

Reinforcement Learning and Sequential Decision Making Section 01

CS 272

Fall 2024 3 Unit(s) 08/21/2024 to 12/09/2024 Modified 07/18/2024

Course Description and Requisites

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

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

Letter Graded

* Classroom Protocols

Communication with the instructor

Students are requested to use the Canvas message function to contact the instructor. Private messages sent to the instructor's email address gets lost due to the large volume of emails received.

The instructor does not write messages after normal business hours, on weekends or holidays.

Reviewing code for the homework and technical trouble-shooting should be done during the office hours.

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 machine 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 (Details will be given in class)

Course Requirements and Assignments

Item	Percentage
Exam 1	15%
Exam 2	15%
Programming Assignments (1 - 5)	50% (10% Each)
Final Presentation	10%
Wednesday Quizzes	10% (1% Each)

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.

Programming Assignments (PA)

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

Final Presentation

Students are expected to present their projects. Details will be posted on Canvas.

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/answer or True-False questions and do not involve complicated calculations.

✓ Grading Information

Extra-credits and Reworks

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

Late Submission

Late submissions within 24 hours will be deducted 10% of its final grade. Submissions over 24 hours late will have 20% grade deducted. Late submissions over 2 days will not 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) (<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) (<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
8/21	Overview		
8/26	What is Learning?		
8/28	Markov Decision Processes	Chap 3	
9/2	Labor Day; No class		
9/4	Policies and Value Functions		wq1: Learning Concepts
9/9	Dynamic Programming	Chap 3	
9/11	Coding: DP	Chap 4	wq2: MDP; PA1 due
9/16	Try and Error: MAB	Chap 2	
9/18	MC	Chap 5	wq3: Value functions and DP
9/23	TD	Chap 6	

Date	Topic	Reference	Note
9/25	Policy Gradient Methods 1	Chap 13	wq4: MC; PA2 due
9/30	Policy Gradient Methods 2		
10/2	Exam Review		wq5: TD
10/7	Exam 1		
10/9	Coding Workshop: Gymnasium		PA3 due
10/14	Coding Workshop: Custom Environment		
10/16	Approximation	Chap 10	wq6: PG
10/21	Linear Approximation Implementation		
10/23	PA4 Discussion		wq7: Approximation; Project doc due
10/28	Deep Learning		PA4 due
10/30	Deep RL		wq8: DL
11/4	Implmentation aspects of Deep RL		
11/6	Coding Workshop: Ray rllib		wq9: DRL
11/11	Veterans Day; No class		
11/13	Integrating Learning and Planning	Chap 8	
11/18	Integrating Learning and Planning		
11/20	Advanced topics on RL		wq10: Learning and planning
11/25	Coding Workshop		
11/27	Fall break; No class		
12/2	Final presentation		PA5 due
12/4	Exam Review		
12/9	Exam 2		