CSCI-B 659: Reinforcement learning for Artificial Intelligence
tentative course outline

Adam White, School of Informatics and Computing, Indiana University, Bloomington

Course content: Reinforcement learning is a framework for modeling the an autonomous agent's interaction with an unknown world. The agent's objective is to learn the effects of it's actions, and modify its policy in order to maximize future reward. The study of Reinforcement learning emphasizes a learning approach to artificial intelligence. Unlike supervised learning, the agent is not explicitly told the correct answers (labels), rather an RL agent must learn only from reward and trial and error interaction with the world. This general framework has been used to optimize helicopter flight, schedule elevators, and achieve super-human level performance in many games (e.g., Backgammon, GO, and Atari). Ideas from reinforcement learning has also be used to explain learning in animals, and model dopamine activity in the human brain.

This course provides an introduction to some of the foundational ideas on which modern reinforcement learning is built, including Markov decision processes, value functions, Monte Carlo estimation, dynamic programming, temporal difference learning, eligibility traces, and function approximation. This course will develop an intuitive understanding of these concepts (taking the agent's perspective), while also focusing on the mathematical theory of reinforcement learning. Programming assignments and projects will require implementing and testing complete decision making systems.

The objective of this course is twofold. The first is to prepare you for conducting research in reinforcement learning. The second is to provide you with the required knowledge to apply reinforcement learning techniques to novel applications.

Topics will include:

- Course overview and introduction to reinforcement learning
- Bandit problems and online learning
- Markov decision processes
- Returns, and value functions
- Solution methods: Dynamic Programming
- Solution methods: Monte Carlo learning
- Solution methods: Temporal difference learning learning
- Eligibility traces
- Models and planning (table lookup case)
- Value function approximation (function approximation)
- Frontiers of RL research (last year this included: policy gradient methods, off-policy learning, least squares methods)

- RL ins Psychology and Animal Learning
- Case studies: successful examples of RL systems

Prerequisites: This course will rely on basic statistics (e.g., probability distributions and expected values), basic linear algebra (e.g., inner products), and basic calculus (e.g., computing simple derivatives). I will not be covering these topics, rather I will assume you know them. You will need to be able to program in C. We will all use the same RL-glue framework for all experiments and assignments—this framework is written in C. No exceptions will be made for other programming languages.

Grading:

- 50% Assignments. The assignments will include both programming and pen and paper problems.
- 10% Thought questions.
- 5% Midterm quiz
- 35% Final exam or Project (project for Phd students and students with permission only).

Text and resources:

- **Required:** Reinforcement Learning: An Introduction, by Richard S. Sutton and Andrew G. Barto

We will be using the online 2016 Sept edition: http://incompleteideas.net/sutton/book/bookdraft2016sep.pdf