Bayes-Optimal Optimization

Caspar Österheld

Optimization is one of the most general problems in computer science and mathematics. In its most general form, the problem consists of finding an input x to some mathematical objective function f such that the corresponding output f(x) is as big as possible. I also assume the function to be treated as a blackbox and that optimization is done by an iterative process of trial and error. The problem of optimization has received wide attention and a lot of algorithms can be found in the literature. Although the no-free-lunch theorems show that optimization algorithms can only work well under certain, usually probabilistic, assumptions, general optimization algorithms usually don't make these assumptions explicitly nor are they parametrized by them. Overall, they tend to be proven in practice rather than being based on solid theoretical foundations.

In my Bachelor's thesis I develop an optimization algorithm, which uses a prior probability distribution over objective functions and can be shown to always choose the Bayes-optimal trials. I analyze the algorithm's time and space complexity to show that while it is not applicable to the vast majority of realworld optimization problems, it can be used for small dummy problem classes to benchmark other optimization algorithms. I then go on to compare the performances of the Bayes-optimal with a Bayesian optimization algorithm called expected improvement maximization as well as a simple hill climbing heuristic on various problem classes. Such comparisons are hoped to reveal insights into why and when heuristics work, both in optimization and artificial intelligence in general.