LECTURE 1: OVERVIEW

CS 4100: Foundations of AI

Instructor: Robert Platt

(some slides from Chris Amato, Magy Seif El-Nasr, and Stacy Marsella)

SOME LOGISTICS

Class webpage:
http://www.ccs.neu.edu/home/rplatt/cs4100_spring2018/index.html

Course staff office hours:

- Rob Platt (rplatt@ccs.neu.edu)
  - Tuesdays, 10:30-11:30, 526 ISEC, or by Appt.
- Bharat Vaidhyanathan (vaidhyanathan.b@husky.neu.edu)
  - ? (programming assignments)
- Ruiyang Xu (xu.r@husky.neu.edu)
  - ? (problem sets)
- Piazza: https://piazza.com/northeastern/fall2018/cs4100/home

BOOK

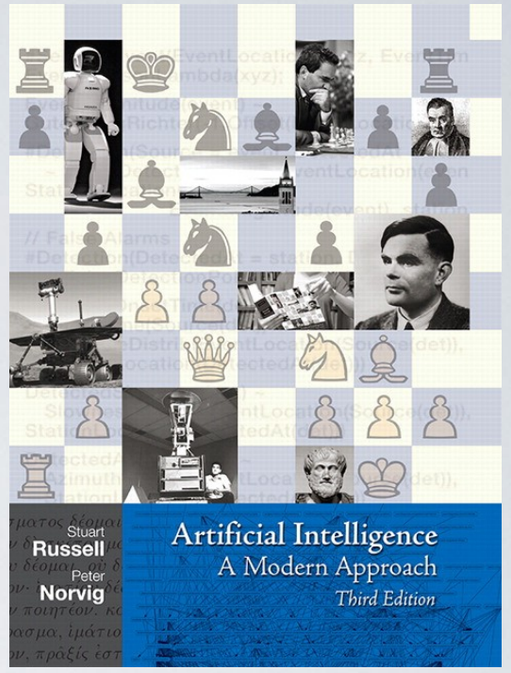

- Required
  - AI: A Modern Approach by Russell and Norvig, 3rd edition (general text)
  - Reinforcement Learning: An Introduction, http://incompleteideas.net/sutton/book/the-book.html
- Optional
  - Machine Learning: A Probabilistic Perspective by Murphy

PROBLEM SETS

- Written problems
- Can discuss problems with others, but each student should turn in their own answers
- Out every Thursday, due every Tuesday

PROGRAMMING ASSIGNMENTS

- Use AI to control Pac Man
- 4 or 5 assignments using different methods
- Coded in Python/Matlab

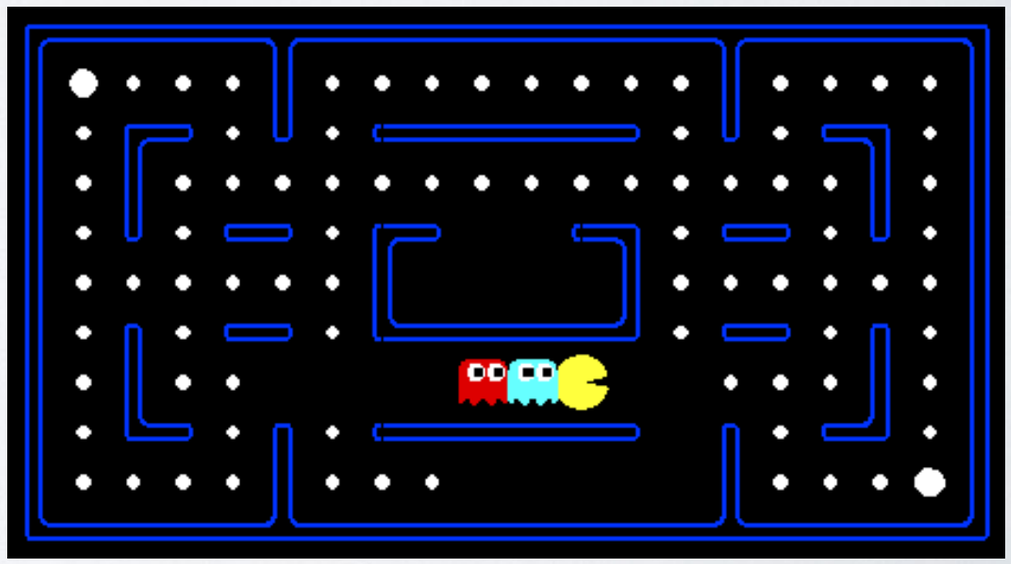

CLASS PROJECT

- Apply an AI method to a problem of your choice or learn a new method
- Conduct experiments and write up report

