Artificial Intelligence for Engineers

ECE 562

Autumn 2023

Administrative Details

- Instructor: Linda Shapiro, 634 CSE, shapiro@cs.washington.edu
- TA: Nishat Khan, nkhan51@uw.edu
- TA: Saygin Seyfioğlu, msaygin@uw.edu

- Course Home Page: http://homes.cs.washington.edu/~shapiro/EE562
- Text: Artificial Intelligence: A Modern Approach, Russell and Norvig, 3rd or 4th edition, see Canvas Discussion Board.

This Lecture

- What is AI all about, roughly from Chapters 1 and 2.
- Overview the Python language we will use.

What is intelligence?

- What capabilities should a machine have for us to call it intelligent?

Turing's Test

- If the human cannot tell whether the responses from the other side of a wall are coming from a human or computer, then the computer is intelligent.

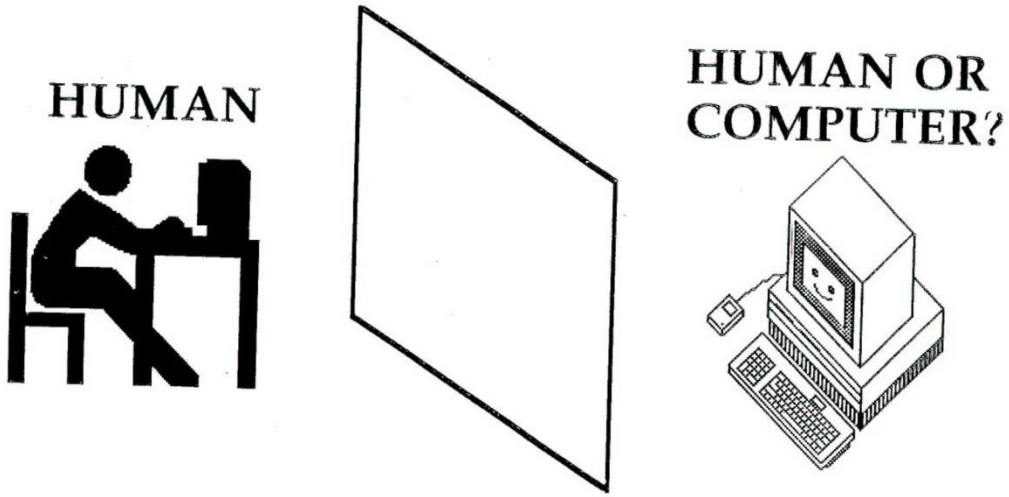

Performance vs. Humanlike

- What is more important: how the program performs or how well it mimics a human?
- Can you get a computer to do something that you don’t know how to do? Like what?
- What about creativity?

Mundane Tasks

- Perception
  - Vision
  - Speech
- Natural Language
  - Understanding
  - Generation
  - Translation
- Reasoning
- Robot Control

Formal Tasks

- Games
  - Chess
  - Checkers
  - Kalah, Othello
- Mathematics
  - Logic
  - Geometry
  - Calculus
  - Proving properties of programs

Expert Tasks

- Engineering
  - Design
  - Fault Finding
  - Manufacturing planning
- Medical
  - Diagnosis
  - Medical Image Analysis
- Financial
  - Stock market predictions

What is an intelligent agent?

- What is an agent?
- What does rational mean?
- Are humans always rational?
- Can a computer always do the right thing?
- What can we substitute for the right thing?

Intelligent Agents

- What kinds of agents already exist today?

