Introduction to Robust Optimization

Wolfram Wiesemann
Imperial College London, UK

Abstract

Traditionally, uncertainty-affected decision problems are solved by modelling the uncertain problem data as random variables and subsequently discretizing the outcomes of these random variables. Although this is a very natural approach, it has several shortcomings: it requires the exact specification of the underlying stochastic process (which is rarely available in practice), and it results in a curse of dimensionality for dynamic (multi-stage) problems, which implies that the computation times grow exponentially with problem size. In this lecture, we review the rapidly growing literature on robust and distributionally robust optimization, which aims to alleviate the aforementioned shortcomings. A robust optimization problem specifies an uncertainty set that contains all possible values for the uncertain problem parameters, and it seeks the best decision in view of the worst parameter realization. A distributionally robust optimization problem, on the other hand, specifies an ambiguity set that contains all possible probability distributions that could govern the uncertain problem parameters, and it seeks the best decision in view of the worst probability distribution. Topics covered include the reformulation and solution of static and dynamic (distributionally) robust optimization problems as well as discrete robust optimization.

Note: To participate, a working knowledge of linear programming is required.