

CS 188: Artificial Intelligence

Fall 2007

Lecture 1: Introduction

8/28/2007

Dan Klein – UC Berkeley
Many slides over the course adapted from
either Stuart Russell or Andrew Moore

Administrivia

<http://inst.cs.berkeley.edu/~cs188>

CS 188: Artificial Intelligence, Fall 2007

[Announcements](#) [Course Info](#) [Lectures](#) [Projects](#) [FAQ](#) [Section](#) [Demos](#) [Newsgroup](#)



Announcements

- [8/27/07] Waitlist: we've asked for a new room to accommodate the waitlisted students. Until we get one, space will be tight. Stay tuned.
- [8/27/07] There will be a Python tutorial lab Friday, August 31st from 10am-3pm, in Soda 275. You can arrive and leave as you like; the tutorial should only take an hour or two. You must complete the tutorial, but you may [complete it at home](#) if you choose.
- [8/27/07] No sections the first week.
- [8/13/07] Welcome to CS188! Check back for updates in the next few weeks.

Course Staff

- Course Staff

Professor

Dan Klein



GSIs



John
DeNero



Aria
Haghighi



Simon
Lacoste-
Julien



Adam
Pauls

Course Details

- Book: Russell & Norvig, AI: A Modern Approach, 2nd Ed.
- Prerequisites:
 - (CS 61A or B) and (Math 55 or CS 70)
 - There will be a lot of statistics and programming
- Work and Grading:
 - Four assignments divided into checkpoints
 - Programming: Python, groups of 1-2
 - Written: solve together, write-up alone
 - 5 late days
 - Mid-term and final
 - Participation
 - Fixed Scale
 - Academic integrity policy

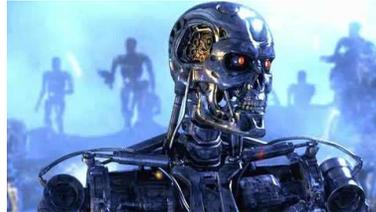
Announcements

- Waitlist:
 - We are likely but not certain to get a larger room
 - We'll know sometime this week
 - If yes, all will probably get in the course, if not, few will
- Important stuff:
 - No section this week, section 103 permanently cancelled
 - Python lab this Friday 10am-3pm in 275 Soda Hall
 - Get your account forms (in front after class)
 - First assignment on web soon
- Questions?

Today

- What is AI?
- Brief History of AI
- What can AI do?
- What is this course?

Sci-Fi AI?



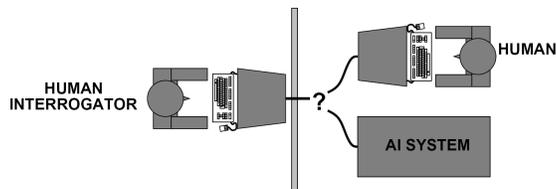
What is AI?

The science of making machines that:

Think like humans	Think rationally
Act like humans	Act rationally

Acting Like Humans?

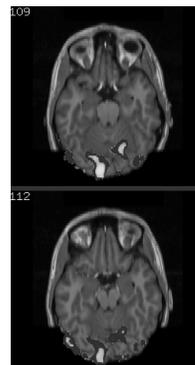
- Turing (1950) "Computing machinery and intelligence"
 - "Can machines think?" → "Can machines behave intelligently?"
 - Operational test for intelligent behavior: the *Imitation Game*



- Predicted by 2000, a 30% chance of fooling a lay person for 5 minutes
- Anticipated all major arguments against AI in following 50 years
- Suggested major components of AI: knowledge, reasoning, language understanding, learning
- Problem: Turing test is not reproducible or amenable to mathematical analysis

Thinking Like Humans?

- The Cognitive Science approach:
 - 1960s "cognitive revolution": information-processing psychology replaced prevailing orthodoxy of behaviorism
- Scientific theories of internal activities of the brain
 - What level of abstraction? "Knowledge" or "circuits"?
 - **Cognitive science:** Predicting and testing behavior of human subjects (top-down)
 - **Cognitive neuroscience:** Direct identification from neurological data (bottom-up)
 - Both approaches now distinct from AI
 - Both share with AI the following characteristic:
The available theories do not explain (or engender) anything resembling human-level general intelligence
- Hence, all three fields share one principal direction!



Images from Oxford fMRI center

Thinking Rationally?

