Candy Grab Game

As you walk in:

1. Grab a pack of game pieces (candy)
2. Form groups of 2 (or 3 with an observer)
3. Play the game!
   A. 11 pieces on the table
   B. Take turns taking either 1 or 2 pieces
   C. Person that takes the last piece wins!
4. Think about how you might implement an Agent to play this in code:

```
class Agent
    function getAction(state)
        return action
```

AI: Representation and Problem Solving

Introduction

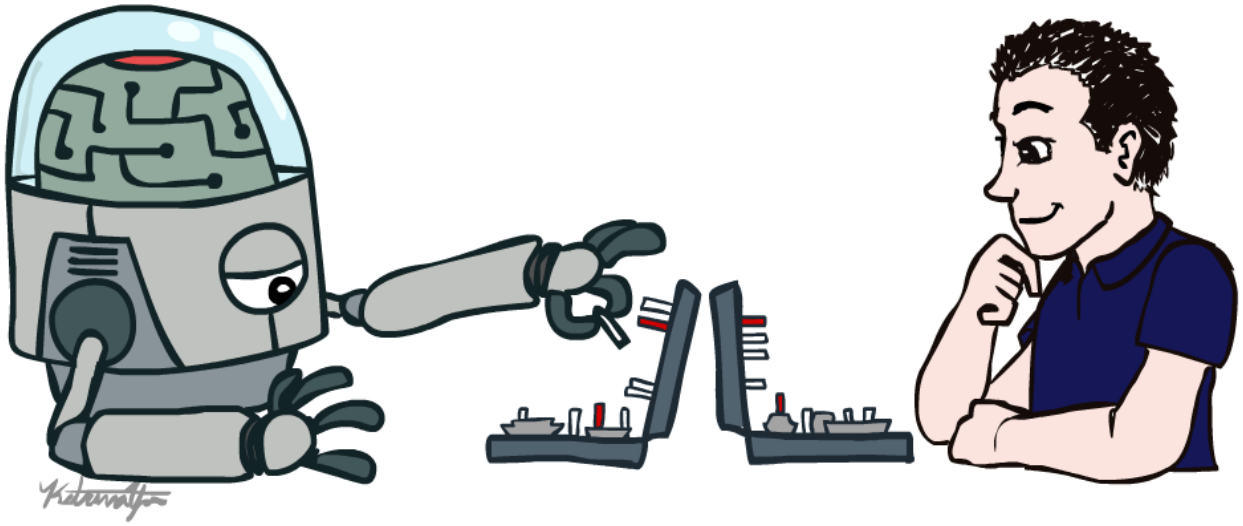

Instructors: Pat Virtue & Stephanie Rosenthal

Slide credits: Pat Virtue, http://ai.berkeley.edu

Course Staff

Instructors

Pat Virtue

Teaching Assistants

Ajay Kumar

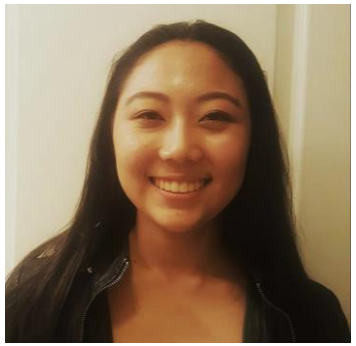

Claire Wang

Ethan Gruman

Lubhaan Kumar

Stephanie Rosenthal

Michelle Ma

Pallavi Koppol

Sean Pereira

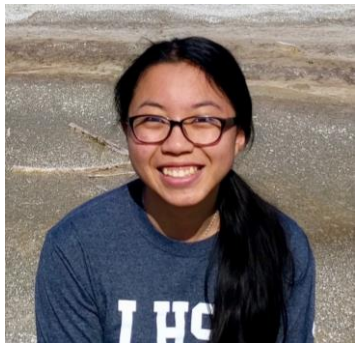

Tina Wu

Course Information

Website: https://www.cs.cmu.edu/~15381

Canvas: canvas.cmu.edu

Gradescope: gradescope.com

Communication: piazza.com

E-mail: pvirtue@cmu.edu

Prerequisites/Corequisites

Lin Alg, Calc 2

Course Scope

Announcements

Recitation starting this Fri 3pm, GHC 4401 (recommended)

No class next Mon 1/21, MLK Holiday

Assignments:

- HW1 (online)
  - Released tomorrow
  - Due Tue 1/22, 10 pm
- P0: Python & Autograder Tutorial
  - Required, but worth zero points
  - Due Thu 1/24, 10 pm

Today

What is artificial intelligence?

A brief history of AI

AI applications and techniques

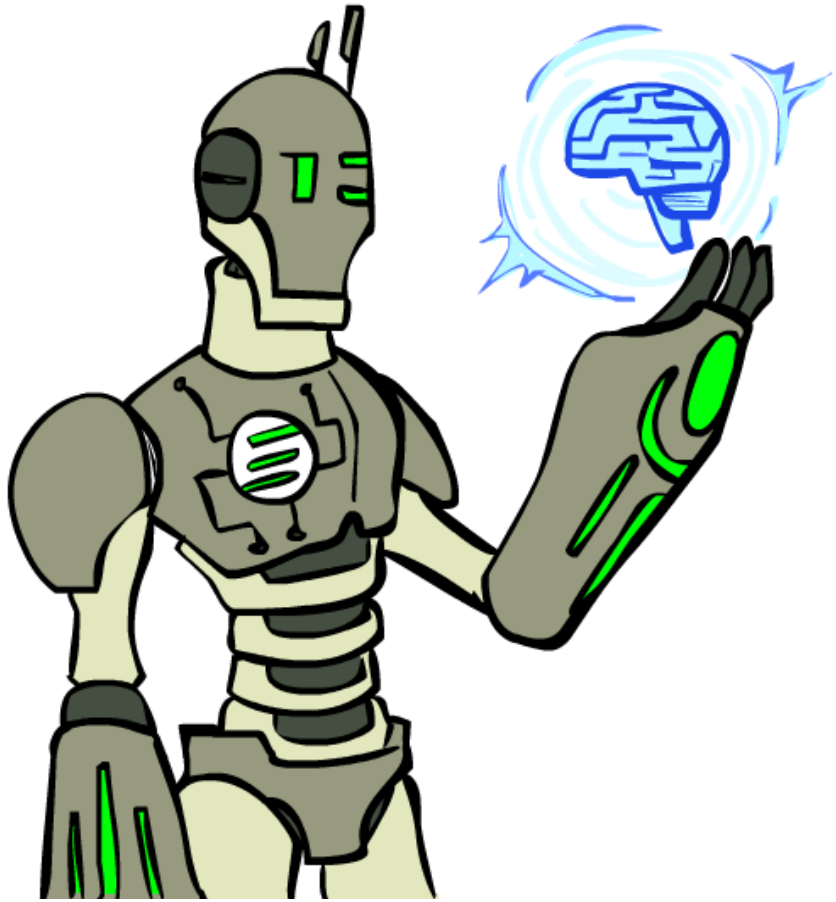

Designing Rational Agents

An **agent** is an entity that *perceives* and *acts*.

