CS 7642, Reinforcement Learning and Decision Making
Spring 2025

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
Christopher Serrano, cserrano7@gatech.edu

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

Office Hours:
Check Ed Discussion for weekly announcements.

Living Schedule (*subject to change*): cs7642_schedule_spring25Links to an external site.

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.** Our primary texts for the course are Sutton and Barto's *Reinforcement Learning* (see: http://www.incompleteideas.net/book/the-book-2nd.htmlLinks to an external site.) and Albrecht et al.'s *Multi-Agent Reinforcement Learning* (see: https://www.marl-book.com/Links to an external site.). Additionally, we will use a selection of research papers, which will be provided to you. As an optional resource, you might find *Grokking Deep Reinforcement Learning* useful, a book written by one of your instructors (see: https://www.manning.com/books/grokking-deep-reinforcement-learningLinks to an external site.). This is purely supplementary and not required for the course.
- **Computing.** You will have access to CoC clusters for your projects, but you will unlikely 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 codeLinks to an external site.). 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.

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, assignments correspond to the reading material, and it will be challenging to do those without at least passing familiarity with the material.

Assignments and assessments

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

- **Projects.** There will be four project assignments involving programming and analysis. These are designed to help you dig deep into the algorithmic challenges of reinforcement learning and develop the knowledge to apply it

to real-world scenarios. Each of the four projects will consist of a write-up and submission of your code (Python is required).

- **Quizzes.** At the end of each major section of the course, you will complete a quiz administered through Socratic Mind, an AI-powered oral assessment tool. These four quizzes are designed to evaluate your understanding of reinforcement learning concepts and your ability to apply them in various contexts. They aim to help you stay on track, gauge your grasp of the material, and prepare effectively for the final exam.
- **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 projects are due by the time and date indicated on Canvas. We 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) **immediate notification** of the problem when it arises (2) a **letter from the Office of the Dean of Students**. Please contact the Dean of Students with the appropriate documentation, such as a doctor's note for an incapacitating illness or family emergency. Documentation must be provided on letterhead with the signature of a physician, supervisor, or another appropriate official to the Dean of Students. Please do not send this documentation through me. Fill out the form you will find at https://studentlife.gatech.edu/request-assistanceLinks to an external site.. 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.

Grading

Component	Weight
Projects (4)	68%
Quizzes (4)	4%
Exams (1)	28%

A priori, the course will initially follow traditional grading scale:

Letter grade	Weighted points cutoffs

A	>=90
B	>= 80 and < 90
C	>= 70 and < 80
D	>= 60 and < 70
F	< 60

However, final grades will be curved based on a statistical analysis of the class's overall performance. This approach is designed to benefit students and will never result in a grade worse than what would be assigned using the traditional grading scale. For example, achieving a grade of 90 or higher guarantees an "A," and similarly, cutoffs for other grades (e.g., 80 for a "B") serve as the baseline minimums. The curve may lower these thresholds based on the class median and standard deviations to reflect the distribution of grades.

After the curve has been applied and the thresholds have been potentially lowered, note that you will only achieve a passing grade if you have attempted every assignment, quiz, and exam, and achieved a non-zero score.

Extra credit opportunities will also be available throughout the semester, and these points will be added after curving is complete. This ensures that extra credit enhances your final grade without influencing the curve itself.

You are welcome here!

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, nationality, 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.

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.

Course Summary:

Date	Details	Due
Mon Jan 6, 2025	Calendar Event Lectures and Readings Week 1	12am
Thu Jan 9, 2025	Calendar Event [CS7642] Open Office Hours	7pm to 8pm
Mon Jan 13, 2025	Calendar Event Lectures and Readings Week 2	12am
	Assignment Setup	due by 7am
Thu Jan 16, 2025	Calendar Event [CS7642] Open Office Hours	7pm to 8pm
Mon Jan 20, 2025	Calendar Event Lectures and Readings Week 3	12am
Thu Jan 23, 2025	Calendar Event [CS7642] Open Office Hours	7pm to 8pm
Mon Jan 27, 2025	Calendar Event Lectures and Readings Week 4	12am
Thu Jan 30, 2025	Calendar Event [CS7642] Open Office Hours	7pm to 8pm
Mon Feb 3, 2025	Calendar Event Lectures and Readings Week 5	12am
	Assignment Project 1	due by 7am
Thu Feb 6, 2025	Calendar Event [CS7642] Open Office Hours	7pm to 8pm
Mon Feb 10, 2025	Calendar Event Lectures and Readings Week 6	12am
	Assignment Q1: Reinforcement Learning Foundations	due by 7am
Thu Feb 13, 2025	Calendar Event [CS7642] Open Office Hours	7pm to 8pm
Mon Feb 17, 2025	Calendar Event Lectures and Readings Week 7	12am

Date	Details	Due
Thu Feb 20, 2025	Calendar Event [CS7642] Open Office Hours	7pm to 8pm
Mon Feb 24, 2025	Calendar Event Lectures and Readings Week 8	12am
Thu Feb 27, 2025	Calendar Event [CS7642] Open Office Hours	7pm to 8pm
Mon Mar 3, 2025	Calendar Event Lectures and Readings Week 9	12am
	Assignment Project 2	due by 7am
Thu Mar 6, 2025	Calendar Event [CS7642] Open Office Hours	7pm to 8pm
Mon Mar 10, 2025	Calendar Event Lectures and Readings Week 10	12am
	Assignment Q2: Deep Reinforcement Learning	due by 7am
Thu Mar 13, 2025	Calendar Event [CS7642] Open Office Hours	7pm to 8pm
Mon Mar 17, 2025	Calendar Event Lectures and Readings Week 11	12am
Thu Mar 20, 2025	Calendar Event [CS7642] Open Office Hours	7pm to 8pm
Mon Mar 24, 2025	Calendar Event Lectures and Readings Week 12	12am
Thu Mar 27, 2025	Calendar Event [CS7642] Open Office Hours	7pm to 8pm
Mon Mar 31, 2025	Calendar Event Lectures and Readings Week 13	12am
	Assignment Project 3	due by 7am
Thu Apr 3, 2025	Calendar Event [CS7642] Open Office Hours	7pm to 8pm

Date	Details	Due
Mon Apr 7, 2025	Calendar Event Lectures and Readings Week 14	12am
	Assignment Q3: Multi-Agent Reinforcement Learning	due by 7am
Thu Apr 10, 2025	Calendar Event [CS7642] Open Office Hours	7pm to 8pm
Mon Apr 14, 2025	Calendar Event Lectures and Readings Week 15	12am
Thu Apr 17, 2025	Calendar Event [CS7642] Open Office Hours	7pm to 8pm
Mon Apr 21, 2025	Calendar Event Lectures and Readings Week 16	12am
	Assignment Project 4	due by 7am
	Assignment Q4: Advanced Topics	due by 7am
Thu May 1, 2025	Quiz Final Exam	due by 8am