

CS 7642, Reinforcement Learning and Decision Making

Spring 2021

Instructor of Record:

Charles Isbell, isbell@cc.gatech.edu
259, College of Computing Building

Creators of Online Material:

Prof. Charles Isbell, isbell@cc.gatech.edu
Prof. Michael Littman, littman@cs.brown.edu

Head TAs:

Tim Bail, timbail@gatech.edu
Miguel Morales, mimoralea@gatech.edu

Piazza:

Piazza will be our primary source of communication and discussion.

Office Hours:

Check Piazza 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, that area of Artificial Intelligence that is concerned with computational artifacts that modify and improve their performance through experience. This course focuses on automated computational decision making through a combination of classic papers and more recent work. It examines

efficient algorithms, where they exist, for single-agent and multiagent planning as well as approaches to learning near-optimal decisions from experience. Topics include Markov decision processes; stochastic and repeated games; partially observable Markov decision processes; reinforcement learning; and interactive reinforcement learning. The class is particularly interested in issues of generalization, exploration, and representation.

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 to build RLDM systems
- To develop the basic skills necessary 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 more computational aspects of developing decision-making systems. Having said that, our concern with research is expressed by having students replicate results in published papers in the area.

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 not strictly necessary, you will find that the lectures make constant call-backs 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. I say that every semester and in every course, but it's true. In the end, it will be your own motivation to understand the material that gets you through it more than anything else. If you are not sure whether this class is for you, please talk to me.

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, ask.
- **Web.** We will use Canvas Announcements and Piazza 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, I feel that it would be impolite to harp on cheating, so I won't. You are all adults, more or less, and are expected to follow the university's code of academic conduct (you know, [the honor code](#)). Furthermore, at least some of you are researchers-in-training, and I expect that you understand proper attribution and the importance of intellectual honesty.

This 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, I do not allow copies of my exams out in the ether (so there should not be any out there for you to use anyway). Just as you are not to use the previous material you are not to share current material—including lecture material—with others either now or in the future. My 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. I also have several... friends... from Texas who will help me personally deal with you. They are on retainer from my Machine Learning course and they've tasted blood.

Readings and Lectures

The online lectures are meant to summarize the readings and stress the important points. You are expected to critically read any assigned material. Your active participation in the material, the lectures, and various forums are crucial in making the course successful. This is less about my teaching than about your learning. My role is to merely assist you in the process of learning more about the area.

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

Grading

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

- **Homework.** There will be six short homework assignments involving programming. You will be provided Jupyter Notebooks and will submit your solution to Gradescope.
- **Projects.** Students will be asked to replicate results from relevant papers from the literature. 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 written, closed-book final exam scheduled for our class's final exam. Although I'm told I have a reputation for creative exams, these exams are meant to be a walk in the park if you follow and read the material.

When you upload files on Canvas, make sure that all the answers are clearly visible and the files shown are the ones you want to be graded. Upon submitting you acknowledge that you are aware that illegible or incorrect PDFs will receive 0 and you will not be able to submit for a regrade.

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: a **note** from the appropriate authority and **immediate notification** of the problem when it arises. Naturally, your excuse must be acceptable. If an alien parasite that thrives on electronic assignments gets into your computer and erases all copies of your work from existence, I will need a signed note from the relevant galactic authorities who have investigated... in English. We only accept submissions 1 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 a makeup. 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 reading the material and still earn an A. Similarly, one can't blow off a project either. Not that you would do either of those things, but it's all about incentives, people.

Disclaimer

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