- The “Laws of Thought” approach
 - What does it mean to “think rationally”?
 - Normative / prescriptive rather than descriptive
- Logician tradition:
 - Logic: notation and rules of derivation for thoughts
 - Aristotle: what are correct arguments/thought processes?
 - Direct line through mathematics, philosophy, to modern AI
- Problems:
 - Not all intelligent behavior is mediated by logical deliberation
 - What is the purpose of thinking? What thoughts should I (bother to) have?
 - **Logical systems tend to do the wrong thing in the presence of uncertainty**



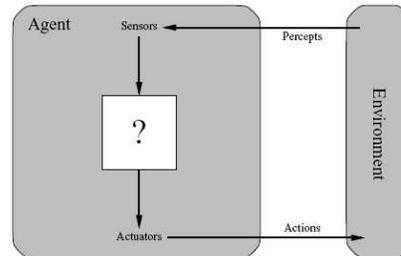
Acting Rationally

- Rational behavior: doing the “right thing”
 - The right thing: that which is expected to maximize goal achievement, given the available information
 - Doesn't necessarily involve thinking, e.g., blinking
 - Thinking can be in the service of rational action
 - Entirely dependent on goals!
 - Irrational \neq insane, irrationality is sub-optimal action
 - Rational \neq successful
- Our focus here: rational agents
 - Systems which make the best possible decisions given goals, evidence, and constraints
 - In the real world, usually lots of uncertainty
 - ... and lots of complexity
 - Usually, we're just approximating rationality
- “Computational rationality” a better title for this course

Rational Agents

- An **agent** is an entity that perceives and acts (more examples later)
- This course is about designing rational agents
- Abstractly, an agent is a function from percept histories to actions:

$$\mathcal{P}^* \rightarrow \mathcal{A}$$



- For any given class of environments and tasks, we seek the agent (or class of agents) with the best performance
- Computational limitations make perfect rationality unachievable
- So we want the best program for given machine resources

[demo: pacman]

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—88: 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"
- 1988—: Statistical approaches
 - Resurgence of probability, focus on uncertainty
 - General increase in technical depth
 - Agents, agents, everywhere... "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?
- ✓ 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 complex surgical operation?
- ✗ Unload a dishwasher and put everything away?
- ✓ Translate spoken English into spoken Swedish 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**

- Automatic speech recognition (ASR)
- Text-to-speech synthesis (TTS)
- Dialog systems



- **Language processing technologies**

- Machine translation:

Aux dires de son président, la commission serait en mesure de le faire .
According to the president, the commission would be able to do so .

Il faut du sang dans les veines et du cran .
We must blood in the veins and the courage .

- Information extraction
- Information retrieval, question answering
- Text classification, spam filtering, etc...



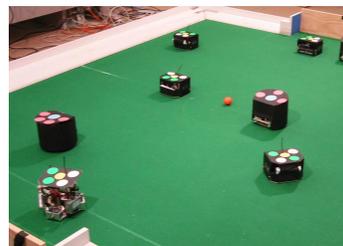
Vision (Perception)



Images from Jitendra Malik

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 stanfordracing.org, CMU RoboCup, Honda ASIMO sites

Logic

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

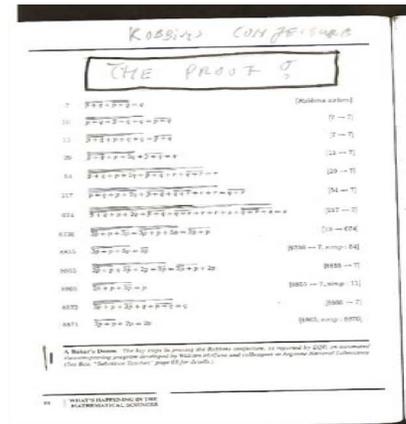


Image from Bart Selman

Game Playing

- **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 big 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."



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

Decision Making

- **Many applications of AI: decision making**
 - Scheduling, e.g. airline routing, military
 - Route planning, e.g. mapquest
 - Medical diagnosis, e.g. Pathfinder system
 - Automated help desks
 - Fraud detection
- ... the list goes on.

Course Topics

- Part I: Search and Plans
 - Fast search
 - Constraint satisfaction
 - Adversarial and uncertain search
- Part II: Uncertainty and Beliefs
 - Reinforcement learning
 - Bayes' nets
 - Decision theory
- Throughout: Applications
 - Natural language
 - Vision
 - Robotics
 - Games

Course Projects

- Pacman
- Robot control
- Battleship
- Spam / digit recognition