Popular science summary

Anton Åkerman

June 2023

Optimization is used in everything from machine learning and finance to signal processing, and it revolves around finding optimal choices for problems, given some criteria and limitations.

For solving optimization problems, optimization algorithms usually start at an initial choice and improve it, one iteration at a time. This is repeated until we hopefully have found a choice that is close to optimal. Important questions to ask are: how much can we guarantee that our choice improves in each iteration, and how can we tune our optimization algorithm so that we approach the optimal solution as much as possible in each iteration, in the worst case? These are questions we need to answer in order to know which algorithms to use for which problems, and in order to have secure guarantees for how long optimization algorithms will need to run or how close to the optimal solution we will end up after a certain number of iterations with our algorithm.

This master's thesis tries to combine two different approaches from the field of optimization, in order to tune the parameters of optimization algorithms. The method we propose uses a recent framework for finding these convergence guarantees called the Performance Estimation Problem, and combines it with derivatives. By calculating derivatives, we can find in which directions to modify the parameters we are allowed to tune. This means the tuning can be made faster or more accurately, at least if there are many parameters to tune, and although tuning parameters is a challenging problem, the method works decently for some simpler examples. A main advantage of this method is that it scales well and can handle optimization algorithms with many parameters to tune, but testing for these algorithms with more parameters is left for future work.

We also present a minor result regarding the possible structure of a certain type of optimization algorithms.