

CSE 573:

Intro to Artificial Intelligence

Hanna Hajishirzi



slides adapted from
Dan Klein, Pieter Abbeel ai.berkeley.edu
And Dan Weld, Luke Zettlemoyer

Website

- Website

- tentative schedule
- lecture slides
- course policies, etc.

- <https://courses.cs.washington.edu/courses/cse573/21wi/>

Course Staff

Hanna Hajishirzi	hannaneh@cs	Mondays	Zoom
Aida Amini	amini91@cs	Fridays	Zoom
Josh Gardner	jpgard@cs	Tuesday/Thursdays	Zoom

- Office hours
 - Schedule on the website
 - TAs: concepts, projects, homework
 - Hanna: concepts, high level guidance, homework
- Introductions?

Logistics

- Canvas: grades, submitting assignments:
 - private matters – private messages
 - if your message is not answered promptly enough, use the staff email:

- Ed: Discussion board: ask and answer questions; announcements

Course Format

- Programming Assignments
 - 4 projects
 - Python
 - Autograded
 - Give you hands-on experience with the algorithms
 - I expect you to get 100% on projects
- Written homeworks
 - 2 written homeworks
 - Gives you a more conceptual understanding of the material

Course Format (continued)


- Paper report
 - Learn how to read and criticize research papers
- Final Project:
 - Encourage to pick a project related to your research
 - We will provide recommendations for picking projects
 - There will be a proposal day.

Prerequisites

- Data Structure or Equivalent:
CSE 332
- Math:
 - Basic exposure to probability and data structures
- Programming – Familiar with Python
 - There is a 0th project (P0)

Textbook

- Not required, but for students who want to read more we recommend
 - Russell & Norvig, AI: A Modern Approach, 3rd Ed.




- Warning: Not a course textbook, so our presentation does not necessarily follow the presentation in the book.

Course Policies



- Grade:
 - Your **grade** will be: 5% class participation, 5% paper reports, 30% programming assignments, 30% homeworks, and 30% project.
- Assignments should be done individually unless otherwise specified.
- Late Policy: **Six** penalty-free late day for the whole quarter; maximum 4 days per assignment. No late day for the final.

Today

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



AI







TUG
CAUTION
MAY CONTAIN
CHEMOTHERAPY DRUG

CAUTION
MAY CONTAIN
CHEMOTHERAPY DRUG



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




Maximize Your Expected Utility



Maximize Your Expected Utility




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



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



Topics in This Course

- Part I: Intelligence from Computation
 - Fast search
 - Adversarial and uncertain search
- Part II: Reasoning under Uncertainty
 - Bayes' nets
 - Decision theory
 - Machine learning
- Throughout: Applications
 - Natural language, vision, robotics, games, ...

AI

Machine Learning

[learning decisions;
sometimes independent]

Robots

[physically
embodied]

Rational Agents

[decisions]


Human-AI Interaction

NLP


Computer Vision

Today

- Course overview
- What is artificial intelligence (AI)?
- **What can AI do?**
- What is this course?





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—2012: Statistical approaches
 - Resurgence of probability, focus on uncertainty
 - General increase in technical depth
 - Agents and learning systems... "AI Spring"?
- 2012— present: Excitement: Look, Ma, no hands!
 - Big Data, big compute, neural networks
 - Some re-unification of subfields
 - AI is being used in industry.



What Can AI Do?


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

- Play a decent game of Jeopardy?
- Win against any human at chess?
- Win against the best humans at Go?
- Play a decent game of tennis?
- Grab a particular cup and put it on a shelf?
- Unload any dishwasher in any home?
- Drive safely along the highway?
- Drive safely along University Avenue?
- Buy a week's worth of groceries on the web?
- Buy a week's worth of groceries at QFC?
- Discover and prove a new mathematical theorem?
- Perform a surgical operation?
- Unload a known dishwasher in collaboration with a person?
- Translate spoken Chinese into spoken English in real time?
- Write an intentionally funny story?
-
-



Unintentionally Funny Stories


- One day Joe Bear was hungry. He went to visit his friend Irving Bird where some honey was. In the oak tree, there was a beehive in the oak tree. He ate the bees.
- Henry Squirrel was thirsty. He went to the river bank where his good friend Henry slipped and fell in the river. The End.
- Once upon a time there was a fox sitting in his tree, holding a piece of cheese. He became angry at the piece of cheese. He became angry to the crow. The End.



crow. One day the crow was sitting in his tree, holding a piece of cheese. He noticed that he was holding the piece of cheese. The fox walked over

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



"It is impossible for journalists to enter Tibetan areas"

Philip Bruno, correspondent 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...

