

## Course Syllabus

[Edit](#)

# CS 8801, Summer 2026

## Special Topics Seminar: Deep Reinforcement Learning and Humanoid Robotics

Lecture Hours: 1 hour 15 minutes weekly (Online Synchronous)

### Class Sessions:

1 Credit Hour

All lectures will be conducted online synchronously.

Lectures will be recorded and shared for students who cannot attend live sessions.

### Instructor:

- Sehoon Ha (sehoonha@gatech.edu)
- 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:

1. Develop a strong understanding of the state of the art in reinforcement learning.
2. Identify limitations of existing methods, recognize open research problems, and formulate meaningful research directions.
3. 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.

Assignments: None

Quizzes and Exams: None

## Textbooks & Resources:

There is no official textbook for the class.

A number of the supporting readings will come from:

1. Reinforcement Learning: An Introduction, Sutton and Barto, 2nd Edition.

This is available for free [here](http://incompleteideas.net/book/the-book-2nd.html) and references will refer to the final pdf version available [here](http://incompleteideas.net/book/RLbook2018.pdf).

Some other additional references that may be useful are listed below:

1. Reinforcement Learning: State-of-the-Art, Marco Wiering & Martijn van Otterlo, Eds. [\[link\]](https://link.springer.com/book/10.1007%2F978-3-642-27645-3)
2. Artificial Intelligence: A Modern Approach, Stuart J. Russell and Peter Norvig. [\[link\]](http://aima.cs.berkeley.edu/)
3. Deep Learning, Ian Goodfellow, Yoshua Bengio, and Aaron Courville. [\[link\]](http://www.deeplearningbook.org/)

Additional Resources from similar courses.

1. RL Course from UofT - Animesh Garg [\[link\]](https://www.pair.toronto.edu/csc2621-w20/)
2. RL Course from UW - Byron Boots [\[link\]](https://homes.cs.washington.edu/~bboots/RL-Fall2020/)

3. RL Course from Stanford - Emma Brunskill [[link ↗](http://web.stanford.edu/class/cs234/index.html) (<http://web.stanford.edu/class/cs234/index.html>)]
4. RL Course from University of Alberta - Martha White [[link ↗](https://marthawhite.github.io/rlcourse/schedule.html) (<https://marthawhite.github.io/rlcourse/schedule.html>)]
5. RL course at ASU/MIT, Dimitry Bertsekas -- [[Link ↗](http://web.mit.edu/dimitrib/www/RLbook.html) (<http://web.mit.edu/dimitrib/www/RLbook.html>)]
6. David Silver's course on Reinforcement Learning [[link ↗](http://www0.cs.ucl.ac.uk/staff/D.Silver/web/Teaching.html) (<http://www0.cs.ucl.ac.uk/staff/D.Silver/web/Teaching.html>)]
7. Deep RL Course from Berkeley - Sergey Levine [[link ↗](http://rail.eecs.berkeley.edu/deeprlcourse/resources/#courses) (<http://rail.eecs.berkeley.edu/deeprlcourse/resources/#courses>)]

## Prerequisites:

You need to be comfortable with: introductory machine learning concepts (such as from CS7641 or equivalent), linear algebra, basic multivariable calculus, intro to probability. Note: if you don't meet all the prerequisites above please contact the instructor by email. Optional, but recommended: experience with neural networks, such as from 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 and Student](#)

## 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](https://disabilityservices.gatech.edu/) policies.

## Course Summary:

Date	Details	Instructor
Week 1	Introduction to Reinforcement Learning	Sehoon Ha and Donghoon Baek
Week 2	RL Methods overview	Sehoon Ha
Week 3	Policy Gradient Methods	Sehoon Ha
Week 4	Actor-Critic Algorithms	Sehoon Ha
Week 5	Value Based Methods	Sehoon Ha
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	Character Animation	Donghoon Baek

Week 13  
Week 14  
Week 15

Sim-to-real  
Vision-Language-Action Models  
Research Trends in Humanoid Robotics

Donghoon Baek  
Donghoon Baek  
Donghoon Baek

---