Problem Solving

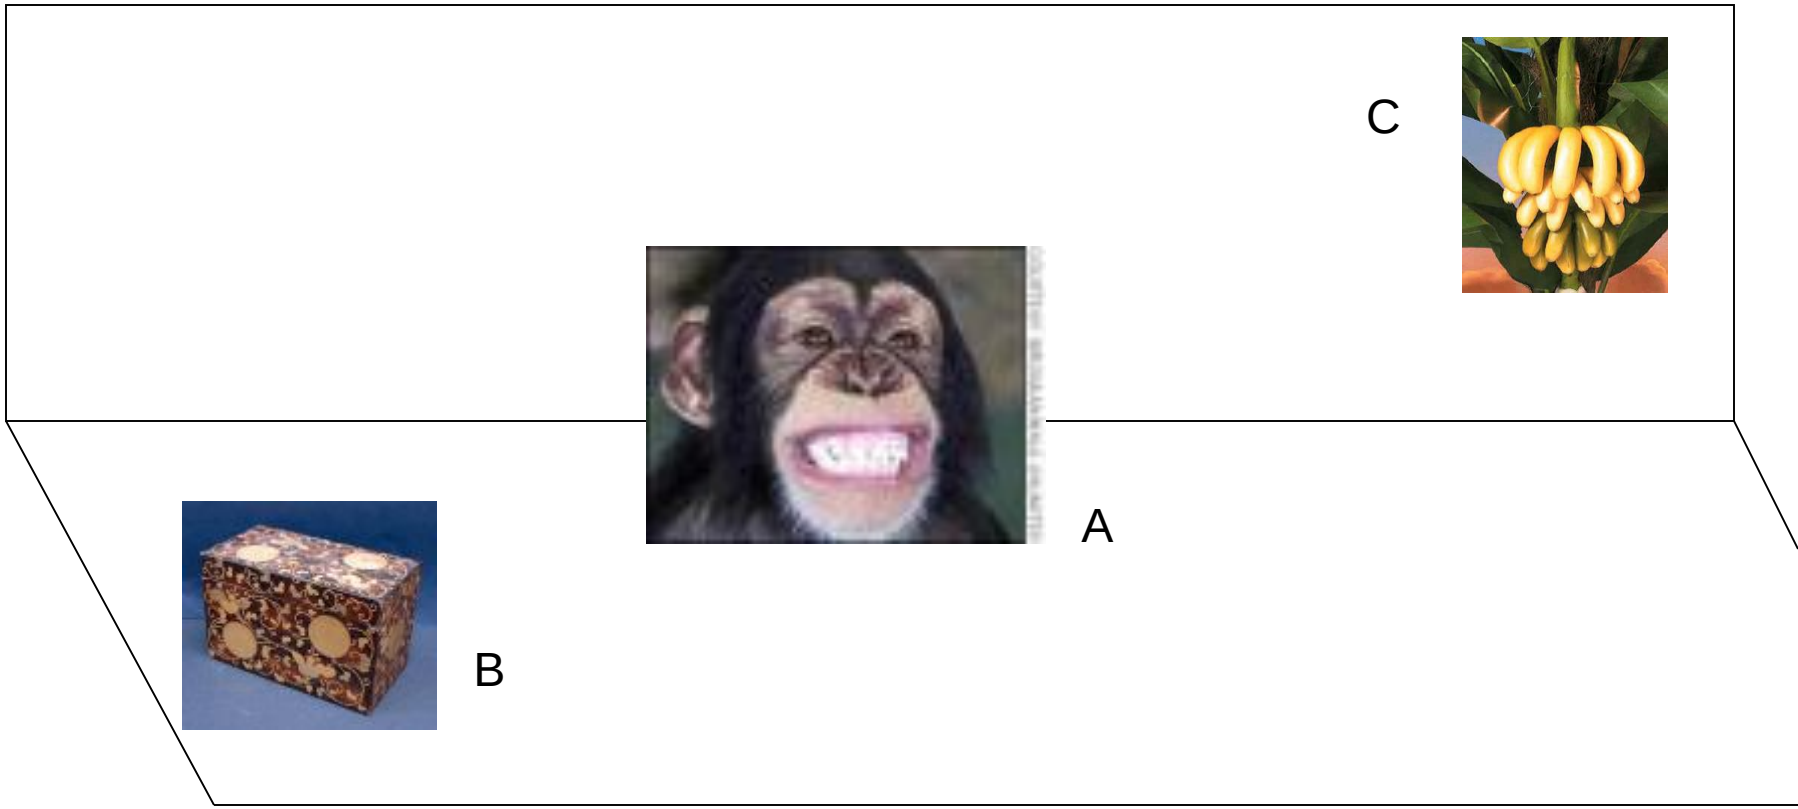

Find a sequence of operations to produce the desired situation from the initial situation.

Game Playing

- Given:
  - An initial position in the game
  - The rules of the game
  - The criteria for winning the game
- WIN!

