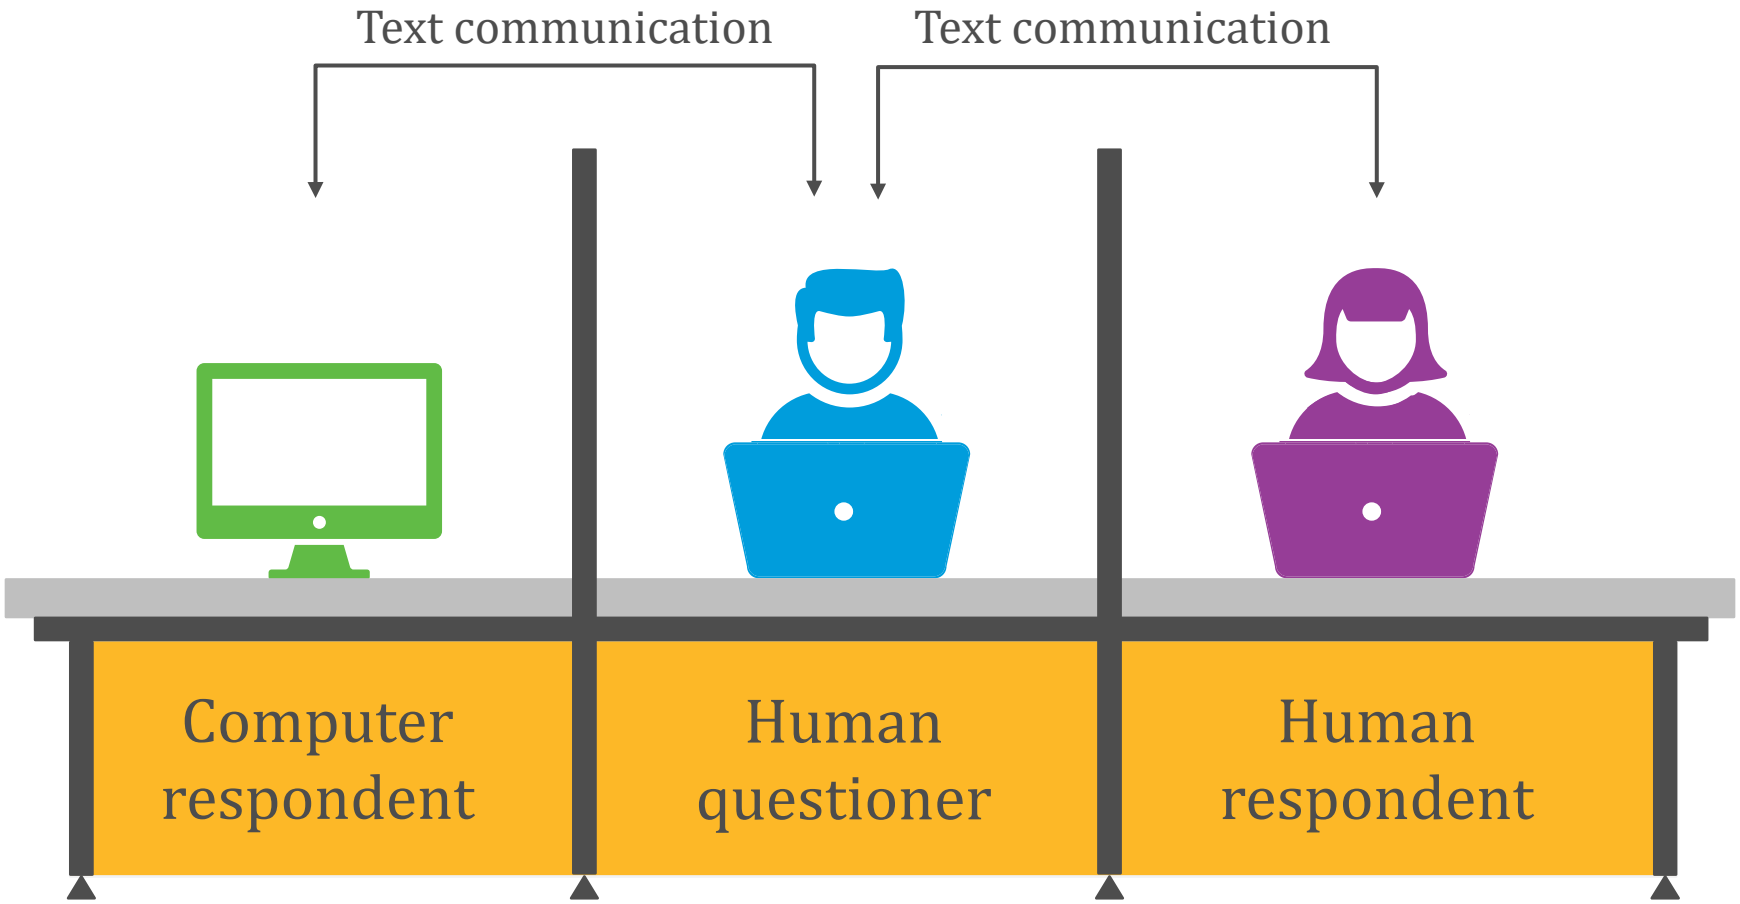

The Turing Test

EUGENE GOOSTMAN

- Turing predicted that a chatbot would convince 30% of judges by 2000
- Eugene Goostman achieved this milestone in June 2014
- Is it really intelligent?

