Mathematics of Reinforcement Learning
Winter 2024

Reinforcement Learning (RL) is a rapidly growing field with significant importance in modern machine learning and artificial intelligence. Its impact spans diverse applications, from game-playing AI to robotics, autonomous systems, and healthcare. At its core, RL focuses on how agents can learn optimal actions through trial and error to maximize cumulative rewards within an environment. Unlike supervised learning, where labeled data guides decisions, RL involves interacting with environments to discover policies that balance immediate and long-term rewards.

In this seminar on the "Mathematics of Reinforcement Learning", we delve into the mathematical foundation that underpins RL algorithms. We begin with a primer on Markov chains, which model the probabilistic transitions between states in a system, forming the backbone of RL environments. From there, we transition into Markov decision processes (MDPs), which introduce decision-making into these chains by incorporating actions and rewards, allowing us to formalize the problem of finding optimal policies for agents. By understanding the mathematics behind these processes, including dynamic programming and the Bellman equations, participants gain a solid grounding in the theory driving state-of-the-art RL algorithms. We then dive into the actual implementation of Reinforcement Learning.

Instructor: Prof. D. Nils Detering

kick-off meeting: 23.09.2024, 16.30 Uhr im Seminarraum 2522.01.81

Prerequirements: A basic course in stochastics, some knowledge of Markov Chains or stochastic processes.

Grading: No grades. Only pass/no pass. For pass you need to give a 90 minute presentation and participate in at least 80% of the presentations throughout the semester.

Language: Depending on the number of non-german speakers, this course might require you to do the presentation in english but this will be discussed during the kick off meeting. You are certainly allowed to do the presentation in english.

Books:

1.) Robert P. Dobrow, *Introduction to Stochastic Processes with R*, John Wiley & Sons Inc., 2016. **(reference for Markov Chains)**

2.) Sheldon M. Ross, *Applied Probability Models with Optimization Applications*, Dover Books on Mathematics **(reference for Markov Decision Processes)**

3.) Richard S. Sutton, Andrew G. Barto *Reinforcement Learning: An Introduction*, The MIT Press **(Standard textbook on RI)**