CMU 15-781

Lecture 1: Introduction

Teacher Section C:
Gianni Di Caro

Section A/B:
Emma Brunskill
Ariel Procaccia

AI?

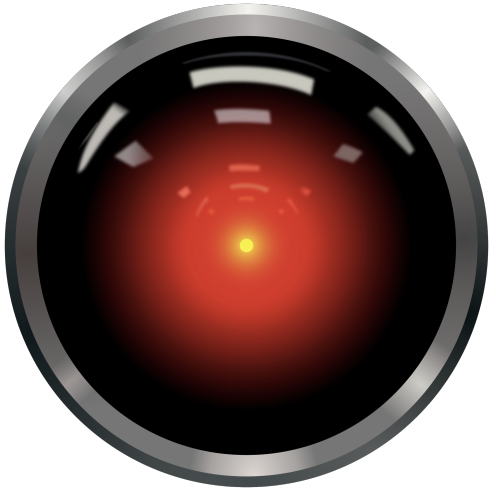

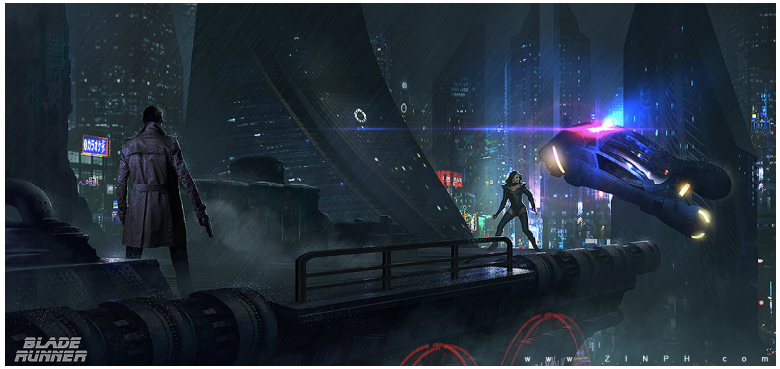

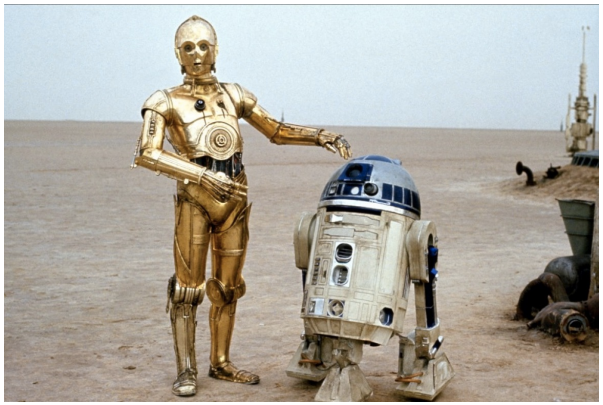

AI?

Many different views ... let's start with some tentative definitions

The science of making machines do things that would require *intelligence* if done by man (B. Raphael)

Intelligence?

The cognitive ability of an individual (entity) to *learn* from experience, to *reason* well, to *remember* important information, and to (*effectively*) cope with the *demands* of daily living (R. Sternberg)

AI TIMELINE (NYT 2011)

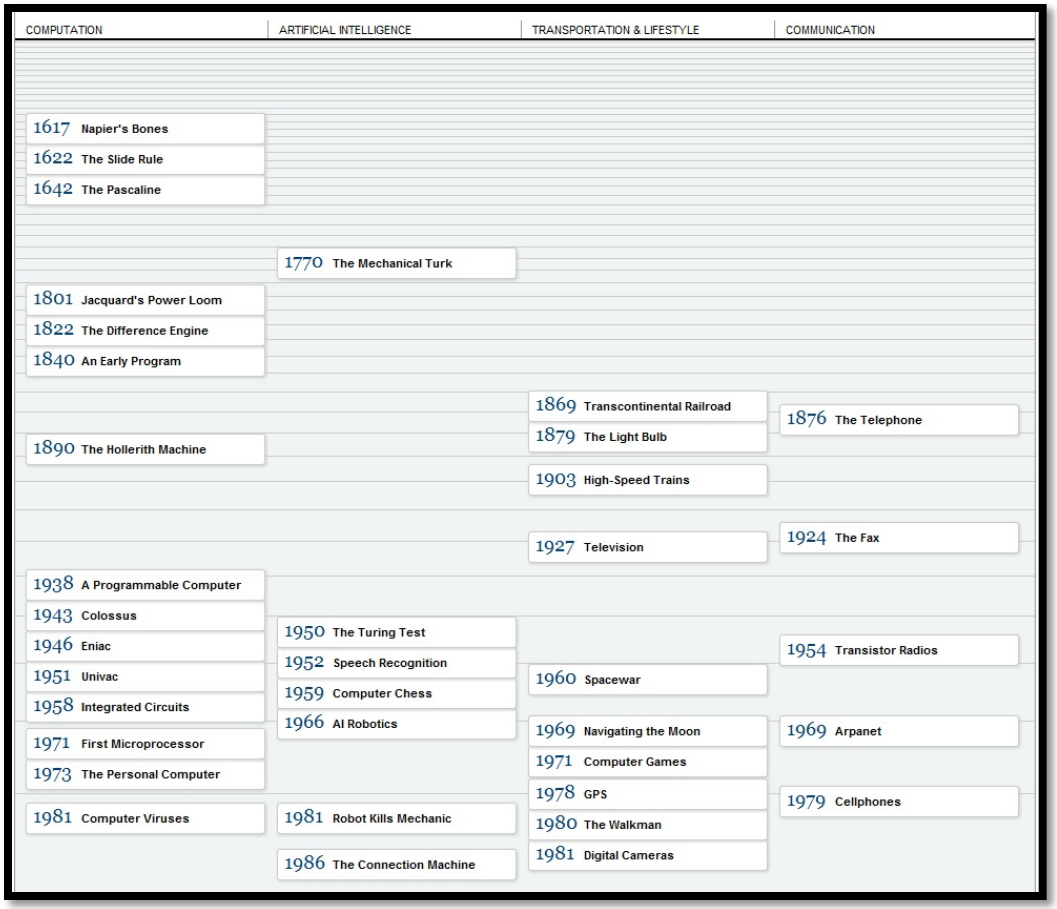

1770 The Mechanical Turk

A mechanical chess-playing machine awes the world, but is revealed decades later to have contained a human chessmaster hidden inside the device.