Computer Vision


- Object Recognition
- Scene Classification
- Image Segmentation
- Human Activity Recognition

Scene Segmentation



(b)

(c)




Object Recognition

- window
- tree
- sky
- road
- pole
- car
- building
- unlabeled




<https://pjreddie.com/darknet/yolo/>

Google Goggles



Smile Detection



Leaf Snap



Image captioning: What begins to work



The flower was so
vivid and attractive.



Blue flowers are **running rampant** in my garden.

We sometimes do well: 1 out of 4 times, machine captions were preferred over the original Flickr captions:



Spring in a white dress.

Blue flowers have no scent. Small white flowers have no idea what they are.



Scenes around the lake on my bike ride.

This horse walking along the road as we drove by.



But many challenges remain (better examples of when things go awry)



The couch is definitely bigger than it looks in this photo.



My cat laying in my duffel bag.



Yellow ball suspended in water.



A high chair in the trees.




Game Agents

- 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
- 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.”



Game Agents

- Reinforcement learning




Pong



Enduro



Beamrider



Q*bert


2016



AlphaGo deep RL defeats Lee Sedol (4-1)


Simulated Agents

Iteration 0



Robotics

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



Images from UC Berkeley, Boston Dynamics, RoboCup, Google

Robots




Robocup

- https://www.youtube.com/watch?v=_PC-V5GJP6Q


Robocup








Watch later



Tools for Predictions & Decisions




Decision Making

- Applied AI in many kinds of automation:
 - Scheduling, airline routing
 - Route planning
 - Medical diagnosis
 - Web search
 - Spam classification
 - Automated help desks
 - Smarter devices, like cameras
 - Fraud detection
 - Product recommendation
 - ... Lots more!


Today

- Course overview
- What is artificial intelligence (AI)?
- What can AI do?
- **What is this course?**




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





Maximize Your Expected Utility




Topics in This Course

- Part I: Intelligence from Computation
 - Fast search
 - Adversarial and uncertain search
- Part II: Reasoning under Uncertainty
 - Bayes' nets
 - Decision theory
 - Machine learning
- Throughout: Applications
 - Natural language, vision, robotics, games, ...

Pac-Man as an Agent



Assignments: Pac-man




Originally developed at UC Berkeley:

<http://www-inst.eecs.berkeley.edu/~cs188/pacman/pacman.html>

This course vs. others

- CSE 515 – Stat methods
- CSE 517 – NLP
- CSE 546 – ML
- CSE 571 – Robotics
- CSE 576,7 – Vision
- Advanced RL




PS1: Search

Goal:

- Help Pac-man find his way through the maze

Techniques:

- Search: breadth-first, depth-first, etc.
- Heuristic Search: Best-first, A*, etc.




PS2: Game Playing

Goal:

- Play Pac-man!

Techniques:

- Adversarial Search: minimax, alpha-beta, expectimax, etc.




PS3: Ghostbusters

Goal:

- Help Pac-man hunt down the ghosts

Techniques:

- Probabilistic models: HMMS, Bayes Nets
- Inference: State estimation and particle filtering




PS4: Reinforcement Learning

Goal:

- Help Pac-man learn about the world

Techniques:

- Planning: MDPs, Value Iterations
- Learning: Reinforcement Learning



Important This Week

- Important this week:
 - **Check out** canvas--- our main resource for assignments and grades
 - **Check out** website – for schedule and slides
 - **Check out** Ed – for discussions; we are going to add everyone to Ed
 - **P0: Python tutorial** is out
- Also important:
 - **Office Hours** start next week.