

From the knapsack problem to drawing metro maps - mixed integer linear programming in R

Dirk Schumacher

`mail@dirk-schumacher.net`

Keywords: optimization, modelling.

Combinatorial optimization problems can be found in many places: finding the optimal seating plan for you and your coworkers, designing a conference schedule or setting up facilities in an emergency situation. In general, these problems can be modelled using an objective function and a set of constraints that define valid solutions.

For example, we want to find the optimal conference schedule while taking into account time preferences of the speakers and room capacities based on predicted attendance. Given an implicit set of all possible conference schedules, we try to find one schedule that maximizes our objective function (e.g. tries to satisfy speaker preferences as much as possible). We also would like to be certain that the found solution is indeed optimal. It turns out, many of these (real world) optimization problems can be naturally formulated as a special problem class, called a mixed-integer linear program (MILP). As the name suggests, the aim is to optimize a linear objective function, subject to a set of linear inequalities with some of the variables being integer valued. Once able to formulate the problem as a MILP, we can use specialized open-source and commercial solvers that have been developed over the past decades to efficiently solve it to optimality.

My talk will be about modelling and solving combinatorial optimization problems with R in general and about mixed integer-linear programming in particular. As such I will introduce the theory of mixed-integer linear programming from an applied perspective and review the packages currently available to model and solve these optimization problems. In addition, we will also look further into applications and other modelling and optimization techniques offered by R packages.

With this talk I hope I can show that R is a great language to model and solve combinatorial optimization problems while being able to draw from state of the art methods in statistics, machine learning, graphics and reproducible research. With the right tools, it is surprisingly straight forward to solve real world problems. However, as the members of the operations research community, from where many applied developments come from, have not been traditional R users, there are a lot of exciting opportunities for future developments!