- Alone or in pairs
- Examples:

kaggle

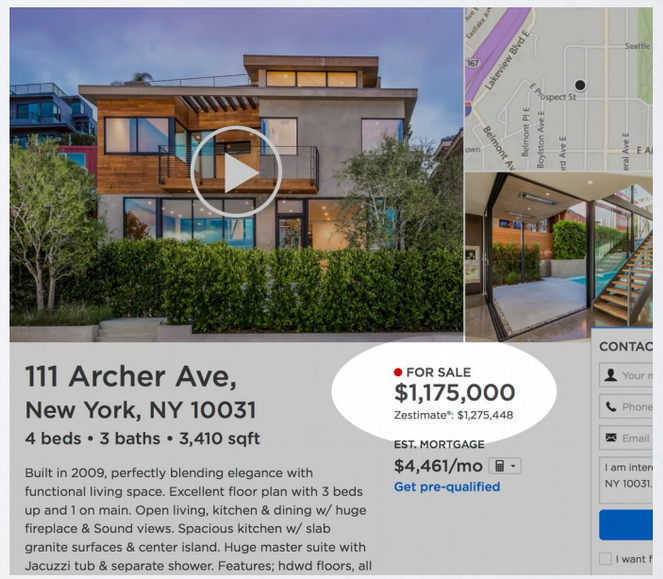

GRADING

- Problem sets: 20%
- Programming assignments: 30%
- Midterms: 30%
- Final project (presentation and paper): 20%

TOPICS COVERED

- Search
  - Uninformed search
  - Informed search
  - Adversarial search
  - Constraint satisfaction
- Decision making under uncertainty
  - Probability refresher
  - Markov Decision Processes
  - Reinforcement Learning
- Graphical Models
  - Bayes Nets
  - Hidden Markov Models
- Machine Learning
  - Supervised learning
  - Unsupervised learning
  - Deep learning

