EC 710: Dynamic Programming and Reinforcement Learning
Spring 2025

Instructor Information

Name: Alex Olshevsky
Office: PHO 419
Email: alexols@bu.edu
Office Hours: TBD

Class Information

Time: MW 10:10-11:55
Classroom: EMB 105

Course Description

Reinforcement learning is a subfield of artificial intelligence which deals with learning from repeated interactions with an environment. Reinforcement learning is the basis for state-of-the-art algorithms for playing strategy games such as Chess, Go, Backgammon, and Starcraft, as well as a number of problems throughout robotics, operations research, and other fields of engineering. In this course, we will study the fundamental algorithms of reinforcement learning. Our goal will be to understand the mathematical principles underlying these algorithms and to test them on several popular benchmarks.

This class will be math-heavy.

Prerequisites: Knowledge of fundamentals of probability and linear algebra.

Course Outline

1. Markov Decision Processes. Dynamic Programming.
2. Linear Quadratic Regulator.
3. Value Iteration and Policy Iteration.
4. Monte Carlo, Temporal Difference Methods, and Q-learning.
5. Function Approximation in RL. Neural Networks and Nonlinear Function Approximation.
6. Policy Gradient and Actor Critic Methods.

Grading

(a) A robot hand solves a Rubik's cube.

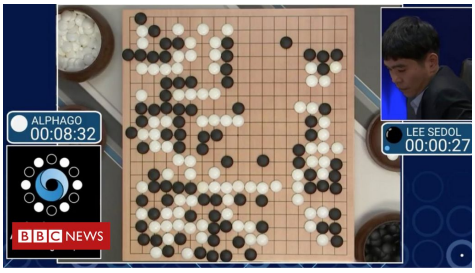

(b) AlphaGo plays the (human) world champion.

(c) Spectators watch AlphaStar play Starcraft.

(d) A quadruped navigates a rocky terrain.

Figure 1: Four systems built with reinforcement learning.

- Homeworks: 20%
- Midterm: 40%
- Final Project: 40%

Your lowest homework grade will be dropped.

With an acceptable written excuse, a missed exam will be dropped from the computation of the final grade.

Academic Policy

BU takes academic integrity very seriously. Academic misconduct is conduct by which a student misrepresents his or her academic accomplishments, or impedes other students's opportunities of being judged fairly for their academic work. Knowingly allowing others to represent your work as their own is as serious an offense as submitting another's work as your own. More information on BU's Academic Conduct Code, with examples, may be found at http://www.bu.edu/academics/policies/academic-conduct-code

Collaboration Policy

In this class you may use any textbooks or web sources when completing your homework and

programming exercises. You may also use human collaborators from class subject to the following strictly enforced conditions:

- You must clearly acknowledge all your sources (including your collaborators) on the top of your homework.
- You must write all answers in your own words and write your own code.
- You must be able to fully explain your answers upon demand.
- You may not use any human resource outside of class (including web-based help services, outside tutors, etc) in doing your homeworks or programming exercise.
- Obviously, you may not collaborate with anyone on quizzes.

Failure to meet any of the above conditions would constitute plagiarism and will be considered cheating in this class. If you are not sure whether something is permitted by the course policy, ASK ME! (it's much more awkward to explain your actions after the fact to the college disciplinary committee). The penalty for academic misconduct at BU is severe.

AI Policy

You can use any AI service (e.g., ChatGPT, Claude) on the homework in any capacity as long as you acknowledge it in your submission.

Grade Scale

Final grades will be assigned according to the following scale:

A	93 – 100	C+	77 – 79.99
A–	90 – 92.99	C	73 – 76.99
B+	87 – 89.99	C–	70 – 72.99
B	83 – 86.99	D	60 – 69.99
B–	80 – 82.99	F	0 – 59.99

However, I may curve up the grades on individual assignments or quizzes.

Incomplete Grades

Incomplete grades will not be given to students who wish to improve their grade by taking the course in a subsequent semester. An incomplete grade may be given for medical reasons if a physician's note is provided. The purpose of an incomplete grade is to allow a student who has a legitimate interruption in the course to complete the remaining material in another semester. In particular, students will not be given an opportunity to improve their grade by doing "extra work."

Drop Date

Students are responsible for being aware of the drop dates for the current semester. Drop forms will not be back-dated.