CSE 473: Artificial Intelligence

Hanna Hajishirzi

https://courses.cs.washington.edu/courses/cse473/19au

Several slides from Luke Zettlemoyer, Dan Klein, Dan Weld, Stuart Russell, Andrew Moore

AI









Today

Course Format

• What is artificial intelligence (AI)?

• What can AI do?

• What is this course?

Course Staff

Hanna Hajishirzi	hannaneh@cs	Mondays 11-12	CSE654
Aida Amini	amini91@cs	Mondays 4-5	CSE386
Chris Clark	csquared@cs	Tue 2-3pm	Allen 220
Xinyue Chen	chenxy20@cs	Thu 2-3pm	Allen 220
Andrey Ryabtsev	ryabtsev@cs	Wed 5-6pm	Allen 220
Alyssa La Fleur	lafleur1@cs	Fri 11 -12	Allen 021
Svetoslav Kolev	swetko@cs	Wed 12-1pm	Allen 220

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

Website

• Website

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

o <u>https://courses.cs.washington.edu/courses/cse473/19au</u>

Canvas

• Communication, grades, submitting assignments:

Discussion board: ask and answer questions; announcements
 private matters - private messages

 \circ if your message is not answered promptly enough, here is the staff email:

Course Format

Programming Assignments

- ° 4 projects
- Python
- Autograded

- \circ Give you hands-on experience with the algorithms
- \circ I expect you to get 100% on projects

Written homeworks

- ° 2 written homeworks
- \circ Gives you a more conceptual understanding of the material

Course Format (continued)

• Exams

Midterm: Nov. 4th
Final: Dec. 9th
Both take home

- Exam Review Sessions
- Late days do not apply to midterm, final

Prerequisites

 Data Structure or Equivalent: CSE 332

• Math:

 $\circ\,$ Basic exposure to probability and data structures

• Programming - Familiar with Python

 \circ There is a 0th project (P0)

Textbook

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

Today

• Course overview

• What is artificial intelligence (AI)?

• What can AI do?

• What is this course?

What is AI?

The science of making machines

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

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

Pac-Man as an Agent

Pac-Man is a registered trademark of Namco-Bandai Games, used here for educational purposes

Machine Learning

AI

[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—: 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 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? X Drive safely along the highway? **P**Drive safely along University Avenue? **X** 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? **X** Write an intentionally funny story?

Unintentionally Funny Stories

3:24 PM

44 Tell me a story 99

- One day Joe Bear was history in the coak tree. He ate the
- Henry Squirrel was thirs river bank where his good Henry slipped an The End.
- Once upon a time there the crow was sitting in k He noticed that he was and swallowed the chee

riend d him lked to

o the s sitting. ravity drowned.

%

and a vain crow. One day ece of cheese in his mouth. cheese. He became hungry, over to the crow. The End.

[Shank, Tale-Spin System, 1984]

Natural Language

• Speech technologies (e.g. Siri)

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

$\circ\,$ 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

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

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

Google Goggles

Smile Detection

Image captioning: What begins to work

The flower was so vivid and attractive.

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

Blue flowers are running rampant in my garden.

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)

<u>The couch</u> 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.

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!

Robots

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
- $\circ~$ 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

Enduro

E

Q*bert

2016

AlphaGo deep RL defeats Lee Sedol (4-1)

Simulated Agents

Iteration 0

[Schulman, Moritz, Levine, Jordan, Abbeel, ICLR 2016]

Robotics

• Robotics

- Part mech. eng.
- Part AI
- Reality much harder than simulations!

• Technologies

- \circ 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

Robocup

Robocup (Stockholm '99)

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

Types of Environments

- Fully observable *vs.* partially observable
- Single agent *vs.* multiagent
- Deterministic *vs*. stochastic
- Static *vs*. sequential
- Discrete *vs*. continuous

Fully observable vs. Partially observable

VS.

Can the agent observe the complete state of the environment?

Single agent vs. Multiagent

VS.

Is the agent the only thing acting in the world?

Deterministic vs. Stochastic

Is there uncertainty in how the world works?

Static vs. Sequential

Does the agent take more than one action?

VS.

Discrete vs. Continuous

Is there a finite (or countable) number of possible environment states?

Utility?

Clear utility function

Not so clear utility function

Topics in This Course

• Part I: Making Decisions

• Fast search

Adversarial and uncertain search

• Part II: Reasoning under Uncertainty

- Bayes' nets
- Decision theory
- Machine learning

Throughout: Applications

Natural language, vision, robotics, games, ...

Assignments: Pac-man

Originally developed at UC Berkeley:

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

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 discussion and communication
 - Check out website- for schedule and slides
 - PO: Python tutorial is out
 - Mark exam dates in your calendars
- Also important:
 - Office Hours start next week.