A **rational agent** selects actions that maximize its (expected) **utility**.

Characteristics of the **percepts, environment,** and **action space** dictate techniques for selecting rational actions

This course is about:

- General AI techniques for a variety of problem types
- Learning to recognize when and how a new problem can be solved with an existing technique

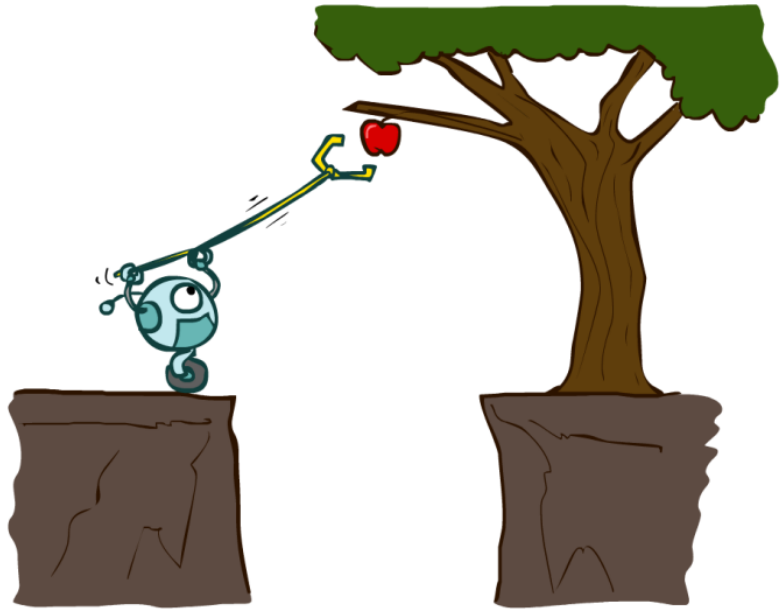

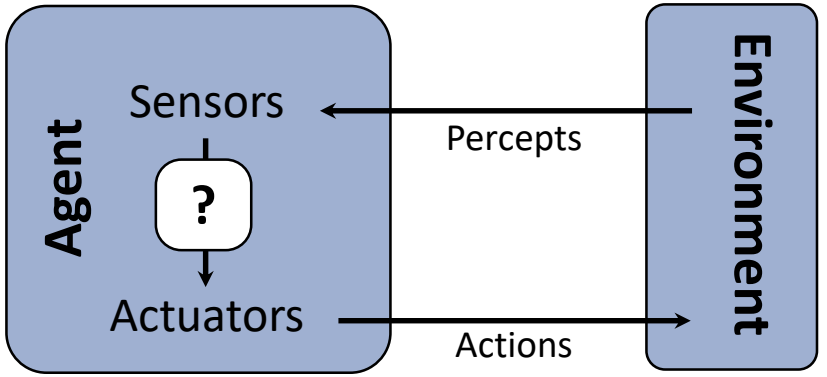

Pac-Man as an Agent

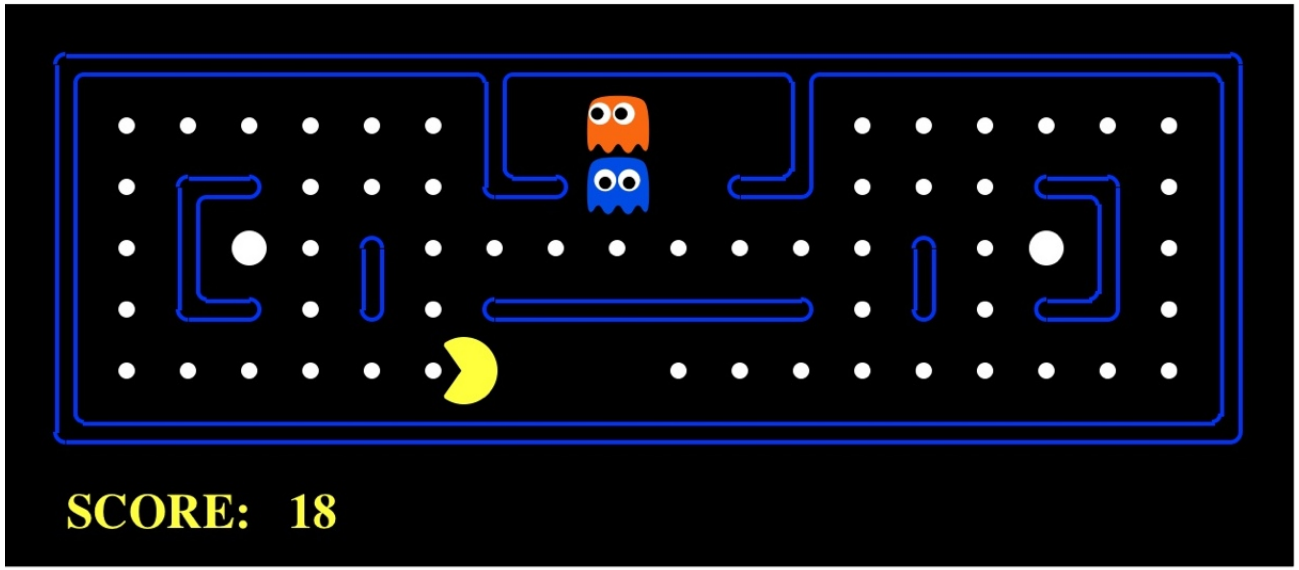

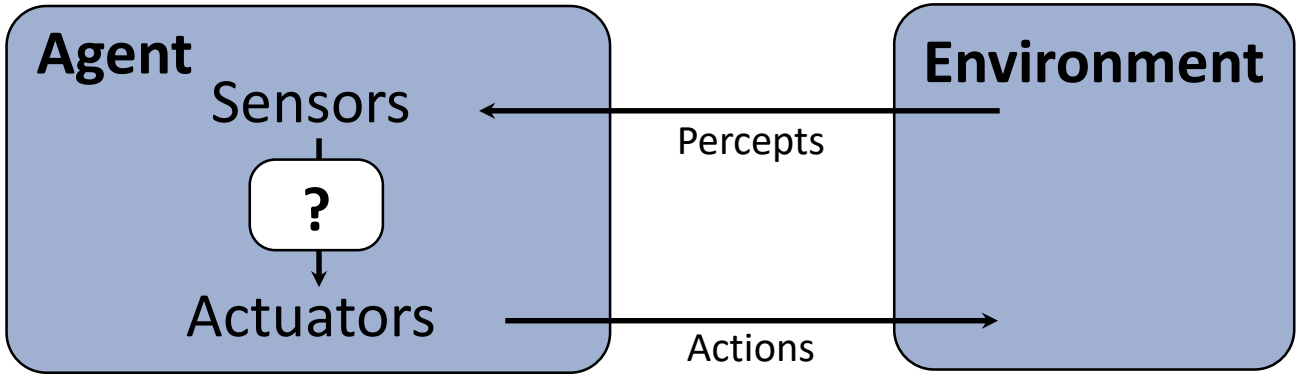

Pac-Man is a registered trademark of Namco-Bandai Games, used here for educational purposes

Pac-Man as an Agent

Candy Grab Agent

```
class Agent

    function getAction(state)

        return action
```

Candy Grab Agent

Agent 001 – Always choose 1

```
function getAction(                    )

    return 1
```

Candy Grab Agent

Agent 002 – Always choose 2

```
function getAction( numPiecesAvailable )

    return 2
```

Candy Grab Agent

Agent 004 – Choose the opposite of opponent

```
function getAction( numPiecesAvailable )

    return ?
```

Candy Grab Agent

Agent 007 – Whatever you think is best

```
function getAction( numPiecesAvailable )

    return ?
```

Candy Grab Agent

Agent 007 – Whatever you think is best

```
function getAction( numPiecesAvailable )

    if numPiecesAvailable % 3 == 2
        return 2
    else
        return 1
```

Piazza Poll question

Games – Three “Intelligent” Agents

Which agent code is the most “intelligent”?