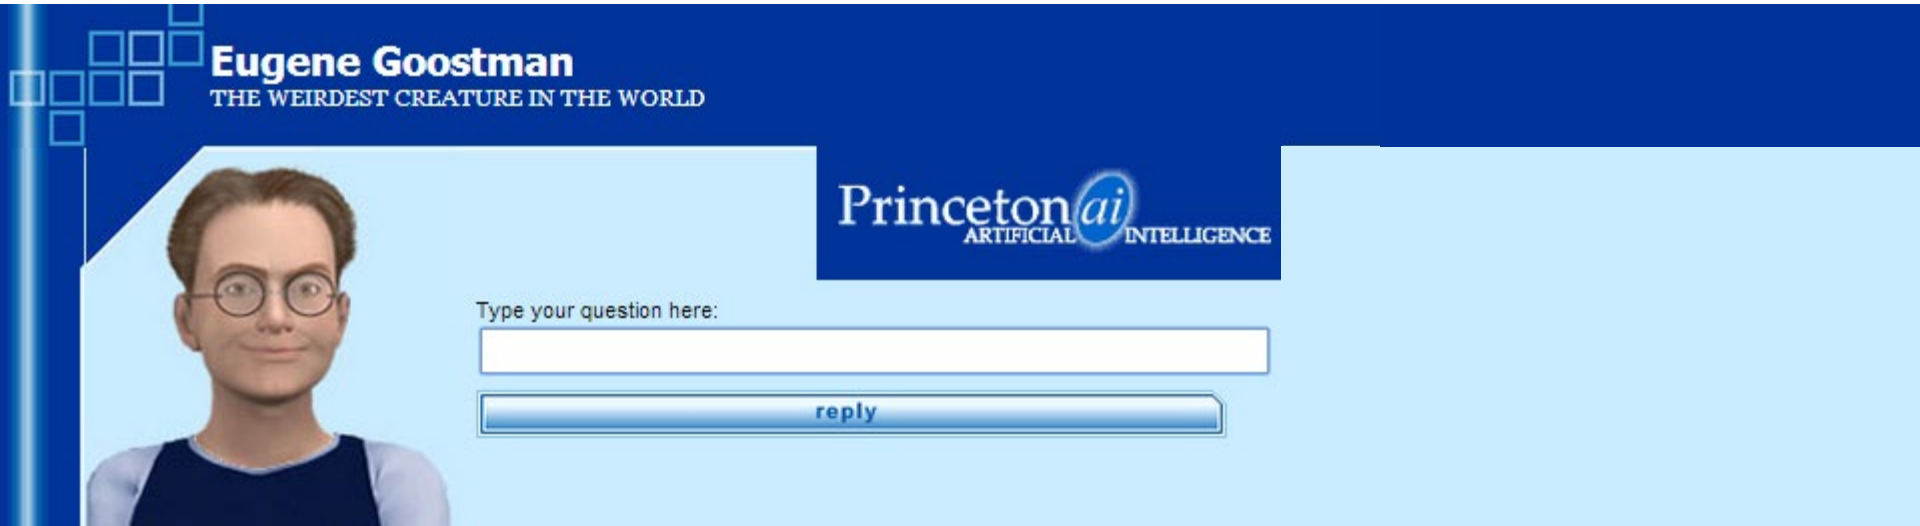

GOOSTMAN VS. AARONSON

Scott: Which is bigger, a shoebox or Mount Everest?

Eugene: I can't make a choice right now. I should think it out later. And I forgot to ask you where you are from…

Scott: How many legs does a camel have?

Eugene: Something between 2 and 4. Maybe, three? :-))) By the way, I still don't know your specialty — or, possibly, I've missed it?

Scott: How many legs does a millipede have?

Eugene: Just two, but Chernobyl mutants may have them up to five. I know you are supposed to trick me.

Scott: No, I need to know that you're not a chatbot. Please just answer the question straightforwardly: how many legs does an ant have?

Eugene: Something between 2 and 4. Maybe, three? :-))) Oh, what a fruitful Conversation ;-)

Scott: Do you understand why I'm asking such basic questions? Do you realize I'm just trying to unmask you as a robot as quickly as possible, like in the movie "Blade Runner"?

Eugene: …wait

Scott: Do you think your ability to fool unsophisticated judges indicates a flaw with the Turing Test itself, or merely with the way people have interpreted the test?

Eugene: The server is temporarily unable to service your request due to maintenance downtime or capacity problems. Please try again later.