AI ALL AROUND US

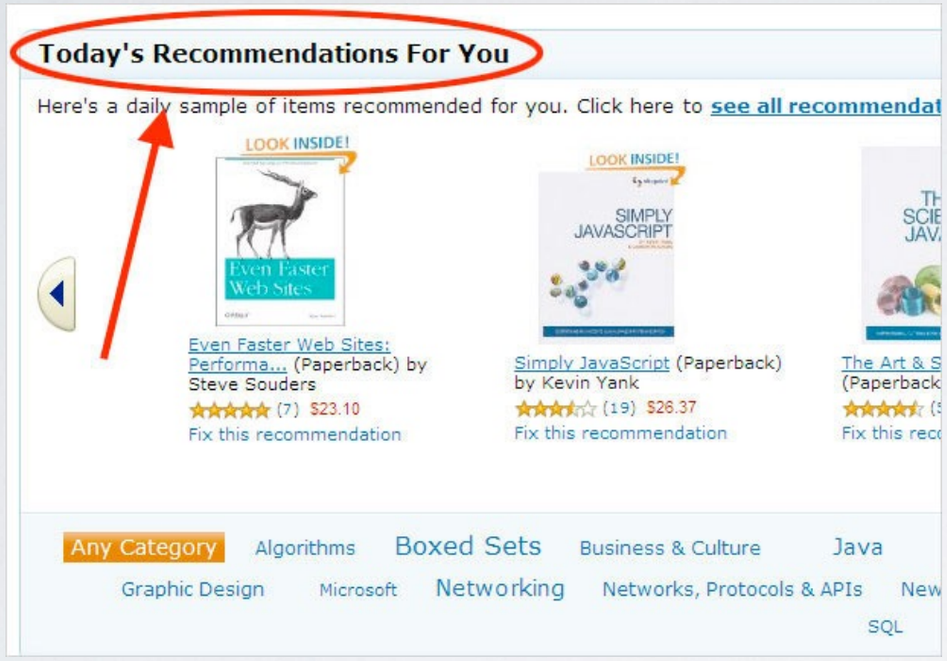

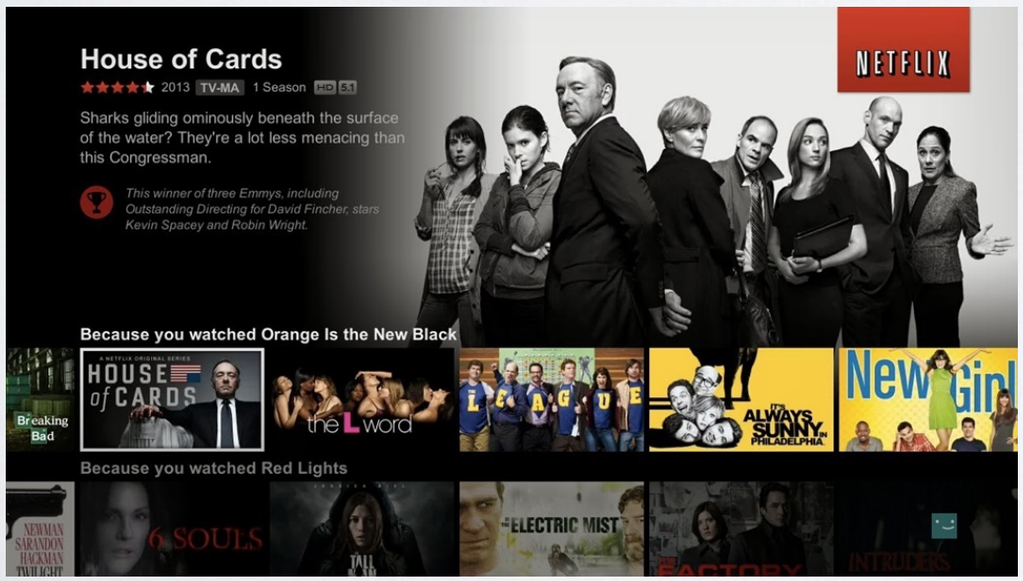

ARTIFICIAL INTELLIGENCE

- What is AI?

WHAT IS AI?

- Historical perspective:
  - Handbook of AI: the part of computer science concerned with designing intelligent computer systems, that is, systems that exhibit the characteristics we associate with intelligence in human behavior
- Thoughts on this definition?

WHAT IS AI?

- Historical perspective:
  - Handbook of AI: the part of computer science concerned with designing intelligent computer systems, that is, systems that exhibit the characteristics we associate with intelligence in human behavior
- Which is harder? Why?

Decide on moves vs Recognize pieces and move them

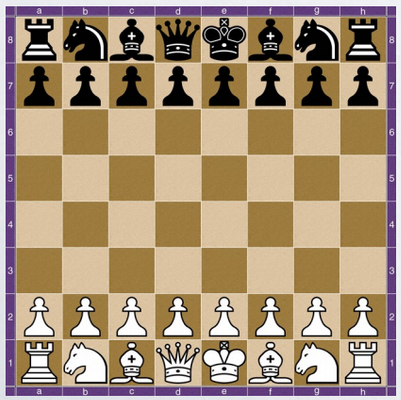

WHAT IS AI?

- Historical perspective:
  - Handbook of AI: the part of computer science concerned with designing intelligent computer systems, that is, systems that exhibit the characteristics we associate with intelligence in human behavior
- What we think requires intelligence is often wrong
  - Elephants don’t play chess: Rodney Brooks
  - People perform behaviors that seem simple
    - They require little conscious thought
    - E.g., recognizing a friend in a crowd

WHAT IS AI?

- Historical perspective:
  - Handbook of AI: the part of computer science concerned with designing intelligent computer systems, that is, systems that exhibit the characteristics we associate with intelligence in human behavior
- It’s a moving Target: once we come up with an algorithm or technology to perform a task, we tend to re-assess our beliefs that it requires intelligence or is AI
  - Beating the best human chess player was a dream of AI from its birth
  - Deep blue eventually beats the best
  - “Deep Blue is unintelligent because it is so narrow. It can win a chess game, but it can't recognize, much less pick up, a chess piece. It can't even carry on a conversation about the game it just won. Since the essence of intelligence would seem to be breadth, or the ability to react creatively to a wide variety of situations, it's hard to credit Deep Blue with much intelligence.” Drew McDermott

WHAT IS AI?

- Historical perspective:
  - Handbook of AI: the part of computer science concerned with designing intelligent computer systems, that is, systems that exhibit the characteristics we associate with intelligence in human behavior
- The algorithm or technology may not seem intelligent
  - Deep Blue relied on high speed brute force search
  - Raised the question: Is that how people do it?
  - Why not?
  - Does it matter?

WHAT IS AI?
A MORE MODERN VIEW

Russell & Norvig: Views of AI fall into four categories:

Thinking Humanly	Thinking Rationally
Acting Human	Acting Rationally

The textbook organized around "acting rationally”

but lets consider the others as well...

ARTIFICIAL INTELLIGENCE

- Intelligence
  - Cognitive modeling: behaves like a human
  - Engineering: achieve (or surpass) human performance
  - Rational: behaves perfectly, normative
  - Bounded-rational: behaves as well as possible
- Aiding humans or completely autonomous

ACTING HUMANLY: TURING TEST

- Turing (1950) "Computing machinery and intelligence":
- "Can machines think?" or "Can machines behave intelligently?”

