



CS 188x: Artificial Intelligence Introduction



Dan Klein, Pieter Abbeel
University of California, Berkeley

Today

- What is artificial intelligence?
- What can AI do?
- What is this course?



Sci-Fi AI?



What is AI?

The science of making machines that:

Think like people



Think rationally

Act like people



Act rationally

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 (Short) History of AI



A (Short) 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"
- 2000-: 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 along Telegraph Avenue?
- Buy a week's worth of groceries on the web?
- Buy a week's worth of groceries at Berkeley Bowl?
- 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?



Unintentionally Funny Stories

- One day Joe Bear was hungry. He asked his friend Irving Bird where some honey was. Irving told him there was a beehive in the oak tree. Joe walked to the oak tree. He ate the beehive. The End.
- Henry Squirrel was thirsty. He walked over to the river bank where his good friend Bill Bird was sitting. Henry slipped and fell in the river. Gravity drowned. The End.
- Once upon a time there was a dishonest fox and a vain crow. One day the crow was sitting in his tree, holding a piece of cheese in his mouth. He noticed that he was holding the piece of cheese. He became hungry, and swallowed the cheese. The fox walked over to the crow. The End.



[Shank, Tale-Spin System, 1984]

Natural Language

- Speech technologies (e.g. Siri)
 - Automatic speech recognition (ASR)
 - Text-to-speech synthesis (TTS)
 - Dialog systems



- Language processing technologies
 - Question answering
 - Machine translation




- Web search
- Text classification, spam filtering, etc...



Vision (Perception)

- Object and face recognition
- Scene segmentation
- Image classification



Images from Erik Sudderth (left), wikipedia (right)

Robotics

- Robotics
 - Part mech. eng.
 - Part AI
 - Reality much harder than simulations!
- Technologies
 - Vehicles
 - Rescue
 - Soccer!
 - Lots of automation...
- In this class:
 - We ignore mechanical aspects
 - Methods for planning
 - Methods for control



Images from UC Berkeley, Boston Dynamics, RoboCup, Google

Logic

- Logical systems
 - Theorem provers
 - NASA fault diagnosis
 - Question answering
- Methods:
 - Deduction systems
 - Constraint satisfaction
 - Satisfiability solvers (huge advances!)




Image from Bart Selman

Game Playing

- Classic Moment: May, '97: Deep Blue vs. Kasparov
 - First match won against world champion
 - "Intelligent creative" play
 - 200 million board positions per second
 - Humans understood 99.9 of Deep Blue's moves
 - Can do about the same now with a PC cluster
- Open question:
 - How does human cognition deal with the search space explosion of chess?
 - Or: how can humans compete with computers at all??
- 1996: Kasparov Beats Deep Blue
 - "I could feel --- I could smell --- a new kind of intelligence across the table."
- 1997: Deep Blue Beats Kasparov
 - "Deep Blue hasn't proven anything."
- Huge game-playing advances recently, e.g. in Go!



Text from Bart Selman, image from IBM's Deep Blue pages


Decision Making

- Applied AI involves many kinds of automation
 - Scheduling, e.g. airline routing, military
 - Route planning, e.g. Google maps
 - Medical diagnosis
 - Web search engines
 - Spam classifiers
 - Automated help desks
 - Fraud detection
 - Product recommendations
 - ... Lots more!





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

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, ...