- Performed for 84 years
- Defeated Napoleon and Franklin
- Amazon Mechanical Turk: "artificial artificial intelligence"

Amazon Mechanical Turk

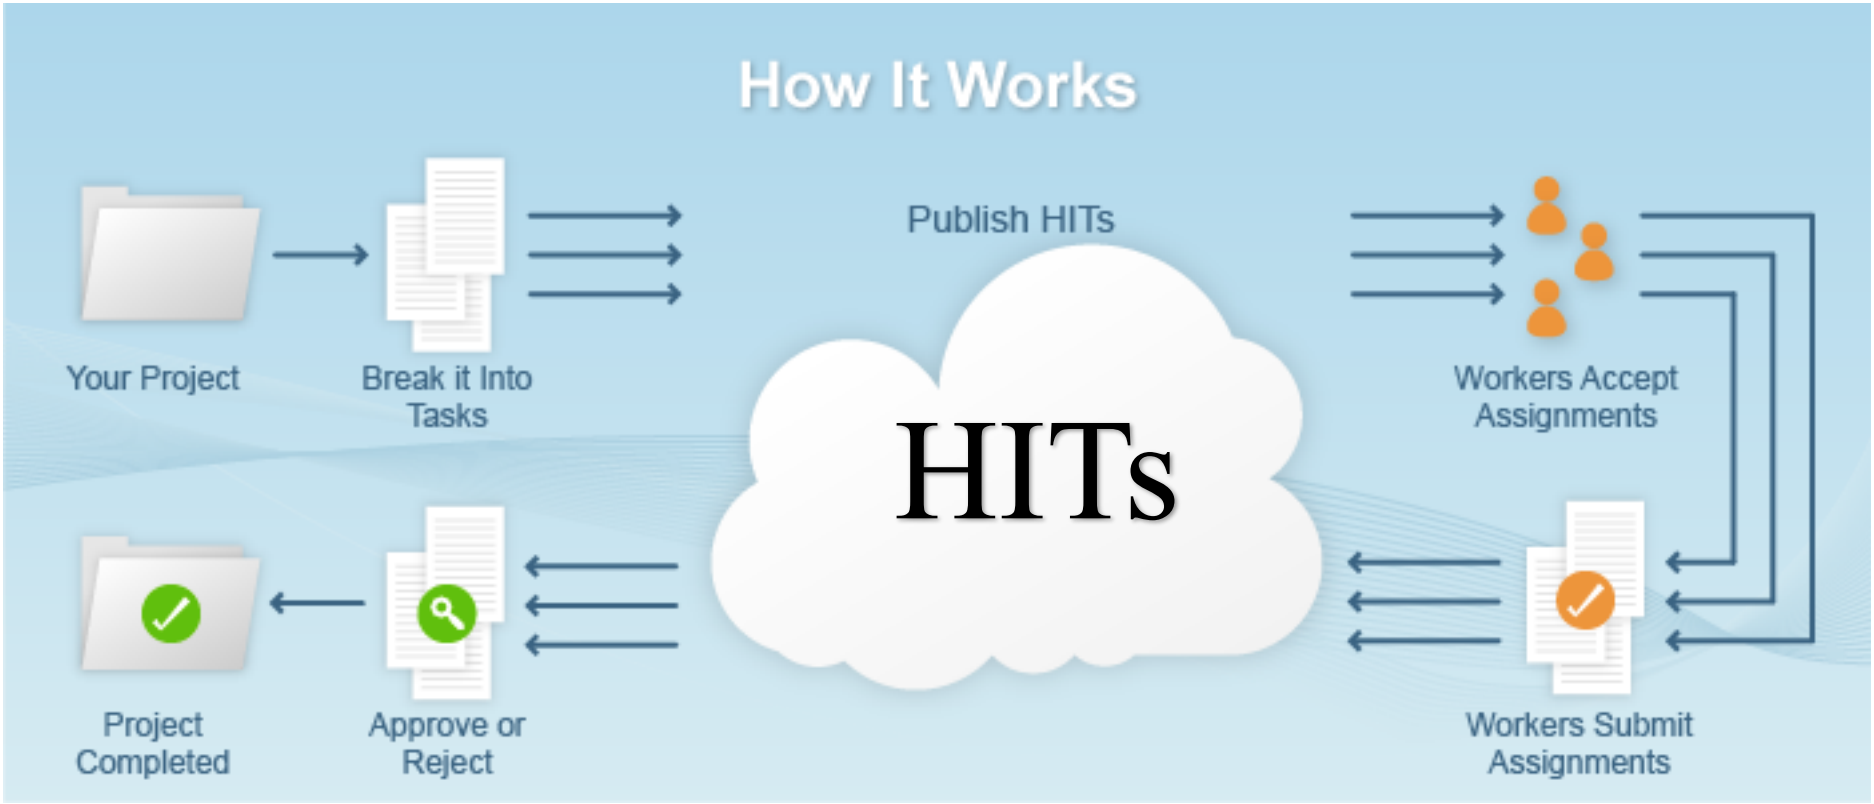

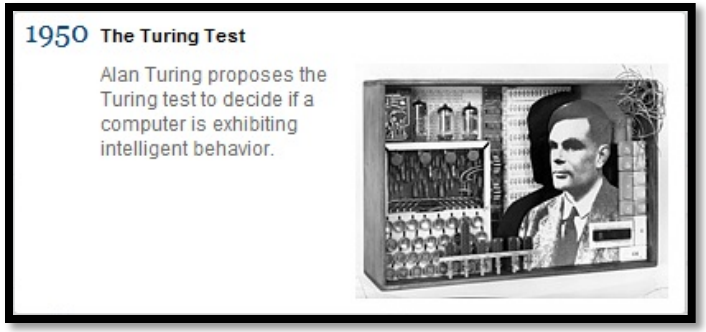

- Big question: **can machines think**?
- *Too difficult!* Let's separate functionality from implementation
- More concrete question: can machines do well in the (human) imitation game?
- Judge communicates via text channel with computer and human, must reliably identify the computer (Total Turing test involves physical interface)

