CMPUT 609 TOPICS IN COMPUTING SCIENCE (LEC A1 Fa16)

Course Outline

CMPUT 609/499 Reinforcement Learning in Artificial Intelligence Course Outline

General Information

Term: Fall 2016
Date and Time: TR, 3:30-4:50. First class Sept 1. Last class Dec 6.

Location: ETLC E1-008

Number of credits: 3
Instructor: Rich Sutton, rsutton@ualberta.ca, http://www.cs.ualberta.ca/~sutton)

Office: Athabasca 3-13

Overview

This course will provide a comprehensive introduction to reinforcement learning as an approach to artificial intelligence, emphasizing the design of complete agents interacting with stochastic, incompletely known environments. Reinforcement learning has adapted key ideas from machine learning, operations research, psychology, and neuroscience to produce some strikingly successful engineering applications. The focus is on algorithms for learning what actions to take, and when to take them, so as to optimize long-term performance. This may involve sacrificing immediate reward to obtain greater reward in the long-term or just to obtain more information about the environment. The course will emphasize the development of intuition relating the mathematical theory of reinforcement learning to the design of human-level artificial intelligence.

Objectives

After this you will be prepared to conduct research in reinforcement learning and be able to apply reinforcement learning techniques to novel applications

be able to apply reinforcement learning techniques to novel applications.

Pre-requisites

Interest in learning approaches to artificial intelligence; basic probability theory; basic linear algebra; computer programming ability. You should be comfortable with statistical ideas such as probability distributions and expected values.

Course Topics

The course will cover Markov decision processes, dynamic programming, temporal-difference learning, Monte Carlo reinforcement learning methods, eligibility traces, the role of function approximation, and the integration of learning and planning.

Course Work and Evaluation

Grading will be on the basis of (with relative weighting):

- 5 Written exercises (5)
- Mid-term exam (4)
- 3 Programming projects (4)
- Final project (4)

Course Materials

Textbook: *Reinforcement Learning: An Introduction*, second edition in progress, by Richard S. Sutton and Andrew G. Barto

This and other course materials will be distributed via the course dropbox

Academic Integrity

The University of Alberta is committed to the highest standards of academic integrity and honesty. Students are expected to be familiar with these standards regarding academic honesty and to uphold the policies of the University in this respect. Students are particularly urged to familiarize themselves with the provisions of the Code of Student Behaviour and avoid any behaviour which could potentially result in suspicions of cheating, plagiarism, misrepresentation of facts and/or participation in an offence. Academic dishonesty is a serious offence and can result in suspension or expulsion from the University.

Last modified: Thursday, 1 September 2016, 12:20 PM