CS 669 - Reinforcement Learning Syllabus

Course Modality:
This is an online course, which will be conducted fully online, asynchronously via Canvas. For more information on using Canvas and other supported learning tools, visit the IST Service Desk Knowledgebase.

Course Workload: This course values your time and effort and aims to provide a rewarding learning experience. You can expect to dedicate approximately **7.5 hours** to the course per week. This estimate includes, but is not limited to, time spent on readings, watching course videos, completing assignments, participating in discussions, and reviewing feedback.

Instructor Information

Instructor	Email	Office Hours
Jing Li	jingli@njit.edu	Available by appointment by Zoom.

General Information

Course Description

This course covers current topics, key concepts, and classic and modern algorithms in reinforcement learning and contains both theory and applications. The topics include but are not limited to, Markov Decision Processes, exploration and exploitation, planning, value-based learning, policy gradient, etc. Students will present recent papers in reinforcement learning, and work on written and programming assignments and do a reinforcement learning project. After completing this course, students will be able to start using reinforcement learning for real world problems that can be specified as Markov Decision Processes.

Extended Course Description

Reinforcement learning (RL) is a rapidly growing area, with applications across fields like autonomous driving, robotics, optimization, psychology, and neuroscience. Unlike other machine learning paradigms, RL emphasizes learning through direct interaction with the

environment, enabling agents to learn from trial and error without relying on supervision. This makes RL one of the closest forms of machine learning to human learning, and it has become a major research focus in artificial intelligence and neural networks. The field's practical relevance is underscored by its impressive industry applications.

In this course, you will explore both the theoretical foundations and practical applications of reinforcement learning. Topics will include essential concepts like Markov Decision Processes (MDPs), value-based learning, Monte Carlo and Temporal Difference methods, policy gradients, and more. Through a combination of lectures, paper presentations, programming assignments, and a hands-on project, you will gain the skills to apply RL to real-world problems that can be framed as MDPs. To get started, familiarize yourself with the course syllabus and plan your approach to the upcoming assignments and project work.

Prerequisites/Co-requisites

Students should have learned linear algebra, basic probability, basic calculus, computer programming, OR approval of the instructor. Experience with machine learning, artificial intelligence, or deep learning (e.g., CS 675, CS 670, CS 677) is recommended.

Course Learning Outcomes

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

1. Identify and define key concepts such as agents, environments, states, actions, rewards, and policies.
2. Differentiate and compare related concepts in RL, such as planning vs. learning, exploration vs. exploitation, model-based vs. model-free approaches, and different reinforcement learning paradigms vs. other types of machine learning.
3. Analyze and apply the Markov Decision Process (MDP) framework and Bellman equation to real-world problems.
4. Implement and demonstrate the basic concepts of dynamic programming techniques, Monte-Carlo Methods, and Temporal-Difference learning.
5. Evaluate the need for function approximation and analyze key features of advanced RL algorithms such as Deep Q-Networks (DQN), Policy Gradients, Actor-Critic methods.
6. Construct and implement in code classical RL algorithms using standard API and environments in RL libraries.
7. Analyze and critique the theoretical and empirical performance of different RL solutions and assess the trade-offs between their computational efficiency and solution optimality.
8. Design and execute experiments to systematically investigate the effects of varying parameter settings on the performance of RL algorithms in real-world applications.
9. Synthesize and evaluate findings the findings from experiments to develop insights and recommendations for tuning various parameter settings of different RL approaches.

Required Materials

Reinforcement Learning: An Introduction by Rich Sutton and Andrew Barto (open educational resource)

Course Work

Quizzes: (10 points) There will be frequent multiple-choice quizzes throughout the course. They are meant to help you practice course concepts.

Discussions: (10 points) You are expected to participate in the discussion forums on Canvas. Active participation in these discussions fosters an engaging learning environment that enhances your understanding of the material, allows you to learn from your peers, and contributes to your success in the class. Please post your response to the prompt and/or reply to your classmates depending on the instructions by Sunday at 11:59 pm of the module they are listed.

Programming Assignments: (45 points) Programming assignments will be given to give you an opportunity to apply the learned reinforcement learning algorithms on different gaming environments. These assignments are designed to give you hands-on practice and empirical experience of how different algorithms work and how hyperparameters affect the performance. They will also prepare you for the projects.

Projects: (**35 points**) There will be projects with regular milestones, including presentations and written reports. You can work on the course project in a team of up to 3 members. You will have opportunities to iterate and revise your work based on peer and instructor feedback. You will present the project progress to the class in the project proposal, midterm demo, project paper presentation, and final project demo. The final project outcome will be in the form of a final report, describing the design, analysis, implementation and experimentation efforts, together with the code submission. The report must acknowledge and document in detail all contributions that anyone made to the work.

