Ph.D. Seminar on Advanced Dynamic Programming

This six-ECTs course will be taught by Prof. Xie (jingui.xie@tum.de) at TUM Campus Heilbronn. We will be leveraging hybrid teaching so that Ph.D. students from Munich may join online. This course takes place in the summer semester of 2021.

1. Objectives

Dynamic programming is an optimization approach that has been widely applied in operations management. In this course, the students will study the advanced models in dynamic programming systematically. At the end of this course, students will be able to understand the mathematical derivation of dynamic programming, model the sequential decision-making problems, and find the optimal solution proficiently.

2. Structure

There will be an intense 180 minutes session discussion each week. During this time, we are going to discuss a specific topic in dynamic programming in the greatest depth. In the kickoff meeting, the topics will be assigned to the student. One student will present the topic and lead the discussion. Each student must read, study, and prepare each session thoroughly. The following topics will be covered:

- Dynamic programming
- Approximate dynamic programming
- Markov decision process
- Partially observable Markov decision process
- Reinforcement learning
- Robust dynamic programming
- Other topics related to advanced dynamic programming

The textbooks include and are not limited to the following:

- Dimitri P. Bertsekas. Dynamic Programming and Optimal Control (Vol. I and II, 4th Edition), Athena Scientific (2012)
- Martin L. Puterman. Markov Decision Processes. Wiley-Interscience (2005)
- Warren B. Powell. Approximate Dynamic Programming: Solving the Curses of Dimensionality, (2nd Edition), Willy (2011)
- Vikram Krishnamurthy. Partially Observed Markov Decision Processes. Cambridge University Press (2016)
- Sutton R.S., Barton A.G. Reinforcement learning: An introduction, The MIT Press (2018)

3. Subject area and scope

This seminar is open for all Ph.D. students at TUM School of Management who are interested in advanced dynamic programming, Markov decision process, reinforcement learning and other relevant topics. Literature, including textbooks and papers, might be discussed.

4. Prerequisites

The module requires solid knowledge in advanced mathematics, especially the knowledge of probabilities, Markov chains, and so on. Mathematical maturity and the ability to write down precise and rigorous arguments are also important.

5. Grading

The students need to host at least one session to present the topic and lead the discussion. Everyone should read relevant literature after class and participate actively in the discussion. Students who show up unprepared and/or unable to contribute actively may be considered absent.