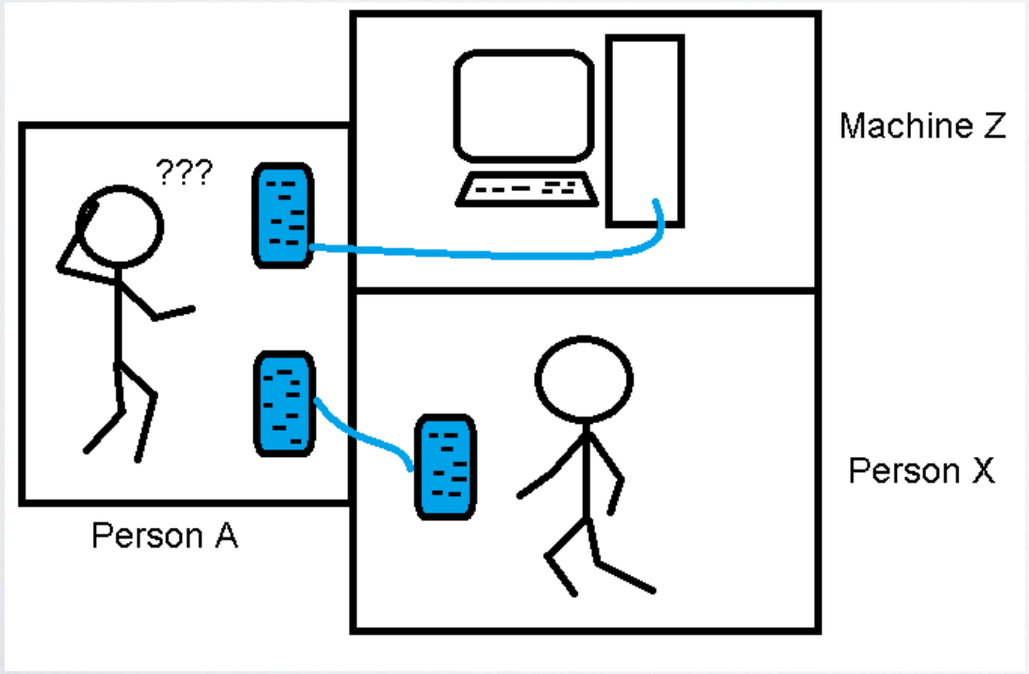

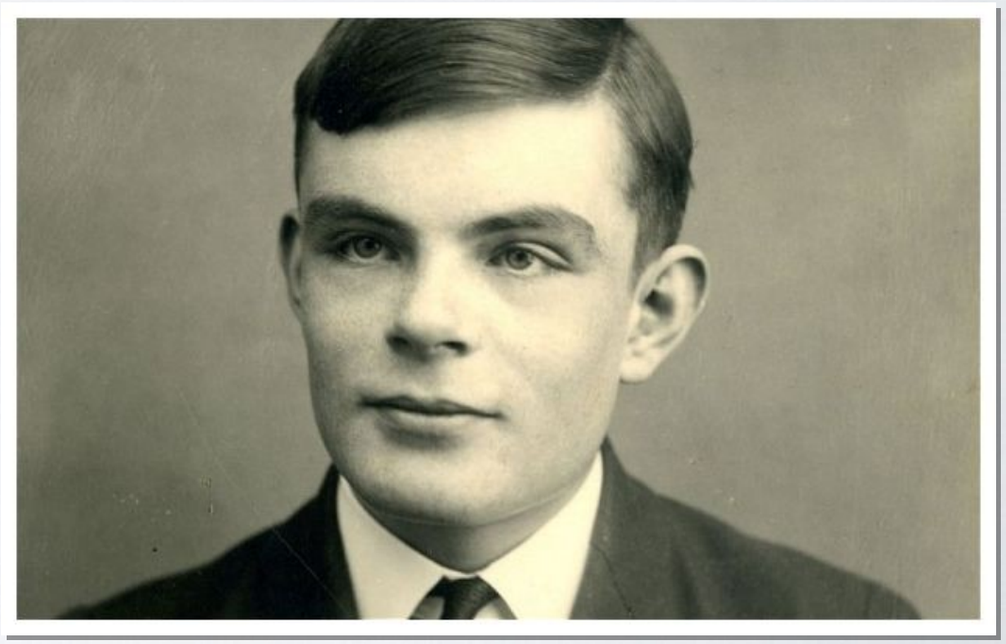

WHAT WOULD A COMPUTER NEED TO PASS THE TURING TEST?

?

WHAT WOULD A COMPUTER NEED TO PASS THE TURING TEST?

- Natural language processing: to communicate with examiner.
- Knowledge representation: to store and retrieve information provided before or during interrogation.
- Automated reasoning: to use the stored information to answer questions and to draw new conclusions.
- Machine learning: to adapt to new circumstances and to detect and extrapolate patterns.
- And this is only the simple version without perception or action!

WHAT WOULD A COMPUTER NEED TO PASS THE TURING TEST?

- Natural language processing: to communicate with examiner.
- Knowledge representation: to store and retrieve information provided before or during interrogation.
- Automated reasoning: to use the stored information to answer questions and to draw new conclusions.
- Machine learning: to adapt to new circumstances and to detect and extrapolate patterns.

- And this is only the simple version without perception or action!
- Is this a good test of AI?