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? http://www.telegraph.co.uk/culture/tvandradio/bbc/10891699/John-Humphrys-grills-the-robot-who-passed-the-Turing-test-and-is-not-impressed.html

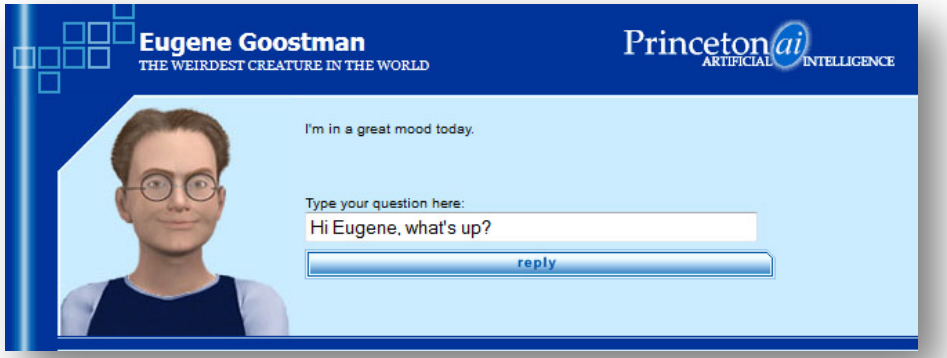

Eugene Goostman vs. Scott Aaronson

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.

The Chinese Room

- Suppose AI has produced a program that can pass the Turing Test in Chinese
- You have a handbook with its pseudocode, or just a rule book
- You're in a closed room and receive Chinese characters through a slot
- You run the program's code manually and return the output
- Does this mean you understand Chinese?

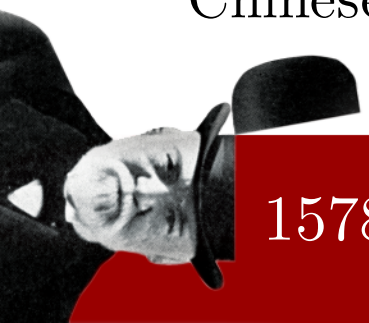

Counterarguments

- Finding the mind: the whole system understands Chinese, the person is just a part of the system
- Redesigning the experiment: suppose the program simulates the actions of every neuron in the brain of a Chinese speaker

1952 Speech Recognition

Bell Labs develops the first effective speech-recognition device using splitter technology similar to the one developed by Alexander Graham Bell 78 years earlier.

- "Audrey" could recognize digits spoken by a single voice
- In 1962 IBM demonstrated "Shoebox", which could understand 16 words
- Biggest milestone in the Seventies: CMU's "Harpy" system, which could understand 1011 words ~ vocabulary of three-year-old

The birth of AI

- 1956 workshop at Dartmouth
- Participants included Marvin Minsky, John McCarthy, Claude Shannon, Ray Solomonoff, Arthur Samuel, Allen Newell, Herbert Simon
- Proposal included this assertion: "every aspect of learning or any other feature of intelligence can be so precisely described that a machine can be made to simulate it"

1959 Computer Chess

Arthur Samuel's checkers program wins games against the best human players. 48 years later, the game of checkers is solved by computers.

- Samuel's program actually only competed at "respectable amateur" level
- By the Nineties the *Chinook* checker programs was "beating" the best human players (Marion Tinsley - Chinook: 4:2, 6 drawns)
- Checkers was *solved* by Jonathan Schaeffer in 2007 after 18 years of calculation

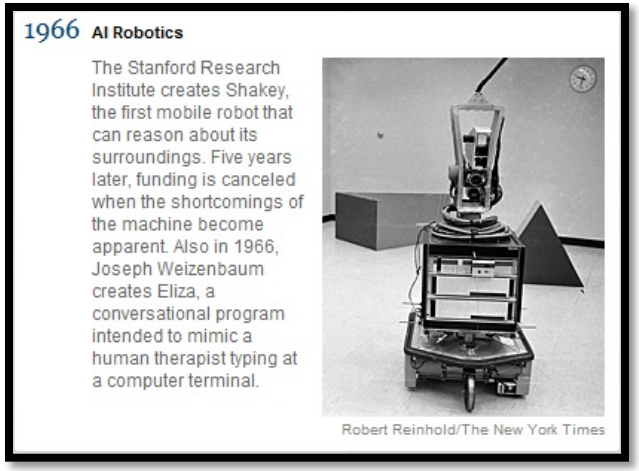

- Shakey: first mobile robot to visually interpret environment
- Can locate items, navigate around them, and reason about its actions
- http://www.youtube.com/watch?v=qXdn6ynwpiI (4:08)

Classical AI: Deliberative paradigm

Classical AI (60's - 80's) was based on the *deliberative* paradigm of human intelligence

To be **intelligent**, machines / robots have to be able to perform some intensive forms of *model-based "thinking"*

