CS394R: Reinforcement Learning: Theory and Practice -- Fall 2019

Instructors: Scott Niekum and Peter Stone
Department of Computer Science

Tuesday, Thursday 9:30-11:00am
GDC 2.216

IMPORTANT: The first reading assignment is due at 5pm on Wednesday (Aug. 28th). Please read chapter one of the textbook (available online) and email a reaction to the reading to the professors and TAs by that time. Section 1.7 is optional. Details are available on the assignments page and below. For this first week, leniency will be given if the assignment isn't done by that time, but after that please make sure you finish before class.

Jump to the assignments page.
Jump to the resources page.
Jump to the textbook page.

IMPORTANT: Due to seating limitations in the classroom, there is a hard limit on the number of students who can register. Though it is listed as "closed," there may still be room in the class. If you are interested in registering for the course, please come to the first class session. If there is physical room in the room, more seats may be opened.

Instructor Contact Information

Scott Niekum
office hours: TBA and by appointment
office: GDC 3.404
phone: 232-74741
fax: 471-8885
email: sniekum@cs.utexas.edu

Peter Stone
office hours: TBA and by appointment
office: GDC 3.508
phone: 471-9796
fax: 471-8885
email: pstone@cs.utexas.edu

Teaching Assistants

Ishan Durugkar
office hours: TBA and by appointment
office:
email: ishand@cs.utexas.edu

Wonjoon Goo
office hours: TBA and by appointment
office:
email: wonjoon@cs.utexas.edu

Bo Liu
office hours: TBA and by appointment
office:
email: bliu@cs.utexas.edu

Faraz Torabi
office hours: TBA and by appointment
office:
email: faraztrb@utexas.edu

Yifeng Zhu
office hours: TBA and by appointment
office:
email: yifengz@cs.utexas.edu

Course Description

"The idea that we learn by interacting with our environment is probably the first to occur to us when we think about the nature of learning. When an infant plays, waves its arms, or looks about, it has no explicit teacher, but it does have a direct sensori-motor connection to its environment. Exercising this connection produces a wealth of information about cause and effect, about the consequences of actions, and about what to do in order to achieve goals. Throughout our lives, such interactions are undoubtedly a major source of knowledge about our environment and ourselves. Whether we are learning to drive a car or to hold a conversation, we are all acutely aware of how our environment responds to what we do, and we seek to influence what happens through our behavior. Learning from interaction is a foundational idea underlying nearly all theories of learning and intelligence."

"Reinforcement learning problems involve learning what to do --- how to map situations to actions --- so as to maximize a numerical reward signal. In an essential way these are closed-loop problems because the learning system's actions in uence its later inputs. Moreover, the learner is not told which actions to take, as in many forms of machine learning, but instead must discover which actions yield the most reward by trying them out. In the most interesting and challenging cases, actions may affect not only the immediate reward but also the next situation and, through that, all subsequent rewards. These three characteristics --- being closed-loop in an essential way, not having direct instructions as to what actions to take, and where the consequences of actions, including reward signals, play out over extended time periods --- are the three most important distinguishing features of the reinforcement learning problem."

These two paragraphs from chapter 1 of the course textbook describe the topic of this course. The course is a graduate level class. There will be assigned readings and class discussions and activities. The exact content of

the course will be guided in part by the interests of the students. It will cover at least the first 13 chapters of the (2nd edition of the) course textbook. Beyond that, we will either continue with the text or move to more advanced and/or recent readings from the field with an aim towards focussing on the practical successes and challenges relating to reinforcement learning.

There will be at least one exam, some problem sets, and also a programming component to the course. Students will be expected to be proficient programmers.

Prerequisites

Some background in artificial intelligence and strong programming skills are recommended.

Text

The course textbook is:
Reinforcement Learning: An Introduction.
By Richard S. Sutton and Andrew G. Barto.
MIT Press, Cambridge, MA, 1998.
Note that the book is available on-line, though if you take the course, it's probably a book you'll want for your bookshelf.

Assignments

Reading, written, and programming assignments will be updated on the assignments page. A tentative schedule for the entire semester is posted. But the readings and exercises may change up until the Wednesday of the week before they are due (1 week in advance).

Resources

Slides from class and other relevant links and information are on the resources page. If you find something that should be added there, please email it to the instructors and/or TAs.

