Logistics and Introduction

CS 480
Intro to Artificial Intelligence

Course Info

CS 480-001 Introduction to Artificial Intelligence
When: Tuesdays/Thursdays, 10:30-11:45am
Where: Horizon hall, room 2009
Course syllabus: https://cs.gmu.edu/~hrolenok/teaching/cs-480-spring2025/index.html

Instructor: Brian Hrolenok
Instructor office hours: Date/Time TBD, ENGR 4708

Topics

- Search
- Bayes Nets
- Filtering
- Markov Decision Processes and Reinforcement Learning
- Local Search / Optimization (maybe)
- Regression
- Neural Networks
- Decision Trees

Readings and tentative schedule on the syllabus

Textbook

“Artificial Intelligence: A Modern Approach” by Stuart Russell and Peter Norvig, Third edition.

ISBN: 978-0136042594

Web: http://aima.cs.berkeley.edu/

Fourth (purple) and Second (green) editions should be roughly equivalent, page numbers may have changed

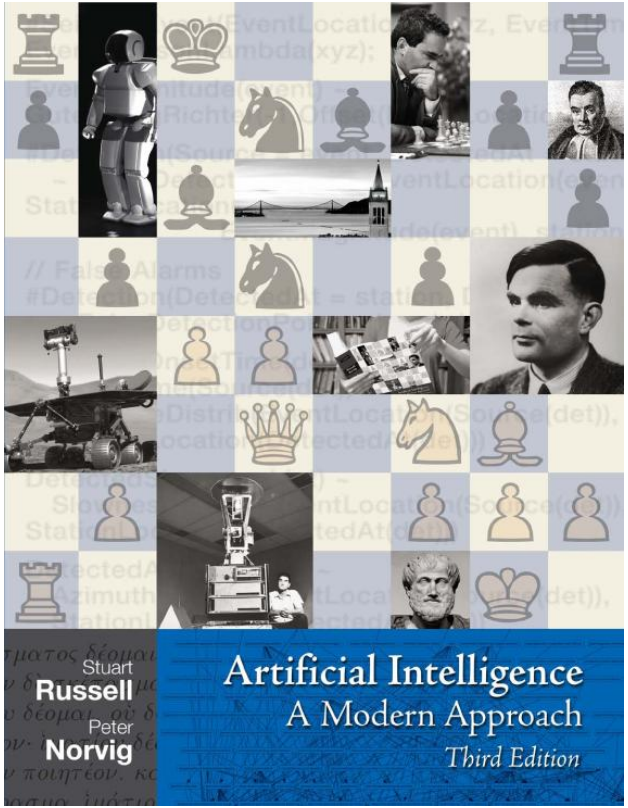

Grading

- 7 problem sets (10% total, on Gradescope now)
- 4 homeworks (30% total, on Gradescope (first one available now))
- 1 midterm (25%, roughly week 9)
- 1 final (35%, scheduled by the registrar)

Midterm will be closed-book with 1 page of notes, individual work only.

Final will be closed-book with 2 pages of notes, individual work only.

Assignments can be done in groups of **at most 3**. The homework assignments are autograded, and the autograder is available to you. Problem sets are graded by hand.

See the syllabus for details

Prerequisites

Official prereq's: CS 310, CS 330

Useful programming skills to have

- Python
- Objects and Classes
- Data structures
- Algorithms and Complexity

Useful math to know

- Probability
- Calculus
- Linear Algebra
- Discrete

If you are unsure, check out the first assignments.

What is Artificial Intelligence?

What is intelligence?

- Thinking like a human
  - Cognitive Science
  - Introspection
  - Difficult to experimentally measure/ test
- Acting like a human
  - Turing Test (1950!)
  - Narrow criteria for success
  - Humans make mistakes
  - Do we really understand what “intelligence” means?
- Thinking Rationally
  - “Laws of thought”
  - natural computer implementation
  - Informal vs formal knowledge?
  - Inefficient?
- **Acting rationally**
  - **Agent** based approach
  - Easy* to define success and experimentally validate
  - Allows for novel artificial solutions (wheels vs feet)

The Turing Test

- In room 1: a human at a terminal (interrogator)
- In room 2: Two subjects, A and B
  - A human at a terminal
  - An AI program running on a computer
- Only communication allowed via terminal
- Interrogator goal: pick which of A or B is the human and which the computer
- AI goal: fool the interrogator into believing it is the human