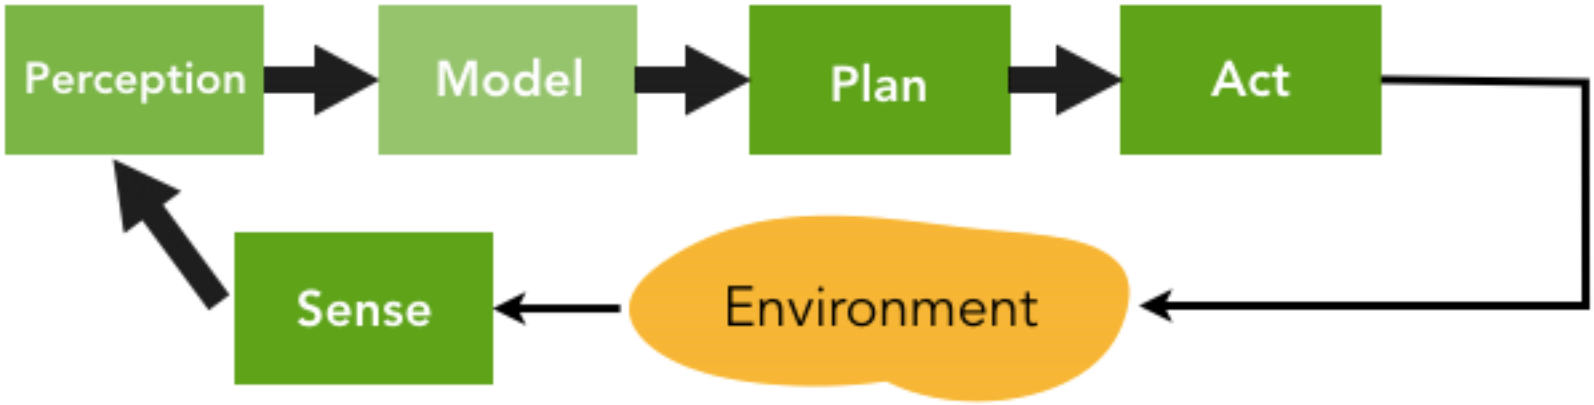

Reactive paradigm: Don't think, React!

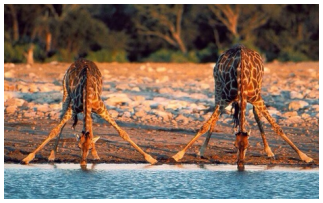

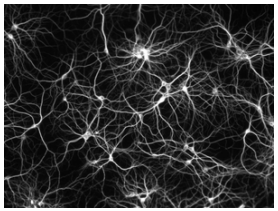

Bio-inspired, sub-symbolic, neural models ...

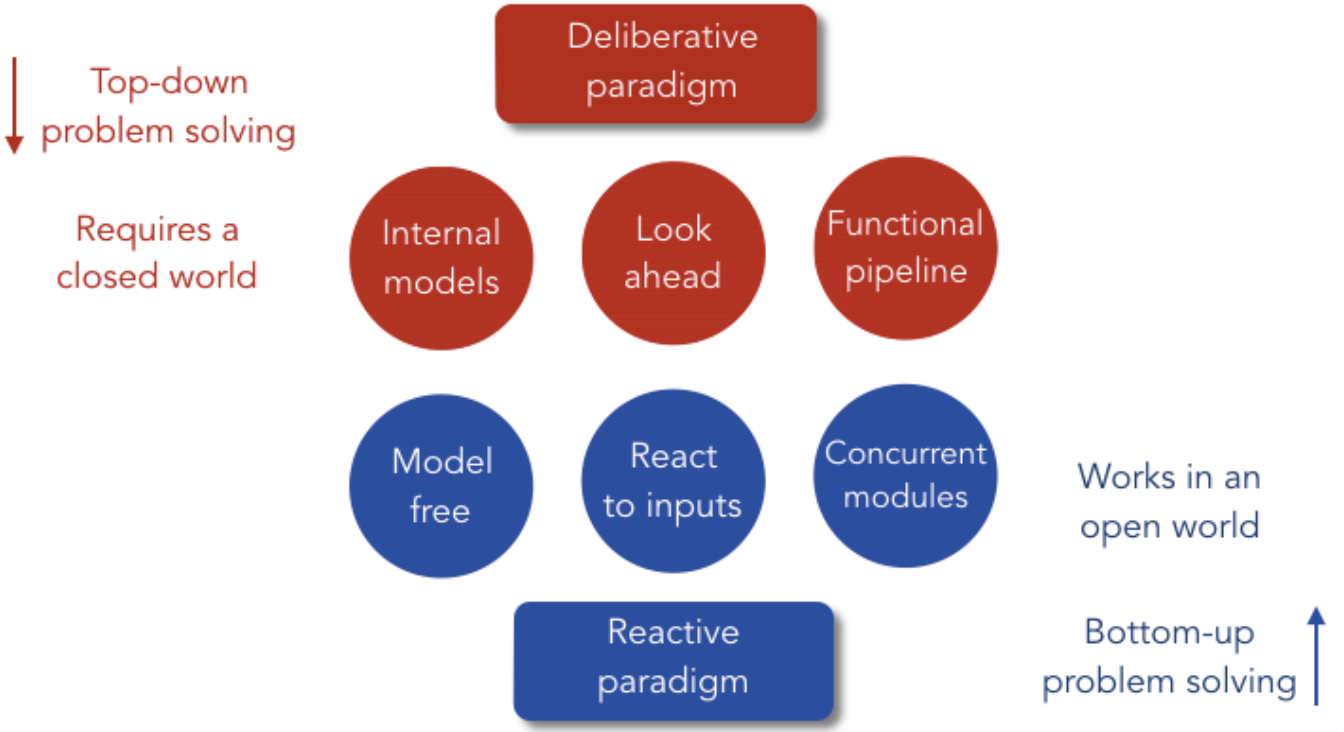

- Started as “ChipTest” at CMU, followed by “Deep Thought”
- After graduation, developers were hired by IBM
- Defeated Kasparov 3.5-2.5 in 1997
- Kasparov played anti-computer opening moves to get Deep Blue out of its opening book
- Kasparov accused IBM of cheating

- **A**dvanced **S**tep in **I**nnovative **Mo**bility (resemblance to Asimov is a coincidence)
- Can recognize moving objects, postures, gestures, its surrounding environment, sounds and faces, which enables it to interact with humans
- http://www.youtube.com/watch?v=NZngYDDDfW4

DARPA Urban Challenge

- 96 km urban area course, to be completed < 6 hours, took place in 2007
- Tartan Racing (CMU+GM) claimed the $2 million prize
- Challenge involves mission planning, motion planning, behavior generation, perception, world modeling
- http://www.youtube.com/watch?v=lULl63ERek0

- Watson defeated the two greatest-ever Jeopardy! champions
- Involves natural language processing, information retrieval, knowledge representation and reasoning, and machine learning
- http://www.youtube.com/watch?v=oUj9AzSE_9c

Go and deep learning

- In March 2016, AlphaGo beat the 9-dan player Lee Sedol 4-1
- It is based on deep learning and reinforcement learning
- Closer to general AI than Deep Blue or Watson

