Western Science

Department of Computer Science

Course Outline
Reinforcement Learning (9670/9671/9170)

1. Course Information

Course Information

Reinforcement Learning (9670/9671/9170), Winter 2023

2. Instructor Information

Instructors	Email	Office	Phone	Office Hours
Dr. Apurva Narayan	apurva.narayan@uwo.ca	MC 368	519 661-2111 x81138	

Students must use their Western (@uwo.ca) email addresses when contacting their instructors. Please start the subject as 'Reinforcement Learning'.

Office hours will be via Zoom and set up on a need basis via email.

3. Course Syllabus, Schedule, Delivery Mode

This course will provide a broad introduction to the foundational concepts and algorithms of reinforcement learning, one of the largest and most active areas in machine learning. The main focus will be on fundamental algorithms and their applications, and will end with an introduction to deep reinforcement learning. Knowledge of probability theory, logic, expectation, and basic machine learning principles (e.g., gradient descent) will be very helpful.

Upon successful completion of this course, students will be able to:

- Define the core features of reinforcement learning, and explain how RL differs from other artificial intelligence / machine learning approaches.
- Determine if a given problem should be approached as a reinforcement learning problem. Compare different algorithms to select the most appropriate for a particular application/problem space.
- Implement (in code) various common/classic reinforcement algorithms from scratch in Python.
- List and define the various criteria for evaluating reinforcement learning algorithms (e.g., regret, sample efficiency, sample complexity, …).
- List and define the major weaknesses and complications that come with a reinforcement learning approach (e.g., sample in efficiency).

In-person

Mon: AHB 1B02, 1:30pm – 3:30pm
Wed: AHB 1B06, 2:30pm – 3:30pm

Classes begin: January 9, 2023
Fall Reading February 18 – 26, 2023
Classes end: April 10, 2023
Exam period: April 13 – 30, 2023

Table of Contents and Schedule

Week	Topics	Reading	Notes
1	Introduction and Motivation	[SutBar] Chapter 1	
2	Markov Processes and MDPs	[SutBar] Chap. 3	
3	Value Iteration and Policy Iteration	[SutBar] Sec. 4.1, 4.4, [SutBar] Sec. 4.3	
4	Introduction to RL	[SutBar] Sec. 5.1-5.3, 6.1-6.3, 6.5	
5	Deep Neural Networks		
6	Deep Q-Networks	[SutBar] Sec. 9.4, 9.7	
7	Policy Gradient	[SutBar] Sec. 13.1-13.3, 13.7	
8	Actor-Critic Method	[SutBar] Sec. 13.4-13.5	
9	Revisit – Multi-Armed Bandit	[SutBar] Sec. 2.1-2.7	
10	Model Based RL	[SutBar] Chap. 8	
11	Imitation RL	Ho, J., & Ermon, S. (2016). Generative adversarial imitation learning. In NeurIPS (pp. 4565-4573). Torabi, F., Warnell, G., & Stone, P. (2018). Behavioral cloning from observation. In IJCAI (pp. 4950-4957).	
12	Inverse RL	Ziebart, B. D., Bagnell, J. A., & Dey, A. K. (2010). Modeling interaction via the principle of maximum causal entropy. In ICML. Finn, C., Levine, S., & Abbeel, P. (2016). Guided cost learning: Deep inverse optimal control via policy optimization. In ICML (pp. 49-58).	

13	Inverse RL		
14	Project Presentations		

Contingency plan for an in-person class pivoting to 100% online learning
In the event of a COVID-19 resurgence during the course that necessitates the course delivery moving away from face-to-face interaction, affected course content will be delivered entirely online, either synchronously (i.e., at the times indicated in the timetable) or asynchronously (e.g., posted on OWL for students to view at their convenience). The grading scheme will **not** change. Any remaining assessments will also be conducted online as determined by the course instructor.

4. Course Materials

Sutton, R. S., & Barto, A. G. (2011). Reinforcement learning: An introduction.
This is the text on reinforcement learning, written by Richard Sutton and Andrew Barto from the University of Alberta and made available for free online. It is roughly split into two parts: Part 1 covers the fundamentals of reinforcement learning, namely core concepts, multi-armed bandits, dynamic programming, temporal difference (TD) learning, Q-learning, Monte-Carlo techniques, and the classic SARSA and Dyna algorithms. Part 2 delves into function approximation, eligibility traces, and policy gradient methods, while Part 3 covers the conceptual cross-field foundations of reinforcement learning with chapters on psychology, neuroscience, and classic case studies.
Additional readings will be provided through OWL as they arise.

Students are responsible for checking the course OWL site (http://owl.uwo.ca) on a regular basis for news and updates. This is the primary method by which information will be disseminated to all students in the class.

All course material will be posted to OWL: http://owl.uwo.ca.

If students need assistance with the course OWL site, they can seek support on the OWL Help page. Alternatively, they can contact the Western Technology Services Helpdesk. They can be contacted by phone at 519-661-3800 or ext. 83800.