ACTING PIGEONLY

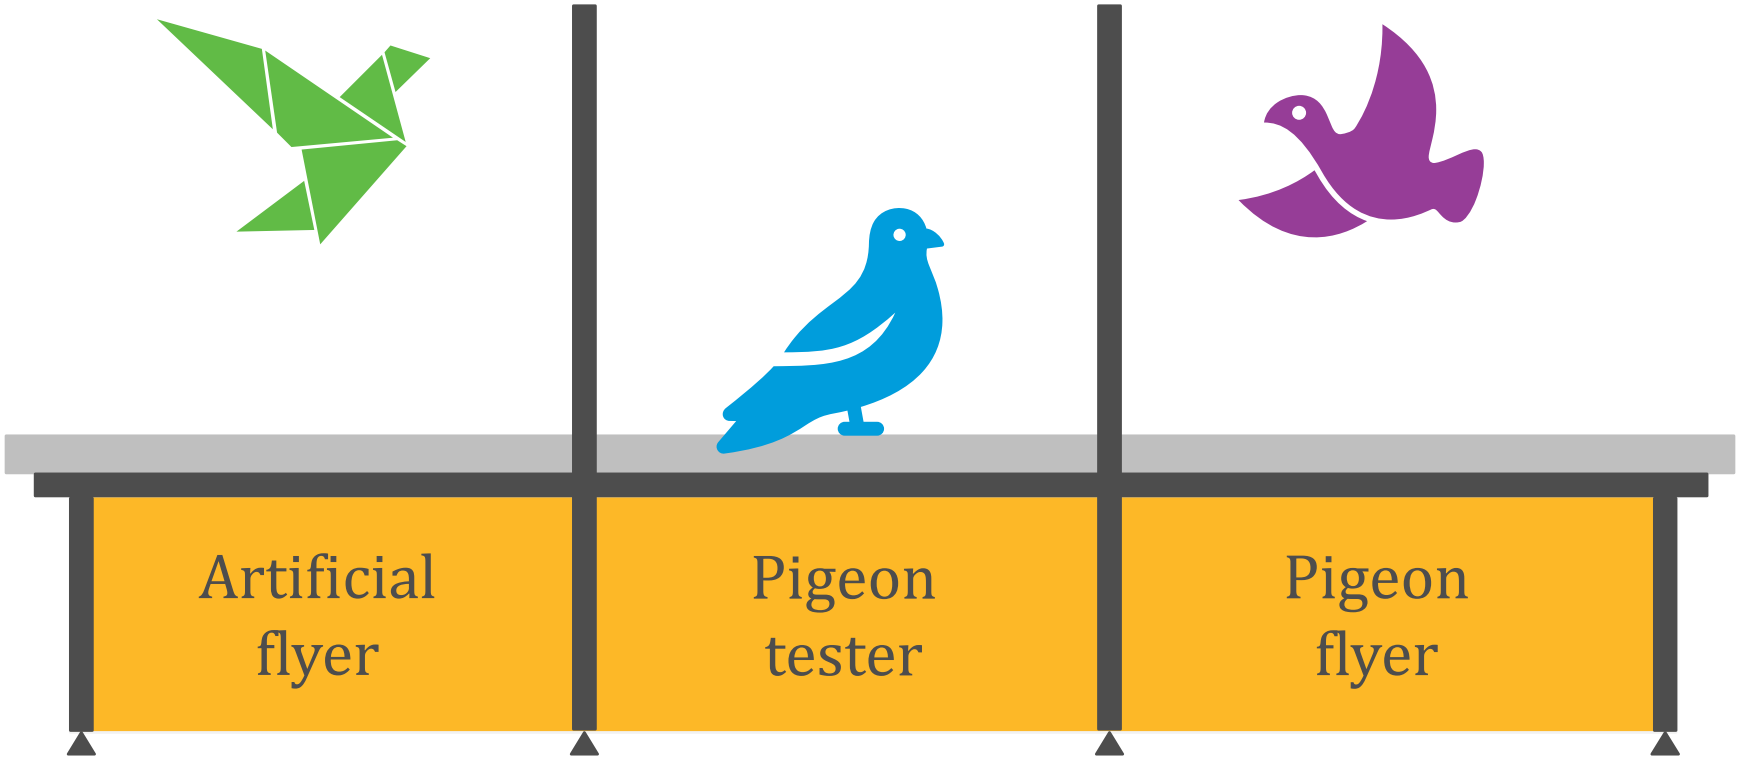

R&N: Aeronautical engineering texts do not define the goal of their field as making "machines that fly so exactly like pigeons that they can fool even other pigeons."

