From benchmarking optimization heuristics to dynamic algorithm configuration

Vermetten, D.L.

Summary

Within the context of optimization, benchmarking is sometimes viewed as a bridge between theory and practice. Its empirical nature allows for the creation of large-scale experiments to investigate a wide variety of questions, which can be intractable using theoretical approaches. At the same time, benchmarking does not have to fit neatly into a specific application domain but can cover a broader range of fundamental questions. This combination allows benchmarking studies to provide a range of new insights into the strengths and weaknesses of different optimization heuristics.

Given the potential benefits of robust benchmarking setups, it is critical that tools are created which lower the barrier to entry for its use. These setups should be flexible enough to fit the wide variety of questions being asked while providing rigorous frameworks for the rest of the benchmarking pipeline. In this thesis, we focused on the IOHprofiler framework, which aims to be such a tool for the wider research community. Using a modular design structure, IOHprofiler can be integrated with many commonly used tools for black-box optimization and benchmarking. These integrations form the backbone for the experiments used throughout this thesis, as the corresponding performance and algorithm behavior data can be easily reused. This way, we can go beyond the common 'competitive' benchmarking practice, where we only care about the algorithm with the best average performance, to gaining insights about the complementarity between different algorithms.

One of the key ways in which algorithm complementarity can be beneficial is in the context of algorithm selection and configuration. Instead of relying on a single algorithm for a wide set of problems, we use a set of features to determine which algorithm or algorithm configuration to use. This meta-learning task benefits from variety between algorithms, which can be achieved by looking at different types of algorithms, but also within an algorithm family by considering a wide range of modifications proposed over the last decades. We show how two modularizations of popular evolutionary algorithms, CMA-ES and Differential Evolution, can lead to new insights about combinations of algorithmic ideas, resulting in improved performance over previous hand-designed versions of these same algorithms.

While algorithm complementarity exists on the problem-level, looking at benchmark data from a variety of sources also reveals a complementarity in performance within individual functions. While some algorithms are great at finding promising regions, others excel at fast convergence once this region is found. As such, the notion of dynamic algorithm selection, where we can switch between different algorithms during the optimization procedure, is studied in detail. We highlight the inherent potential in this approach, while simultaneously highlighting the aspects of this switching approach which need to be further developed to create reliably dynamic algorithm combinations.

Dynamic algorithm selection is a promising field of research within optimization, but current studies into its performance are still somewhat limited. This is largely a result of the difficulty of creating reliable benchmarking setups for this scenario. While collections of problems for benchmarking optimization algorithms are widely available, they are not particularly suited for meta-learning scenarios, as there is no truly fair way to create the commonly required train-test set distinction used by the required machine learning techniques. For this reason, we propose a benchmark problem generator based on commonly used black-box optimization problems, which can be used to generate arbitrary amounts of training and testing problems to benchmark these meta-learning mechanisms.