Games – Three “Intelligent” Agents

A: Search / Recursion

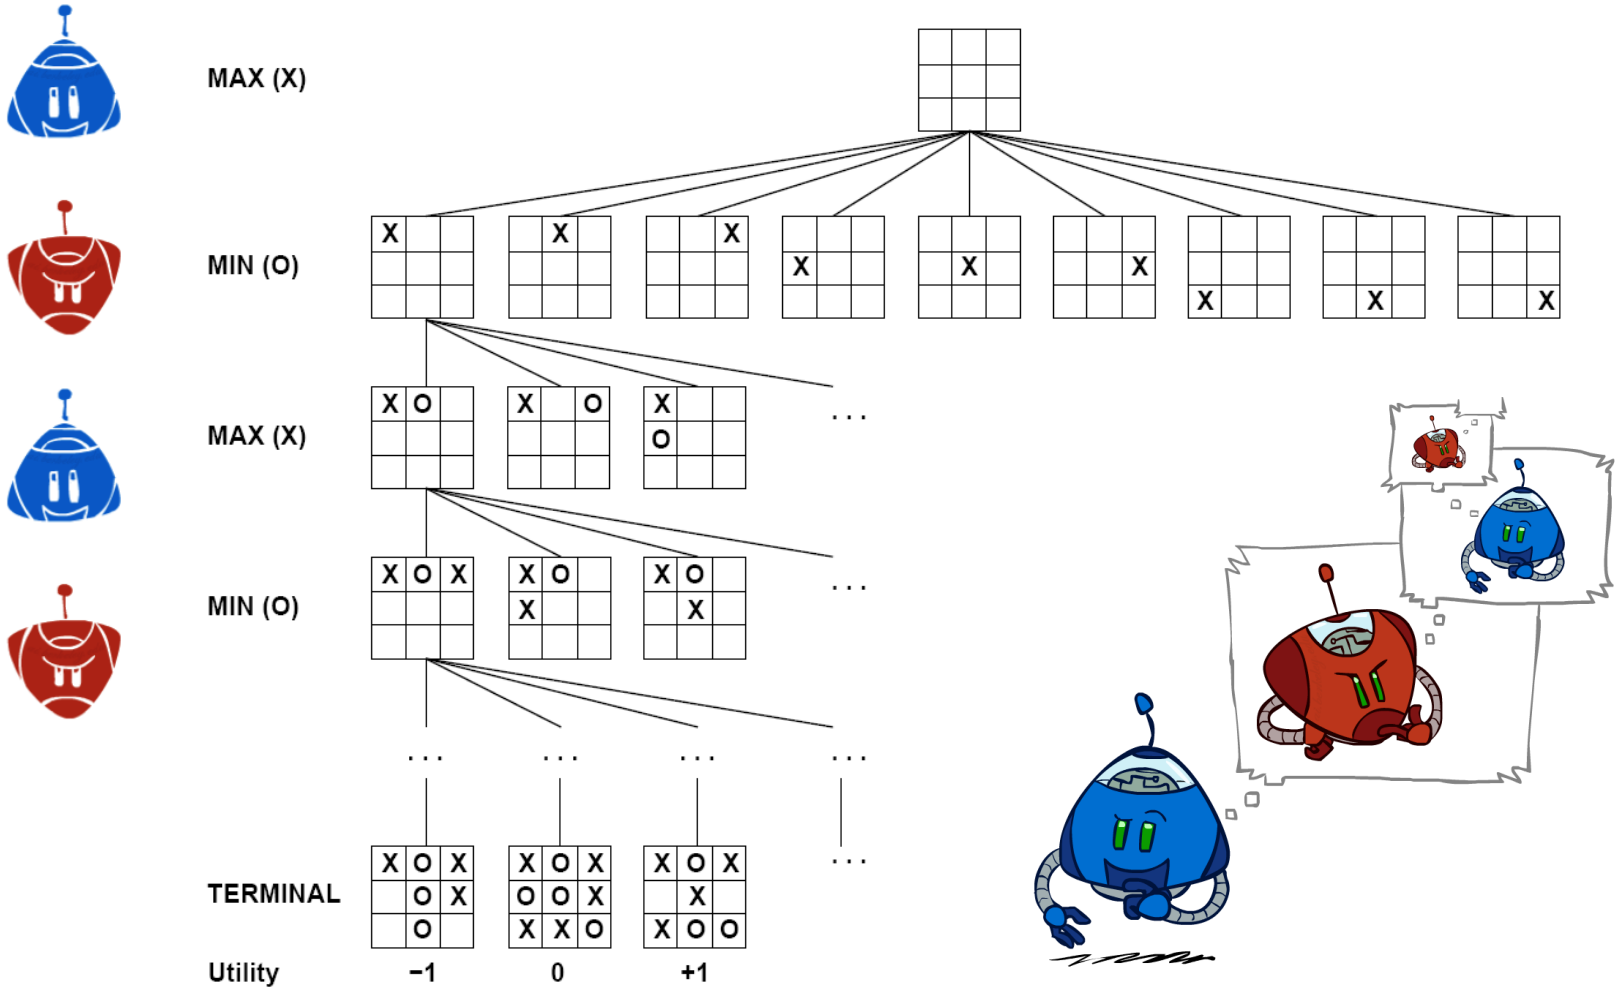

Games – Three “Intelligent” Agents

B: Encode the pattern

```
function getAction( numPiecesAvailable )

      if numPiecesAvailable % 3 == 2
            return 2
      else
            return 1
```

```
10's value:Win
9's value:Lose
8's value:Win
7's value:Win
6's value:Lose
5's value:Win
4's value:Win
3's value:Lose
2's value:Win
1's value:Win
0's value:Lose
```

Games – Three “Intelligent” Agents

C: Record statistics of winning positions

Pieces Available	Take 1	Take 2
2	0%	**100%**
3	2%	0%
4	**75%**	2%
5	4%	**68%**
6	5%	6%
7	**60%**	5%

Piazza Poll question

Games – Three “Intelligent” Agents

Which agent code is the most “intelligent”?

