

Large-Scale and Data-Driven Markov Decision Processes

Professor Wolfram Wiesemann

Imperial College Business School

Monday 9 October 2023

1300hrs

LR7, IEB

Abstract:

Markov decision processes (MDPs) constitute one of the predominant modelling and solution paradigms for dynamic decision problems affected by uncertainty. MDPs model the dynamics of a system through a random state evolution that generates rewards over time. The decision maker aims to select actions that influence this state evolution so as to maximize rewards.

In this talk, we review recent advances in MDPs along two directions: *(i)* the construction of data-driven policies that combine the (traditionally separated) tasks of estimating the system's behaviour and selecting actions that maximize rewards in the estimated system, and *(ii)* the exploitation of structure to solve large-scale problems. In view of *(i)*, we will show how the consideration of data-driven policies naturally leads to the study of robust MDPs, where the decision maker combats overfitting by hedging against the worst system dynamics that are plausible under some given training data. We will also discuss how alternative models of robustness offer different trade-offs between the competing goals of out-of-sample performance and complexity of the involved policies and computations. As for *(ii)*, we will review two types of structure that allow us to alleviate the well-known curse of dimensionality: weakly coupled MDPs that combine a potentially large number of MDPs via a small number of linking constraints, and factored MDPs whose states are represented by assignments of values to state variables that evolve and contribute to the system's rewards largely independently.