Course Syllabus

Jump to Today | Edit

CS 7642, Reinforcement Learning and Decision Making

Spring 2024

Course Instructors:
Miguel Morales, mimoralea@gatech.edu
Rodrigo Borela, rborelav@gatech.edu

Creators of Online Material:
Charles Isbell
Michael Littman

Head TAs:
Tim Bail, timbail@gatech.edu
Taka Hasegawa, taka@gatech.edu
Chris Serrano, cserrano7@gatech.edu

Ed Discussion:
Ed Discussion will be our official source of communication and discussion.

Office Hours:
Check Ed Discussion for weekly announcements.

General Information

Reinforcement Learning and Decision Making is a three-credit course on, well, Reinforcement Learning and Decision Making. Reinforcement Learning is a subarea of Machine Learning concerned with computational artifacts that modify and improve their performance through experience. One key distinction of Reinforcement Learning is the data used to train the model typically comes in the form of trial-and-error experiences often collected by the model itself. This course focuses on algorithms that can learn control policies programmatically, through a combination of classic papers and more recent work. It examines efficient algorithms, where they exist, for single-agent and multi-agent planning as well as approaches to learning near-optimal decisions from experience. Topics include Markov decision processes; dynamic programming methods; value-based methods; partially observable Markov decision processes; policy-based methods; stochastic and repeated games; decentralized partially observable Markov decision processes; and multi-agent methods. The class is particularly interested in issues of generalization, exploration, representation, and multi-agent systems.

Objectives

There are four primary objectives for the course:

- To provide a broad survey of approaches and techniques in RLDM
- To develop a deeper understanding of several major topics in RLDM

- To develop the design and programming skills that will help you build RLDM systems
- To develop the necessary skills to pursue research in RLDM

As you will see in the next section, we assume that you are already familiar with machine learning techniques and have some comfort with doing empirical work in machine learning. As a result, we emphasize the computational aspects of developing decision-making systems.

Prerequisites

The official prerequisite for this course is an introductory course in machine learning at the graduate level. While having taken such a course is optional, you will find that the lectures make constant callbacks to material covered in graduate machine learning courses (and the course offered by the creators of this material in particular). Of course, having said all that, the most important prerequisite for enjoying and doing well in this class is your interest in the material. We say this in every semester and every course, but it's true. In the end, it will be your motivation to understand the material that gets you through it more than anything else. If you are unsure whether this class is for you, please get in touch with the instructors.

Resources

- **Readings.** We use research paper readings, and those will be provided for you. We also use Sutton and Barto's Reinforcement Learning book (see: http://www.incompleteideas.net/book/the-book-2nd.html).
- **Computing.** You will have access to CoC clusters for your assignments, I suppose, but you won't need them. You are required to use Python for all assignments, and you can leverage many of the libraries available to you. However, you are **not allowed** to use any reinforcement learning library. All reinforcement learning related code must be your own. If in doubt, it is your responsibility to ask.
- **Web.** We will use Ed Discussion to post last-minute announcements, so check it early and often. You are responsible for keeping up with class announcements.

Statement of Academic Honesty

At this point in your academic careers, we feel it would be impolite to harp on cheating, so we won't. You are all adults and are expected to follow the university's code of academic conduct (honor code). Some of you are researchers-in-training, and we expect that you understand proper attribution to integrity of intellectual honesty.

We should also point out that "proper attribution" does not absolve the writer of the "intellectual honesty" that comes from original writing. While it is definitely the case that copying text without attribution is considered plagiarism, it is also the case that copying too much text even with attribution is a violation of our policy. In particular, more than three quotes longer than two sentences will be considered plagiarism and a terminal lack of academic originality. Do not press this issue, and we will all have fun.

Some of you have taken CS 7641 with us, so let me point out that this course is not CS 7641. Do not assume anything you read on that syllabus applies to this in any way, shape, or form. Note that unauthorized use of any previous semester course materials, such as tests, quizzes, homework, projects, videos, and any other coursework, is prohibited in this course. You are not to use code from

previous or current students. You must submit your own work. Using these materials will be considered a direct violation of academic policy and will be dealt with according to the GT Academic Honor Code.

Furthermore, we do not allow the distribution of copies of exams outside the course. Just as you are not to use the previous material, you are not to share current material with others either now or in the future. Our policy on that is strict. If you violate the policy in any shape, form, or fashion, you will be dealt with according to the GT Academic Honor Code.