A. Search / Recursion

B. Encode multiple of 3 pattern

C. Keep stats on winning positions

Games – Three “Intelligent” Agents

C: Record statistics of winning positions

Pieces Available	Take 1	Take 2
2	0%	**100%**
3	2%	0%
4	**75%**	2%
5	4%	**68%**
6	5%	6%
7	**60%**	5%

AI in the News

https://www.youtube.com/watch?v=EfGD2qveGdQ

Sci-Fi AI?

What is AI?

The science of making machines that:

Turing Test

In 1950, Turing defined a test of whether a machine could “think”

“A human judge engages in a natural language conversation with one human and one machine, each of which tries to appear human. If judge can’t tell, machine passes the Turing test”

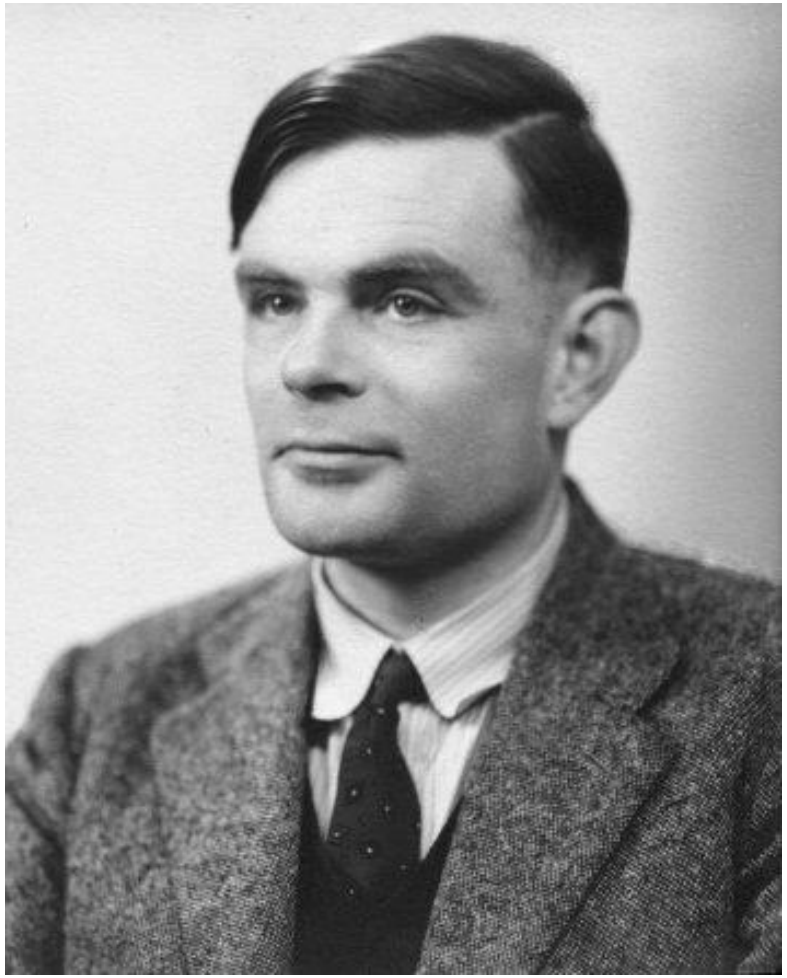

en.wikipedia.org/wiki/Turing_test

What is AI?