Discussion Forum

While the Professor and the TA would be glad to answer any questions you have, you would frequently find your peers to be an equally important resource in this class.

Course Requirements

Grades will be based on:

Written responses to the readings and other class participation (10%):

By 5pm on the afternoon before a class with a new reading assignment due, everyone must submit a brief question or comment about the readings **as an email in plain ascii text**. Please send it in the body of the email, rather than as an attachment. Please use the subject line "class readings for [due date]". In some cases, specific questions may be posted along with the readings. But in general, it is free form. Credit will be based on evidence that you have done the readings carefully. Acceptable responses include (but are not limited to):

- Insightful questions;
- Clarification questions about ambiguities;
- Comments about the relation of the reading to previous readings;
- Solutions to problems or exercises posed in the readings;
- Critiques;
- Thoughts on what you would like to learn about in more detail;
- Possible extensions or related studies;
- Thoughts on the paper's importance; and
- Summaries of the most important things you learned.

Example successful responses from a previous class are available on the sample responses page.

These responses will be graded on a 10-point scale with a grade of 9 being a typical full-credit grade. Responses will be due by 5pm on Monday. No late responses will be accepted.

This deadline is designed both to encourage you to do the readings before class and also to allow us to incorporate some of your responses into the class discussions.

Students are expected to be present in class having completed the readings and participate actively in the discussions and activities.

Multiple choice and short answer exercises (30%):

There will be a series of multiple choice and/or short answer questions to complete most weeks on EdX.

Programming exercises (30%):

Each student will be required to complete a series of minor programming assignments. These exercises will not involve extensive or elaborate programs. The emphasis is to be on empirically analyzing various learning algorithms and reporting on the results. They will likely be auto-graded. Details TBA.

Final Exam (30%):

There will be a final exam during the normal final exam period for the class as scheduled by the university.

Extension Policy

If you turn in your assignment late, expect points to be deducted. No exceptions will be made for the written responses to readings-based questions (subject to the ``notice about missed work due to religious holy days'' below). For other assignments, TBA.

The greater the advance notice of a need for an extension, the greater the likelihood of leniency.

Academic Dishonesty Policy

You are encouraged to discuss the readings and concepts with classmates. But all written work must be your own. And programming assignments must be your own except for 2-person teams when teams are authorized. All work ideas, quotes, and code fragments that originate from elsewhere must be cited according to standard academic practice. Students caught cheating will automatically fail the course. If in doubt, look at the departmental guidelines and/or ask.

Notice about students with disabilities

The University of Texas at Austin provides upon request appropriate academic accommodations for qualified students with disabilities. To determine if you qualify, please contact the Dean of Students at 471-6529; 471-4641 TTY. If they certify your needs, I will work with you to make appropriate arrangements.

Notice about missed work due to religious holy days

A student who misses an examination, work assignment, or other project due to the observance of a religious holy day will be given an opportunity to complete the work missed within a reasonable time after the absence, provided that he or she has properly notified the instructor. It is the policy of the University of Texas at Austin that the student must notify the instructor at least fourteen days prior to the classes scheduled on dates he or she will be absent to observe a religious holy day. For religious holy days that fall within the first two weeks of the semester, the notice should be given on the first day of the semester. The student will not be penalized for these excused absences, but the instructor may appropriately respond if the student fails to complete satisfactorily the missed assignment or examination within a reasonable time after the excused absence.

Related Courses

- Previous offering of this course (Fall 2016)
- The one before that (Spring 2013)
- The one before that (Spring 2011)
- The one before that (Fall 2007)
- The one before that (Fall 2004)
- Emma Brunskill's course at Stanford
- Michael Littman's course at Brown
- Sergey Levine's course at Berkeley
- David Silver's course at UCL
- Yishay Mansour's course at Tel Aviv University

UTCS Reinforcement Learning Reading Group

- The UTCS Reinforcement Learning Reading Group is a student run group that meets bi-weekly to discuss papers related to reinforcement learning. The RL Reading Group web page also provides a repository of past readings.
- Here's An RL reading list from Shivaram Kalyanakrishnan.
- Csaba Szepesvari's list of RL applications

[Back to Department Homepage]

Page maintained by Peter Stone
Questions? Send me mail