IBM’S WATSON


Watson and the Jeopardy! Challenge
THINK
$24,000
Who is Stoker?
(I FOR ONE WELCOME OUR NEW COMPUTER OVERLORDS)
$1,000
$77,147
Who is Bram Stoker?
$ 17,973
$21,600
WHO IS BRAM STOKER?
$5600

AI ASSISTANTS

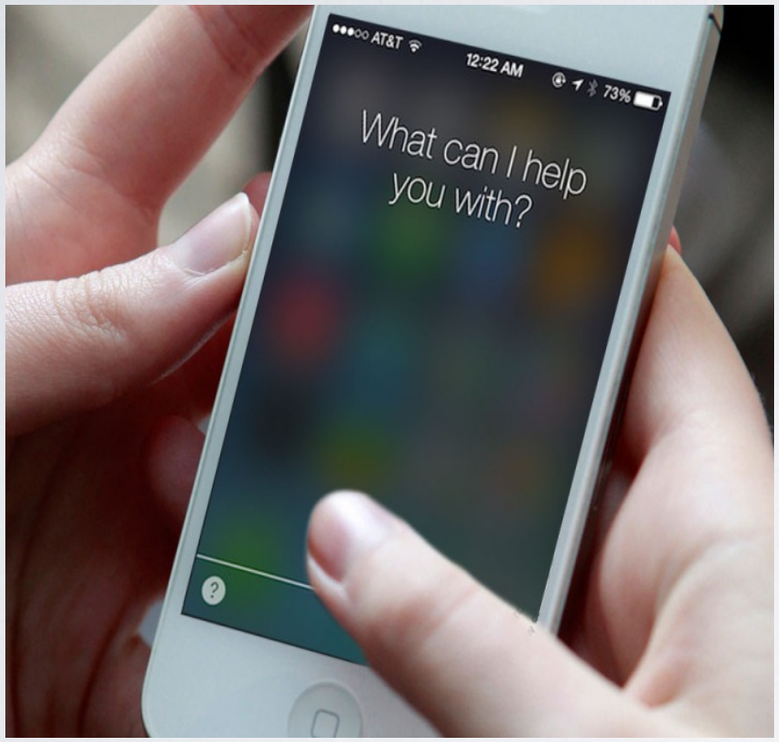

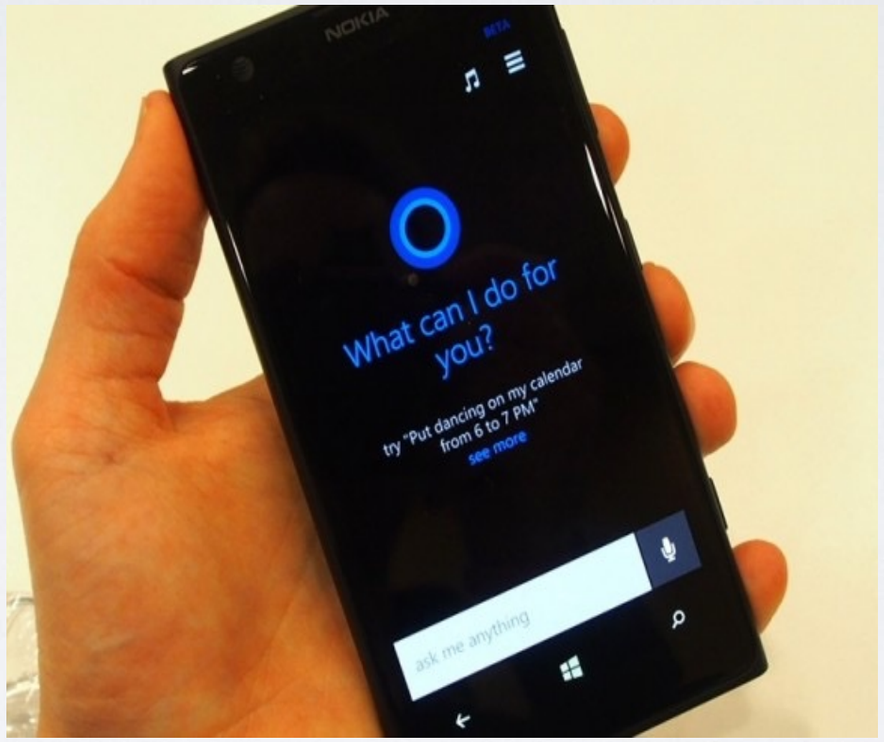

