COMP 5801W
Reinforcement Learning
Course Outline

Instructor: Junfeng Wen (junfeng.wen [AT] carleton.ca)

Winter 2025
School of Computer Science
Carleton University

Course Information

Instructor: Junfeng Wen

Contact: junfeng.wen [AT] carleton.ca

Classroom: Please check Carleton Central or Brightspace for the room location

Lectures: Tuesdays 2:35 pm - 5:25 pm

Office Hours: By appointment

Brightspace access for University of Ottawa Students: Please see information here. uOttawa OCICS students can find the room location on Brightspace.

Important dates and deadlines can be found here, including class suspension for fall, winter breaks, and statutory holidays.

Course Description

Reinforcement learning (RL) focuses on learning by interacting with a complicated environment through trial and error. In this research-oriented course, students will first learn the fundamental concepts of RL, including Markov decision processes, value prediction, and optimal control. Then we will investigate several advanced topics in the recent RL literature, such as offline/batch RL, distributional RL and RL applications. Priority will be given to OCICS joint institute students and students in thesis-based research programs.

Precludes Any additional credit for COMP 5900 Introduction to Reinforcement Learning.

Prerequisites Students are expected to be familiar with linear algebra, calculus, basic statistics and Python programming.

Recommended Textbooks

- *Reinforcement Learning: An Introduction* (2nd edition), Sutton Barto
- *Algorithms for Reinforcement Learning*, Szepesvari
- *Markov Decision Processes: Discrete Stochastic Dynamic Programming*, Puterman

Students are not required to purchase textbooks or other learning materials for this course.

Topics Covered and Learning Outcomes

This course will cover the following (tentative) topics

- Markov decision process
- Dynamic programming methods
- Monte-Carlo methods
- Temporal difference learning
- Prediction and control with function approximation
- Policy gradient and actor-critic methods
- Common deep RL algorithms
- Offline/Batch RL
- Distributional RL
- Exploration
- RL applications

Upon completion, students should be able to

- Develop a solid understanding of the fundamental concepts and principles in reinforcement learning
- Understand a wide range of reinforcement learning algorithms, their applicability, strengths and weaknesses
- Design and implement reinforcement learning algorithms for real-world problems, and evaluate their performance

Assessment Scheme

Two assignments 40% (20% each)

- Done individually
- For assignments, you have three excused days **throughout the term** (rounded up to the nearest day) to account for any unforeseeable difficulties. After that no late submission will be accepted without proper justifications
- Submissions are handled electronically (i.e., through Brightspace). Technical problems do not exempt you from late policy, so if you wait until the last minute and then have issues with your connection, it will still count as a late submission. Consequently, you are advised to
    1. Periodically upload you progress
    2. Attempt to submit your final submission early (e.g., at least one hour in advance of the due date and time) and
    3. Download the submitted files to make sure they are correct
- You have **one week** to change the assignment grade after posted if there is any issue. After that there will be no regrade

Paper presentation 10%

Project 50%

- Proposal (0%)
- In-class demo (10%)
- Final report (40%)

Graduate Academic Advisors

The Graduate Advisors for the School of Computer Science are available in Room 5302 HP; or by email at grad.scs@carleton.ca. The graduate advisors can assist with understanding your academic audit and the remaining courses required to meet graduation requirements.

Intellectual Property

All materials created for this course (including, but not limited to, lecture notes, in-class examples, tutorial exercises, assignments, examinations, and posted solutions) remain the intellectual property of the instructor. These materials are intended for the personal and non-transferable use of students registered in the current offering of the course. Reposting, reproducing, or redistributing any course materials, in part or in whole, without the written consent of the instructor, is strictly prohibited.

Sharing assignment or quiz specifications or posting them online (to sites like Chegg, CourseHero, OneClass, etc.) is considered academic misconduct. You are never permitted to post, share, or upload course materials without explicit permission from your instructor. Academic integrity offences are reported to the office of the Dean of Science. Penalties for such offences can be found on the ODS webpage.

Academic Accommodations and Regulations

Academic Accommodations. Carleton is committed to providing academic accessibility for all individuals. You may need special arrangements to meet your academic obligations during the term. The accommodation request processes are outlined on the Academic Accommodations website.

ChatGPT/Generative AI Usage. As our understanding of the uses of AI and its relationship to student work and academic integrity continue to evolve, students are required to discuss their use of AI in any circumstance not described here with the course instructor to ensure it supports the learning goals for the course. Students can use AI tools for checking grammar/spell, or rewriting part of the submitted documents for better explanation/exposition. However, students cannot use AI tools to generate coding solutions or large volumes of text. Any AI-generated content should be clearly identified/highlighted.

Academic Integrity. Every student should be familiar with the Carleton University student academic integrity policy. A student found in violation of academic integrity standards may be awarded penalties which range from a reprimand to receiving a grade of F in the course or even being suspended or expelled from the program or University. Examples of punishable offences include: plagiarism and unauthorized co-operation or collaboration. Any such reported offences will be reviewed by the office of the Dean of Science. More information on this policy may be found on the ODS Academic Integrity page.

Student Rights & Responsibilities. Students are expected to act responsibly and engage respectfully with other students and members of the Carleton and the broader community. See the 7 Rights and Responsibilities Policy for details regarding the expectations of non-academic behaviour of students. Those who participate with another student in the commission of an infraction of this Policy will also be held liable for their actions.

Unauthorized Collaboration. Senate policy states that "to ensure fairness and equity in assessment of term work, students shall not co-operate or collaborate in the completion of an academic assignment, in whole or in part, when the instructor has indicated that the assignment is to be completed on an individual basis".