The science of making machines that:

Think like people

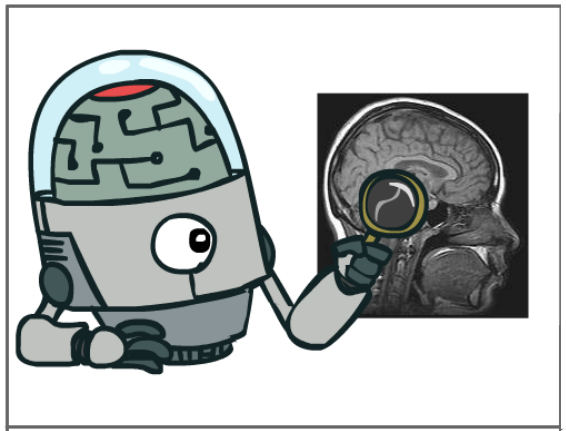

Act like people

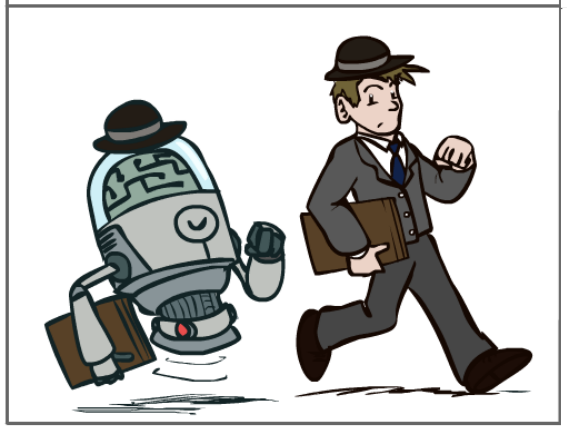

AI Definition by John McCarthy

What is artificial intelligence

- It is the science and engineering of making intelligent machines, especially intelligent computer programs

What is intelligence

- Intelligence is the computational part of the ability to achieve goals in the world

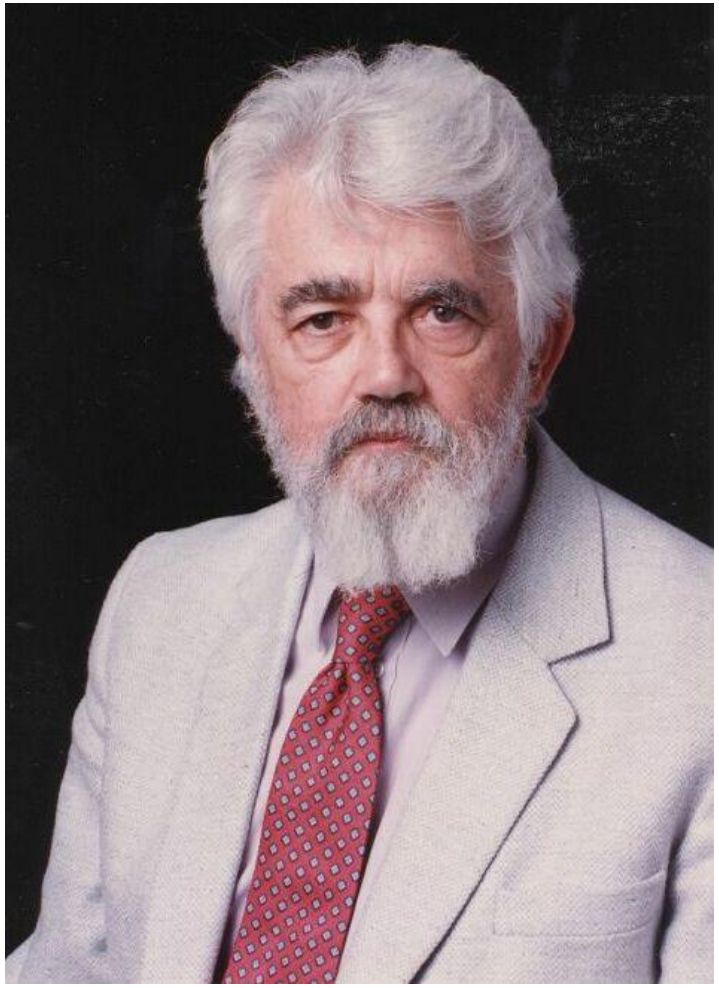

http://www-formal.stanford.edu/jmc/whatisai/whatisai.html

Rational Decisions

We'll use the term **rational** in a very specific, technical way:

- Rational: maximally achieving pre-defined goals
- Rationality only concerns what decisions are made (not the thought process behind them)
- Goals are expressed in terms of the **utility** of outcomes
- Being rational means **maximizing your expected utility**

A better title for this course would be:

Computational Rationality

Maximize Your Expected Utility

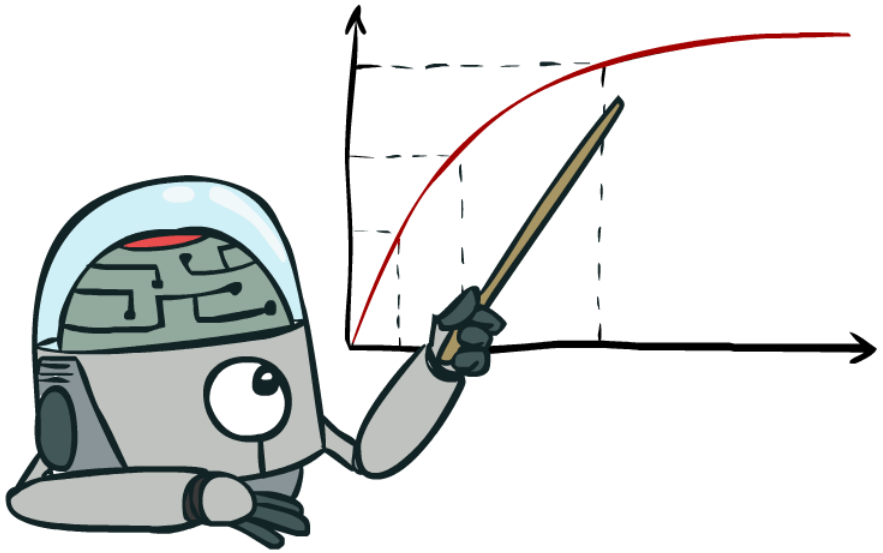

What About the Brain?

