CS 188: Artificial Intelligence

Lecture 1: Introduction

Pieter Abbeel – UC Berkeley
Many slides from Dan Klein.

Course Information

http://inst.cs.berkeley.edu/~cs188/sp12

This semester's website will be live by Friday midnight.

- Communication:
  - Announcements on webpage
  - Questions? Try piazza!
  - If not suitable for piazza, staff email: cs188-staff@lists.berkeley.edu
  - Office hours: see website for schedule

Course Staff

- **Course Staff**

  Professor
  Pieter Abbeel

  GSIs

  Yangqing Jia   Jonathan Long   Arjun Singh

Course Information

- Book: Russell & Norvig, AI: A Modern Approach, 3rd
- Slides
- Prerequisites:
  - (CS 61A or B) and (Math 55 or CS 70)
  - Strongly recommended: CS61A, CS61B and CS70
  - There will be a lot of math and programming
  - Self diagnostic
- Work and Grading:
  - 5 programming projects: Python, groups of 1-2
    - 5 late days, 2 per project
  - Regular assignments --- details forthcoming!
  - 2 midterms
  - 1 final
  - Participation
  - Fixed scale
  - Academic integrity policy
- Contests!

Waiting List

- Will take a while to sort out. We don't control enrollment. Contact Michael-David Sasson (msasson@cs) with any questions on the process.

Today

- What is artificial intelligence?

- What can AI do?

- What is this course?

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
    - E.g., generate plan for driving to the airport
  - 1966: Weizenbaum's Eliza / Turing test

Herb Simon, 1957

It is not my aim to surprise or shock you---but the simplest way I can summarize is to say that there are now in the world machines that think, that learn and that create. Moreover, their ability to do these things is going to increase rapidly until---in a visible future---the range of problems they can handle will be coextensive with the range to which human mind has been applied.

More precisely: within 10 years a computer would be chess champion, and an important new mathematical theorem would be proved by a computer.

Harder than originally thought

- Herb Simon's prediction came true, but after roughly 40 years instead of after 10
- Eliza:
  - " … mother …" → "Tell me more about your family"
  - "I wanted to adopt a puppy, but it's too young to be separated from its mother." → ???
- 1957: Sputnik
  - Automatic Russian → English translation
  - Famous example:
    - "The spirit is willing but the flesh is weak."
    - E → R → E: "The vodka is strong but the meat is rotten."

Observations

- Need some understanding about the world

- Computational tractability, NP-completeness, exponential scaling.

A (Short) History of AI (ctd)

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

- Language processing technologies
    - Machine translation

Compétitivité : pourquoi l'écart se creuse entre la France et l'Allemagne

Competitiveness: why the gap between France and Germany

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

[demos: language]

Vision (Perception)

[videos: vision]

- Object and character recognition
- Scene segmentation
- 3D reconstruction
- Image classification

## Robotics

[videos: 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

- Scheduling, e.g. airline routing, military
- Route planning, e.g. google maps
- Medical diagnosis
- Automated help desks
- Fraud detection
- Spam classifiers
- Web search engines
- Movie and book recommendations

- … Lots more!

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 would be:

Computational Rationality

Maximize Your Expected Utility

Designing Rational Agents

- 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.

- 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

Pacman as an Agent

SCORE: 18

What About the Brain?

- Brains (human minds) are very good at making rational decisions (but not perfect)

- "Brains are to intelligence as wings are to flight"

- Brains aren't as modular as software

- Lessons learned: **prediction** and **simulation** are key to decision making

Course Topics

- Part I: Making Decisions
  - Fast search
  - Constraint satisfaction
  - Adversarial and uncertain search

- Part II: Modeling Uncertainty
  - Bayes' nets
  - Decision theory

- Part III: Machine learning
  - Perceptron, kernels

- Throughout: Applications
  - Natural language, vision, robotics, games