A THOUGHT EXPERIMENT

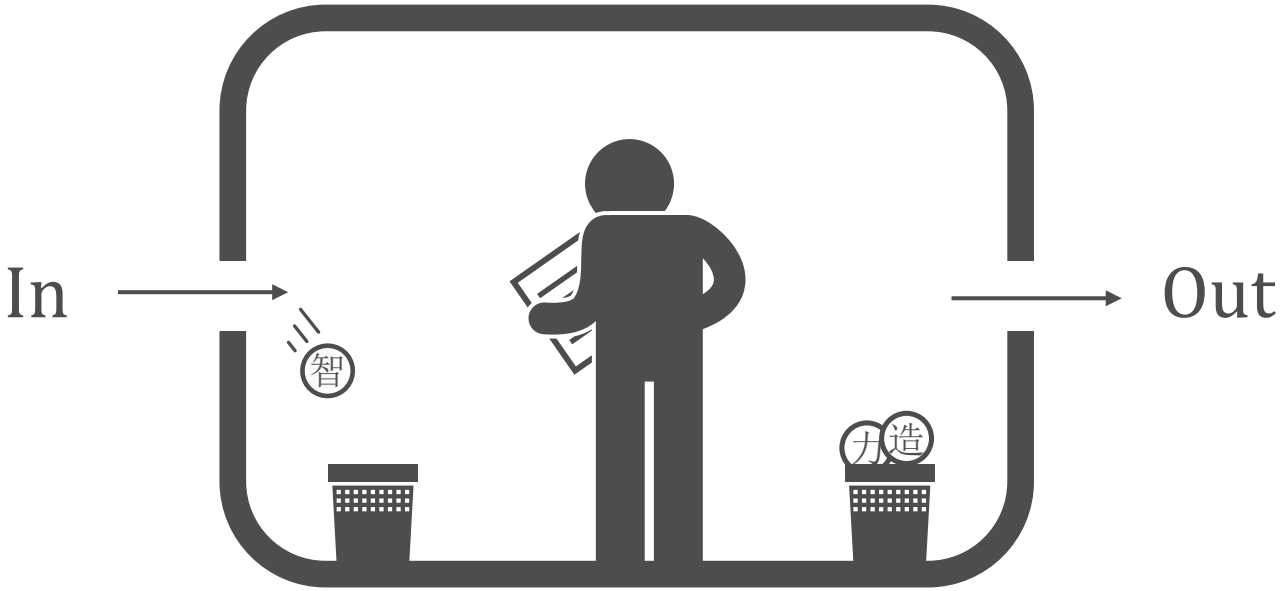

The Chinese Room Argument
[Searle 1980]

ACTING RATIONALLY

- A rational agent acts to achieve the best outcome or “do the right thing”
- This can be based on thinking rationally, but not necessarily
- This approach is especially concrete because the quality of an outcome can be rigorously defined
- The same approach can also be found in control theory, operations research, statistics and economics

TEACHING PHILOSOPHY

Key course objective:
Learn to represent problems

TEACHING PHILOSOPHY

AI courses are too shallow. I should teach AI with as many theorems and as much math as possible!

I should focus on representation. Hey, where did my hair go?

Procaccia 2011

Procaccia 2021

SYLLABUS

Problem solving

- Uninformed search
- Informed search
- Motion planning
- Constraint satisfaction problems
- Convex optimization
- Integer programming

Multi-agent systems

- Game theory
- AI game playing
- Wildlife protection (Tambe)
- Social choice

Reasoning with uncertainty

- Bayesian networks
- Hidden Markov Models
- Markov decision processes

Machine learning

- Reinforcement learning
- Decision trees
- Linear classification
- Neural networks
- Language models (Alvarez-Melis)

Ethics

- Fairness
- Value alignment
- Embedded EthiCS

CS 181 (S21) VS. 182

Problem solving
Uninformed search
Informed search
Motion planning