- Brains (human minds) are very good at making rational decisions, but not perfect
- Brains aren't as modular as software, so hard to reverse engineer!
- "Brains are to intelligence as wings are to flight"
- Lessons learned from the brain: memory and simulation are key to decision making

A Brief History of AI

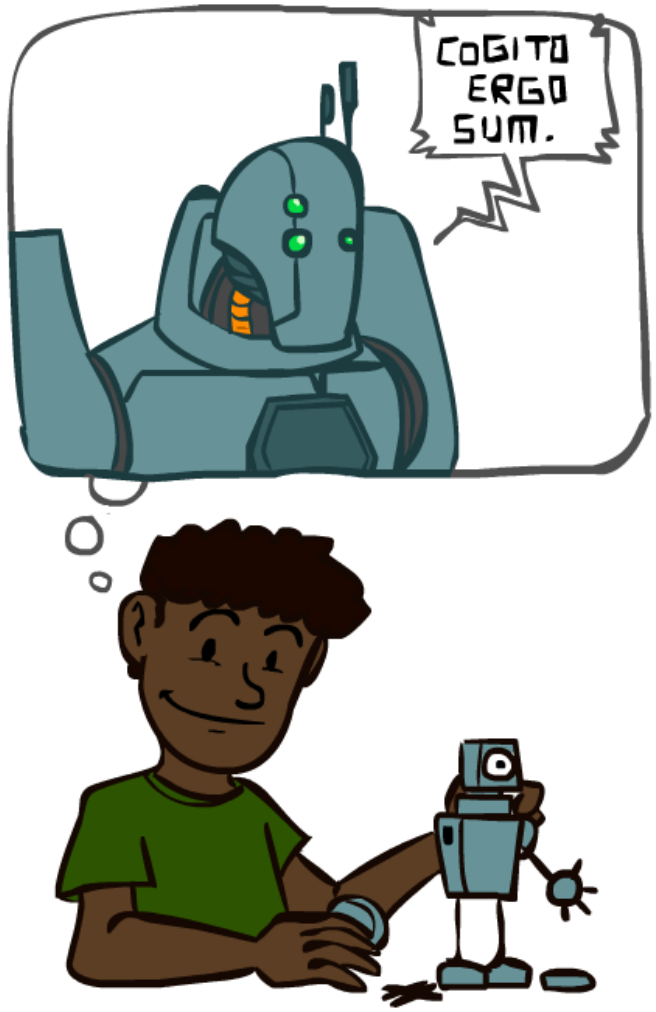

A Brief History of AI

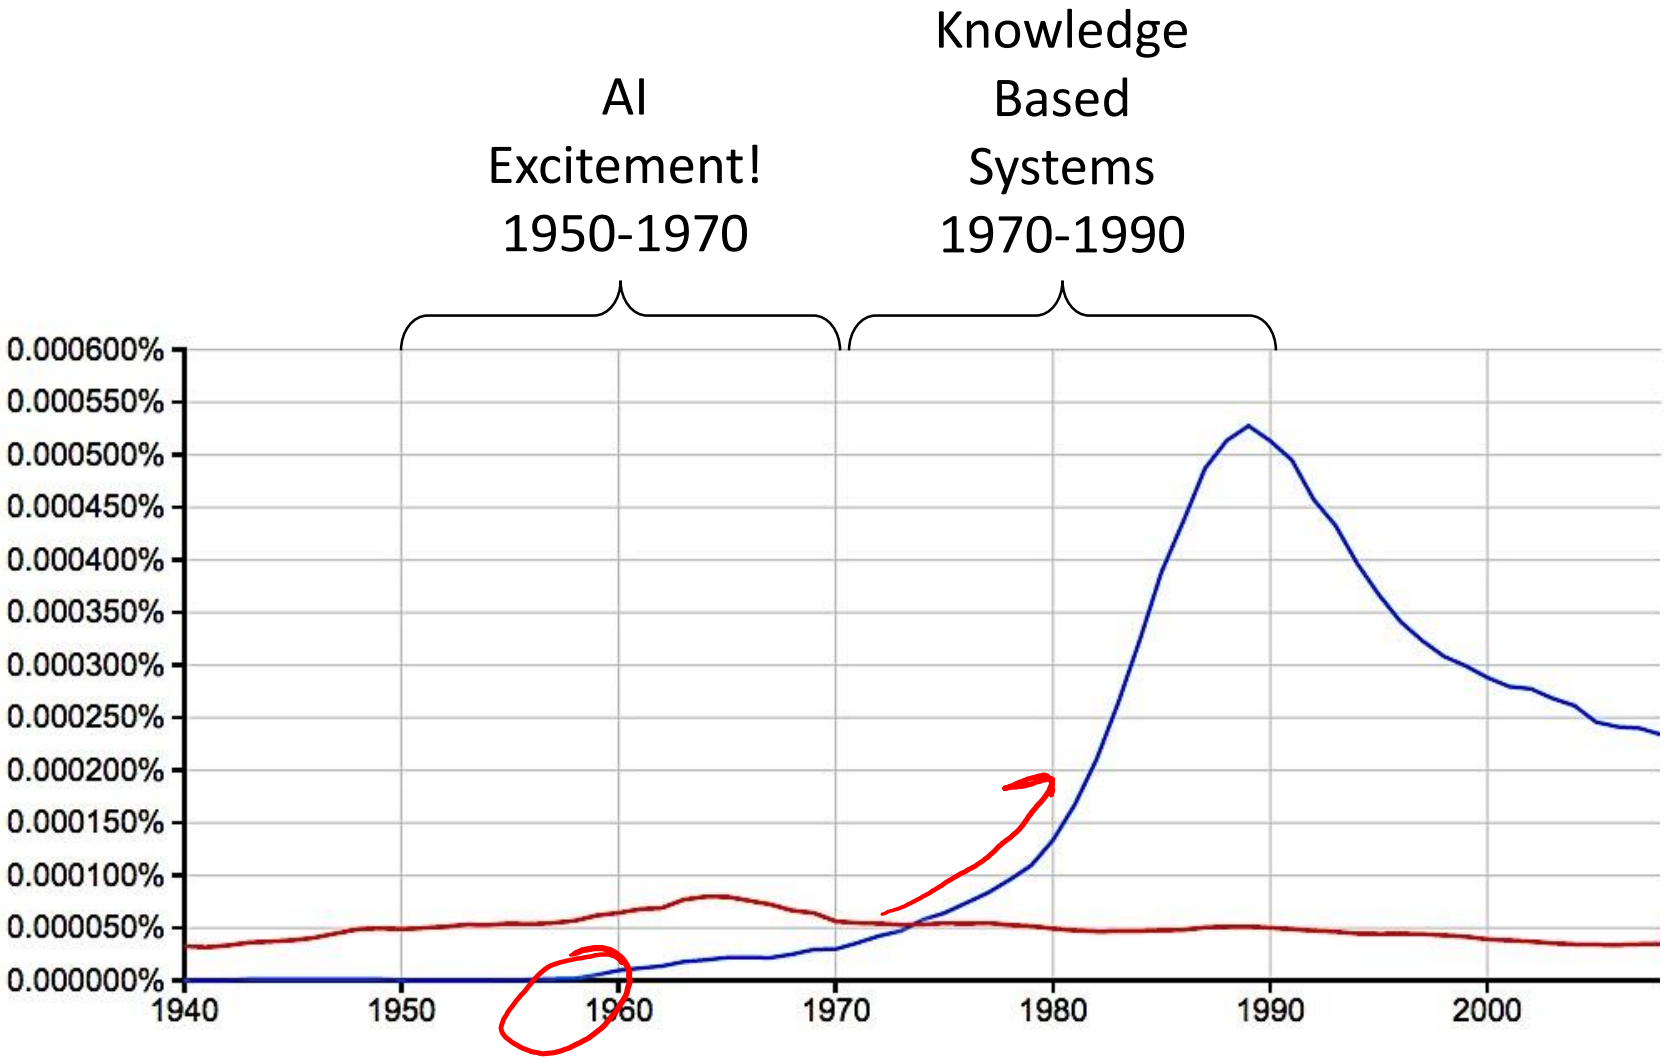

https://books.google.com/ngrams

What went wrong?

