Reinforcement Learning and AI @ Yonsei (Summer 2026)

Profs. Chang-hee Won (Temple University, USA), Jongbaeg Kim (Yonsei University, Korea)

Course Overview

Reinforcement learning, a subset of artificial intelligence, is a computational approach to decision-making that models the cause-and-effect relationships between actions and rewards. It is based on the idea that a system can learn through trial and error. Reinforcement learning is one of the three primary paradigms in machine learning, alongside supervised learning and unsupervised learning. This framework provides a powerful approach to solving optimization problems, in which an agent interacts with its environment and progressively refines its policy over time. Closely related to both optimal control and sequential decision-making areas. This course explores the fundamental principles of reinforcement learning and its integration into control applications. Topics include reinforcement learning algorithms, Markov decision processes, dynamic programming, and other reinforcement learning techniques. Applications of reinforcement learning span a wide range of fields, including robotics, manufacturing, finance, healthcare, operations, natural language processing, and game-playing systems such as AlphaGo.

The objective of this short course is to introduce reinforcement learning to advanced undergraduate and graduate students in mechanical engineering and related disciplines concerned with sequential decision-making processes. The central goal of the course is to understand how physical systems can be represented using the framework of reinforcement learning. We will begin by introducing the fundamentals of reinforcement learning and its relationship to artificial intelligence. We will then discuss how to incorporate machine learning principles and constraints into optimal control methods through techniques such as Markov decision processes, dynamic programming, model-free prediction, and model-free control. Students are expected to have a background in linear algebra, probability, and computer programming.

Schedule

Day 1. Introduction to Reinforcement Learning and AI

Lecture 1: Onboarding & course logistics, introduction to RL (2 hours)

Lecture 2: Reinforcement Basics (2 hours)
- The RL Problem
- Inside RL Agent
- Problems with RL

Day 2. Modeling and Solving RL

Lecture 3: Markov Decision Process (2 hours)
- Markov Processes
- Markov Reward Processes
- Markov Decision Processes

Lecture 4: Planning by Dynamic Programming (2 hours)
- Policy Evaluation
- Policy Iteration
- Value Iteration

Day 3. Model-Free RL Prediction and Control

Lecture 5: Model-Free Prediction (2 hours)
- Monte-Carlo Learning
- Temporal-Difference Learning

Lecture 6: Model-Free Control and Concluding Remarks (2 hours)
- On-Policy Monte-Carlo Control
- On-Policy Temporal-Difference Learning
- Off-Policy Learning
- Concluding Remarks

Required Technical Resources

Richard Sutton and Andrew G. Barto, *Reinforcement Learning*, First or Second Edition, The MIT Press, 2020. ISBN 978-0262039246

You are recommended to bring a laptop computer with either:
- Python and relevant development environment configured; OR
- Internet connection to access Google Colab (https://colab.research.google.com). Chrome web browser is recommended.

Instructor Information

Prof. Chang-hee (Andy) Won, Ph.D.

Professor, CSNAP Lab Director
Department of Electrical and Computer Engineering
Temple University, Philadelphia, USA

Email: cwon@temple.edu
Website: https://sites.temple.edu/csnap

Prof. Jongbaeg Kim, Ph.D.

Professor, Department of Mechanical Engineering
Yonsei University

Email: kimjb@yonsei.ac.kr
Website: https://ntl.yonsei.ac.kr/