

Course Syllabus



CS 7642 Reinforcement Learning

Fall 2025

Course Instructors:

Miguel Morales, mimoralea@gatech.edu (<mailto:mimoralea@gatech.edu>)

Rodrigo Borela, rborelav@gatech.edu (<mailto:rborelav@gatech.edu>)

Creators of Online Material:

Charles Isbell

Michael Littman

Miguel Morales

Rodrigo Borela

Head TAs:

Tim Bail, timbail@gatech.edu (<mailto:timbail@gatech.edu>)

Taka Hasegawa, taka@gatech.edu (<mailto:taka@gatech.edu>)

Christopher Serrano, cserrano7@gatech.edu (<mailto: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_fall25.xlsx](#)  (https://gtvault-my.sharepoint.com/:x/g/personal/rvalente6_gatech_edu/EXhg6bl8cv1Jv0LOsX7rASMBf875gQiyZxa42tomQkGvYw?e=5g2sML).

General Information

Reinforcement Learning (RL) 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

Upon successful completion of this course, you should be able to design, implement, and critically evaluate reinforcement learning systems, integrating theoretical principles with practical experiments across single-agent, multi-agent, and domain-specific contexts.

The specific learning objectives are as follows:

- Formulate RL problems in terms of agents, environments, states, actions, and rewards, and determine appropriate solution strategies.
- Implement core reinforcement learning algorithms and apply function approximation methods to solve control problems in various environments.
- Design and evaluate algorithms for multi-agent reinforcement learning settings, addressing cooperation and competition.
- Integrate reinforcement learning methods into domain-specific applications such as autonomous systems.

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 original creators of this material in particular). 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.html>  (<http://www.incompleteideas.net/book/the-book-2nd.html>)) and Albrecht et al.'s *Multi-Agent Reinforcement Learning* (see: <https://www.marl-book.com/>  (<https://www.marl-book.com/>)). 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-learning> ➡ <https://www.manning.com/books/grokking-deep-reinforcement-learning>)). This is purely supplementary and not required for the course.

- **Online forum.** We will use Ed Discussion to post announcements and clarifications, so make sure to check it frequently. It will also serve as a space for you to discuss the content, brainstorm solutions when you get stuck and receive feedback from the instruction team.
- **Computing.** You will have access to CoC clusters for your projects, but you will likely only need them for the final project. You are required to use Python for all projects, and you can leverage many of the libraries available to you for visualization, data analysis, function approximation (neural networks). However, you are **NOT ALLOWED** to use any single- or multi-agent reinforcement learning library. All reinforcement learning related code must be written by you in its entirety. If in doubt, it is your responsibility to ask.

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](https://osi.gatech.edu/index.php/students/honor-code) ➡ <https://osi.gatech.edu/index.php/students/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.

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](https://osi.gatech.edu/index.php/students/honor-code) ➡ <https://osi.gatech.edu/index.php/students/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](https://osi.gatech.edu/index.php/students/honor-code) ➡ <https://osi.gatech.edu/index.php/students/honor-code>).

Academic dishonesty, including plagiarism, unauthorized collaboration, use of prohibited resources, falsification, or misrepresentation, will not be tolerated. Any suspected violations will be reported to the [Office of Student Integrity](https://osi.gatech.edu/index.php/students/honor-code) ➡ <https://osi.gatech.edu/index.php/students/honor-code>). Maintaining integrity in your academic work is essential to your success at Georgia Tech and in your professional career.

Collaboration with AI

We treat AI-based assistance (e.g. ChatGPT, 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.

With regards to code writing: All code you submit must be written entirely by you, except for any template code provided for the project. Do not include in your repository any code you did not write yourself. Implementing algorithms from scratch is essential to developing a deep understanding of their mechanics and building the skills needed for you to innovate in the future.

With regards to report writing: You may use AI tools to check or improve spelling, grammar, or sentence phrasing, but you must declare this use in your report in the "Acknowledgements" section. **The use of AI to generate multi-sentence portions of text is not permitted.** All ideas, explanations, analyses and conclusions must be entirely your own. If the instructional team suspects a violation, you may be asked to provide evidence of authorship (e.g., Overleaf edit history). Failure to provide convincing evidence will result in a zero for the project and a report to the Office of Student Integrity.

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 instructional team is here to assist you in learning and preparing to participate in the growing field of reinforcement learning.

To help you pace yourself, we have provided a nominal schedule (check the living schedule above) 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, projects and quizzes 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 challenges of reinforcement learning problems and develop the knowledge to apply them to real-world scenarios. Each of the projects will consist of a write-up and submission of your code (Python is required).
- **Quizzes.** Throughout each major section of the course, you will complete quizzes administered through Socratic Mind, an AI-powered oral assessment tool. These quizzes are designed to engage you in discussions about the course concepts their applicability to 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-assistance> (<https://studentlife.gatech.edu/request-assistance>). 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.

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.

Grading

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

A priori, the course will initially follow a 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 receive a passing grade if you have attempted and submitted every project and exam, and earned a non-zero score on each.** Missing any required submission (e.g., not turning in a project) will result in a failing grade, regardless of the your overall average.

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

Regrade Policy

Once project grades are released, there will be a 24-hour review period during which regrade requests will not be accepted. Use this time to carefully review grader comments. Regrade requests will only be considered if there is a clear mismatch between the assignment expectations and the grading; for example, if your report presents results/analyses/information that were overlooked. **Requests must cite specific passages from your submission;** we cannot consider unstated knowledge. We will not revisit deductions that are based on valid issues noted by the grader, as all grading follows the same rubric, consistently applied to all students. Please note that if a re-evaluation results in a lower grade, that grade will stand.

Communication and Student-Faculty Expectations

We strive for a professional and respectful learning environment by following the [Georgia Tech Student-Faculty Expectations](https://catalog.gatech.edu/rules/21/)  (<https://catalog.gatech.edu/rules/21/>). Professional behavior is expected at all times, both in-person and online. This includes clear, courteous, and constructive communication with peers, teaching assistants, and instructors in all contexts, including email, forums, meetings, and class discussions.

A positive classroom community depends on respectful dialogue and a commitment to mutual responsibility between students and faculty. Inappropriate tone, hostile language, or disrespect directed toward instructional staff or fellow students, particularly in matters related to assignments, will not be tolerated. Any student who communicates disrespectfully about an assignment may receive a zero on that assignment, regardless of its point value or context.

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.