Constraint satisfaction problems
Convex optimization
Integer programming
Multi-agent systems
Game theory
AI game playing
Wildlife protection (Tambe)
Social choice

Reasoning with uncertainty
Bayesian networks
Hidden Markov Models
Markov decision processes
Machine learning
Reinforcement learning
Decision trees
Linear classification
Neural networks
Language models (Alvarez-Melis)
Ethics
Fairness
Value alignment
Embedded EthiCS

MIT 6.036 (F20) VS. 182

Problem solving
Uninformed search
Informed search
Motion planning
Constraint satisfaction problems
Convex optimization
Integer programming
Multi-agent systems
Game theory
AI game playing
Wildlife protection (Tambe)
Social choice

Reasoning with uncertainty
Bayesian networks
Hidden Markov Models
Markov decision processes
Machine learning
Reinforcement learning
Decision trees
Linear classification
Neural networks
Language models (Alvarez-Melis)
Ethics
Fairness
Value alignment
Embedded EthiCS

TEACHING FELLOWS

Wenqi Chen

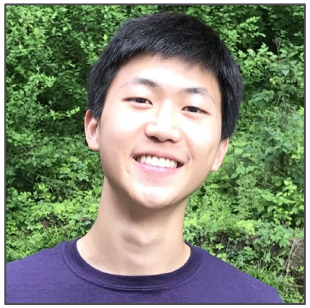

Max Guo

Kavya Kopparapu

Eric Lin

Shuvom Sadhuka

Amir Shanehsazzadeh

Zuzanna Skoczylas

Meiling Thompson

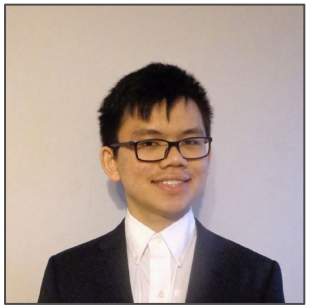

William Zhang

REQUIREMENTS

- Five homework assignments
  - Theory and programming
  - 10% each
- Midterm and final exams
  - 15% and 25%, respectively
- Attendance
  - 10%
  - Exceptions for Covid and other reasons
  - Tracked through polls