Turing's point: if the program can win this game, we **can't** say the computer is **not** intelligent*

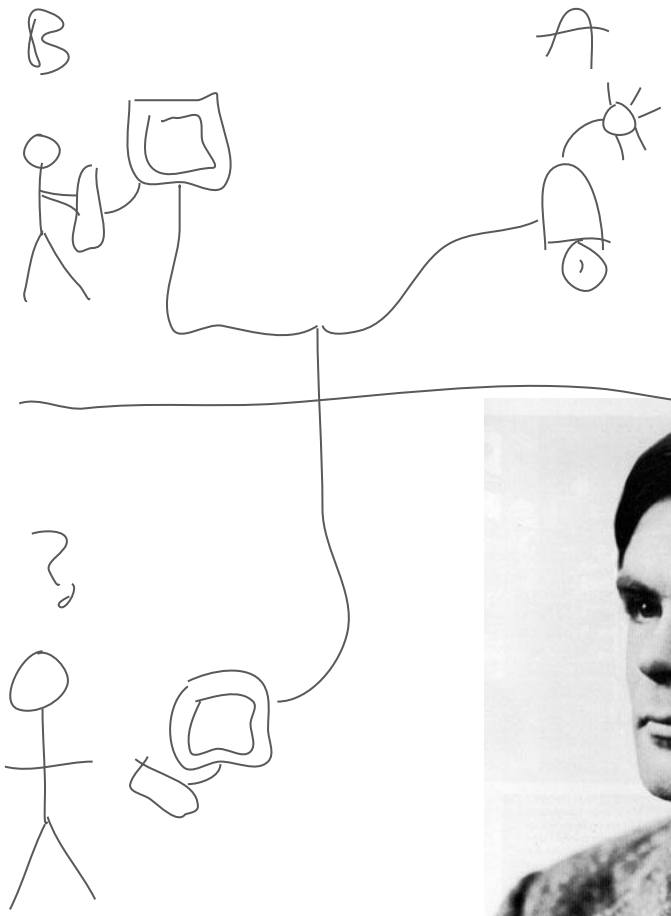

A brief and incomplete history of AI

- 1950 - Turing Test
- ‘52 - Checkers (Samuel)
- ‘56 - The Dartmouth Summer Workshop
- ‘58 - Lisp (McCarthy)
- ‘62 - Perceptrons (Rosenblatt, Block)
- ‘65 - First order logic (Robinson)
- ‘67 - Algebra word problems (Bobrow)
- ‘68 - SHRDLU (Winograd)
- ‘69 - Limits of perceptrons (Minsky)*
- ‘70-75 Blocks World (Winston, Sussman, Waltz)
- ‘74 - Lighthill report, SUR*
- ‘80-87 - Expert Systems, commercial success
- ‘81 - Fifth Generation project launched
- ‘87 - SCI cutbacks, LISP machine market collapse*
- ‘91 - Fifth Generation project canceled*
- ‘97 - LSTMs
- ‘98 - ConvNets for MNIST (LeCun)
- 2009 - Pre-training for DNNs
- 2012 - AlexNet wins ImageNet
- 2015 - AlphaGo, DQN
- 2014-2019 - Google, Facebook, Uber, Amazon, Apple, Netflix...

*AI “winters”

Where are we today? (1)

