Optimizing solvers for real-world expensive black-box optimization with applications in vehicle design

Long, F.X.

Chapter 6

Conclusions and Future Work

The contributions of this thesis serve as a significant step in the long journey towards our ultimate vision, i.e., developing an automated optimization pipeline for efficiently solving real-world expensive BBO problems. A summary of this thesis is provided in Section 6.1, followed by a list of future work in Section 6.2.

6.1 Conclusions

Efficiently solving real-world BBO problems w.r.t. real-world constraints, such as wall-clock time and computational cost, has gained increasing attention over the years. Unlike solving benchmark functions, such as the popular BBOB functions, the function evaluation budget available for real-world applications is typically rather limited, particularly when the function evaluation is expensive. In the automotive industry, for instance, the quality of vehicle designs is nowadays commonly evaluated using time-consuming and costly FE simulation runs. Following this, the development of effective optimization approaches for real-world applications is becoming a pressing issue, such as finding better optimization solutions, within a shorter time duration, and/or using less computational resources.

Motivated to fill the gap, we propose an automated optimization pipeline for (near) optimally solving real-world expensive BBO problems within the scope of this thesis. In this regard, we propose considering a set of scalable and cheap-to-evaluate representative functions that belong to the same optimization problem classes as the BBO problems for the fine-tuning of optimization configurations w.r.t. some real-world constraints. Subsequently, optimization configurations can be properly fine-tuned at a relatively low cost, prior to the optimization runs using expensive function evaluations. Eventually, the optimal configurations identified can be applied for an efficient solving of the expensive BBO problems. Throughout this thesis, we consider automotive crashworthiness optimization as a representative real-world expensive BBO problem, which requires dealing with complex functions and expensive function evaluations.

Firstly, we begin our investigation by analyzing the optimization landscape characteristics of real-world expensive BBO problems in terms of ELA features in Chapter 3, which is still lacking according to the best of our knowledge. Based on our investigations on 20 automotive crash problem instances of different crash scenarios and problem dimensionalities, we show that the optimization landscape characteristics of these automotive crash problems are different from those of the BBOB functions. In other words, **the automotive crash problems belong to optimization problem classes that are different from the BBOB functions (*RQ1*)**. Subsequently, the automotive crash problems are insufficiently represented by the BBOB functions. Beyond that, we investigate a feature-free approach to characterize the optimization landscape of BBO problems based on some latent representations, which are computed using deep NN model like VAE. Interestingly, our results reveal that such latent space representations could potentially complement the classical ELA features in better capturing the optimization landscape of BBO problems.

Since the BBOB functions are insufficiently representative for real-world expensive BBO problems, we shift our focus towards generating test functions that belong to the same optimization problem classes in Chapter 4. **Using a tree-based random function generator, RGFs with similar optimization landscape characteristics in terms of ELA features can be identified for the automotive crash problems (*RQ2*)**. More importantly, **these similar RGFs can be exploited for estimating the actual performance of optimization configurations, and thus, for identifying optimal optimization configurations on expensive BBO problems (*RQ3*)**. Following this, we propose to consider RGFs as scalable and cheap-to-evaluate representations of real-world expensive BBO problems for HPO purposes. Given that a similar RGF could not be identified in some cases, e.g., for some of the BBOB functions, we additionally investigate the potential of GP in guiding the function generation towards specific optimization problem classes based on ELA features. **While some potential of this GP-based function generator can be observed, this approach is not as straightforward as expected (*RQ4*)** and substantial work is necessary to vigorously guide the function evolution, such as an improved feature selection strategy.

Based on the findings in previous chapters, we evaluate the performance of the proposed optimization pipeline for real-world expensive BBOB problems in Chapter 5. When tested on 24 BBOB functions in 20-d, we show that better optimization configurations can be indeed identified based on some representative functions, which can outperform the default configuration on most BBOB functions and even compete against the SBS in some cases. Moreover, the proposed optimization approach has a high flexibility, as it can work well with different combinations of BBO algorithms, e.g., ModCMA, ModDE, and BO, and optimizers for HPO, e.g., TPE and SMAC. More crucially, our results show that **the performance of BO can be improved using optimal configurations identified by our optimization approach, when solving a real-world automotive crashworthiness optimization problem using two load cases, namely a single pole impact and a multi-pole impact (*RQ5*)**. Especially for solving the multi-pole problem, a significant improvement in optimization performance can be achieved using our optimal BO configurations, which clearly outperform the state-of-the-art RSM and SRSM. Subsequently, we are confident that our optimization approach can be applied for an optimal solving of real-world expensive BBOB problems that require dealing with complex functions, while using a limited function evaluation budget. Apart from that, we evaluate the performance of predictive models in identifying optimal configurations for real-world expensive BBO problems, using deep NN models and RGFs as training dataset, similar to a landscape aware HPO context. Based on our investigation on the BBOB functions, **the trained NN models can perform quite well in identifying configurations that can most of the time outperform the default configuration (*RQ6*)**.

In summary, the proposed automated optimization pipeline has a motivating potential for an optimal solving of real-world expensive BBO problems, such as automotive crashworthiness optimization problems. Considering that the proposed approach can perform rather well on the BBOB suite, which covers a wide range of optimization problem classes, we believe that our approach can generalize to other BBO domains that are sufficiently represented by the BBOB functions. For instance, ship design problems in the maritime industry [25] and turbomachinery designs in the aerospace industry [103] are two potential real-world expensive BBO problems that are attractive for an application of our approach with appropriate modifications and extensions.

6.2 Future Work

During our investigations, nevertheless, the following weaknesses and limitations in our proposed optimization approach have been identified, where further analysis and improvements could be worthwhile:

Outlook 1: Within the scope of this thesis, we focus on optimally solving unconstrained single objective optimization problems. To improve the overall applicability of our proposed optimization approach for real-world applications, a proper extension towards constraint handling and/or multi-objective optimization would be valuable;

Outlook 2: Since some DoE samples of real-world expensive BBO problems are required as input, it is essential to investigate an optimal trade-off between the DoE sample size and effectiveness of our approach w.r.t. the problem dimensionality. Ideally, expensive function evaluations are performed only as much as necessary, but as little as possible to minimize computational effort. Meanwhile, an application of the proposed optimization approach for solving real-world BBO problems like automotive crash problems using a full vehicle FE model could be interesting, to properly estimate its value for industrial applications;

Outlook 3: In this thesis, we focus on evaluating the potential of fine-tuning optimization configurations using HPO based on some representative functions. Subsequently, another attractive research direction could be extending our approach towards CASH or AAC, where both the choice of optimization algorithms and their hyperparameters are properly fine-tuned;

Outlook 4: Based on our investigations, we show that BO indeed can perform well in solving real-world expensive BBO problems using a limited function evaluation budget. To further improve the performance of BO, for instance, the development of a configurable BO framework that is similar to ModCMA and ModDE could offer a great functional flexibility for real-world applications. In this regard, the BO hyperparameter search space could be expanded, e.g., including the hyperparameters of internal optimizer in BO and/or combining different BO variants (Section 2.1.1); and

Outlook 5: While a diverse set of RGFs can be generated using the tree-based random function generator (Section 4.1), the optimization problem classes that are being covered is primarily limited by the predefined pool of mathematical operands and operators. Subsequently, further effort could be invested to improve the diversity of RGFs that can be generated, to cover a broader spectrum of optimization problem classes. The discontinuity in real-world problems, for instance, could be potentially included by considering step functions, as briefly mentioned in Section 4.1.2.