College of Science · Computer Science

Reinforcement Learning and Sequential Decision Making Section 01

CS 272

Fall 2023 3 Unit(s) 08/21/2023 to 12/06/2023 Modified 08/21/2023

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 simply stop by my office.

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.

Never send your entire code for an assignment to the instructor. The instructor will not fix all 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 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

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. All submissions will 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 a paper on state-of-the-art RL techniques and related topics. A list of suggested papers will be provided during the class.

It is recommended to form a group of two students. I may approve exceptions (individual or group of three) with a valid reason.

Wednesday Quizzes

There will be a small quiz at the beginning of (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

Exam 1	15%
Exam 2	15%

Programming Assignment 1	10%
Programming Assignment 2	10%
Programming Assignment 3	10%
Programming Assignment 4	10%
Programming Assignment 5	10%
Final Presentation	10%
Wednesday Quizes	10% (1% Each)

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
92% to 96%	A
90% to 91%	A minus
87% to 89%	B plus
82% to 86%	B
80% to 81%	B minus
77% to 79%	C plus
72% to 76%	C
70% to 71%	C minus
67% to 69%	D plus
62% to 66%	D
60% to 61%	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
8/21	Overview		
8/23	What is Learning?		
8/28	Coding: Python	https://docs.python.org/3/library/unittest.html	PA1 posted
8/30	Markov Decision Processes	Chap 3	wq1: Learning Concepts
9/4	No class; Labor Day		
9/6	Policies and Value Functions	Chap 3	wq2: MDP
9/11	Dynamic Programming	Chap 4	
9/13	Coding: DP		wq3: Value functions and DP; PA2 posted
9/18	MC	Chap 5	
9/20	TD	Chap 6	wq4: MC
9/25	Try and Error: MAB	Chap 2	
9/27	Gymnasium Implementation	https://gymnasium.farama.org/	wq5: TD
10/2	Coding: Model-free control		PA3 posted
10/4	Exam 1		
10/9	Approximation	Chap 10	
10/11	Linear Approximation Implementation		wq6: Approximation
10/16	Deep Learning		
10/18	NN Implementation		wq7: Deep Learning
10/23	Deep RL		
10/25	Coding: Deep RL		wq8: Deep RL: PA4 posted
10/30	Ray rllib Implementation		
11/1	Policy Gradient Methods	Chap 13	wq9: Coding Deep RL

11/6	Coding: Policy based		
11/8	Offline RL	https://arxiv.org/abs/2005.01643	wq10: Policy-based RL
11/13	Integrating Learning and Planning	Chap 8	
11/15	Scalability: Representation Learning (GNN)	https://distill.pub/2021/gnn-intro/	PA5 posted
11/20	Implementation Aspects of State Representation Learning		
11/22	No class; Fall break		
11/27	Final presentation		
11/29	Final presentation		
12/4	Exam Review		
12/6	Exam 2		