TEST POLL

- We will test the hypothesis that the SEC floor is statistically indistinguishable from the floor of my garage
- Poll: Select all the photos that show the SEC floor (the others are from my garage)

Option 1 ✓

Option 2

Option 3 ✓

PREREQUISITES

- Students must have previously taken Statistics 110 (Probability) or an equivalent course
- Experience with Python programming and a good understanding of time complexity (including big O notation) are assumed
- A certain degree of “mathematical maturity” is crucial

COURSE TEXTBOOK

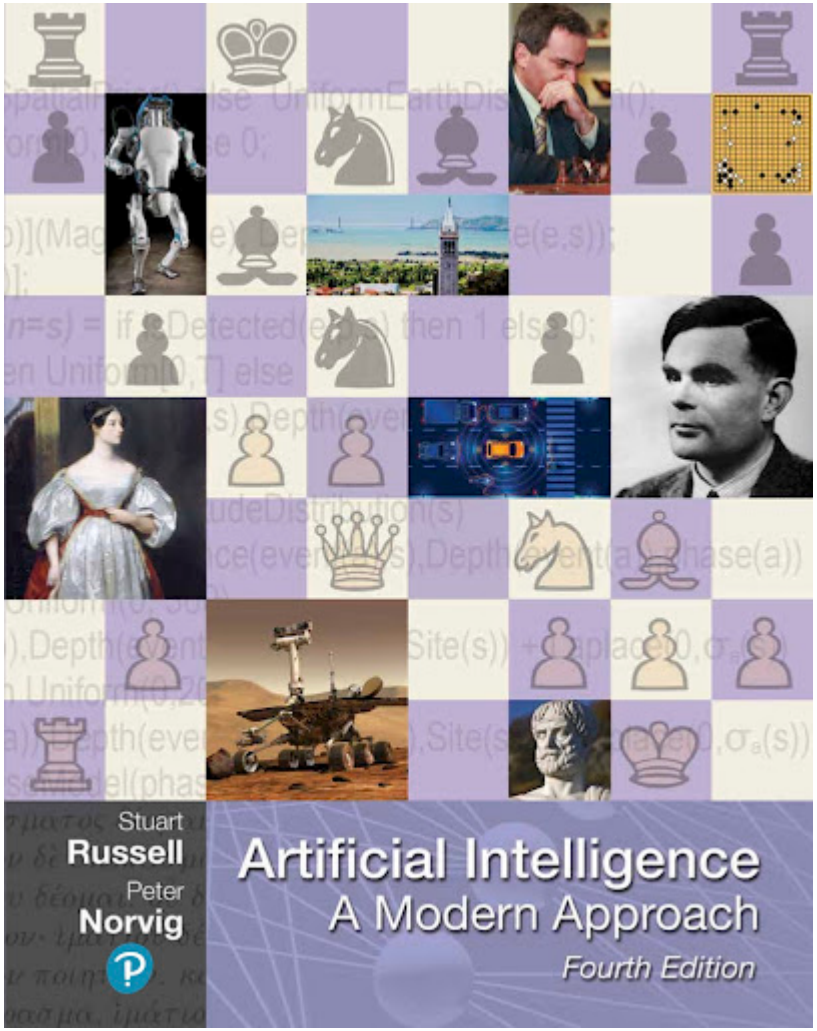

SECTION STRUCTURE

- We are trying out more personalized sections
- TFs will provide to students in their section a method for suggesting topics/problems and will base the section on these suggestions
- Sections will be held in person
- Attendance is optional

OTHER COMMENTS

- Detailed info is available on the course website (Google “Harvard CS 182 fall 2021”)
- Lecture slides will be posted to the website ahead of time
- Lectures are recorded and Panopto links will be posted to the website
- Sign up for Ed Discussion