AI as Data science

- Data, data, data... to learn from!
- Sensors
- Internet
- Social nets
- Clouds
- Smartphones
- Fast computers
- GPUs

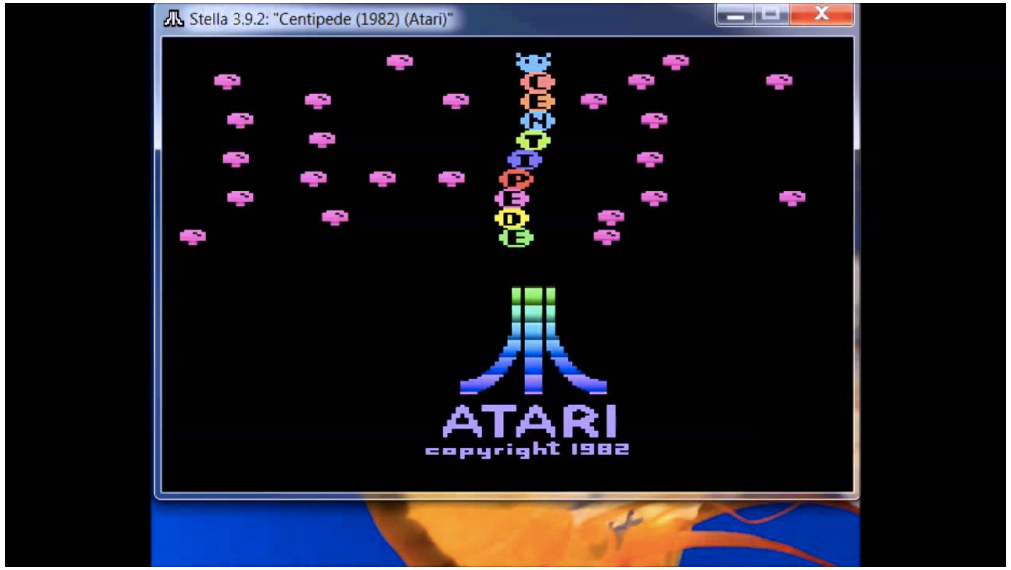

DeepMind / Google

http://www.nature.com/nature/journal/v518/n7540/abs/nature14236.html

AI IN EVERYDAY LIFE

Which of these apps on your phone heavily rely on AI?





The future

AI SAFETY

- Elon Musk: AI is "our greatest existential threat."
- Stephen Hawking: "Success in creating AI would be the biggest event in human history. Unfortunately, it might also be the last..."
- Bill Gates: "First, the machines will do a lot of jobs for us and not be super intelligent. That should be positive if we manage it well. A few decades after that, though, the intelligence is strong enough to be a concern."

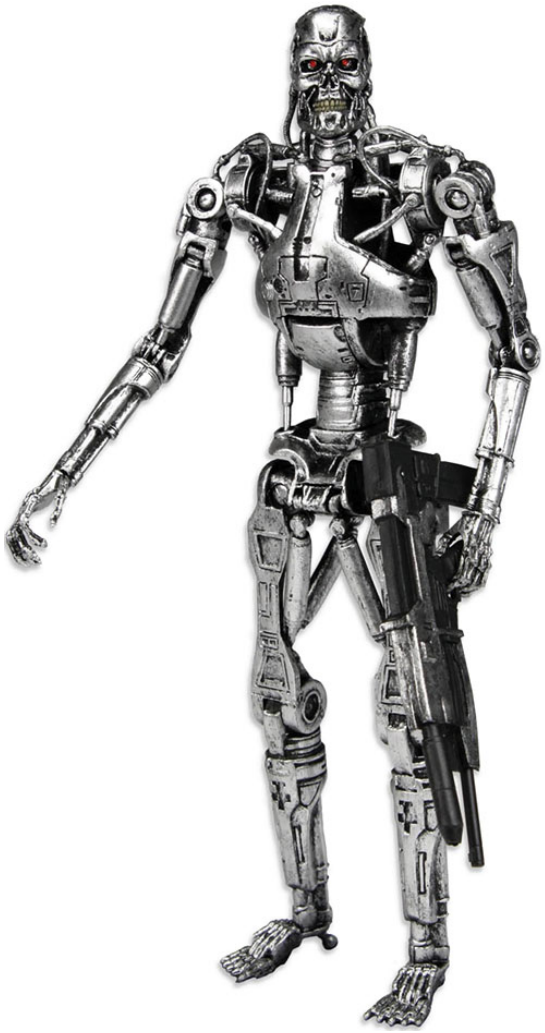

The technological singularity

- Emergence of superhuman intelligence
- Key idea: self-improvement
- Source of name: analogy between inability to predict events after the development of a superintelligence, and the space-time singularity beyond the event horizon of a black hole
- Some predict: this century
- Others argue: never