Constraint Satisfaction

Example: Map Coloring

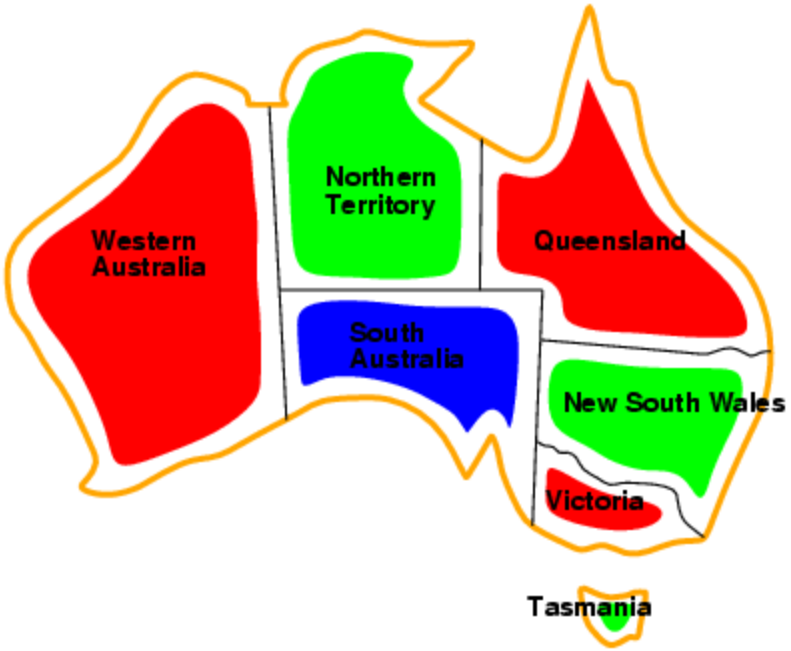

Reasoning

- Given:
  - ∀x (human(x) -> animal(x))
  - ∀x (animal(x) -> (eats(x) & drinks(x)))
- Prove:
  - ∀x (human(x) -> eats(x))

Learning

- Example: Neural Network

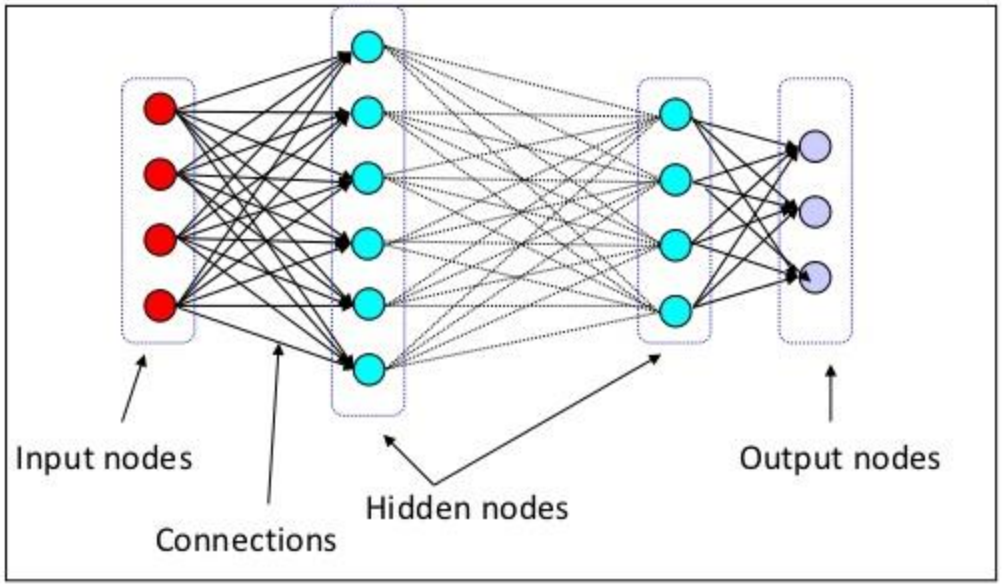

$$Output: \quad y_i = f(w_i^1 x_1 + w_i^2 x_2 + w_i^3 x_3 + \cdots + w_i^m x_m) = f(\sum_j w_i^j x_j)$$

Natural Language Understanding

- Pick up a big red block.
- OK.
- While hunting in Africa, I shot an elephant in my pajamas.
- I don't understand.