Statement on the use of AI

We treat AI-based assistance, such as ChatGPT and Copilot, the same way we treat collaboration with other people: you are welcome to talk about your ideas and work with other people inside of the class, as well as with AI-based assistants.

However, all work you submit must be your own. You should never include in your assignment anything that was not written directly by you without proper citation (including quotation marks and in-line citation for direct quotes).

Including anything you did not write in your assignment without proper citation will be treated as an academic misconduct case. If you are unsure where the line is between collaborating with AI and copying AI, we recommend the following heuristics:

Heuristic 1: Never hit “Copy” within your conversation with an AI assistant. You can copy your own work into your own conversation but do not copy anything from the conversation back into your assignment.

Instead, use your interaction with the AI assistant as a learning experience, then let your assignment reflect your improved understanding.

Heuristic 2: Do not have your assignment and the AI agent open simultaneously. Like the above, use your conversation with the AI as a learning experience, then close the interaction down, open your assignment, and let your assignment reflect your revised knowledge.

This heuristic includes avoiding using AI directly integrated into your composition environment: just as you should not let a classmate write content or code directly into your submission, so you should also avoid using tools that directly add content to your submission.

Deviating from these heuristics does not automatically qualify as academic misconduct; however, following these heuristics guarantees your collaboration will not cross the line into misconduct.

Diversity and inclusion

Our course cannot be successful without appreciating the diversity of our students. In this class, we aim to create an environment where all voices are valued, respecting the diversity of gender, sexuality, age, socioeconomic status, ability, ethnicity, race, and culture. We always welcome suggestions that can help us achieve this goal.

Students with disabilities: your access to this course is extremely important to us. The institute has policies regarding disability accommodation, which are administered through the Office of Disability Services. Please request your accommodation letter as early in the semester as possible so we can arrange your approved academic accommodation.

Readings and Lectures

The online lectures are meant to summarize the readings and stress the critical points. You are expected to read any assigned material critically. Your active participation in the material, the lectures, and office hours are crucial in making the course successful. The more you put into the material, the more you will get out. The entire teaching staff is to assist you in learning and growing in this exciting growing field of reinforcement learning and decision-making.

To help you pace yourself, we have provided a nominal schedule (check the Calendar page in Canvas) that tells you when we would be covering material if we met once a week for three hours during the term. Try to keep that pace. More to the point, weekly assignments correspond to the reading material, and it will be challenging to do those without at least passing familiarity with the material.

Grading

Your final grade is divided into homework, projects, and a final exam.

- **Homework.** There will be six short homework assignments involving programming. These are designed to keep you engaged in the course and help you understand the material more intimately through hands-on experimentation. These are straightforward and concise by design. You will be provided Jupyter Notebooks and will submit your solution to Gradescope.
- **Projects.** There will be three project assignments involving programming and analysis. These are designed to help you dig deeper into the algorithmic challenges of reinforcement learning. Each of the three projects will consist of a short write-up and submission of your code (Python is required).
- **Exams.** There will be one closed-book, multiple-choice question final exam. The final exam will cover everything you learned during the semester, so keep notes of all that you're learning. They will come in handy as you prepare for the final.

Due Dates

All graded assignments are due by the time and date indicated on Canvas. We do not accept late submissions for *homework* assignments. No exceptions whatsoever. We do accept late *project* assignments for a 20-point per-day penalty, a max of 5 days, or a 0 grade. The only exceptions to late project assignment penalties will require (1) a **note** from the appropriate authority and (2) **immediate notification** of the problem when it arises. With each notification, we need a proper explanation. We are here to work with you all; please do not try to abuse the system as it will not work. Also, we only accept approved late submissions one full week after the due date, including any exceptional cases. After that week, you will automatically get a 0 for that assignment, with no change for makeup. This policy is for your own good to help you keep pace with the coursework. For cases that require longer than a week, we suggest dropping the course or asking for an incomplete semester.

Numbers

Component	
Homework (6)	30%
Projects (3)	45%
Exams (1)	25%

In the spirit of mechanism design, the grading scheme is set up so that one can't blow off any component and still earn an A. One can't blow off reading the material, that would be a bad idea anyway, one can't blow off a project either. Not that you would do any, but it's all about incentives and rewards, people.

Disclaimer

We reserve the right to modify any of these plans as needed during the class at any time; however, we won't do anything capriciously, anything we do change won't be too drastic, and you'll be informed as far in advance as possible.

Enjoy the semester, we wish you learn a lot and have fun.