# CS7642 Syllabus

## Course Information

**Course Prefix and Number: CS7642**
**Course Name: Reinforcement Learning**
**Instructor: Borela Valente, Rodrigo**

## Course Description

*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.

## Course Learning Outcomes

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.

## Required Course Materials

Our primary texts for the course are Sutton and Barto's Reinforcement Learning (see: https://www.incompleteideas.net/book/the-book-2nd.html) and Albrecht et al.'s Multi-Agent Reinforcement Learning (see: https://www.marl-book.com). Additionally, we will use a selection of research papers, which will be provided to you.

## Grading Policy

Four individual projects account for 68% of the overall course grade. Four quizzes account for 4% of the overall course grade. One final exam accounts for 28% of the overall course grade.

The class is graded on a letter scale with:

- The cutoff for an A will be at most 90%.
- The cutoff for a B will be at most 80%.
- The cutoff for a C will be at most 70%.
- The cutoff for a D will be at most 60%.
- Anything below corresponds to not passing the class.

Overall course numerical grades will be curved based on a statistical analysis of the class's 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 described above. 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 the final exam, and earned a non-zero score on each.** Missing any required project submission or the final exam will result in a failing grade, regardless of the your overall average.

## Description of Graded Components

- **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 short quizzes. 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.

## Attendance Policy

Students enrolled into this online class will learn individually and in their own pace (yet having to meet set deadlines for assignments, quizzes, and exams). As such, no formal attendance policy for class session applies.

## Academic and Research Honesty/Integrity Statement

Georgia Tech aims to cultivate a community based on trust, academic integrity, and honor. Students are expected to act according to the highest ethical standards. Review the Georgia Tech's Honor Code and the Student Code of Conduct.

Any student suspected of cheating or plagiarism on a quiz, exam, or project will be reported to the Office of Student Integrity, who will investigate the incident and identify the appropriate penalty for violations.

## Artificial Intelligence Use Policy

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.

## Core IMPACTS

Not applicable.

## Accommodations for Students with Disabilities

If you are a student with learning needs that require special accommodation, as soon as possible to make an appointment to discuss your special needs and to obtain an accommodations letter. Please also e-mail me as soon as possible in order to set up a time to discuss your learning needs contact the Office of Disability Services

## Student-Faculty Expectations Agreement

At Georgia Tech, we believe that it is important to strive for an atmosphere of mutual respect, acknowledgement, and responsibility between faculty members and the student body. This summarizes my expectations for you and what you can expect from me. In the end, simple respect for knowledge, hard work, and cordial interactions will help build the environment we seek. Therefore, I encourage you to remain committed to the ideals of Georgia Tech while in this class. Expectations of Advisors and Advisees