CS 8801, Fall 2026

Special Topics Seminar: Deep Reinforcement Learning and Humanoid Robotics

Lecture Hours: Tuesdays, 2:00–3:00 PM Eastern Time (Online Synchronous)

Credit Hours: 1 Credit Hour

Zoom Link: https://gatech.zoom.us/j/94117821133?pwd=5ynM5VMwVCqIsxTb6MiUDWfAyC7Jb9.1

Password: 20260519

All lectures will be conducted online synchronously. Lectures will be recorded and shared for students who cannot attend live sessions.

Instructor:

- Donghoon Baek (dbaek36@gatech.edu)

Contact:

Please contact Donghoon Baek by email for all course-related questions and communication.

Logistics:

This seminar will be graded on a Pass/Fail basis.

Description:

This seminar focuses on deep reinforcement learning (DRL) and its applications to humanoid robotics. We begin with an introduction to RL foundations and a broad overview of learning methods, extending these concepts to deep neural network-based policy learning. The curriculum covers essential approaches including policy gradients, value-based methods, and actor-critic algorithms. We will also delve into advanced techniques such as TRPO, PPO, and Deep RL with Q-functions, establishing a strong theoretical toolkit for continuous robot control.

Building on this foundation, the seminar explores how DRL enables complex humanoid behaviors like locomotion, whole-body coordination, and dexterous manipulation. We will address practical challenges such as sim-to-real transfer and cross-embodiment generalization, alongside complementary paradigms like imitation learning, diffusion policies, and vision-language-action (VLA) models. Ultimately, the course provides a comprehensive understanding of achieving robust, adaptive, and scalable humanoid robot intelligence.

Learning Objectives:

By the end of this course, students will be able to:

- Develop a strong understanding of the state of the art in reinforcement learning.
- Identify limitations of existing methods, recognize open research problems, and formulate meaningful research directions.

- Critically evaluate research papers and communicate their ideas.

Grading & Evaluation:

This seminar will be graded on a Pass/Fail basis, primarily through attendance and participation. (We will NOT strictly check attendance, but we strongly encourage you to attend.)

- Assignments: None
- Quizzes and Exams: None

Textbooks & Resources:

There is no official textbook for the class.

Primary Reference:

- Reinforcement Learning: An Introduction, Sutton and Barto, 2nd Edition. (free PDF)

Additional References:

- Reinforcement Learning: State-of-the-Art, Marco Wiering & Martijn van Otterlo, Eds. (link)
- Artificial Intelligence: A Modern Approach, Stuart J. Russell and Peter Norvig. (link)
- Deep Learning, Ian Goodfellow, Yoshua Bengio, and Aaron Courville. (link)

Additional Resources from Similar Courses:

- RL Course from UofT – Animesh Garg (link)
- RL Course from UW – Byron Boots (link)
- RL Course from Stanford – Emma Brunskill (link)
- RL Course from University of Alberta – Martha White (link)
- RL Course at ASU/MIT – Dimitry Bertsekas (link)
- David Silver's Course on Reinforcement Learning (link)
- Deep RL Course from Berkeley – Sergey Levine (link)

Prerequisites:

You need to be comfortable with:

- Introductory machine learning concepts (e.g., CS7641 or equivalent)
- Linear algebra
- Basic multivariable calculus
- Intro to probability

Note: If you don't meet all the prerequisites, please contact the instructor by email.

Optional but recommended: Experience with neural networks (e.g., CS 4644/7643 or equivalent), introductory-level familiarity with reinforcement learning and control.

Academic Integrity:

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.

Learning Accommodations:

If needed, we will make accommodations for students with documented disabilities. These accommodations must be arranged in advance and in accordance with the Office of Disability Services policies.

Course Summary:

Week	Topic	Instructor
Week 1	Introduction to Reinforcement Learning	Donghoon Baek
Week 2	RL Methods Overview	Donghoon Baek
Week 3	Policy Gradient Methods	Donghoon Baek
Week 4	Actor-Critic Algorithms	Donghoon Baek
Week 5	Value Based Methods	Donghoon Baek
Week 6	Deep RL with Q-Function	Donghoon Baek
Week 7	Advanced Policy Gradients	Donghoon Baek
Week 8	Model Based Reinforcement Learning	Donghoon Baek
Week 9	Learning from Data	Donghoon Baek
Week 10	Offline RL	Donghoon Baek
Week 11	Humanoid and RL	Donghoon Baek
Week 12	RL for Whole-body Humanoid Control	Donghoon Baek
Week 13	Sim-to-Real	Donghoon Baek
Week 14	Vision-Language-Action Models	Donghoon Baek
Week 15	Research Trends in Humanoid Robotics	Donghoon Baek