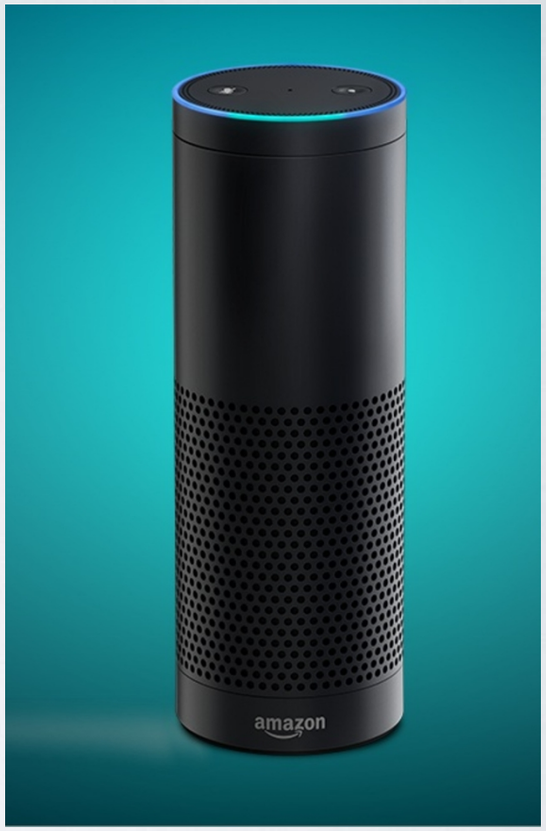

AI ASSISTANTS

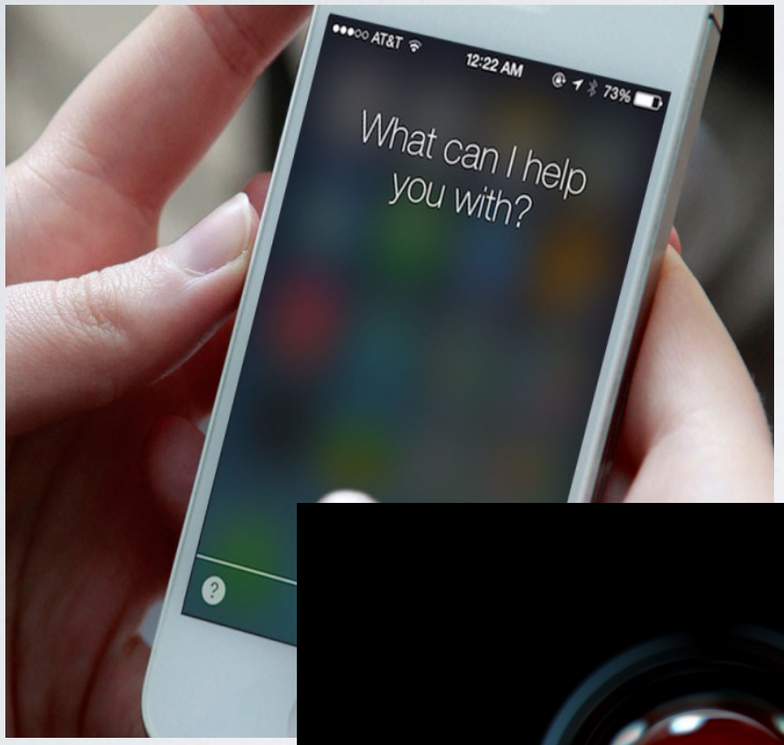

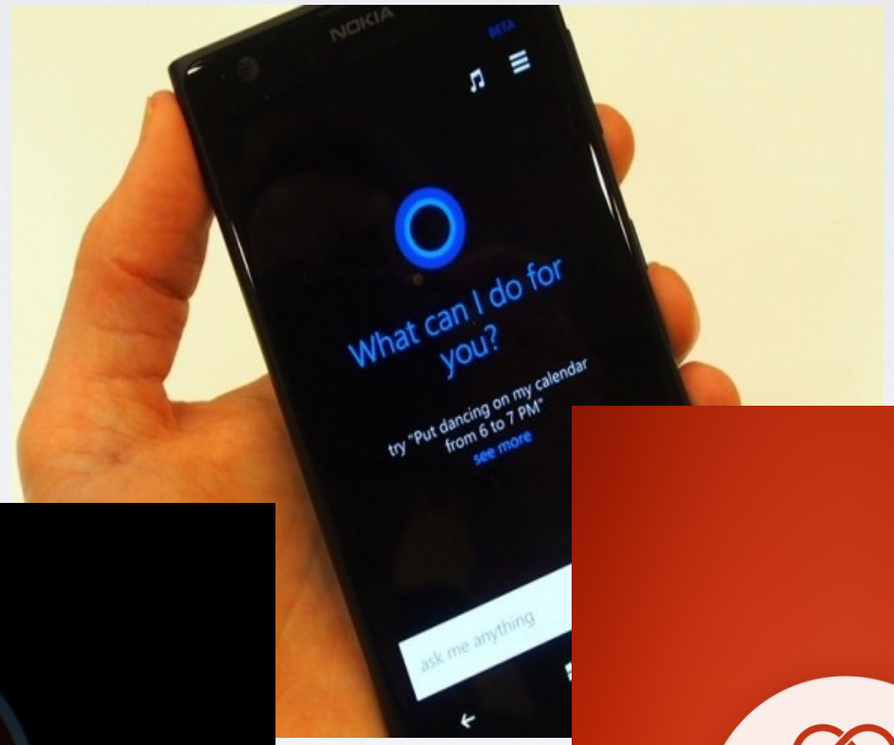

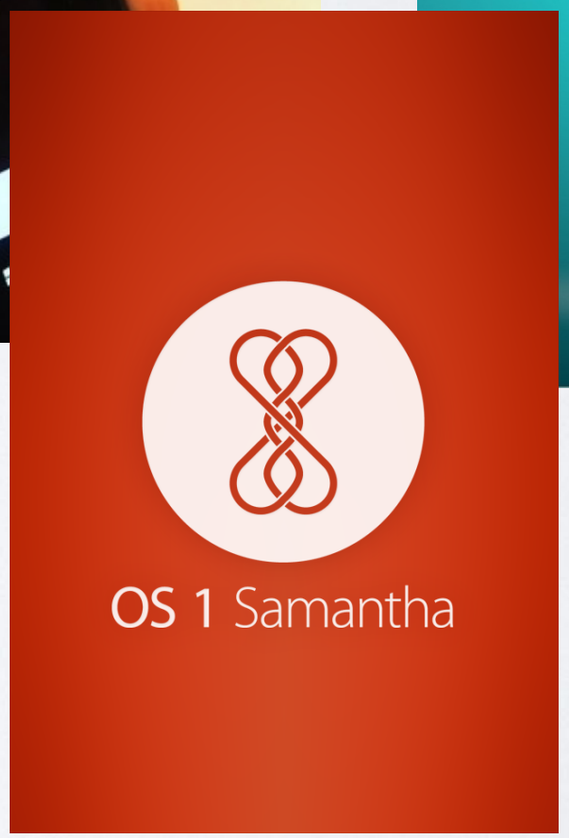

AUTONOMOUS CARS

AUTONOMOUS CARS

- Google, Tesla, Audi, BMW, GM, Ford, Uber, Lyft, Apple, nuTonomy...