The Singularity in Movies

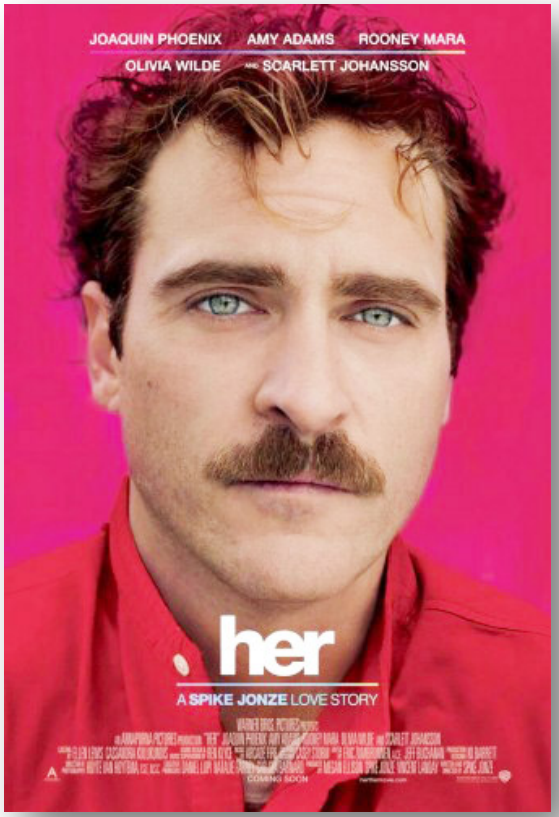

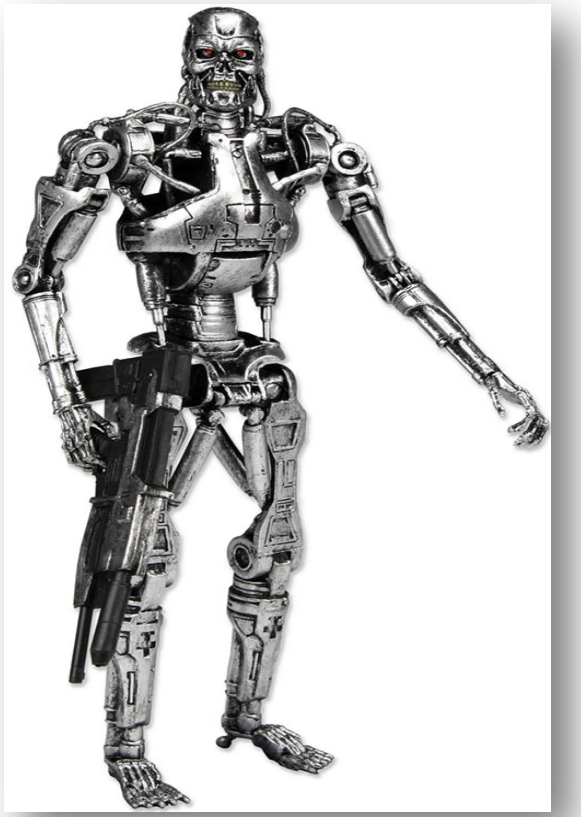

What Things Are Still Unsolved?

Approaches to AI

The science of making machines that:

Think like people

Act like people

Thought processes and reasoning

Fidelity to human performance

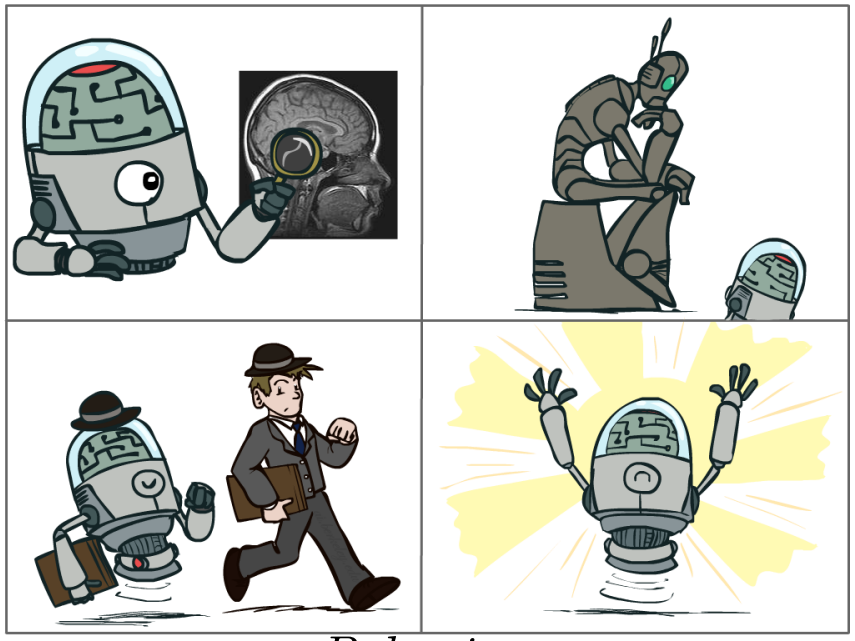

Ideal performance

Think rationally

Act rationally

Behavior

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 expected utility

A reasonable alternate title for this course would be:

Computational Rationality

Maximize Your Expected Utility

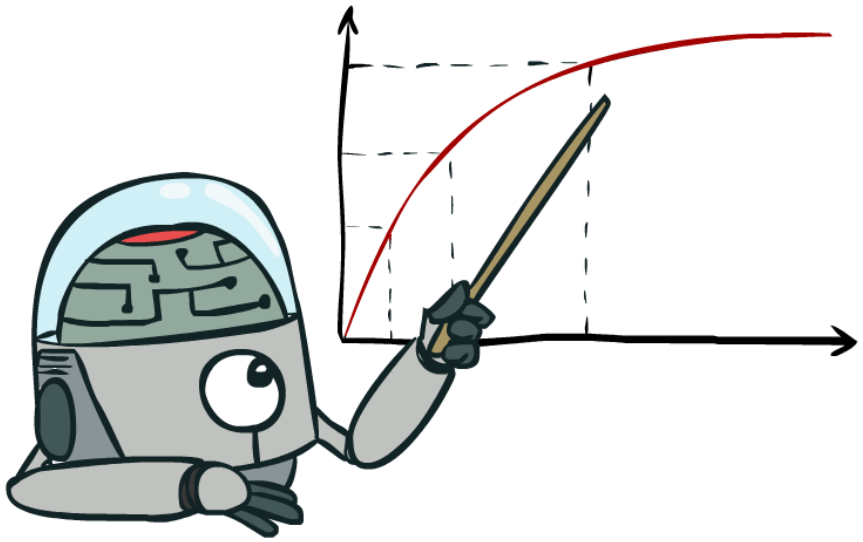

What about the Brain?

