


CSCI 446: Artificial Intelligence

Introduction



Instructor: Michele Van Dyne

Montana Tech

New This Year!


- Using resources from the UC Berkeley AI class
 - Slides, lectures, videos, assignments
- Why?
 - Machine learning will be taught as a separate class, so AI course material has been re-worked
- How does this help me?
 - Framework for programs are provided so you only do the AI parts
 - Autograding so you can 1) check you answers and 2) get feedback faster
 - Access to lecture videos for review
- What's the downside?
 - All programs must be written in Python
 - There are more assignments, but they are more applicable

New This Year!


- First, register for an account at:
 - <https://edge.edx.org/login>
- Second, enroll in the course:
 - https://edge.edx.org/courses/BerkeleyX/CS188x-8/Artificial_Intelligence/about
- Third, you will need to get to the latest version of the course:
 - Click the link in the message above the Homework list
 - You will see an outline listing by week, not by homework

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:

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


- **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?


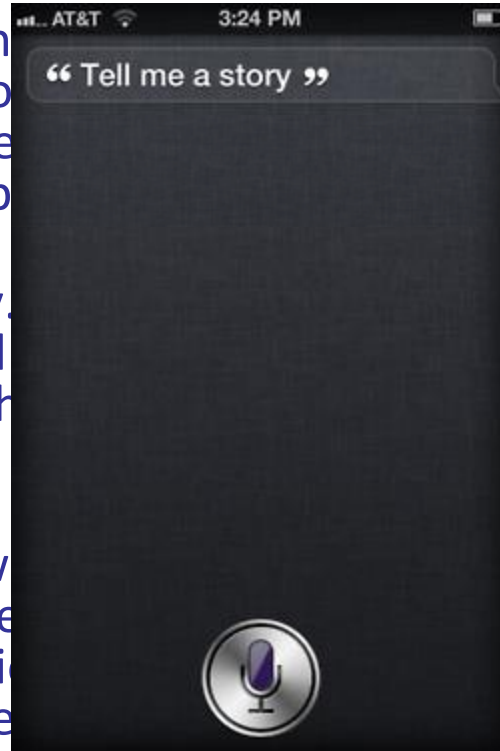
Quiz: Which of the following can be done at present?

- ✓ Play a decent game of table tennis?
- ✓ Play a decent game of Jeopardy?
- ✓ Drive safely along a curving mountain road?
- ? Drive safely down Park St. as the bars are closing?
- ✓ Buy a week's worth of groceries on the web?
- ✗ Buy a week's worth of groceries at Henessey Market?
- ? 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 hunting Irving Bird where some honey was. There was a beehive in the tree above the oak tree. He ate the bees.
- Henry Squirrel was thirsty. He went to the river bank where his good friend was. Henry slipped and fell in the water. The End.
- Once upon a time there was a fox and a vain crow. One day the crow was sitting in his tree holding a piece of cheese in his mouth. He noticed the fox was hungry, and swallowed the cheese. The fox walked up and ate the cheese. The End.



...a vain crow. One day the crow was sitting in his tree holding a piece of cheese in his mouth. He noticed the fox was hungry, and swallowed the cheese. The fox walked up and ate the cheese. The End.


Natural Language

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



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



"Il est impossible aux journalistes de rentrer dans les régions tibétaines"

Bruno Philip, correspondant du "Monde" en Chine, estime que les journalistes de l'AFP qui ont été expulsés de la province tibétaine du Qinghai "n'étaient pas dans l'illégalité".



Les faits Le dalaï-lama dénonce l'"enfer" imposé au Tibet depuis sa fuite, en 1959

Vidéo Anniversaire de la rébellion tibétaine en Chine

"It is impossible for journalists to enter Tibetan areas"

Philip Bruno, correspondant for "World" in China, said that journalists of the AFP who have been deported from the Tibetan province of Qinghai "were not illegal."





Facts The Dalai Lama denounces the "hell" imposed since he fled Tibet in 1959

Video Anniversary of the Tibetan rebellion: China on guard

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

Vision (Perception)

- Object and face recognition
- Scene segmentation
- Image classification



Demo1: VISION – lec_1_t2_video.flv

Demo2: VISION – lec_1_obj_rec_0.mpg

Robotics

Demo 1: ROBOTICS – soccer.avi



Demo 4: ROBOTICS – laundry.avi

Demo 2: ROBOTICS – soccer2.avi

Demo 5: ROBOTICS – petman.avi

Demo 3: ROBOTICS – gcar.avi

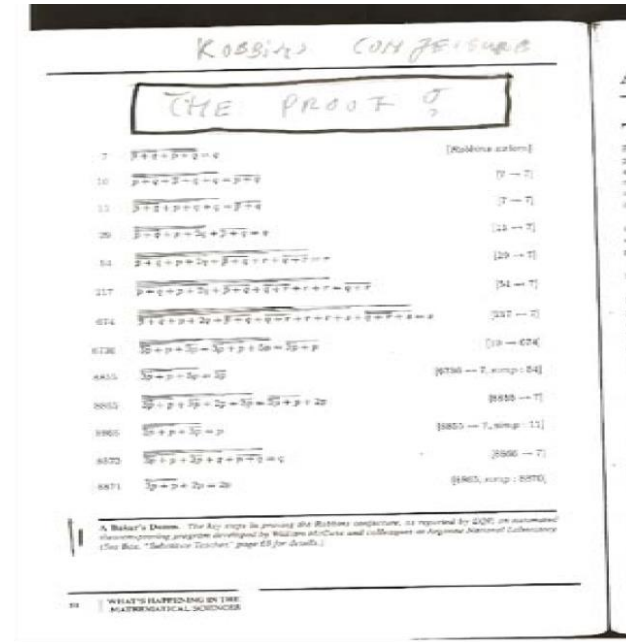
- Robotics
 - Part mech. eng.
 - Part AI
 - Reality much harder than simulations!
- Technologies
 - Vehicles
 - Rescue
 - Soccer!
 - Lots of automation...
- In this class:
 - We ignore mechanical aspects
 - Methods for planning
 - Methods for control



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!)



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!**



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



Course Topics

- Part I: Making Decisions
 - Fast search / planning
 - Constraint satisfaction
 - Logic
 - Adversarial and uncertain search
- Part II: Reasoning under Uncertainty
 - Bayes' nets
 - Decision theory
 - Machine learning
- Throughout: Applications
 - Natural language, vision, robotics, games, ...



Summary

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

