

EC 418: Introduction to Reinforcement Learning

Fall 2023

Instructor Information

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

Class Information

Time: MW 4:30-6:15
Classroom: LSE B03

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 principles underlying these algorithms and to test them on several popular benchmarks.

Prerequisites: Multi-variable calculus, Linear Algebra, and Probability.

Course Outline


1. Probability Review
2. Dynamic Programming and Markov Decision Processes, Value Iteration and Policy Iteration.
3. Monte Carlo, Temporal Difference Methods, and Q-learning.
4. Function Approximation in RL, Neural networks.
5. Recurrent Neural Networks for Function Approximation (time permitting).
6. Policy Gradient and Actor Critic (time permitting).

A good resource to read up on these topics are the the lecture videos for CS 285 at Stanford

<https://www.youtube.com/watch?v=FgzM3zpZ55o&list=PLoROMvodv4rOSOPzutgyCTapiGLY2Nd8u>



(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.

All of the programming assignments in this class will be done in Python, and you will benefit from previous exposure. The beginning of the class will contain a fast-paced introduction to Python for students who have not encountered it before. You will probably be fine if you have lots of programming experience in a different language.

Textbook

We will not be using a textbook. Lecture slides will be posted on blackboard. The closest textbook is *Reinforcement Learning: An Introduction* by Sutton and Barto, which we will follow for some parts of the class. It is available for a free download at <http://incompleteideas.net/book/the-book.html>

Place in the BU Curriculum

There is a some overlap between this class and EC 414, as both classes cover neural networks. However, about 95% of this class is different from EC 414, and if you are interested in machine learning, you will benefit from taking both classes.

Grading

- Homework: 5%
- Coding exercises: 5%
- Midterm 1: 35%

- Midterm 2: 35%
- Final Project: 20%

Your lowest homework grade will be dropped, along with your lowest programming assignment grade.

Midterms are closed-book and closed-notes. Calculators, computing, and communication devices are neither needed nor permitted. However, you are allowed to bring one 8.5×11 -inch sheet of handwritten notes (both sides).

There will be no make-up exams, even in the case of an emergency. A missed exam counts as a zero unless a valid excuse from a physician is presented. 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. However, this is subject to the following strictly enforced conditions:

- You must clearly acknowledge all your sources (including your collaborators, or any AI models used) on the top of your homework.
- 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 exams.

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.

Grade Scale

Final grades will be assigned according to the following scale:

A	93 – 100	C+	77 – 79
A–	90 – 92	C	73 – 76
B+	87 – 89	C–	70 – 72
B	83 – 86	D	60 – 69
B–	80 – 82	F	0 – 59

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.