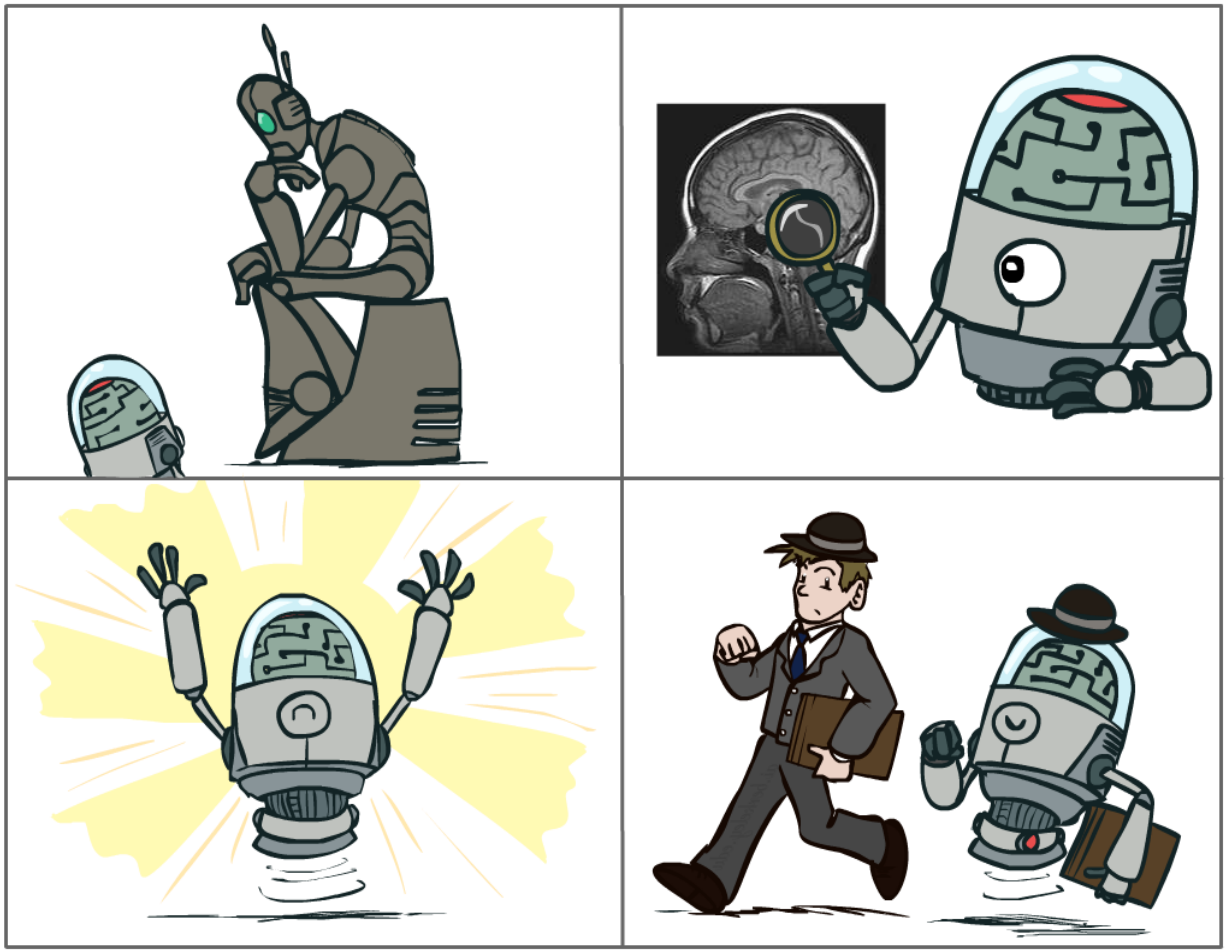

Slide adapted from Abbeel&Klein

Designing Rational Agents

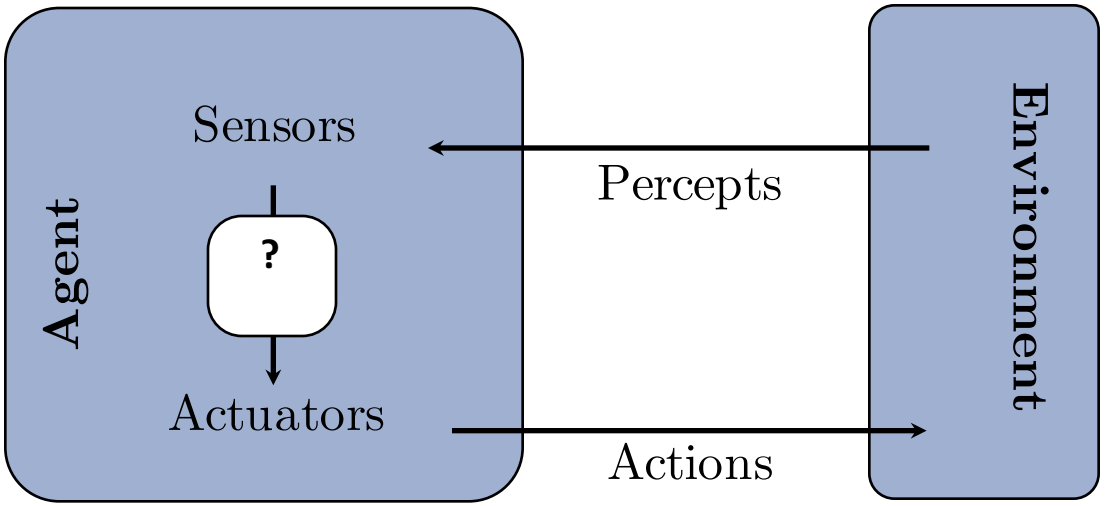

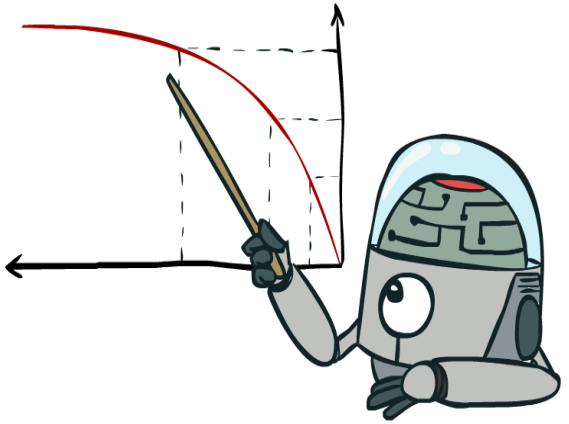

- This course is about:
  - General AI techniques for a variety of problem types
  - Learning to recognize when and how a new problem can be solved with an existing technique