- Project Milestone 0: Topic Exploration
- Project Milestone 1.1: Project Proposal (Presentation Slides Submission)
- Project Milestone 1.2: Project Proposal (Presentation Video and Discussion)
- Project Milestone 2.1: Project Midterm Demo (Presentation Slides Submission)
- Project Milestone 2.2: Project Midterm Demo (Presentation Video and Discussion)
- Project Milestone 3.1: Project Paper Presentation (Presentation Slides Submission)
- Project Milestone 3.2: Project Paper Presentation (Presentation Video and Discussion)
- Project Milestone 4.1: Project Final Demo (Presentation Slides Submission)
- Project Milestone 4.2: Project Final Demo (Presentation Video and Discussion)
- Project Milestone 5: Final Project Report

Feedback

Feedback will be provided on each assignment and project milestone using the comments feature in Canvas. I will typically respond to direct communications, such as email, within 48 hours. Allow up to 2 weeks for feedback on submitted assignments.

Grading Policy

NJIT Grading Legend

Final Grade Calculation

Final grades for all assignments will be based on the following percentages:

Category	Total Amount of Points
Quizzes (10 quizzes = 1 point each)	**10**
Discussion Forums (5 forums)	**10**
Programming Assignments (Assignment 0 = 3, Assignment 1 = 12, Assignment 2 = 15, Assignment 3 = 15)	**45**
Projects (Milestone 0 = 2, Milestone 1 = 5, Milestone 2 = 5, Milestone 3 = 8, Milestone 4 = 5, Milestone 5 = 10)	**35**
Course total:	**100**

Letter to Number Grade Conversions

Letter	A	B+	B	C+	C	F
Number	91-100	86-90	76-85	71-75	61-70	0-60

Exam Information and Policies

This course does not have any exams. Per the NJIT Online Course Exam Proctoring Policy, this course will use authentic assessment, meaning you will be assessed and graded on your ability to deliver real-world outputs as well as your participation and feedback to other students.

Policy for Late Work

For programming assignments and project milestones, solutions submitted up to 72 hours late will be given a 20% penalty; solutions submitted less than a week late will be given a 50% penalty; solutions submitted after a week late will not be given credit. Late submission for quizzes and discussions will not be allowed and will receive no grade.

Academic Integrity

"Academic Integrity is the cornerstone of higher education and is central to the ideals of this course and the university. Cheating is strictly prohibited and devalues the degree that you are working on. As a member of the NJIT community, it is your responsibility to protect your educational investment by knowing and following the NJIT academic code of integrity policy.

Please note that it is my professional obligation and responsibility to report any academic misconduct to the Dean of Students Office. Any student found in violation of the code by cheating, plagiarizing or using any online software inappropriately will result in disciplinary action. This may include a failing grade of F, and/or suspension or dismissal from the university. If you have any questions about the code of Academic Integrity, please contact the Dean of Students Office at dos@njit.edu"

Generative AI

The usage of generative artificial intelligence (AI) tools is permitted in this course. If and when you use artificial intelligence in this course, you must clearly document (e.g., via submitting a screenshot of your interaction with the generative AI tool) how you used it, what answers were provided by the tool, how you integrated the provided answers into your solutions, and why.

Netiquette

Throughout this course, you are expected to be courteous and respectful to classmates by being polite, active participants. You should respond to discussion forum assignments in a timely manner so that your classmates have adequate time to respond to your posts. Please respect opinions, even those that differ from your own, and avoid using profanity or offensive language.

Weekly Expectations

This course is organized by weekly modules. Each week, students must watch lecture videos, complete a quiz, programming assignment, and/or project milestone.

Course Schedule

Week	Topic	Reading	Assignment	Due Dates
1	Introduction to Reinforcement Learning	*Reinforcement Learning: An Introduction (RL)* Section 1 *Reinforcement Learning with LLM*:	1. Module 1 Quiz: RL Introduction 2. Discussion: Introduce Yourself!	1. Due by Sunday 11:59pm. 2. Post your introduction by Thursday 11:59PM. Post your responses by Sunday 11:59PM.