5. Methods of Evaluation

The overall course grade will be calculated as listed below:
Assignments (4) 20%
Midterm Test 20%
Project Presentation 10%
Project 50%
NOTE: NO FINAL EXAM

6. Student Absences

If you are unable to meet a course requirement due to illness or other serious circumstances, please follow the procedures below.

Late assessments without illness self-reports will be subject to a late penalty 5%/day

- Late assessments with illness self-reports should be submitted within 24 hours of submission of the last illness self-report
- An assessment cannot be submitted after it has been returned to the class; an alternate assessment will be assigned.
- A make-up test will be offered
- If a make-up assessment is missed, the student will receive an INC and complete the task the next time the course is offered
- If permission to waive the requirement that students receive evaluation on work totaling 10 % of their final grade at least three days prior to the deadline for withdrawal without academic penalty has been obtained from the Dean's Office, a statement to this effect must be made.

For work totalling 10% or more of the final course grade, you must provide valid medical or supporting documentation to the Academic Counselling Office of your Faculty of Registration as soon as possible. For further information, please consult the University's medical illness policy at

https://www.uwo.ca/univsec/pdf/academic_policies/appeals/accommodation_medical.pdf.

The Student Medical Certificate is available at

https://www.uwo.ca/univsec/pdf/academic_policies/appeals/medicalform.pdf.

6. Accommodation and Accessibility

Religious Accommodation
When a course requirement conflicts with a religious holiday that requires an absence from the University or prohibits certain activities, students should request accommodation for their absence in writing at least two weeks prior to the holiday to the course instructor and/or the Academic Counselling office of their Faculty of Registration. Please consult University's list of recognized religious holidays (updated annually) at

https://multiculturalcalendar.com/ecal/index.php?s=c-univwo.

Accommodation Policies
Students with disabilities are encouraged to contact Accessible Education, which provides recommendations for accommodation based on medical documentation or psychological and cognitive testing. The policy on Academic Accommodation for Students with Disabilities can be found at:

https://www.uwo.ca/univsec/pdf/academic_policies/appeals/Academic Accommodation_disabilities.pdf.

7. Academic Policies

The website for Registrarial Services is http://www.registrar.uwo.ca.

In accordance with policy,

https://www.uwo.ca/univsec/pdf/policies_procedures/section1/mapp113.pdf,

the centrally administered e-mail account provided to students will be considered the individual's official university e-mail address. It is the responsibility of the account holder to ensure that e-mail received from the University at their official university address is attended to in a timely manner.

Computers are allowed to be used during the exams.

Scholastic offences are taken seriously and students are directed to read the appropriate policy, specifically, the definition of what constitutes a Scholastic Offence, at the following Web site:

http://www.uwo.ca/univsec/pdf/academic_policies/appeals/scholastic_discipline_undergrad.pdf.

All required papers may be subject to submission for textual similarity review to the commercial plagiarism detection software under license to the University for the detection of plagiarism. All papers submitted for such checking will be included as source documents in the reference database for the purpose of detecting plagiarism of papers subsequently submitted to the system. Use of the service is subject to the licensing agreement, currently between The University of Western Ontario and Turnitin.com (http://www.turnitin.com).

Tests and examinations in this course may be conducted using a remote proctoring service. By taking this course, you are consenting to the use of this software and acknowledge that you will be required to provide **personal information** (including some biometric data) and the session will be **recorded**. Completion of this course will require you to have a reliable internet connection and a device that meets the technical requirements for this service. More information about this remote proctoring service, including technical requirements, is available on Western's Remote Proctoring website at:

https://remoteproctoring.uwo.ca.

8. Support Services

Please visit the Science & Basic Medical Sciences Academic Counselling webpage for information on adding/dropping courses, academic considerations for absences, appeals, exam conflicts, and many other academic related matters: https://www.uwo.ca/sci/counselling/.

Students who are in emotional/mental distress should refer to Mental Health@Western (https://uwo.ca/health/) for a complete list of options about how to obtain help.

Western is committed to reducing incidents of gender-based and sexual violence and providing compassionate support to anyone who has gone through these traumatic events. If you have experienced sexual or gender-based violence (either recently or in the past), you will find information about support services for survivors, including emergency contacts at

https://www.uwo.ca/health/student_support/survivor_support/get-help.html.

To connect with a case manager or set up an appointment, please contact support@uwo.ca.

Please contact the course instructor if you require lecture or printed material in an alternate format or if any other arrangements can make this course more accessible to you. You may also wish to contact Accessible Education at

http://academicsupport.uwo.ca/accessible_education/index.html

if you have any questions regarding accommodations.

Learning-skills counsellors at the Student Development Centre (https://learning.uwo.ca) are ready to help you improve your learning skills. They offer presentations on strategies for improving time management, multiple-choice exam preparation/writing, textbook reading, and more. Individual support is offered throughout the Fall/Winter terms in the drop-in Learning Help Centre, and year-round through individual counselling.

Western University is committed to a thriving campus as we deliver our courses in the mixed model of both virtual and face-to-face formats. We encourage you to check out the Digital Student Experience website to manage your academics and well-being: https://www.uwo.ca/se/digital/.

Additional student-run support services are offered by the USC, https://westernusc.ca/services/.