Markov Decision Processes and Reinforcement Learning

This version: (First official draft)

Course instructors

Name: **G.P. Kiesmüller**
Mail: office.cdt@mgt.tum.de

Course aims and objectives

Many real life systems are subject to uncertainty and should therefore be modelled with stochastic models. In this course, we focus on the theory and the application of Markov Decision Processes and Semi Markov Decision Processes. The students should gain knowledge about these models such that they are able to construct these models and apply them to solve real life problems. For illustration, we use among others, models of inventory systems, manufacturing systems and maintenance systems. We practice to derive the Bellmann equation for these systems and show how an optimal solution can be computed numerically. Besides the traditional solution approaches, we also discuss approaches based on reinforcement learning.

Content

We first define all elements of a discrete time Markov Decision Process and show how dynamic systems can be modelled with the Bellman equation. We discuss the existence of an optimal stationary policy and show how it can be computed numerically. We distinguish between the expected total discounted reward criteria and long-run average costs. We discuss policy iteration, value iteration, linear programming, and reinforcement learning algorithms to derive optimal stationary policies.

Application procedure

Prerequisits

Participants should have basic knowledge in probability theory and Markov Chains (Random variables, discrete and continuous distribution functions, conditional distributions, moments of random variables, discrete time Markov Chains, Continuous Time Markov Chains). Additionally, students should have basic programming experience.

Application process

PhD Students interested in the course can send an e-mail to office.cdt@mgt.tum.de

Preliminary schedule

The course will take place in September/October 2023

Kick-off meeting 90 minutes

1 days traditional solution methods for Markov Decision Processes

1 days reinforcement learning approaches for Markov Decision Processes

1 day project presentation and discussion

Core readings

A first course on stochastic models. Chapter 6 and 7. (2003). HC Tijms. Wiley

Markov Decision Processes. (2005). ML Putermann. Wiley

Reinforcement learning. (2020) . RS Sutton and AG Barto

Assessment

Students have to participate actively in the exercises and have to present a project related to the course topics.

Workload

3 ECTS (22.5 hours lectures, 90 hours total workload)