Tuning the global optimization solver BARON using derivative-free optimization algorithms

Nikolaos Ploskas, Jianfeng Liu, Nikolaos V. Sahinidis

Optimization solvers provide users with numerous options that control different algorithmic aspects. These options can have a significant impact on solver performance. Therefore, tuning solver options is often necessary and leads to significant performance improvements. The aim of this work is to identify option settings that result in the best solver performance in terms of execution time and solution quality. Tuning options can be regarded as an optimization problem. This problem is hard to solve for two reasons. First, the relationship between the parameters and solver performance is not explicit. Second, some of the options may take discrete values, so the objective function is complex and non-smooth. Hence, the solver must be treated as a black-box system, whose input is values for the different options and output is a performance metric, such as the execution time. Derivative-free optimization algorithms (DFO) are attractive for this tuning problem since they do not require explicit functional representations of the objective function. We perform a computational study over a set of 126 problems from GLOBALLib and MINLPLib collections in order to identify optimal values for each one of the problems and also find a single set of options that can improve the performance of BARON across the entire test collection. A total of 27 DFO algorithms are used for this reason. We present extensive computational results.