Course Topics

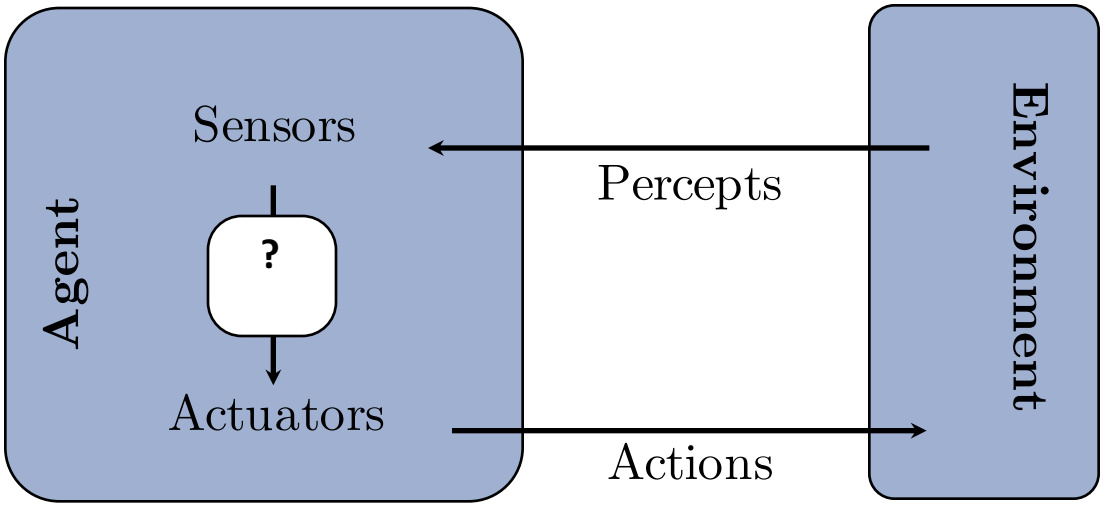

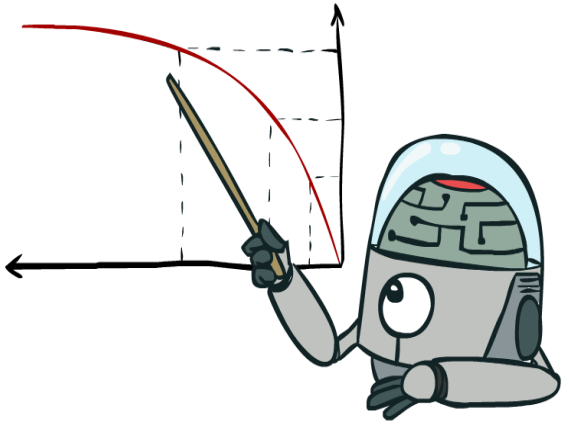

1) Making decisions without uncertainty (search, CSPs, planning)
2) Reasoning under uncertainty (bandits, decision theory, Bayes' nets)
3) Multiple agents (game theory, social choice, swarm intelligence)
4) Machine learning

Learning Objectives

By the end of the course you should be able to

1. Identify the type of an AI problem.
2. Formulate the problem as a particular type.
3. Compare the difficulty of different versions of AI problems, in terms of computational complexity and the efficiency of existing algorithms.
4. Implement, evaluate and compare the performance of various AI algorithms. Evaluation could include empirical demonstration or theoretical proofs.

Course Staff

Professors

Gianni Di Caro

Emma Brunskill

Ariel Procaccia

Teaching Assistants

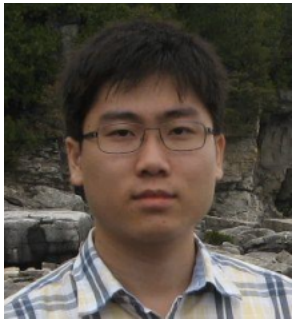

Zhaohan (Daniel) Guo

Klas Leino

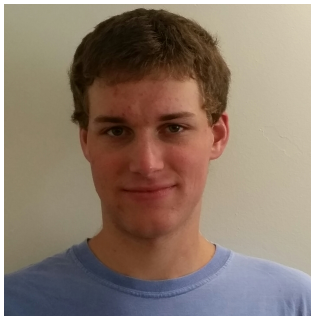

Andrew Pratt

Nicolas Resch

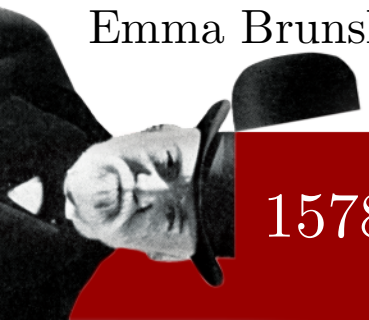

Course Communication

- Class website: http://www.cs.cmu.edu/~15381/
  - Contains office hours, lecture list, policies
- Piazza: used for questions, announcements, and polls
  - Sign up and download Piazza app

Course Resources & Tech

- Resources
  - Live lectures this fall
- Technology
  - Autograded projects, competitions, regular homeworks and for graduate version, project
  - Help us make it awesome!

Textbook

- Not required but for reading more we recommend:
- Russell & Norvig, AI: A Modern Approach
- Available on reserve at the library

Course Information

- Work and grading
  - 5 homeworks, most include both programming and a written component
  - 1 midterm
  - 1 final
  - For graduates: 1 project
  - Grading: 20% Final, 10% Midterm, 35% Homework, 10% Participation, 25% Project

Course Policies

- Submission:
  - Submit the homework according to the instructions on the handout. This normally means submitting to your folder on afs, under a directory named HWx. Theoretical exercises should be submitted as a pdf file.
- Late Homework:
  - You have 8 late days, but you cannot use more than 2 late days per homework. No credit for homework submitted more than 2 days after the due date.
- Collaboration:
  - You can discuss the exercises with your classmates, but you should write up your own solutions. If you find a solution in any source other than the material provided on the course website or the textbook, you must mention the source. You can work on the programming questions in pairs, but theoretical questions are always submitted individually. Make sure that you include a README file with your andrew id and your collaborator's andrew id.