

CSE 473: Artificial Intelligence Spring 2018

Introduction & Agents

Course Staff:



This presentation includes slides from :
Dieter Fox, Dan Weld, Dan Klein, Stuart Russell, Andrew Moore, Luke Zettlemoyer

Selected Texts and Authors

- Earl Hunt (UW) A.I. 1975
- S. Tanimoto (UW) E.A.I. 1987-95
- Stuart Russell (Berkeley) and Peter Norvig (Google) A.I.A.M.A. 1995-2003



Course Logistics

Textbook:



Artificial Intelligence: A Modern Approach, Russell and Norvig (3rd ed)

Prerequisites:

- Data Structures (CSE 332)
- Understanding of probability, logic, algorithms, complexity

Work:


Readings (text & papers)
Programming assignments / hw (40%),
Midterm (20%)
Final (30%)
Class participation (10%)



Pacman, autograder

Today

- What is (AI)?
- Agents
- What is this course?




Brain: Can We Build It?



10¹¹ neurons
10¹⁴ synapses
cycle time: 10⁻³ sec

vs.

10⁹ transistors
10¹² bits of RAM
cycle time: 10⁻⁹ sec



What Can AI Do?

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

- Play a decent game of Soccer?
- Play a winning game of Chess? Go? Jeopardy?
- Drive safely along a curving mountain road? University Way?
- Buy a week's worth of groceries on the Web? At QFC?
- Make a car? Bake a cake?
- Discover and prove a new mathematical theorem?
- Perform a complex surgical operation?
- Unload a dishwasher and put everything away?
- Translate Chinese into English in real time?
- Design a company web page?

What is AI?

The science of making machines that:

Think like humans	Think rationally
Act like humans	Act rationally

What is AI?

The science of making machines that:

Think like humans	Think rationally
Act like humans	Act rationally

Rational Decisions

We'll use the term **rational** in a particular way:

- Rational: maximally achieving pre-defined goals
- Rational 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 might be:

Computational Rationality

A (Short) History of AI

Prehistory

- **Logical Reasoning:** (4th C BC+) Aristotle, George Boole, Gottlob Frege, Alfred Tarski



Medieval Times

- **Probabilistic Reasoning:** (16th C+) Gerolamo Cardano, Pierre Fermat, James Bernoulli, Thomas Bayes



1940-1950: Early Days



1942: **Asimov**: Positronic Brain; Three Laws of Robotics

1. A robot may not injure a human being or, through inaction, allow a human being to come to harm.
2. A robot must obey the orders given to it by human beings, except where such orders would conflict with the First Law.
3. A robot must protect its own existence as long as such protection does not conflict with the First or Second Laws.



1943: **McCulloch & Pitts**: Boolean circuit model of brain


1943, 1946: First electronic digital computers -

Colossus (Thomas H. Flowers*), **ENIAC** (John Mauchly & John Presper Eckert, Jr.)

The Turing Test

Turing (1950) "Computing machinery and intelligence"

- "Can machines think?"
- "Can machines behave intelligently?"
- The *Imitation Game*:



- Suggested major components of AI: knowledge, reasoning, language understanding, learning


1950-1970: Excitement

- 1950s: Early AI programs, including
 - Samuel's checkers program,
 - Newell & Simon's Logic Theorist,
 - Gelerter's Geometry Theorem-Proving Machine
- 1956: Dartmouth meeting: "Artificial Intelligence" adopted
- 1965: Robinson's complete algorithm for logical reasoning

"Over Christmas, Allen Newell and I created a thinking machine."

-Herbert Simon

The Thinking Machine (1960's)



16

1970-1980: Knowledge Based Systems

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

The knowledge engineer practices the art of bringing the principles and tools of AI research to bear on difficult applications problems requiring experts' knowledge for their solution.

- Edward Feigenbaum in "The Art of Artificial Intelligence"

1988--: Statistical Approaches




- 1985-1990: Rise of Probability and Decision Theory
Eg, Bayes Nets
Judea Pearl - ACM Turing Award 2011
- 1990-2000: Machine learning takes over subfields:
Vision, Natural Language, etc.

"Every time I fire a linguist, the performance of the speech recognizer goes up"

- Fred Jelinek, IBM Speech Team

1997

"I could feel –
I could smell – a
new kind of
intelligence
across the
table"
-Gary Kasparov



Saying Deep Blue
doesn't really think
about chess is like
saying an airplane
doesn't really fly
because it doesn't
flap its wings.

