



# Reinforcement Learning

ECE 693B

## Instructor Info



Yuanzhang Xiao



Office Hours: TBD or by appointment



POST 201G or on Zoom



yxiao8@hawaii.edu

## Course Info



two days per week (TBD)



1 hour 15 minutes (TBD)



TBD

## Overview

Reinforcement learning (RL) is the technique behind many breakthroughs in artificial intelligence (AI), such as AlphaGo (the first computer program that beats top human players in Go) and training large language models (LLMs). It also finds a variety of applications in wireless communications, smart grids, and so on.

In this class, we will learn the basics of RL, including theories, algorithms, and applications. We will also get our hands dirty through coding, and therefore understand the nuances in applying RL to research projects.

Specifically, the topics include but are not limited to

- Foundations
  - Problem formulation: Markov decision processes
  - Solution methods: dynamic programming, value iteration, policy iteration
  - Learning algorithms: Q-learning, TD learning, etc.
- State-of-the-art algorithms
  - Deep Q-learning
  - Actor-critic methods
  - Policy gradient methods
- Programming
  - RL environments (Gymnasium, building custom environments)
  - Training frameworks (Stable-Baselines3, RL Baselines3 Zoo)
  - Evaluation and benchmarks (rliable, Hugging Face)
- New paradigms
  - Offline RL
  - Multi-agent RL

At the end of this course, the students will obtain solid understandings of RL theories and hands-on experience in applying RL to their research projects.

## Material

- *Reinforcement Learning and Optimal Control* by Dimitri Bertsekas
- various online resources

## Grading Scheme

60% Homework

40% Course Project