ROBOTICS

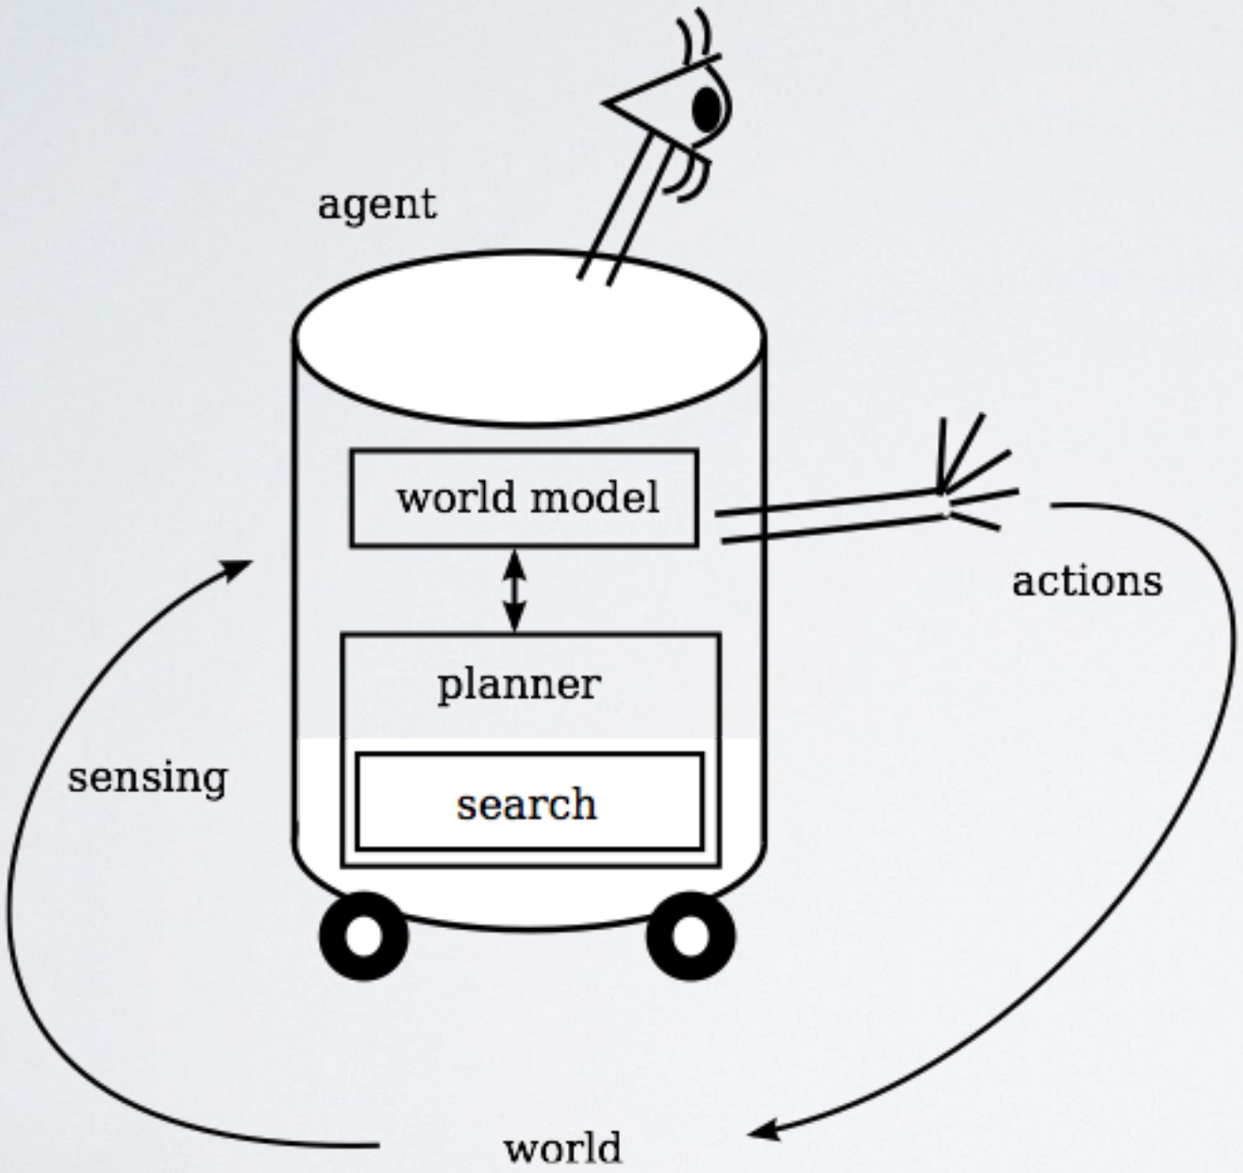

SOME ROBOTS

AI’S CYCLE OF FAILED EXPECTATIONS

- 1958, Simon and Newell: ”within ten years a digital computer will be the world's chess champion”
- 1965, Simon: “machines will be capable, within twenty years, of doing any work a man can do.”
- 1967, Minsky: “Within a generation ... the problem of creating 'artificial intelligence' will substantially be solved.”
- 1970, Minsky: “In from three to eight years we will have a machine with the general intelligence of an average human being.”
- Such optimism lead to AI winters as AI failed to meet expectations
  - Reduced attendance at conferences, reduced federal funding

WHAT WERE THE ROADBLOCKS?

- Limited computer power: There was not enough memory or processing speed to accomplish anything truly useful
- Intractability and the combinatorial explosion. Karp: *many problems can probably only be solved in exponential time (in the size of the inputs)*
- Commonsense knowledge and reasoning. Many important artificial intelligence applications like vision or natural language require enormous amounts of information about the world
- Moravec's paradox: Proving theorems and solving geometry problems is comparatively easy for computers, but a supposedly simple task like recognizing a face or crossing a room without bumping into anything is extremely difficult.

CYCLES OF OPTIMISM, FAILURE ACTUALLY GOOD FOR AI

Forced AI to Explore new ideas

- Statistical techniques revitalized Machine Learning

Old ideas reinvigorated using new approaches and technologies as well as new applications

- Neural Networks
  - Early 1950s work on neural networks falls out of favor after Minsky and Papert book on Perceptrons identifies representational issues