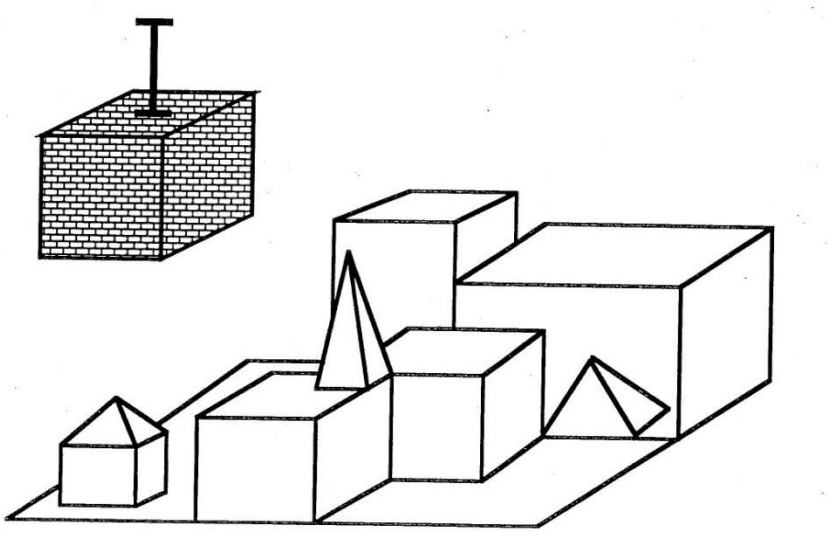

Large Language Models

- ChatGPT from OpenAI

- Write a story about my new AI class
- Make it one paragraph
- Make it about a graduate course in a university.
- Make the professor Linda Shapiro, who is an expert in computer vision

In the hallowed halls of a prestigious university, a cutting-edge graduate course had just debuted: "Advanced Artificial Intelligence and its Impact." Alex, a bright and tech-savvy student, eagerly enrolled, anticipating an exciting journey into the depths of AI research and its societal implications, all under the expert guidance of Professor Linda Shapiro, a renowned authority in computer vision. With Dr. Shapiro at the helm, the class delved deep into the history of AI, explored the intricacies of machine learning and neural networks, and engaged in thought-provoking discussions on the ethical dimensions of AI. Over the semester, Alex collaborated with their friend Mia on a groundbreaking project, devising an AI-driven solution to assist individuals with visual impairments. As the course concluded, they left with a profound sense of responsibility and a determination to leverage AI's potential for positive change in the world, inspired by the wisdom and expertise of Professor Shapiro.

Computer Vision with Machine Learning

Given: Some images and their corresponding descriptions

{trees, grass, cherry trees}

{cheetah, trunk}

{mountains, sky}

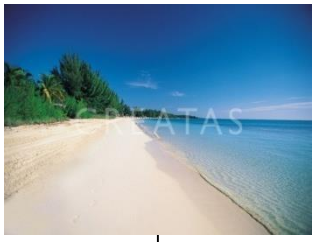

{beach, sky, trees, water}

●●●

To solve: What object classes are present in new images

?

?

?

?

●●●

Groundtruth Data Set: Annotation Samples

tree(97.3), **bush**(91.6), **spring flowers**(90.3), **flower**(84.4), park(84.3), **sidewalk**(67.5), **grass**(52.5), **pole**(34.1)

sky(99.8), **Columbia gorge**(98.8), lantern(94.2), **street**(89.2), house(85.8), bridge(80.8), car(80.5), hill(78.3), boat(73.1), pole(72.3), **water**(64.3), mountain(63.8), **building**(9.5)

sky(95.1), **Iran**(89.3), house(88.6), **building**(80.1), boat(71.7), bridge(67.0), **water**(13.5), **tree**(7.7)

Italy(99.9), grass(98.5), **sky**(93.8), rock(88.8), **boat**(80.1), **water**(77.1), Iran(64.2), stone(63.9), bridge(59.6), **European**(56.3), sidewalk(51.1), **house**(5.3)

Mars Rovers (2003-now)

