Proposed new course: Reinforcement Learning

Lecturer: Bolei Zhou, Computer Science Department

Course Description:

This course aims to teach the fundamentals and the advanced topics of Reinforcement Learning (RL), a computational learning approach where an agent tries to maximize the total amount of reward it receives while interacting with the complex and uncertain environments. The course content includes the introduction of Markov Decision Processes, model-free RL and model-based RL methods, policy optimization, RL distributed system design, as well as the case studies of RL in game playing such as AlphaGo, traffic simulation, autonomous driving, and other machine autonomy applications. Two weeks of lectures will be on the advanced topics of RL such as multi-agent RL, human-in-the-loop method, and imitation learning.

Lectures will be four hours per week, the discussion will be one hour per week, and the outside study will be seven hours per week.

Learning Objectives:

- Students should understand the reinforcement learning foundations such as value-based RL, policy-based RL, actor-critic.
- Students should understand the common RL algorithms such as PPO, SAC, and apply them in their research projects and applications.
- Students should know the strengths and weaknesses of RL and their real-world applications.

Knowledge outcomes:

- Understand the foundation of reinforcement learning and the relevant algorithms
- Be familiar with various reinforcement learning algorithms such as model-based and model-free RL algorithms, on-policy and off-policy algorithms, and their strengths and weaknesses.
- Design and develop applications based on the framework of reinforcement learning

Skills outcomes:

- Set up and run the simulation system environment
- Perform the system design and modeling using reinforcement learning
- Apply and improve the reinforcement learning techniques for practical and academic problems

Weekly Topics:

Week 1: Overview
Lecture 1: Course introduction
Lecture 2: RL basics and coding

Week 2: RL basics
Lecture 3: Markov decision process
Lecture 4: Policy iteration and value iteration

Week 3: Tabular model-free methods
Lecture 5: Model-free prediction
Lecture 6: Model-free control

Week 4: Value-based RL
Lecture 7: Value function approximation
Lecture 8: Deep Q Learning

Week 5: Policy-based RL: basics
Lecture 9: Policy Optimization 1: policy gradient
Lecture 10: Policy Optimization 2: actor-critic algorithms

Week 6: Policy-based RL: advances
Lecture 11: Policy Optimization 3: TRPO, PPO
Lecture 12: Policy Optimization 4: TD3, SAC

Week 7: Model-based RL
Lecture 13: Model-based RL
Lecture 14: Connection to optimal control

Week 8: Advanced topics
Lecture 15: Imitation learning
Lecture 16: Offline RL

Week 9: Advanced topics
Lecture 17: Human-in-the-loop RL
Lecture 18: Distributed computing and RL system design

Week 10: Course summary
Lecture 19: Course project presentation
Lecture 20: Course summary

Course Assignments:

There will be four assignments:

- Assignment 1: RL basics and tabular methods
- Assignment 2: Value iteration, policy iteration, and Q-learning
- Assignment 3: Policy gradient, actor-critic
- Assignment 4: advanced policy optimization, imitation learning, open-ended competitive environments

Grading structure:

Attendance:	10%
Assignments:	30%
Final exam:	40%
Course project:	20%

Reading list:

Textbook:

- Reinforcement Learning: An Introduction 2nd edition by Sutton & Barto
- Reinforcement learning and optimal control by Dimitri P. Bertsekas

Justification:

Reinforcement Learning is one of the most popular machine learning areas, with a wide range of applications like robotics, autonomous driving, and data mining. This course will be interesting to a diverse set of graduate students from various departments within and outside of School of Engineering, such as CS, ECE, Mechanical Engineering, Statistics, and Math. This course is being offered through CS269 this fall quarter, there are more than 65 enrolled students from more than 5 different departments and programs, as well as another 35 students who are registered as listeners to attend the ZOOM auditing. Thus there is a genuine need and popularity to open such a course dedicated to reinforcement learning.

A related graduate course is ECE-239AS at ECE department, which mainly focuses on the theory side of RL, and it is not open regularly. For example, this course is not open for this academic year. I have talked to that lecturer, Prof. Lin Yang, and the ECE chair Prof. Yang, both support the opening and the potential cross-listing between CS and ECE departments.

I have taught this RL course in my previous university for three semesters, and this course is being run through CS269 this fall quarter. The course structure and materials are well-polished and ready for a dedicated new course.