- Drew McDermott

2005



Stanford Car
DARPA Grand Challenge

20

Self-driving car, today



21

2009



More Top Picks for You



Recommendations,
Search result ordering
Ad placement,

22

2011



http://www.youtube.com/watch?v=WFR3lOm_xhE


23

2014 = Momentous Times!

News - Technology

Fooled 33% of judges!

Turing Test breakthrough as super-computer
becomes first to convince us it's human



Eugene Goostman, a computer programme pretending to be a young Ukrainian boy, successfully duped enough humans to pass the iconic test

ANDREW GRIFFIN Sunday 08 June 2014

SHARE TWITTER REDDIT IN SHARE 41K PRINT & A.A.

A programme that convinced humans that it was a 13-year-old boy has become the first computer ever to pass the Turing Test. The test – which requires that computers are indistinguishable



24

Judges were not so smart

Conversation with 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"?

25

Judges were not so smart (cont.)

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.

Scott: Do you think Alan Turing, brilliant though he was, had trouble imagining that the judges of his "imitation game" wouldn't think to ask commonsense questions like the ones above—or that, if they did, they'd actually accept evasion or irrelevant banter as answers?

Eugene: No, not really. I don't think alan turing brilliant although this guy was had trouble imagining that the judges of his imitation game would not consider to Oooh. Anything else?


For more details, see:
<http://www.scottaaronson.com/blog/?p=1858>

26

Robocup (Stockholm '99)



Robocup




What is AI?

The science of making machines that:

Think like humans	Think rationally
Act like humans	Act rationally

Agent vs. Environment

- An **agent** is an entity that *perceives* and *acts*.
- A **rational agent** selects actions that maximize its **utility function**.
- Characteristics of the **percepts, environment, and action space** dictate techniques for selecting rational actions.



Actions? Percepts?



31

Actions? Percepts?

amazon
Prime

More Top Picks for You



Recommender System

32

Types of Environments

- Fully observable *vs.* partially observable
- Single agent *vs.* multiagent
- Deterministic *vs.* stochastic
- Episodic *vs.* sequential
- Discrete *vs.* continuous

Fully observable vs. Partially observable

Can the agent observe the complete state of the environment?



vs.



Single agent vs. Multiagent

Is the agent the only thing acting in the world?



vs.



Aka static *vs.* dynamic

Deterministic vs. Stochastic

Is there uncertainty in how the world works?



vs.



Episodic vs. Sequential

Episodic: next episode doesn't depend on previous actions.




vs.



Discrete vs. Continuous

- Is there a finite (or countable) number of possible environment states?




vs.



Types of Agent

- An **agent** is an entity that *perceives and acts*.
- A **rational agent** selects actions that maximize its **utility function**.
- Characteristics of the **percepts, environment, and action space** dictate techniques for selecting rational actions.



Reflex Agents

- Reflex agents:
 - Choose action based on current percept (and maybe memory)
 - Do not consider the future consequences of their actions
 - Act on how the world **IS**



Goal Based Agents

- Plan ahead
- Ask "what if"
- Decisions based on (hypothesized) consequences of actions
- Must have a model of how the world evolves in response to actions
- Act on how the world **WOULD BE**



Utility Based Agents

- Like goal-based, but
- Trade off multiple goals
- Reason about probabilities of outcomes
- Act on how the world will **LIKELY** be



Pacman as an Agent



Originally developed at UC Berkeley:

<http://www-inst.eecs.berkeley.edu/~cs188/pacman/pacman.html>


Project 1: Search

Goal:

- Help Pac-man find its way through the maze

Techniques:

- Search: breadth-first, depth-first, etc.
- Heuristic Search: Best-first, A*, etc.




Project 2: Game Playing

Goal:

- Play Pac-man!

Techniques:

- Adversarial Search: minimax, alpha-beta, expectimax, etc.




Project 3: Planning and Learning

Goal:

- Help Pac-man learn about the world

Techniques:

- Planning: MDPs, Value Iterations
- Learning: Reinforcement Learning



Project 4: Ghostbusters

Goal:

- Help Pac-man hunt down the ghosts

Techniques:

- Probabilistic models: HMMS, Bayes Nets
- Inference: State estimation and particle filtering



Course Topics

Part I: Making Decisions

- Fast search / planning
- Constraint satisfaction
- Adversarial and uncertain search

Part II: Reasoning under Uncertainty

- Bayes' nets
- Decision theory
- Machine learning

Throughout: Applications

- Natural language, vision, robotics, games, ...

