

CS 47100 – Artificial Intelligence

Time: MW 4:30pm-5:45pm

Location: Lawson B155

Instructor: Elias Bareinboim

TAs: Daniel Kumor and Juan David Correa

Webpage: <http://www.cs.purdue.edu/~eb/>

Summary

The course will provide an introduction to foundational areas of AI and current techniques for building systems that exhibit ‘intelligent’ behavior and can ‘learn’ from experience. We’ll be covering some classic material (e.g., problem-solving, heuristic search techniques, game playing, knowledge representation, logical reasoning, graphical models) as well as more advanced topics (e.g., reinforcement learning, causal inference, counterfactual reasoning).

To give an example of how this material may be attractive to some students, consider the board game “Go”. Go was invented more than 5,000 years ago in China and is much harder than chess given its enormous ‘branching factor’, which is not amenable to exhaustive search. Interestingly, outstanding Go-playing relies on the notion of ‘intuition’ and ‘creativity’, which have been difficult concepts to mathematize and incorporate in traditional AI systems. Recent advances in the field of AI allowed computers to beat one of the greatest Go grandmasters (9-dan professional) alive, Lee Sedol. (The system was trained by a combination of passively observing previously played games (observational learning), playing against other opponents (reinforcement learning), and playing against itself (counterfactual reasoning).) After the games, Lee stated that his eventual loss to the machine was "inevitable" but said that "robots will never understand the beauty of the game the same way that we humans do." Not everyone agreed with this statement, and move 37 (2nd match) was considered a mistake by the Go experts watching the match, but after some time they realized that it was a ‘masterful’ and ‘genius’ move. As in any complex problem, this accomplishment relied on a combination of multiple techniques, which included three main components -- Monte Carlo tree search (TS), Reinforcement learning (RL), and Deep learning (DL). The first two components will be discussed in class.

By and large, it’s difficult to deny that AI systems are starting to become mainstream, being deployed to help in many traditional tasks and having influence in broader contexts than traditional CS. Other examples of this outreach include IBM Watson, which is capable of manipulating “natural language” and beating the best Jeopardy players (and now in med school), and Google self-driving cars. We’ll discuss the different aspects that make AI systems, including representation, inference, and learning. Specifically, we’ll examine the tradeoffs between how much can be learned from passive data, interactions with the environment, and by leveraging background, structured knowledge.

For a more visual presentation of AI systems operating in practice, see:

<https://www.youtube.com/watch?v=V1eYniJ0Rnk>

<https://www.youtube.com/watch?v=YXylqtEQ0tk>

Prerequisites

In order to be successful in this course, you should have a basic knowledge of:

- Discrete Math (search algorithms, proof techniques, and graph theory)
- Calculus (find min/max of functions)
- Statistics (basic probability, modeling, experimental design)
- Some programming experience (with special understanding of complexity analysis)

Material (tentative)

The following is a rough outline of the material we will cover.

Week	Subject	Material	References
1	Introduction, Uninformed Search		RN Ch. 1-3
2	Informed Search, CSPs		RN Ch. 3-4
3	CSPs (cont.), Local Search		RN Ch. 4,6
4	Adversarial Search		RN Ch. 5
5-6	Logic		
7	Planning		
8-10	Probability Review, Bayesian Networks		
11	Uncertainty and Utilities		
12	MDPs		
13	Reinforcement Learning		
14-15	Causal Inference		

References

We will use material selected from different sources, including chapters of the following books:

- [RN] Artificial Intelligence: A Modern Approach (Russell and Norvig, 2009), Pearson.
- [KF] Probabilistic Graphical Models: Principles and Techniques (Koller and Friedman, 2009), MIT Press.
- [Pea] Causality (Pearl, 2000), Cambridge Press.
- [SB] Reinforcement Learning: An Introduction (Sutton and Barto, 1998), Bradford.

RN and KF will be the main references.

Grading

The course will have both a theoretical and practical components. The project will allow students to chose between these two tracks.

- Midterm exam: 30%
- Homeworks: 30%
- Final exam: 40%

Please review the Purdue honor code. While working on assignments in small teams is okay, your homework solutions must be your own.

Course Policies

You are expected to attend lectures and participate in class. Small snacks are allowed. While taking notes on laptops, please make sure you are quiet and respectful of those around you, including those behind you who might be distracted by your snacks/devices. You are expected to come prepared to class, and to participate in class discussion, especially when something is not clear. If you are too shy to ask in class, please post on piazza or attend office hours.