

CS 272 Reinforcement Learning and Sequential Decision Making

Course Information

- Instructor: Genya Ishigaki
 - Telephone: (408) 924-5076
 - Email: genya.ishigaki@sjsu.edu
 - Office Hours:
 - Mondays & Wednesdays 2:00 PM - 3:00 PM
 - Location: MacQuarrie Hall 215
 - You do NOT need to make an appointment for these office hours. You can simply stop by my office.
- Class Days/Time: Mondays & Wednesdays 12:00 PM - 1:15 PM
- Class mode: In-person
- Class Location: MacQuarrie Hall 422
- Prerequisites: CS 157A with a grade of C- or better. Limited to MSCS, MSBI, and MSDS students.

Course Description

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

Course Learning Outcomes (CLO)

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.

Textbook

- Richard S. Sutton and Andrew G. Barto, [Reinforcement learning: An introduction](#) (Second edition), MIT press, 2018.
- Open AI, [Spinning Up in Deep RL](#).
- Michael A. Nielsen, [Neural Networks and Deep Learning](#), Determination Press, 2015.

Other Equipment

- Python development environment
 - [Anaconda](#)
 - [Google Colab](#)
- LaTeX (*for Project Summary)

- [Overleaf](#)

Grading

Exams, Assignments, and Projects

- Four unit exams (Exam 1,2,3,4)
 - The worst-scored exam will be dropped automatically at the end of the semester.
 - All exams are planned to be conducted during the regular class hours.
- Four programming assignments (PA 0,1,2,3)
 - PA0: Coding practice project
 - PA1,2,3: RL projects
- Final project
 - It is recommended to form a group of TWO students. I may approve exceptions (individual or group of three) with a valid reason.
 - Topic Summary: A one-page document describing the project idea
 - Final Presentation: Presentation summarizing the evaluation of selected RL algorithms
 - Project Codes: Implementation of RL algorithms based on selected papers

Item	% in Final Grade
Exam 1,2,3,4	36% (12% each; The worst one will be dropped.)
Programming Assignment (PA) 0	7%
Programming Assignment (PA) 1,2,3	36% (12% each)
Project Topic Summary	5%
Project Final Presentation / Codes	16 %

Grading 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

Total Grade	Letter Grade
62% to 66%	D
60% to 61%	D minus
59% and below	F

Extra-credits and Reworks

The worst-scored exam among the four exam papers will be dropped automatically. No additional 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.

Attendance

I will not take attendance for classes. Students not attending either of the first two classes will be dropped to make room for students on the waiting list. Attempting to get marked as present (by having someone else attend in your place or using technological deceptions) will be considered academic dishonesty and at a minimum will result in you getting dropped from the course.

Grading Policy

The University Policy S16-9, Course Syllabi (<http://www.sjsu.edu/senate/docs/S16-9.pdf>) requires the following language to be included in the syllabus:

"Success in this course is based on the expectation that students will spend, for each unit of credit, a minimum of 45 hours over the length of the course (normally three hours per unit per week) for instruction, preparation/studying, or course related activities, including but not limited to internships, labs, and clinical practica. Other course structures will have equivalent workload expectations as described in the syllabus."

Fall 2022 Announcement: COVID-19 and Monkeypox

Students registered for a College of Science (CoS) class with an in-person component should view [the CoS COVID-19 and Monkeypox Training slides](#) for updated CoS, SJSU, county, state and federal information and guidelines, and more information can be found on the SJSU Health Advisories website. By working together to follow these safety practices, we can keep our college safer. Failure to follow safety practice(s) outlined in the training, [the SJSU Health Advisories website](#), or instructions from instructors, TAs or CoS Safety Staff may result in dismissal from CoS buildings, facilities or field sites. Updates will be implemented as changes occur (and posted to the same links).

University Policies

Per University Policy S16-9, university-wide policy information relevant to all courses, such as **academic integrity, accommodations**, etc. will be available on Office of Graduate and Undergraduate Programs' Syllabus Information web page at <http://www.sjsu.edu/gup/syllabusinfo/>. Make sure to review these policies and resources.

Tentative Schedule and Topics

Date	Topic	Reference	Note
8/22	Overview		
8/24	What is Learning?	Shoham and Leyton-Brown Chap 7, Paper	
8/29	Python Recap		
8/31	Markov Decision Processes	Sutton and Barto Chap 3	
9/5	Labor Day - No class		
9/7	Policies and Value Functions	Sutton and Barto Chap 3	
9/12	Dynamic Programming	Sutton and Barto Chap 4	PA0 Due
9/14	Coding: MDP and DP		
9/19	Exam 1, Coding Q&A		
9/21	Model-free prediction	Sutton and Barto Chap 5, 6	
9/26	Model-free prediction	Sutton and Barto Chap 5, 6	
9/28	Model-free control	Sutton and Barto Chap 5, 6	PA1 Due
10/3	Coding: Model-free control		
10/5	Exam 2, Open AI Gym		
10/10	Approximation	Sutton and Barto Chap 9, 10	
10/12	Linear Approximation Implementation		
10/17	Deep Learning	Nielsen's book	PA2 Due
10/19	Deep RL	OpenAI RL Spinning Up	
10/24	Coding: Deep Learning		
10/26	Exam 3, Project Guidelines		
10/31	MAB and Regret	Sutton and Barto Chap 2	
11/2	Integrating Learning and Planning	Sutton and Barto Chap 8	
11/7	Project Topic Discussion		Project Topic Due
11/9	Policy Gradient Methods	Sutton and Barto Chap 13	PA3 Due
11/14	Actor-Critic Methods	Sutton and Barto Chap 13	
11/16	Policy Gradient Implementation		

Date	Topic	Reference	Note
11/21	Exam 4, Evaluation Guidelines		
11/23	Non-Instructional Day - No class		
11/28	Advanced Topics in RL		
11/30	Final presentation		Final Slides / Project Codes Due
12/5	Final presentation		