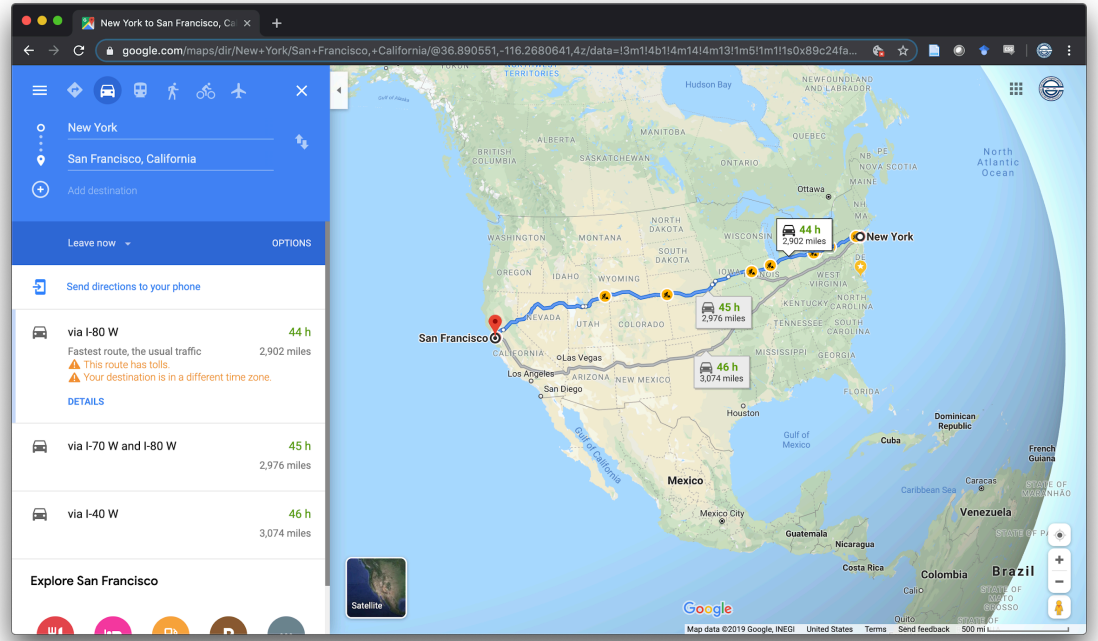

Where are we today? (2)

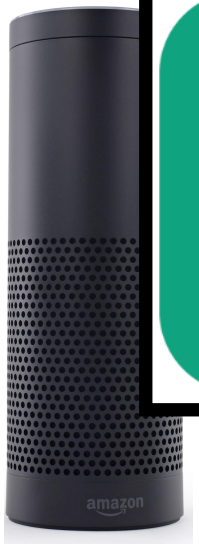

Home speakers

Phone Assistants

Dialog systems

Where are we today? (3)

