DS 669 - Reinforcement Learning
Spring 2024

Class schedule:
Fridays, 06:00 PM-08:50 PM KUPF 205

Instructor:
Jing Li, jingli@njit.edu, GITC 4419

Class Overview:
Reinforcement learning is widely used in many engineering and scientific disciplines, such as autonomous driving, robotics, optimization, psychology, and neuroscience. It emphasizes learning by an agent from direct interaction with its environment, without relying on supervision. Of all the forms of machine learning, reinforcement learning is the closest to the learning that humans do. Thus, it has become one of the most active research areas in machine learning, artificial intelligence, and neural network research. Moreover, the field has developed impressive applications in the industry. This course covers current topics, key concepts, classic and modern algorithms in reinforcement learning and contains both theory and applications. The topics include but are not limited to Markov Decision Process, exploration and exploitation, planning, value-based learning, policy gradient. Students will present recent papers in reinforcement learning, work on written and programming assignments and do a reinforcement learning project. After completing this course, students will be able to start using reinforcement learning for real problems that can be specified as the MDP.

Tentative schedule:
1. Intro to Reinforcement Learning and Sequential Decision Making;
2. Markov Decision Process;
3. Value Functions, Bellman Equations;
4. Value Iteration, Policy Iteration;
5. Temporal Difference Learning;
6. Q-Learning;
7. Neural Networks, DQN;
8. Policy Gradient Methods;
9. Midterm Project Presentation
10. Actor-Critic Methods;
11. Multi-armed Bandits;
12. Exploration and Exploitation;
13. Monte Carlo Tree Search;
14. Final Project Presentation;

Disclaimer: *The schedule of the course is subject to change based on the progress of the class, including test dates after they are announced. Changes will be announced as early as possible.*

Course format:
The course will study principles, challenges and some state-of-the-art solutions in reinforcement learning. It also involves programming assignments and a course project.

You are encouraged to collaborate on assignments. However, you must cite all your collaborators (teammates) and any sources beyond the class materials that you consulted while working on a problem—for example, an "expert" consultant other than the teacher, or another text—must be given a proper scholarly citation, which you should include with your submission.

Programming assignments:
Students can work on the programming assignments in teams of up to 3 members. Teams may change from assignment to assignment, but the sharing of code between teams is strictly prohibited. Solutions submitted on time (as determined by Canvas's receipt time stamp) do not have penalty. Solutions submitted up to 72 hours late will be given a 20% penalty. Solutions submitted less than a week late will be given a 50% penalty. Solutions submitted after a week late will not be given credit.

Course project:
Students can work on the course project in the team of up to 3 members. Students will present their project progress to the class in the project proposal, midterm demo and final project demo. The final project outcome will be in the form of a final report, describing the design, analysis, implementation and experimentation efforts, together with the code submission. The report must acknowledge and document in detail all contributions that anyone has made to the work.

Course materials:
Textbook:
Reinforcement Learning: An Introduction by Rich Sutton and Andrew Barto (available online)

Readings:
Topics in class will be accompanied by some reading materials or classical/recent research papers from top conferences and journals.

Prerequisites:
Students should have learned Linear algebra, basic probability, basic calculus, computer programming, OR approval of instructor. Experience with machine learning, artificial intelligence, or deep learning (e.g., CS 675, CS 670, CS 677) is recommended.

Grading:
Assignments – 60 points
Project – 40 points
Final grades will be curved according to departmental policy.

Office Hours and Contact Information:
Fridays, 04:00 PM - 06:00 PM
Please make appointments by email for individual meeting times.
Jing Li: jingli@njit.edu, GITC 4419

Honor Code:

Academic Integrity is the cornerstone of higher education and is central to the ideals of this course and the university. Cheating is strictly prohibited and devalues the degree that you are working on. As a member of the NJIT community, it is your responsibility to protect your educational investment by knowing and following the academic code of integrity policy that is found at: https://www5.njit.edu/policies/sites/policies/files/NJIT-University-Policy-on-Academic-Integrity.pdf.

Please note that it is my professional obligation and responsibility to report any academic misconduct to the Dean of Students Office. ***Any student found in violation of the code by cheating, plagiarizing or using any online software inappropriately will result in disciplinary action. This may include a failing grade of F, and/or suspension or dismissal from the university.*** *If you have any questions about the code of Academic Integrity, please contact the Dean of Students Office at dos@njit.edu.*

Modifications to syllabus:

The syllabus may be modified at the discretion of the instructor or in the event of extenuating circumstances. Students will be notified in class of any changes to the syllabus.