CS 121: Introduction to AI

Jean-Claude Latombe
ai.stanford.edu/~latombe

cs121.stanford.edu

Required textbook:
S. Russell and P. Norvig.
Artificial Intelligence: A Modern Approach.
3rd edition, Prentice Hall, 2010

Course Assistants

- Jacob Quain
- Nikil Viswanathan

Office Hours and Sections

- **JCL** Mon at 11am-12pm in Gates 135
- **Jacob Quain**
- **Nikil Viswanathan**
- **CA section:**

Today's Agenda

- Introduction to AI
  (Russell and Norvig: Chap. 1 and 2)
- Overview of CS121

What is AI?

an attempt of

- AI is ~~the~~ reproduction of human reasoning and intelligent behavior by computational methods

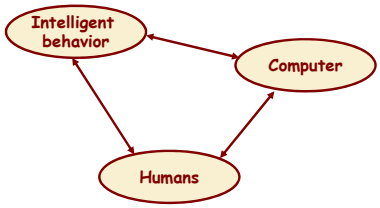

What is AI?
(R&N)

Discipline that systematizes and automates reasoning processes to create machines that:

Act like humans	Act rationally
Think like humans	Think rationally

Act like humans	Act rationally
Think like humans	Think rationally

- The goal of AI is to create computer systems that perform tasks regarded as requiring intelligence when done by humans
- → AI Methodology: Take a task at which people are better, e.g.:
  - Prove a theorem
  - Play chess
  - Plan a surgical operation
  - Diagnose a disease
  - Navigate in a building

  and build a computer system that does it automatically
- But do we want to duplicate human imperfections?

7

Act like humans	Act rationally
Think like humans	Think rationally

- Here, how the computer performs tasks does matter
- The reasoning steps are important
- → Ability to create and manipulate symbolic knowledge (definitions, concepts, theorems, …)
- What is the impact of hardware on low-level reasoning, e.g., to go from signals to symbols?

8

Act like humans	Act rationally
Think like humans	Think rationally

- Now, the goal is to build agents that always make the "best" decision given what is available (knowledge, time, resources)
- "Best" means maximizing the expected value of a utility function
- → Connections to economics and control theory
- What is the impact of self-consciousness, emotions, desires, love for music, fear of dying, etc … on human intelligence?

9

Can Machines Act/Think Intelligently?

"If there were machines which bore a resemblance to our bodies and imitated our actions as closely as possible for all practical purposes, we should still have two very certain means of recognizing that they were not real men. The first is that they could never use words, or put together signs, as we do in order to declare our thoughts to others… Secondly, even though some machines might do some things as well as we do them, or perhaps even better, they would inevitably fail in others, which would reveal that they are acting not from understanding, …"

Discourse on the Method, by Descartes (1598-1650)

10

Can Machines Act/Think Intelligently?

Turing Test:

- http://plato.stanford.edu/entries/turing-test/
- Test proposed by Alan Turing in 1950
- The computer is asked questions by a human interrogator. It passes the test if the interrogator cannot tell whether the responses come from a person
- Required capabilities: natural language processing, knowledge representation, automated reasoning, learning,…
- No physical interaction
- Chinese Room (J. Searle)

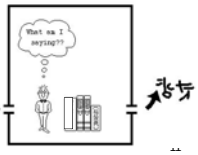

11

An Application of the Turing Test

- CAPTCHA: Completely Automatic Public Turing tests to tell Computers and Humans Apart
- E.g.:
  - Display visually distorted words
  - Ask user to recognize these words
- Example of application: have only humans open email accounts

12

Can Machines Act/Think Intelligently?

- Yes, if intelligence is narrowly defined as information processing

  AI has made impressive achievements showing that tasks initially assumed to require intelligence can be automated

  But each success of AI seems to push further the limits of what we consider "intelligence"

13

Some Achievements

- Computers have won over world champions in several games, including Checkers, Othello, and Chess, but still do not do well in Go
- AI techniques are used in many systems: formal calculus, video games, route planning, logistics planning, pharmaceutical drug design, medical diagnosis, hardware and software trouble-shooting, speech recognition, traffic monitoring, facial recognition, medical image analysis, part inspection, etc...
- Stanford's robotic car, Stanley, autonomously traversed 132 miles of desert
- Some industries (automobile, electronics) are highly robotized, while other robots perform brain and heart surgery, are rolling on Mars, fly autonomously, ..., but home robots still remain a thing of the future


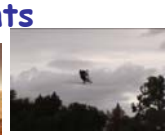


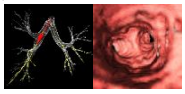


Can Machines Act/Think Intelligently?

- Yes, if intelligence is narrowly defined as information processing

  AI has made impressive achievements showing that tasks initially assumed to require intelligence can be automated

- Maybe yes, maybe not, if intelligence is not separated from the rest of "being human"

15

Some Big Open Questions

- AI (especially, the "rational agent" approach) assumes that intelligent behaviors are only based on information processing? Is this a valid assumption?
- If yes, can the human brain machinery solve problems that are inherently intractable for computers?
- In a human being, where is the interface between "intelligence" and the rest of "human nature", e.g.:
  - How does intelligence relate to emotions felt?
  - What does it mean for a human to "feel" that he/she understands something?
- Is this interface critical to intelligence? Can there exist a general theory of intelligence independent of human beings? What is the role of the human body?

