SJSU SAN JOSÉ STATE UNIVERSITY

College of Science · Computer Science

Reinforcement Learning and Sequential Decision Making

CS 272

Fall 2025 Section 01 In Person 3 Unit(s) 08/20/2025 to 12/08/2025 Modified 08/12/2025

Contact Information

Dr. Genya Ishigaki

Email: genya.ishigaki@sjsu.edu
Office: MH 215
Phone: (408) 924-5076
Website: https://sjsu-interconnect.github.io/ (https://sjsu-interconnect.github.io/)

Office Hours

Monday, Wednesday, 3:00 PM to 4:00 PM, MH 215

You don't need to make an appointment for these office hours. You can stop by my office.

Course Description and Requisites

Introduction to reinforcement learning, deep reinforcement learning, other online learning algorithms, and their applications.

Prerequisite(s):ï¿½CS 157A. Allowed Declared Major: Computer Science MS, Bioinformatics MS, and Data Science MS.ï¿½

Letter Graded

Classroom Protocols

- This course is conducted in person, which requires physical attendance for a significant portion of its delivery.
- Students are requested to use **the Canvas message function** to contact the instructor. Private messages sent to the instructor's email address get lost due to the large volume of emails received.

- The instructor does not write messages after normal business hours, on weekends, or on holidays.
- Reviewing code for the homework and technical troubleshooting should be done during office hours. Never send your entire code for an assignment to the instructor. The instructor will not fix the bugs in your code.

Program Information

Diversity Statement - At SJSU, it is important to create a safe learning environment where we can explore, learn, and grow together. We strive to build a diverse, equitable, inclusive culture that values, encourages, and supports students from all backgrounds and experiences.

Course Learning Outcomes (CLOs)

Upon successful completion of this course, students will be able to:

- Distinguish different types of reinforcement learning algorithms and when to use them.
- Describe the benefits and potential challenges of deep reinforcement learning.
- Apply reinforcement learning algorithms to real-world problems.
- Analyze and evaluate the performance of reinforcement algorithms.
- Create a reinforcement learning project to solve a social or technical issue.

Course Materials

Textbook:

- Richard S. Sutton and Andrew G. Barto, Reinforcement learning: An introduction (Second edition), MIT press, 2018.

Supplemental Textbooks:

- Michael A. Nielsen, Neural Networks and Deep Learning, Determination Press, 2015.
- Marc Peter Deisenroth, A. Aldo Faisal, and Cheng Soon Ong, Mathematics for Machine Learning, Cambridge University Press, 2020.

Technology

- Python development environment with the following packages (It is highly recommended to use an isolated Python workspace, such as Anaconda env.)
  - Gymnasium (https://gymnasium.farama.org/index.html (https://gymnasium.farama.org/index.html))
  - PettingZoo (https://pettingzoo.farama.org/ (https://pettingzoo.farama.org/))
  - Ray RLlib (https://docs.ray.io/en/latest/rllib/index.html (https://docs.ray.io/en/latest/rllib/index.html))
- All demos assume a UNIX-like OS environment.

Course Requirements and Assignments

Item	Percentage
Wednesday Quiz 1 - 10 (wq)	20% (2% Each)
Programming Assignments 1 - 3 (pa)	30% (10% Each)
Exam 1	20%
Exam 2	20%
Final Project	10%

Wednesday Quizzes

There will be a small quiz at the beginning of the class (almost) every Wednesday. (Check the Syllabus on Canvas for the schedule.)

Answers will be discussed in class but will not be published. No retake or makeup opportunity will be given as the solutions will be revealed in class.

The purpose of the Wednesday quizzes is to confirm your understanding of the fundamental concepts in RL right after each class. Most quizzes are multiple-choice or True-False questions and do not involve complicated calculations.

Programming Assignments (PA)

Submissions may be graded based on automated software testing. The specifications will be given in class and in the form of code templates.

Exams

Two exams will be conducted during the regular class hours. A tentative schedule will be given in the course schedule below.

No make-up exams except in case of verifiable emergency circumstances.

Grading Information

Extra-credits and Reworks

No extra-credit assignments or rework opportunities will be given.

Late Submission

No late submissions will be accepted.

Missed Assignments or Exams

When students need to miss an assignment deadline or exam due to health conditions or any other emergency, it should be reported within ONE week after the due date.

Final Grade Table

Total Grade	Letter Grade
97% and above	A plus
93% to 96%	A
90% to 92%	A minus
87% to 89%	B plus
83% to 86%	B
80% to 82%	B minus
77% to 79%	C plus
73% to 76%	C
70% to 72%	C minus
67% to 69%	D plus
65% to 66%	D
60% to 64%	D minus
59% and below	F

University Policies

Per University Policy S16-9 (PDF) (http://www.sjsu.edu/senate/docs/S16-9.pdf), relevant university policy concerning all courses, such as student responsibilities, academic integrity, accommodations, dropping and adding, consent for recording of class, etc. and available student services (e.g. learning assistance, counseling, and other resources) are listed on the Syllabus Information (https://www.sjsu.edu/curriculum/courses/syllabus-info.php) web page. Make sure to visit this page to review and be aware of these university policies and resources.

Course Schedule

Date	Topic	Reference	Note
8/20	Overview		
8/25	What is Learning?		
8/27	Markov Decision Processes	Chap 3	
9/1	Labor Day; No class		
9/3	Policies and Value Functions		wq1: Learning Concepts
9/8	Dynamic Programming	Chap 3	
9/10	Coding: DP	Chap 4	wq2: MDP; PA1 due
9/15	Try and Error: MAB	Chap 2	
9/17	MC	Chap 5	wq3: Value functions and DP
9/22	TD	Chap 6	
9/24	Policy Gradient Methods 1	Chap 13	wq4: MC; PA2 due
9/29	Policy Gradient Methods 2 + coding		
10/1	Exam Review		wq5: TD
10/6	Exam 1		
10/8	Coding Workshop: Gymnasium		
10/13	Approximation		
10/15	Linear Approximation Implementation	Chap 10	wq6: PG
10/20	Coding Workshop: PettingZoo & Custom Environment		PA3 due
10/22	Custom Environment & Deep Learning		wq7: Approximation, Project group formation due
10/27	Deep Learning		
10/29	Deep RL: DQN + Implementation		wq8: Deep Learning
11/3	Deep RL: A3C, PPO		
11/5	Coding Workshop: Ray rllib		wq9: DRL

11/10	Project discussion		
11/12	Integrating Learning and Planning	Chap 8	
11/17	Integrating Learning and Planning		
11/19	Advanced topics in DRL		wq10: Learning and planning
11/24	Exam Review		
11/26	Fall break; No class		
12/1	Project presentation		Final project due
12/3	Project presentation		
12/8	Exam 2		