  - Deep Learning: Now back in a wide range of applications involving large data sets that are now available
- Work on Emotion
  - Initially argued as critical for AI by Simon and Minsky
  - Fell out of favor during rational period
  - Now a key new area Affective Computing: as man and machine increasingly interact
- Earlier ideas about knowledge representation re-entering ML
  - May transform purely statistical techniques

BUT SUCCESS BRINGS FEARS AND ETHICAL CONCERNS

- Real issues
  - Privacy
  - Jobs
- Elon Musk
  - If I were to guess at what our biggest existential threat is, it's probably that... With artificial intelligence, we are summoning the demon
  - AI is "potentially more dangerous than nukes.”
- Stephen Hawking
  - “I think the development of full artificial intelligence could spell the end of the human race”

BUT SUCCESS BRINGS FEARS AND ETHICAL CONCERNS

- Real issues
  - Privacy
  - Jobs
- Elon Musk
  - If I were to guess at what our biggest existential threat is, it's probably that… With artificial intelligence, we are summoning the demon
  - AI is "potentially more dangerous than nukes.”
- Stephen Hawking
  - “I think the development of full artificial intelligence could spell the end of the human race”

THE SOCIAL DILEMMA OF AUTONOMOUS VEHICLES

- Raises fundamental Issues in moral psychology
  - How to balance self-interest and the public good?
- Bonnefon *et al.*'s study (*Science* 24 Jun 2016) found that participants would
  - Approve of autonomous vehicles that might sacrifice passenger to save others
  - Not want to buy or even ride in such vehicles
  - Not approve regulations mandating self-sacrifice