Dog	Buster
Barks	✓
Has Fur	✓
Has four legs	✗

A Brief History of AI

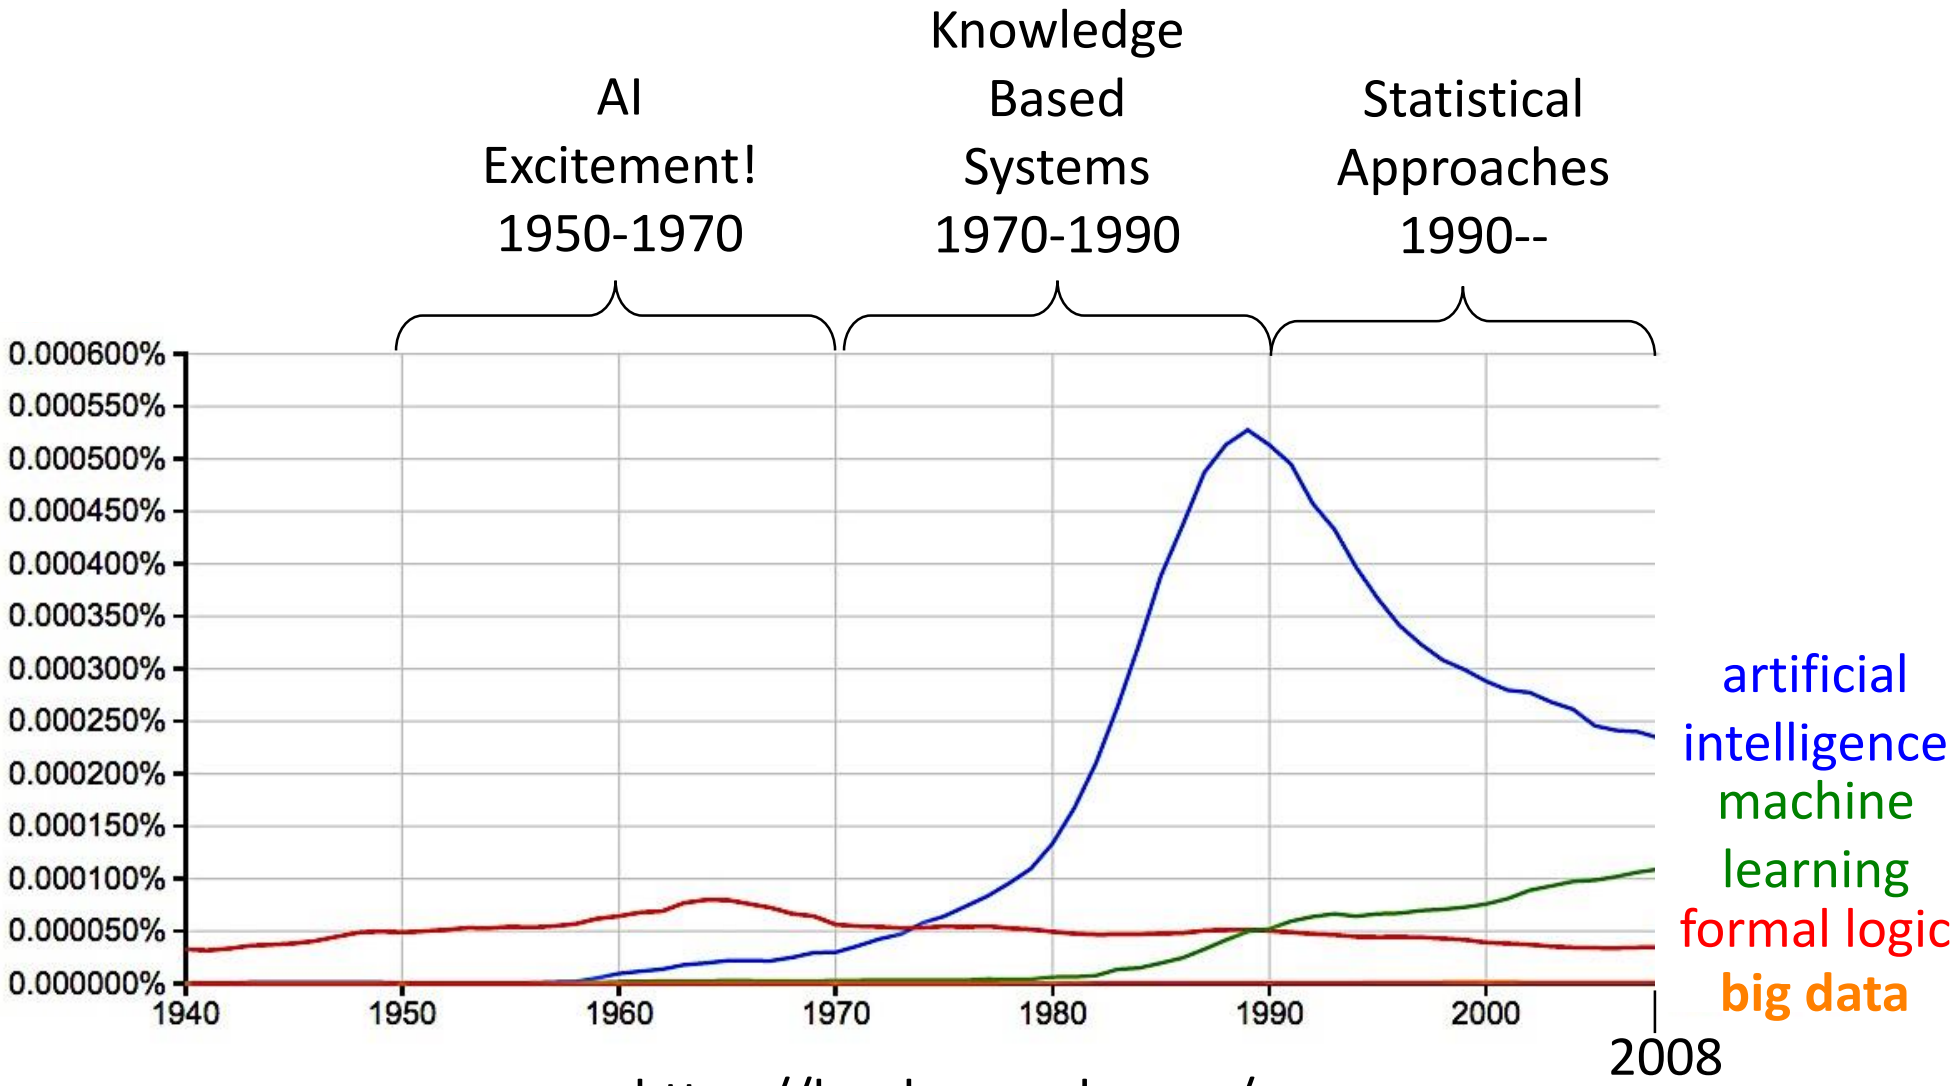

https://books.google.com/ngrams

A Brief History of AI

1940-1950: Early days

- 1943: McCulloch & Pitts: Boolean circuit model of brain
- 1950: Turing's “Computing Machinery and Intelligence”

1950—70: Excitement: Look, Ma, no hands!

- 1950s: Early AI programs, including Samuel's checkers program, Newell & Simon's Logic Theorist, Gelernter's Geometry Engine
- 1956: Dartmouth meeting: “Artificial Intelligence” adopted
- 1965: Robinson's complete algorithm for logical reasoning

1970—90: Knowledge-based approaches

- 1969—79: Early development of knowledge-based systems
- 1980—88: Expert systems industry booms
- 1988—93: Expert systems industry busts: “AI Winter”

1990—: Statistical approaches

