

Statistical Foundation of Reinforcement Learning

CS292F Spring 2021

Instructor: Prof. Yu-Xiang Wang

Lectures: On Zoom. Monday and Wednesday 1:00 - 2:50

Piazza: <https://piazza.com/ucsb/spring2021/cs292/home>
(We use Piazza for Q&A, discussion, and most announcement!)

Gradescope: <https://www.gradescope.com/courses/258384>
(We will be using the gradescope for submitting homeworks and project reports.)

Gauchospace: The Gauchospace is active, but it is only reserved for communications that are restricted to people who are officially registered.

1 Overview

Reinforcement learning is one of the hottest topics in machine learning research lately. It is the main technique behind many recent breakthroughs in artificial intelligence, such as AlphaGo and AlphaZero. RL has shown great promise in addressing real life sequential decision making problems.

This course is about the theory of reinforcement learning. The targeted audience is PhD level students in Computer Science, Statistics, Math and Electrical Engineering who work on (or at least is actively developing an interest in) theoretical research.

2 What you will learn?

1. Markov decision process (MDP) model and analyze the computational complexity of MDP inferences.
2. Various settings of reinforcement learning and their relationships.
3. Ideas and formal arguments behind strategic exploration in (online) Bandits and RL
4. Ideas and formal arguments behind the recent progress in offline reinforcement learning.
5. Background knowledge needed to read and appreciate the recent literature in RL theory.

3 Prerequisites

As this is a graduate level course, there is no hard pre-requisite, but students entering the class are expected of the following.

- Students are expected to be able to follow rigorous mathematical arguments and perform derivation and proofs using notations, definitions and theorems from calculus, linear algebra, probability and statistics.

- Students are expected to have working knowledge of basic algorithms and data structures, and could write simple code in Python / Numpy.
- If you have taken a course in statistical machine learning, linear / convex optimization, control theory, or are familiar with topics such that concentration inequalities, dynamic programming, dynamical systems through other means, then you are very well-prepared to enjoy the course.

4 Textbooks

We will be following the exposition of the recent monograph by Agarwal, Jiang, Kakade, Sun [Link] in many of our topics, but will also draw from research papers and other sources.

5 Assignments and Grades

The grades will be based on the following breakdowns.

- 40% Homework assignments.
- 40% Course project.
- 10% Scribing.
- 10% Lecture attendance and participation.

6 Course project Project

The course project will be in teams of maximum size 3. Individual projects are fine, but it is usually more fun to have someone to talk to. The scope of the project depends on your own interest and bandwidth, and is intended to be flexible. For example, you could leverage your ongoing research work if it is related.

Basic project Read an RL theory paper (or any related subject), reproduce its proofs.

Advanced project Develop an extension of existing results in a paper (to a new setting or a new application).

Special project Apply RL techniques your learned to your own research / existing research.

In all three cases, you will need to submit: a short proposal, a midterm report, and a final report; as well as to present your project to the class towards the end of the quarter. For group projects, each student need to contribute and a section should be included in the midterm and final report on the contribution from each group member.

7 Logistics

The instruction will be online.

Zoom lectures: Only registered students are given access to the real-time lectures over Zoom. You have been emailed instructions on how to connect to the lecture. If you are having trouble, please contact the instructor.

The live Zoom sessions will be recorded for students who may not be able to attend synchronous. By default, your microphone and camera will be muted when you join the session. If you do not want to be included in the recording, simply keep your camera and microphone off. You may ask questions in the chat window. Please refer to the “copyright” section for more details.

Zoom office hours: Office hours will be conducted in a similar manner to the lectures. Details will be announced via email.

Attendance policy: The attendance to lectures is required. It is part of the course evaluation to attend the lectures. Send PM on Piazza if you will have to miss lectures due to other personal businesses.

Late homework policy: Late submissions are allowed, but if you are late more than one full week then there will be a 20% penalty.

8 Policy on Academic Integrity

Please read this section carefully.

