

CSE 4/510 Introduction to Reinforcement Learning

1 Course Description

This course is intended for students interested in artificial intelligence. Reinforcement learning is an area of machine learning where an agent learns how to behave in a environment by performing actions and assessing the results. Reinforcement learning is how Google DeepMind created the AlphaGo system that beat a high-ranking Go player and how AlphaStar become the first artificial intelligent to defeat a top professional player in StarCraft II. We will study the fundamentals and practical applications of reinforcement learning and will cover the latest techniques used to create agents that can solve a variety of complex tasks, with applications ranging from gaming to finance to robotics. The course is comprised of assignments, short weekly quizzes and a final exam.

2 Logistics

Summer 2019

Session: May 28 - Jul 05

Time: T, R 3:00 - 6:15

Room: Knox 104

Course number: CSE4/510 Special Topic - Introduction to Reinforcement Learning

Level: Underdatuates/Graduates

3 Objective

At the end of the course, students are supposed to:

- Learn how to define reinforcement learning (RL) task
- Learn the core principals behind the RL, including policies, value functions, deriving Bellman equations
- Understand the current code standards and libraries used in RL
- Reproduce fundamental RL algorithms
- Understand tabular methods to solve classical control problems
- Learn the policy optimization methods from vanilla to more complex cases
- Know current advance techniques and applications in RL

4 Syllabus*

Foundations

1. Introduction and basics of Reinforcement Learning (RL)
2. Defining RL framework and Markov Decision Process
3. Policies, Value Functions and Bellman Equations
4. Exploration vs. Exploitation
5. Code standards and libraries used in RL (Python/Keras/Tensorflow)

Tabular methods and Q-learning

6. Planning by Dynamic Programming and Monte Carlo
7. Temporal-Difference learning methods (TD(0), SARSA, Q-Learning)
8. Deep Q-networks (DQN): experience replay and backpropagation

Policy optimization

9. Introduction to policy-based methods
10. Vanilla Policy Gradient
11. REINFORCE algorithm and stochastic policy search
12. Actor-critic methods (A2C)
13. Proximal Policy Optimization

Model based RL

14. Model-based RL approach

Recent Advances and Applications

15. Multi-Agent Reinforcement Learning
16. Safety in RL
17. Applying RL for real-world problems

*The preliminary syllabus is subject to change.

5 Reference Materials

- Richard S. Sutton and Andrew G. Barto, "Reinforcement learning: An introduction", Second Edition, MIT Press, 2019
- Li, Yuxi. "Deep reinforcement learning." arXiv preprint arXiv:1810.06339 (2018).
- Lecture notes and other relevant materials will added to UBLearns

6 Evaluation*

50% - Assignments

20% - Short weekly quizzes

30% - Final Exam

*Minor changes to the evaluation scheme might apply before the beginning of the class

7 Projects Overview

Assignments

Assignments will focus on implementing fundamental RL algorithms.

Possible tasks:

- Building custom environment following Gym standards
- Implementing Q-learning / DQN
- Implementing REINFORCE / Proximal policy optimization

8 Prerequisites

CSE 250. Basic knowledge of a object-oriented programming is required (ideally experience in Python). Knowledge of probabilities/statistics is required; calculus and linear algebra are ideal. Experience with machine learning is not required, but some basic familiarity may help. If you have doubts regarding your background, please contact Alina to discuss it.