Abstract

The optimization of computationally expensive black-box problems is a challenge faced in many application fields. Those problems are characterized by the lack of information about their landscape and their computational cost. For instance, in engineering design, the objective function frequently results from complex numerical simulations and only its output is available. We investigate two major ways of dealing with such problems. First, we rely on Machine Learning to build surrogate models that approximate the expensive objective function at a lower computational cost. Second, we use parallel computing to reduce the computational burden of the optimization process. The major research question addressed in this thesis is how surrogate modeling and parallelism can help to efficiently and effectively sample the design space.

We distinguish two main approaches of using the surrogate model inside an optimization process. In Surrogate-Driven Optimization (SDO), the surrogate model actively drives the process through the definition of an Acquisition Function (AF) that evaluates the promisingness of a candidate decision. This is typically the case in Bayesian Optimization (BO), where Gaussian Process (GP) surrogate models are used. The Acquisition Process (AP) points out the most valuable candidates through the optimization of the AF. Alternatively, Surrogate-Assisted Optimization (SAO) uses the surrogate model to discard unpromising candidates and/or to partially replace the objective function. In SAO, the candidates are generated by external operators and filtered out using an Evolution Control (EC) based on the surrogate model.

In this thesis, we focus on the AP of BO Algorithms (BOAs) and its challenge of providing a valuable batch of candidates to be exactly evaluate in parallel. Firstly, the scalability of BO is limited when considering only the parallel evaluation of the candidates. Indeed, the sequential parts of the algorithm tend to be also computationally expensive. Secondly, the effectiveness of the batch acquisition is lesser compared to the sequential selection. An efficient use of parallel computing necessitates an efficient AP and smartly allocating the overall time budget to the AP (including the model fitting) and the simulations. We propose a new approach introducing parallel computing into the AP by leveraging space partitioning. From this, we derive two algorithms, namely Binary Space Partitioning Efficient Global Optimization (BSP-EGO) and Local models BSP-EGO (ℓBSP-EGO). The first one uses a global model to guide the optimization while the second one sets up multiple local models inside the sub-regions. The two developed algorithms are confronted with recent state-of-the-art BOAs using very different APs (*e.g.*, using trust regions, multiple AFs, etc.). We demonstrate the better scalability and batch effectiveness of the BSP-EGO-based algorithms compared to other BOAs.

We also compare the BO approach to other Surrogate-Based Optimization (SBO) algorithms, more precisely to Surrogate-Assisted Evolutionary Algorithms (SAEAs). These latter are usually more time-efficient since the acquisition of the candidates does not require an expensive surrogate model. The experimental protocol involves both benchmark functions and real-world engineering problems. It is designed to identify which approach is the most suitable depending on the operational constraints. The results indicate that BOAs are extremely sample-efficient, providing good outcomes with a few simulations. However, they are generally hampered by their expensive AP (including the model fitting). When the computational budget is higher, either because the simulator is computationally cheap, or because the time frame is large enough, SAEAs are to be preferred. We also demonstrate that both approaches can be combined into a hybrid algorithm benefiting the sample-efficiency of BO and the time-efficiency of SAO.