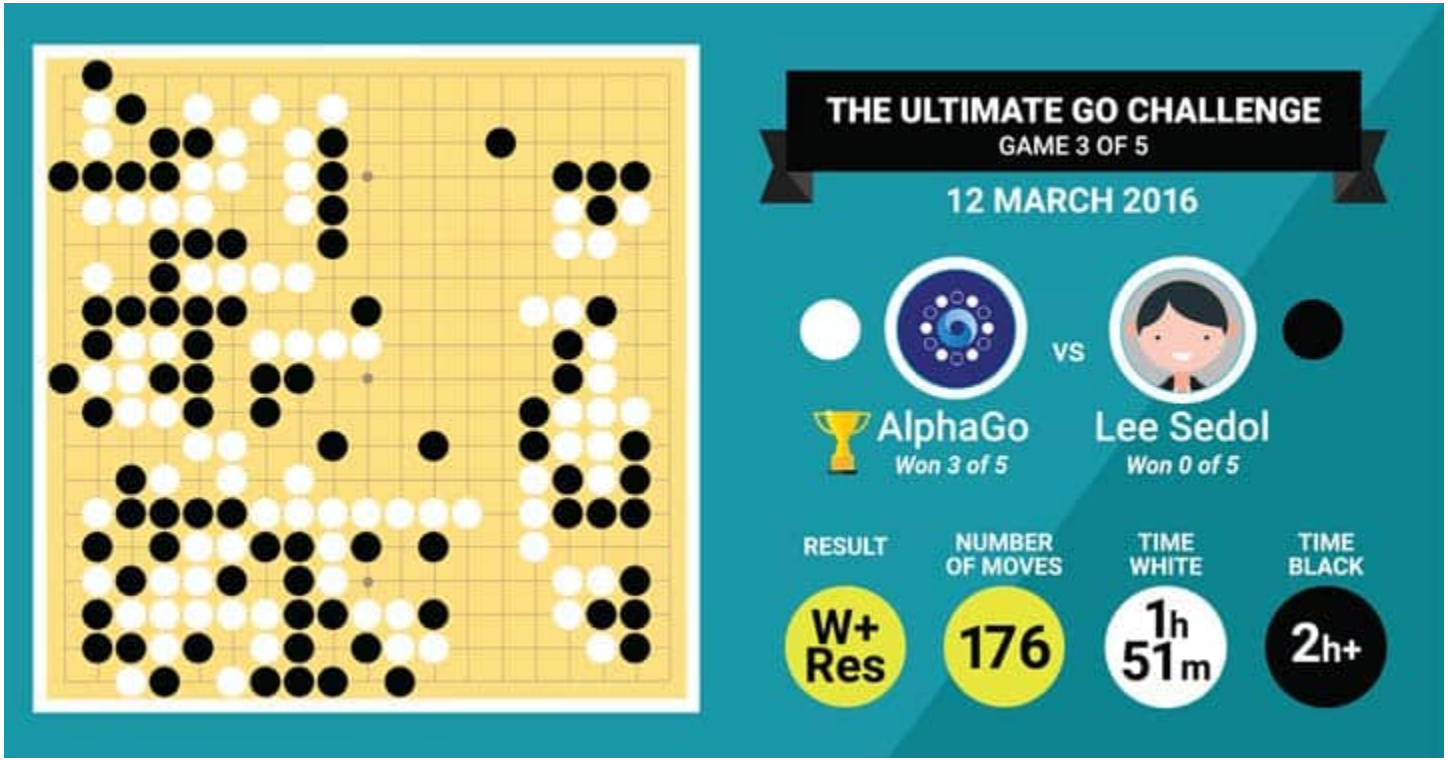

https://www.bgr.in/news/google-deepmind-alphago-vs-lee-sedol-googles-ai-claims-victory-over-go-world-champion/

Where are we today? (4)

https://www.youtube.com/watch?v=5ALIK-z-vUg

An Agent-based Approach

An **agent** is an entity that perceives and acts in an **environment**

SENSE

THINK

ACT

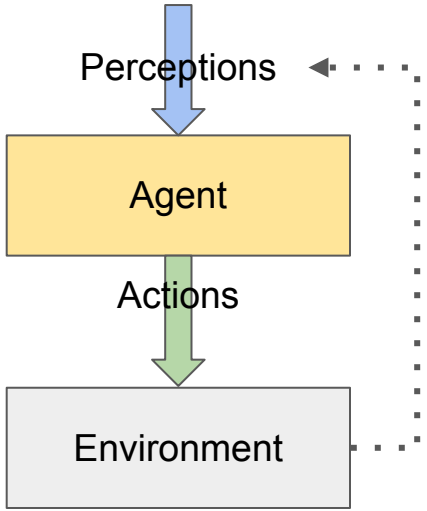

An Agent-based Approach

A mathematical definition of **agent**

$$f : \mathcal{P}^* \rightarrow \mathcal{A}$$

$\mathcal{P}$ perception, $\mathcal{P}^*$ sequences of perceptions, $\mathcal{A}$ actions

How do we design a computer program to compute **f()**?

How do we **efficiently** determine the **best** f() for a given task?

A simple example: Vacuum world

Environment: 2 rooms (cells)

Perceptions: Which room agent is in, whether the room is clean or dirty

Actions: Move left, move right, clean
Moving into a wall does nothing
Cleaning a clean room does nothing

P*: [(A,Clean)], [(A,Dirty)], … [(A,Clean),(B,Dirty), (B,Clean)], …

Objective: Both rooms are clean

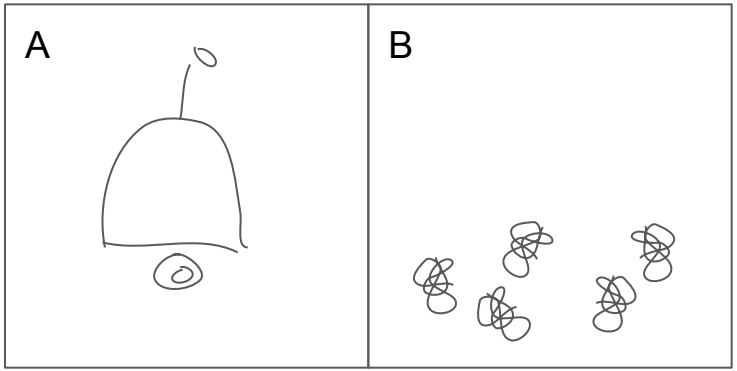

Current P: (A, Clean)

For next time

- Write a program to compute f() for vacuum world
- How simple can we make it?
- What's the **best** f()?
- What kinds of things can we do to make the problem more challenging?
- Do we always need the full history of perceptions?
- State space vs. Perception history