Combinatorial optimization under uncertainty: matroids, matchings and submodular functions

The field of approximation algorithms for combinatorial optimization under uncertainty is extensively studied for more than a decade, and it still brings more attention to it. The most important feature of such problems is the fact that we operate on data that is not fully known in advance.

This interest is mainly due to two reasons. First is the progressing digitization of our lives, which results in collecting vast amount of data. And thanks to it complex systems can be build that have to make nontrivial algorithms that will operate only on predictions infered from such data.

The second reason are internet environments where we have to operate fast on enormously big streams of data. In the internet advertisement millions of requests for an ad impression will come during the day, and we have to make fast and irrevocable decisions without the privilage of waiting for the full set of requests.

Due to the increasing importance of such problems both theoretical and practical research in this area is rapidly progressing. And there is a true need for development of theoretical ideas which later are being used as conceptual building blocks for the design of systems that have to serve in rigorous industrial and business applications.

In this proposal we focus on the theoretical foundations of optimization under uncertainty. Let us first present the most canonic examples of problems from this domain.

Secretary problem In its original form, the task is to select the best out of a set of many secretaries. Secretaries appear one by one in a uniformly random order. Whenever a secretary appears, he can be compared against all previously appeared secretaries. Then, the algorithm has to decide, before the arrival of the next secretary, whether to select the current secretary or not. A well-known classical algorithm from 1963 selects the best secretary with probability atleast $1/e =\approx 0.37$, and this is known to be optimal.

Prophet inequality Imagine a gambler who is going to play a known to him game many times, and hence sees a sequence of outcomes of these games. After seeing a prize of one game, the gambler can either accept the prize and the game ends, or reject the prize and the next prize is presented to him. The classic result from 1978 states that the gambler can obtain at least half of the expected reward that a prophet can make who knows the realizations of the prizes in advance.

Ski rental Imagine we are going skiing for an unknown number of days (you do not know the exact number due to various reasons, e.g., loss of interest, accidents that break your legs, or extremely bad weather). Suppose that renting skis costs \$1 per day and buying skis costs \$10. Every day we have to decide whether to continue renting skis for one more day or buy a pair of skis. If we would know in advance for how many days we would go skiing, we could decide easily whether we should rent or buy — if less than 10 days, then rent; if 10 or more days, then buy. But how can we minimize the cost, when we do not have this knowledge in advance? It is not hard to see that any reasonable strategy will spend at most \$20 (rent for 10 days and then buy), but a surprising result is that a randomized strategy can spend at most 16.

In the proposal we plan to investigate problems of the above nature, but in way more complex scenarios that naturally come from the way computers help our lives nowadays.