EE 597 Reinforcement Learning, Spring 2024

Instructor: Jing Yang
Email: yangjing@psu.edu
Office: 228B EE West, Phone: (814) 863-8732
Lectures: TuTh 12:05PM - 1:20PM, Hammond Bldg 217
Office Hour: TuTh 1:30PM - 2:30PM or by appointment

Course Description

To realize the impact of AI in real-world systems requires autonomous decision making based on historical observations. Reinforcement learning is one powerful paradigm for doing so, and it is relevant to an enormous range of applications, including robotics, game playing, cyber-physical system control and optimization, etc. This class will provide a solid introduction to the field of reinforcement learning, including the basics of reinforcement learning as well as deep reinforcement learning.

Learning Outcomes

By the end of the class students should be able to:

- Understand the key differences between reinforcement learning and non-interactive machine learning.
- Given an application problem (e.g. from robotics, IoT, NLP, etc), decide whether it should be formulated as an RL problem; if yes, be able to define it formally (in terms of the state space, action space, dynamics and reward model), state what algorithm is best suited for addressing it and make it work.
- Implement common RL algorithms in code.
- Describe various performance metrics (e.g. regret, sample complexity, computational complexity, empirical performance, convergence, etc.) for RL algorithms and evaluate algorithms on these metrics.
- Understand the exploration vs exploitation dilemma and existing approaches to addressing this challenge.

Prerequisites

- Proficiency in Python. All coding assignments will be in Python.
- College Calculus, Linear Algebra.
- Basic Probability and Stochastic Processes (EE560 or equivalent).
- Basics in Optimization and Machine Learning.

Course Webpage

- Canvas: http://canvas.psu.edu

References

A collection of notes, relevant papers and materials will be prepared and distributed. Textbooks recommended for further reading are:

- *Reinforcement Learning: An Introduction*, Sutton and Barto, 2nd Edition, The MIT Press, 2018. `http://incompleteideas.net/book/the-book.html`
- *Bandit Algorithms*, Tor Lattimore and Csaba Szepesvari, Cambridge University Press, 2020. `https://tor-lattimore.com/downloads/book/book.pdf`

Course Outline

1. Introduction to RL
2. Finite Markov Decision Processes (MDP)
3. Dynamic Programming
4. Monte Carlo Methods
5. Temporal Difference Learning
6. Value Function Approximation
7. Policy Search Methods
8. Bandit Problems
9. Fast RL
10. Offline Reinforcement Learning
11. Imitation Learning

Course Project

Course project is an opportunity for you to apply what you have learned in class to a problem of your interest in reinforcement learning. The project could be either original research problems in RL or the application of RL in different domains. You will complete it *in groups of two*.

Project deadlines:

- **Week 6 (Sunday, 02/18, 11:59 PM)** Submit a short proposal (no more than 2 pages) stating the topics that you plan to work on. Describe why they are important or interesting, and provide some appropriate references.
- **Week 12 (Sunday, 03/17, 11:59 PM)** Submit a progress report that explicitly refers back to your project proposal: what has been accomplished, what goals should be revised, etc.
- **Weeks 15-16** Presentation and peer evaluation. Deliver a 20-minute presentation with slides in class.
- **Finals week (Monday, 04/29, 11:59 PM)** Submit the final project report.

You are required to use LaTeX for your reports in this course, as LaTeX is a skill you should learn if you haven't already! You should use the official NeurIPS template to structure your reports.

Grading

- Homework assignments (45%), Midterm (20%), Final project (35%).
- There will be three homework assignments. Each homework assignment may contain a written part and a programming part. All programming assignments will be in Python.
- All homework will be submitted to Canvas before the deadline. Late homework will not be accepted. You homework submission should include: 1) a PDF file containing solutions to the written part, 2) the report for the programming part, and 3) your own code. Your code should only use packages from the standard library.
- For written homework assignments, you are welcome to discuss ideas with others, but you are expected to write up your own solutions independently (without referring to another's solutions). For programming homework assignments, you may only share the input-output behavior of your programs. Please remember that if you share your solution with another student, even if you did not copy from another, you are still violating the academic integrity.
- Final project grade consists of three parts: Proposal (5%), Progress report (2%), Oral Presentation (14%), Final Report (14%).
- The final grade will be assigned according to the table below:

Score	[0, 60)	[60, 70)	[70, 77)	[77, 80)	[80, 83)	[83, 87)	[87, 90)	[90, 93)	[93, 100]
Grade	F	D	C	C+	B-	B	B+	A-	A

Academic Integrity

Academic integrity is the pursuit of scholarly activity in an open, honest and responsible manner. Academic integrity is a basic guiding principle for all academic activity at The Pennsylvania State University, and all members of the University community are expected to act in accordance with this principle. Consistent with this expectation, the University's Code of Conduct states that all students should act with personal integrity, respect other students' dignity, rights and property, and help create and maintain an environment in which all can succeed through the fruits of their efforts.

Academic integrity includes a commitment by all members of the University community not to engage in or tolerate acts of falsification, misrepresentation or deception. Such acts of dishonesty violate the fundamental ethical principles of the University community and compromise the worth of work completed by others.

Disability Accommodation Statement

Penn State welcomes students with disabilities into the University's educational programs. Every Penn State campus has an office for students with disabilities. Student Disability Resources (SDR) website provides contact information for every Penn State campus (`http://equity.psu.edu/sdr/disability-coordinator`). For further information, please visit Student Disability Resources website (`http://equity.psu.edu/sdr/`).

In order to receive consideration for reasonable accommodations, you must contact the appropriate disability services office at the campus where you are officially enrolled, participate in an intake interview, and provide documentation: See documentation guidelines (`http://equity.psu.edu/sdr/guidelines`). If the documentation supports your request for reasonable accommodations, your campus disability services office will provide you with an accommodation letter. Please share this letter with your instructors and discuss the accommodations with them as early as possible. You must follow this process for every semester that you request accommodations.

Counseling and Psychological Services Statement

Many students at Penn State face personal challenges or have psychological needs that may interfere with their academic progress, social development, or emotional wellbeing. The university offers a variety of confidential services to help you through difficult times, including individual and group counseling, crisis intervention, consultations, online chats, and mental health screenings. These services are provided by staff who welcome all students and embrace a philosophy respectful of clients' cultural and religious backgrounds, and sensitive to differences in race, ability, gender identity and sexual orientation.

Counseling and Psychological Services at University Park (CAPS)
(`http://studentaffairs.psu.edu/counseling/`): 814-863-0395

Penn State Crisis Line (24 hours/7 days/week): 877-229-6400
Crisis Text Line (24 hours/7 days/week): Text LIONS to 741741

Educational Equity and Reporting Bias

Penn State takes great pride to foster a diverse and inclusive environment for students, faculty, and staff. Acts of intolerance, discrimination, or harassment due to age, ancestry, color, disability, gender, gender identity, national origin, race, religious belief, sexual orientation, or veteran status are not tolerated and can be reported through Educational Equity via the Report Bias webpage (`http://equity.psu.edu/reportbias/`).

Copyright Statement

All course materials students receive or to which students have online access are protected by copyright laws. Students may use course materials and make copies for their own use as needed, but unauthorized distribution and/or uploading of materials without the instructor's express permission is strictly prohibited. University Policy AD 40, the University Policy Recording of Classroom Activities and Note Taking Services addresses this issue. Students who engage in the unauthorized distribution of copyrighted materials may be held in violation of the University's Code of Conduct, and/or liable under Federal and State laws.