Week	Topic	Reading	Assignment	Due Dates
		Introducing ChatGPT		
2	Markov Decision Process	*RL* Sections 3.1 to 3.4	1. Module 2 Quiz: MDP	1. Due by Sunday 11:59pm.
3	Value Functions and Bellman Equations	*RL* Sections 3.3 to 3.5	1. Module 3 Quiz: Bellman Equation 2. Project Milestone 0: Topic Exploration	1. Due by Sunday 11:59pm. 2. Due by Sunday 11:59pm.
4	Policy Evaluation	*RL* Sections 3.6 to 4.1	1. Programming Assignment 0: Warm Up	1. Due by Sunday 11:59pm.
5	Value Iteration and Policy Iteration	*RL* Sections 4.2 to 4.8	1. Module 5 Quiz: Value Iteration and Policy Iteration 2. Project Milestone 1.1: Project Proposal (Presentation Slides Submission) 3. Project Milestone 1.2: Project Proposal (Presentation Video and Discussion)	1. Due by Sunday 11:59pm. 2. Due by Sunday 11:59pm. 3. Due by Sunday 11:59pm.
6	Monte-Carlo Methods	*RL* Sections 5.1 to 5.4	1. Project Milestone 1.2: Project Proposal (Presentation Video and Discussion)	1. Due by Sunday 11:59pm.
7	Temporal-Difference Learning	*RL* 6.1 to 6.3 and 7.1	1. Module 7 Quiz: MC and TD Prediction 2. Programming Assignment 1: Value Iteration and Policy Iteration	1. Due by Sunday 11:59pm. 2. Due by Sunday 11:59pm.
8	Temporal-Difference Control	*RL* Sections 6.4 to 6.9 and 7.2 to 7.7	1. Module 8 Quiz: MC and TD Prediction 2. Project Milestone 2.1: Project Midterm Demo (Presentation Slides Submission) 3. Project Milestone 2.2: Project Midterm Demo (Presentation Video and Discussion)	1. Due by Sunday 11:59pm. 2. Due by Sunday 11:59pm. 3. Due by Sunday 11:59pm.
9	Value Function Approximation	*RL* Sections 9.1 and 9.4	1. Project Milestone 2.2: Project Midterm Demo (Presentation Video and Discussion)	1. Due by Sunday 11:59pm.

Week	Topic	Reading	Assignment	Due Dates
10	Deep Q-Learning	*RL* Sections 10.1, 10.2, and 11	1. Module 10 Quiz: Value Function Approximation 2. Programming Assignment 2: Sarsa and Q Learning	1. Due by Sunday 11:59pm. 2. Due by Sunday 11:59pm.
11	Project Development and Iteration		1. Project Milestone 3.1: Project Paper Presentation (Slides Submission) 2. Project Milestone 3.2: Project Paper Presentation (Video Submission and Discussion)	1. Due by Sunday 11:59pm. 2. Due by Sunday 11:59pm.
12	Policy Gradient	*RL* Sections 13.1 to 13.4	1. Module 12 Quiz: Policy Gradient 2. Project Milestone 3.2: Project Paper Presentation (Video Submission and Discussion)	1. Due by Sunday 11:59pm. 2. Due by Sunday 11:59pm.
13	Planning and Learning	*RL* Sections 8 and 13.5 to 13.8	1. Module 13 Quiz: Planning and Learning 2. Programming Assignment 3: DQN	1. Due by Sunday 11:59pm. 2. Due by Sunday 11:59pm.
14	Project Final Demo	N/A	1. Project Milestone 4.1: Project Final Demo (Slides Submission) 2. Project Milestone 4.2: Project Final Demo (Video and Discussion)	1. Due by Sunday 11:59pm. 2. Due by Sunday 11:59pm.
15	Exploration and Exploitation	*RL* Section 2	1. Module 15 Quiz: Exploration and Exploitation 2. Project Milestone 4.2: Project Final Demo (Video and Discussion) 3. Project Milestone 5: Final Project Report	1. Due by Sunday 11:59pm. 2. Due by Sunday 11:59pm. 3. Due by Sunday 11:59pm.

Additional Information and Resources

Accessibility:

This course is offered through an accessible learning management system. For more information, please refer to Canvas's Accessibility Statement.

Requesting Accommodations:

The Office of Accessibility Resources and Services works in partnership with administrators, faculty, and staff to provide reasonable accommodations and support services for students with disabilities who have provided their office with medical documentation to receive services.

If you are in need of accommodations due to a disability, please contact the Office of Accessibility Resources and Services to discuss your specific needs.

Resources for NJIT Online Students

NJIT is committed to student excellence. To ensure your success in this course and your program, the university offers a range of academic support centers and services. To learn more, please review the “Student Services” page in Canvas, which includes information related to technical support.