Reinforcement Learning

(CSCI 4160/6963, ECSE 4965/6965)

Course Syllabus

Course description

This is an introductory course on the theory and practice of reinforcement learning (RL). We will start with an introduction to linear algebra and probability and will quickly go over standard supervised learning techniques, such as linear regression. After that, we will derive the RL framework, starting from Markov chains and Markov reward processes and building up to Markov decision processes. We will then cover classic RL approaches such as dynamic programming, Monte Carlo methods, Q-learning and policy gradients. In the last part of the course, we will cover deep learning and deep RL.

Learning Outcomes

Students who have successfully completed this course will:

- acquire a solid theoretical understanding of modern reinforcement learning algorithms;
- work on and understand the challenges of classical reinforcement learning benchmarks;
- be able to implement classical reinforcement learning algorithms such as dynamic programming and Q-learning;
- understand the challenges of deep learning and deep reinforcement learning;
- be able to apply reinforcement learning in new domains related to robotics and control.

Instructor: Rado Ivanov

Office: Lally 309
Email: ivanor@rpi.edu

Course Mechanics

Homework submission: Homework assignments will be submitted through LMS. If you have any issues with the system, let me know as soon as possible.

Attendance: While I will not formally collect attendance sheets, please try to attend all lectures, unless you have a good reason not to attend. Some of the assignments will be hard to complete if you don't attend the lectures.

Posting homework solutions online: You are not allowed to post your solutions on github or any other online repository. If you would like to show off your projects to potential employers, it would be much more effective to do so with extracurricular activities such as RCOS and so on.

Covid guidelines: If you have to skip a class due to covid, please let me know as soon as possible. Where possible, I will try to provide lecture recordings for those cases. **Note that lecture recordings will not be provided for any other reason and cannot be used as a substitute for in-person attendance.**

Textbooks

There is no required book for the course since there is no individual book that covers all the course material. All of the necessary material will be included in the lecture slides. I will suggest additional reading material before each lecture. We will follow some of the material in the following books, (most of) which are available for free online:

- Hastie, Trevor, et al. The elements of statistical learning: data mining, inference, and prediction. Vol. 2. New York: springer, 2009. (available online: https://hastie.su.domains/Papers/ESLII.pdf)
- James, Gareth, et al. An introduction to statistical learning. Vol. 112. New York: springer, 2013. (available online: https://www.statlearning.com/)
- Ian Goodfellow, Yoshua Bengio, and Aaron Courville. Deep learning. MIT press, 2016. (available online: https://www.deeplearningbook.org/)
- Sutton, Richard S., and Andrew G. Barto. Reinforcement learning: An introduction. MIT press, 2018. (available online: http://incompleteideas.net/book/RLbook2020.pdf)
- Puterman, Martin L. *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons, 2014.

Grading

Students will be graded on 10 homework assignments. Each assignment is worth 10% of the final grade.

Late submission rule: You can use up to 2 extensions – you don't need to explain why you submitted late. Each extension is for 3 days, midnight to midnight. If your submission is more than 3 days late (or more than 1 minute late in case you have used up your extensions), you will receive a score of 0. **You cannot use both extensions on the same assignment.**

Re-grading: If you think there was a mistake in your grade, you can submit a re-grade request, and we will regrade the entire assignment. You are allowed up to 2 re-grades for the semester.

Collaboration and Academic Honesty

You are expected to work alone on all assignments. In particular:

- Discussion is allowed on homework but submitted work must be your own.
- **YOU ARE RESPONSIBLE FOR ENSURING THAT YOUR HOMEWORKS ARE NOT COPIED.**
- Copying from **anywhere** other than the class notes or your notes is NOT allowed.
- You must write and understand all solutions yourself.

In cases of academic dishonesty, the minimum penalty is a course grade of F, and other institute-mandated protocols may be invoked.