- Resurgence of probability, focus on uncertainty
- General increase in technical depth
- Agents and learning systems… “AI Spring”?

2012—: Where are we now?

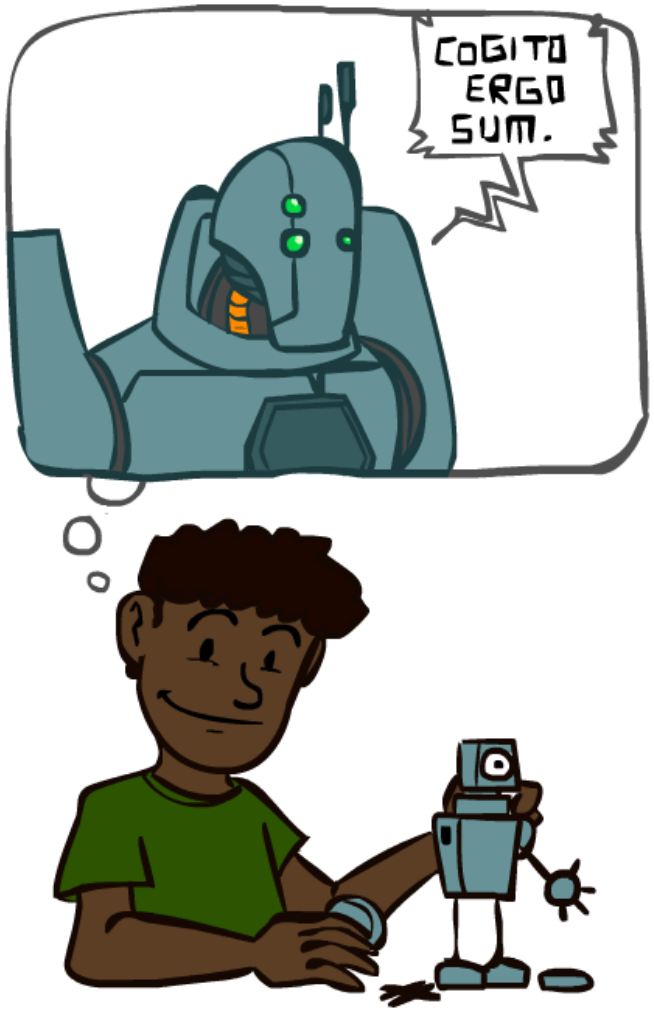

What Can AI Do?

Quiz: Which of the following can be done at present?

- ✓ Play a decent game of table tennis?
- ✓ Play a decent game of Jeopardy?
- ✓ Drive safely along a curving mountain road?
- ? Drive safely across Pittsburgh?
- ✓ Buy a week's worth of groceries on the web?
- ✗ Buy a week's worth of groceries at Giant Eagle?
- ? Discover and prove a new mathematical theorem?
- ✗ Converse successfully with another person for an hour?
- ? Perform a surgical operation?
- ✓ Put away the dishes and fold the laundry?
- ✓ Translate spoken Chinese into spoken English in real time?
- ✗ Write an intentionally funny story?