

COMP 7991 / 8991: ST: Reinforcement Learning – Spring 2021

Basic Information

Time & Place: TR 11:20 AM--12:45 PM

Instructional mode: online

Instructor: Weizi Li, E-mail: wli@memphis.edu

Office Hours: by appointment

Course Overview

Reinforcement learning is a branch of machine learning that learns to make sequential decisions. In recent years, we have seen a surge of breakthroughs based on reinforcement learning, such as AlphaGo. This is made possible by the increased computational power and advances in deep neural networks. As a general framework for sequential decision making, reinforcement learning has the potential to deliver novel solutions in robotics, video games, education, transportation, finance, healthcare, and more. This course will introduce the theory and algorithms of classic and deep reinforcement learning. Students will learn from frontier research results as well as build a reinforcement learning system to solve a problem in a course project of their interest.

Topics include reinforcement learning foundations; bandits; markov decision process; dynamic programming; model-free methods; function approximation; deep neural networks; behavioral cloning; inverse reinforcement learning; deep reinforcement learning; policy gradient methods; actor-critic methods; model-based methods; exploration and exploitation.

Prerequisite

Probability theory, linear algebra, and basic machine learning knowledge

Textbook

No textbook is required.

Tentative Schedule

Unit 1: RL Foundations

Unit 2: Model-free Methods: Basics

Unit 3: Function Approximation

Unit 4: Learning from Demonstrations

Unit 5: Model-free Methods: Policy Gradients

Unit 6: Model-based Methods

Unit 7: Additional Topics

Evaluation

Grades will be computed as follows:

10% Class participation

30% Paper presentation

60% Course project

≥ 87%: A+ or A or A-

≥ 79%: B+ or B or B-

≥ 68%: C+ or C or C-

≥ 60%: D+ or D

< 60%: F

Plus/minus grading will be used based on significant performance differences on the project within each range. Students taking the class at the 8000 level are expected to engage in more encompassing projects, as well as to turn in solutions of a higher quality.

Project Report Policy. Students are required to submit the PDF file of their homework solutions and project reports. A word processing software (e.g., LaTeX or Word) is recommended.

Plagiarism

Plagiarism or cheating behavior in any form is unethical and detrimental to proper education and will not be tolerated. All work submitted by a student (projects, programming assignments, lab assignments, quizzes, tests, etc.) is expected to be a student's own work. The plagiarism is incurred when any part of anybody else's work is passed as your own (no proper credit is listed to the sources in your own work) so the reader is led to believe it is therefore your own effort. Students are allowed and encouraged to discuss with each other and look up resources in the literature (including the internet) on their assignments, but appropriate references must be included for the materials consulted, and appropriate citations made when the material is taken verbatim.

If plagiarism or cheating occurs, the student will receive a failing grade on the assignment and (at the instructor's discretion) a failing grade in the course. The course instructor may also decide to forward the incident to the Office of Student Conduct for further disciplinary action. For further information on U of M code of student conduct and academic discipline procedures, please refer to <https://www.memphis.edu/osa/students/academic-misconduct.php>.

"Your written work may be submitted to Turnitin.com, or a similar electronic detection method, for an evaluation of the originality of your ideas and proper use and attribution of sources. As

part of this process, you may be required to submit electronic as well as hard copies of your written work, or be given other instructions to follow. By taking this course, you agree that all assignments may undergo this review process and that the assignment may be included as a source document in Turnitin.com's restricted access database solely for the purpose of detecting plagiarism in such documents. Any assignment not submitted according to the procedures given by the instructor may be penalized or may not be accepted at all." (Office of Legal Counsel, October 17, 2005).

Disability Notice

Any student who anticipates physical or academic barriers based on the impact of a disability is encouraged to speak with me privately. Students with disabilities should also contact Disability Resources for Students (DRS) at 110 Wilder Tower (901-678-2880). DRS coordinates access and accommodations for students with disabilities (<http://www.memphis.edu/drs/>).