16

Some Big Open Questions

- AI (especially, the "rational agent" approach) assumes that intelligent behaviors are based on information
- **In the movie I, Robot, the most impressive feature of the robots is not their ability to solve complex problems, but how they blend human-like reasoning with other key aspects of human beings (especially, self-consciousness, fear of dying, distinction between right and wrong)**
- Is this interface critical to intelligence? Can there exist a general theory of intelligence independent of human beings? What is the role of the human body?

17

- AI contributes to building an information processing model of human beings, just as Biochemistry contributes to building a model of human beings based on bio-molecular interactions
- Both try to explain how a human being operates
- Both also explore ways to avoid human imperfections (in Biochemistry, by engineering new proteins and drug molecules; in AI, by designing rational reasoning methods)
- Both try to produce new useful technologies
- Neither explains (yet?) the true meaning of being human

18

Main Areas of AI

- Knowledge representation (including formal logic)
- Search, especially heuristic search (puzzles, games)
- Planning
- Reasoning under uncertainty, including probabilistic reasoning
- Learning
- Agent architectures
- Robotics and perception
- Natural language processing

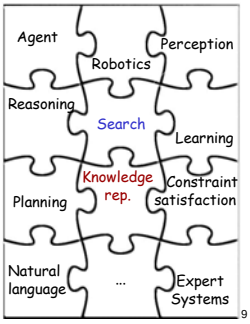

Bits of History

- 1956: The name "Artificial Intelligence" is coined
- 60's: Search and games, formal logic and theorem proving
- 70's: Robotics, perception, knowledge representation, expert systems
- 80's: More expert systems, AI becomes an industry
- 90's: Rational agents, probabilistic reasoning, machine learning
- 00's: Systems integrating many AI methods, machine learning, reasoning under uncertainty, robotics again

Schedule

Date	Topic	HW: Out	Due (Tue.)	Russell & Norvig textbook	Slides (ppt)	Slides (pdf)
1/3	Introduction			Chap. 1 and 2	1	1
1/5	Search problems			Chap. 3, Sections 3.1-2	2	2
1/10	Blind Search	HW1(doc, pdf)		Chap. 3, Sections 3.3-4	3	3
1/12	Heuristic search (1/3)			Chap. 3, Sections 3.5-7	4-5	4-5
1/17	MLK Day (no class)	HW2(doc, pdf)	HW1			
1/19	Heuristic search (2/3)			Chap. 3, Sections 3.5-7	4-5	4-5
1/24	Heuristic search (3/3) + Motion planning (1/2)	HW3(doc, pdf)	HW2	Chap. 4, Section 4.1	6-7	6-7
1/26	Motion planning (2/2)			Chap. 25, Section 25.4	6-7	6-7
1/31	Action planning	HW4(doc, pdf)	HW3	Chap. 10	8	8
2/2	Constraint satisfaction			Chap. 6, Section 6.1	9	9
2/7	Constraint propagation		HW4	Chap. 6, Sections 6.2-5	10	10
2/9	Introduction to uncertainty			Chap. 13, Sections 13.1-2	11	11
2/14	Non-deterministic uncertainty	HW5(doc, pdf)			12	12
2/16	Adversarial Search			Chap. 5	13	13
2/21	Presidents' Day (no class)	HW6(doc, pdf)	HW5			
2/23	Deciding under probabilistic uncertainty			Chap. 16 and 17	14	14
2/28	Bayesian nets	HW7(doc, pdf)	HW6	Chap. 14	15	15
3/2	Inductive learning (1/2)			Chap. 18	16	16
3/7	Inductive learning (2/2)		HW7	Chap. 18	17	17
3/9	Course review by CAs					

- Note that HWs are due on Tuesdays at noon (not on Mondays)
- Final: Wednesday March 16th at 8:30-11:30am

CS121 Web Site

cs121.stanford.edu

ai.stanford.edu/~latombe/cs121/2011/home.htm

(homeworks, exam, grading)

Required textbook:

S. Russell and P. Norvig. Artificial Intelligence: A Modern Approach.

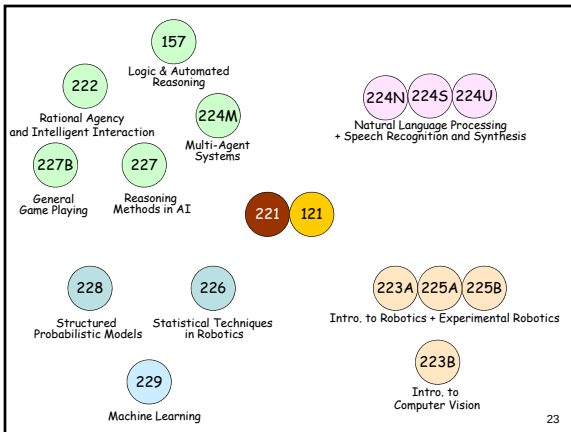

Immediate actions:

1. **Browse cs121.stanford.edu**
2. **Register on AXESS as soon as possible**