## Course prefix, number, and name

CS 8803 - Special Topics Seminar: Deep Reinforcement Learning

## Semester and academic year

Fall 2025

## Course description

Intelligent decision making in interactive and sequential environments is the hall mark of intelligence. These problems appear in varying domains from foundation mod els, animation, robotics, power systems, and control. It is imperative that these systems are built on intelligent and adaptive algorithms. Learning by interaction through reinforcement offers a natural mechanism to postulate these problems.
This graduate-level seminar course will cover topics and new research frontiers in reinforcement learning (RL). Planned topics include model-based and model-free RL, policy search, Monte Carlo Tree Search, off-policy evaluation, temporal abstractio n and hierarchical approaches, inverse reinforcement learning, and imitation learn ing.

## Course objectives / approved course learning outcomes

At the end of this course, students will be able to:
1. Acquire familiarity with the state of the art in reinforcement learning.
2. Articulate limitations of current work, identify open frontiers, and scope rese arch projects.
3. Constructively critique research papers and deliver a tutorial-style presentati on.
4. Work on a research-based project, implement and evaluate experimental results, and discuss future work in a project paper.

## Required course materials

(e.g., textbooks with ISBNs, articles, lab supplies)
There is no official textbook for the class.
Supporting readings may include:
· Reinforcement Learning: An Introduction, Sutton and Barto, 2nd Edition.
· Reinforcement Learning: State-of-the-Art, Marco Wiering and Martijn van Otterlo, eds.
· Artificial Intelligence: A Modern Approach, Stuart J. Russell and Peter Norvig.
· Deep Learning, Ian Goodfellow, Yoshua Bengio, and Aaron Courville.

## Grading policy and weighting

Homework assignments: 60%
Paper presentation and implementation: 30%
Presentation: 10%

Write-up with critique and follow-ups: 5%
Hands-on implementation and notebook: 15%
Quizzes and participation: 10%

## Attendance policy

Regular attendance and active participation are expected. In-class quizzes and activities cannot be made up except in cases of documented excused absences.

## Additional criteria for successful completion of the course

A total of 4 late days are allowed for homeworks. Any late submission will be counted as a minimum of one late day.

## Academic honesty / academic integrity statement

Academic dishonesty will not be tolerated. This includes cheating, lying about course matters, plagiarism, or helping others commit a violation of the Honor Code. Plagiarism includes reproducing the words of others without both the use of quotation marks and citations. Students are reminded of the obligations and expectations associated with the Georgia Tech Academic Honor Code and Student Code of Conduct.

Students are expected to implement the core components of each project on their own, but extra credit opportunities may build on third-party datasets or code. That is acceptable as long as the submitted work clearly cites the third-party source and makes clear what is not the student's own work.
Students should not view or edit anyone else's code. Students should not post code to Ed Discussion, except for starter code or helper code that is not related to the core project. Please refer to GT Honor Code for details.

## Statement about acceptable student conduct

The Georgia Tech community believes that it is important to continually strive for an atmosphere of mutual respect, acknowledgement, and responsibility between faculty members and the student body. Please read GT Student faculty expectations for details.

## Statement about services offered through the Office of Disability Services

If needed, accommodations will be made for students with documented disabilities. These accommodations must be arranged in advance and in accordance with the Office of Disability Services policies. Please read more here.