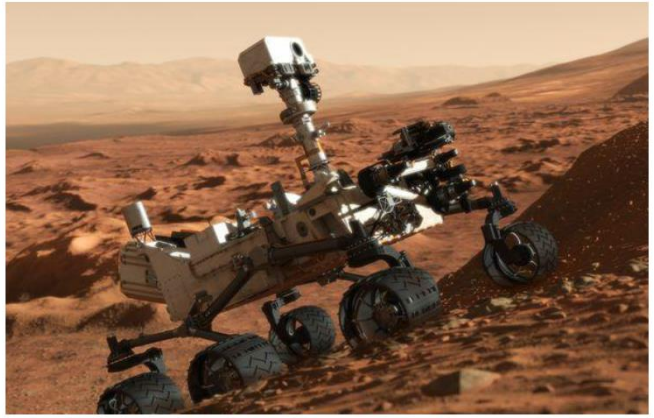

(See NASA website for latest updates) 25

Robots that Learn

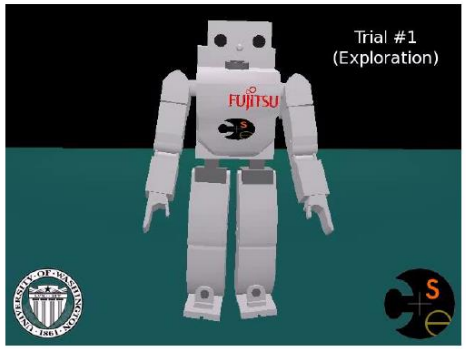

Learning

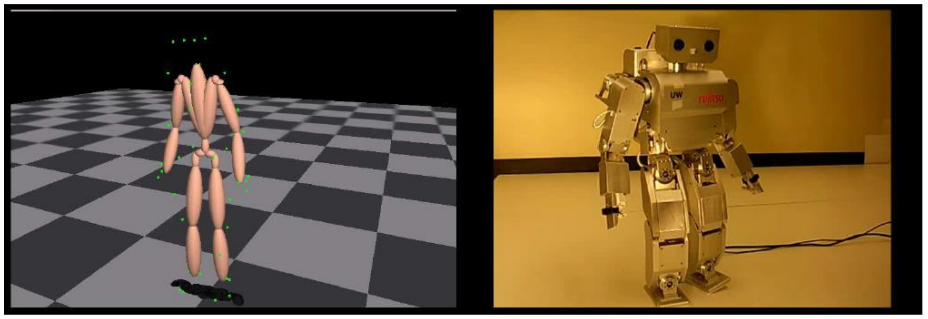

After Learning

(Work by UW CSE PhD David Grimes) 27

Brain-Computer Interfaces

(Work by UW MD-PhD Kai Miller) 29

Stuart Russell's "Potted History of AI"

- 1943 McCulloch & Pitts: neural nets model of the brain
- 1950 Turing's "Computing Machinery and Intelligence"
- 1952-69 Look Ma, no hands
- 1950s Early AI Programs: Logic Theorist, Checker Player, Geom
- 1956 Term "Artificial Intelligence" adopted
- 1965 Robinson's complete algorithm for logical reasoning
- 1966-74 AI discovers computational complexity; neural nets go
- 1969-79 Early development of knowledge-based "expert systems"
- 1980-88 Expert systems boom
- 1988-93 Expert systems bust: "AI Winter"
- 1985-95 Neural networks return
- 1988- AI and Statistics together
- 1995- Agents, agents everywhere
- NOW- PROBABILITY EVERYWHERE!
- NOW- Learning, Learning, Learning
- NOW- DEEP Learning

Overview of Intended Topics

1. Introduction to AI (Chs. 1-2, done)
2. Python (Python as a Second Language, S. Tanimoto)
3. Problem Solving by Search (Ch 3) “Big Chapter”
4. Beyond Classical Search (Ch 4)
5. Adversarial Search (Ch 5) “Game Playing”
6. Constraint Satisfaction Problems (Ch 6)
7. Learning (related to Ch 18 and Ch 20)
8. Computer Vision with Learning (not from book)
9. Deep Learning (not from book)
10. Latest Advances in Deep Learning
11. Other Applications such as Medical, LLMs, Robotics

Tentative Assignments

- Warmup Blind Search
- Heuristic Search
- Game Playing

} Python

- PyTorch Deep Learning Exercise
- Course Project of Your Choice using ML
- Readings and Problem Set