The university, the department, and this instructor all take the issue of academic integrity **very** seriously. A university requires an atmosphere of mutual trust and respect. While collaboration is an integral part of many scholarly activities, it is not always appropriate in a course, and it is never appropriate unless due credit is given to all participants in the collaboration.

Here are some examples:

- Allowed: Discussion of lecture and textbook materials
- Allowed: Discussion of how to approach assignments, what techniques to consider, what textbook or lecture material is relevant
- Allowed: Collaboration on homeworks. You need to declare your collaborators, and describe what you get from the person who helped you, and each student still needs to write their own report / code independently.
- Allowed: Refer to online resources, but you need to cite the exact references.
- Not allowed: Turning in someone else's work as your own, even with that person's permission.
- Not allowed: Allowing someone else to turn in your work as his or her own.
- Not allowed: Turning in work without proper acknowledgment of the sources of the content (including ideas) contained within the work.

For some views on academic integrity at UCSB see the Academic Integrity page of the Office of Judicial Affairs.

Summary: Academic integrity is absolutely required - dishonesty (cheating, plagiarism, etc.) benefits no one and hurts everyone. Violations of these honor codes on academic integrity will be reported to the Office of Student Conduct. If you find yourself in such a situation, please contact the instructor. If you are not sure whether or not something is appropriate, please ask the instructor.

9 Code of conduct

The University of California, Santa Barbara has a general code of conduct for all students published here: http://www.sa.ucsb.edu/docs/default-source/student-conduct/conductofcode2017.pdf?sfvrsn=d3c07d4f_2

The computer science department's commitment to Diversity, Equity and Inclusion is published here: <https://cs.ucsb.edu/content/diversity-equity-and-inclusion>.

As a department, we holds students, staff, and faculty to the following standards:

- Treat all members of the academic community (students, staff, and faculty) with respect regardless of their experiences and background, including (but not limited to) their cultural backgrounds, socioeconomic status, disabilities, age, religion, sexual orientation, neuro(a)typicality, and gender identity.
- Physical or mental harm, sexual harassment, aggression, and derogatory language is not acceptable in any form.
- Respect the personal property of others and University resources. Unauthorized access, use, vandalism, or theft of equipment, computer servers, labs / offices / classrooms, etc. is prohibited.
- The exchange and challenge of ideas are done in a thoughtful, respectful and constructive manner.
- Disruption of departmental activities such as special events, talks, lectures, and meetings is not acceptable.
- Any violation of the given standards should be reported to the Computer Science chair and/or the CS diversity committee (diversity@cs.ucsb.edu). Consequences may include a formal warning, suspension, or expulsion from the University.

10 Students with Disabilities

If you are a student with a disability and would like to discuss special academic accommodations, please contact the instructor. In addition, students with temporary or permanent disabilities are referred to the Disabled Students Program (DSP) at UCSB. DSP will arrange for special services when appropriate (e.g., facilitation of access, note takers, readers, sign language interpreters). Please note that it is the student's responsibility to communicate his or her special needs to the instructor, along with a letter of verification from DSP.

11 Copyright of course materials

My lectures and course materials, including presentations slides, written notes, recorded lectures, homework assignments and similar materials, are protected by U.S. copyright law and by University policy. I am the exclusive owner of the copyright in those materials I create. You may take notes and make copies of course materials for your own use. You may also share those materials with another student who is enrolled in or auditing this course. You may not reproduce, distribute or display (post/upload) lecture notes or recordings or course materials in any other way — whether or not a fee is charged — without my express prior written consent. You also may not allow others to do so.

If you do so, you may be subject to student conduct proceedings under the UC Santa Barbara Student Code of Conduct.

Similarly, you own the copyright in your original papers and exam essays. If I am interested in posting your answers or papers on the course web site, I will ask for your written permission.

The live Zoom session will be recorded for students who may not be able to attend at this time. By default, your microphone and camera will be muted when you join the session. If you do not want to be included in the recording, simply keep your camera and microphone off. You may ask questions in the chat window.