AI for expensive optimization problems in industry

Stein. N. van; Winter, R. de; Bäck, T.H.W.; Kononova, A. V.

2023 IEEE Conference on Artificial Intelligence (CAI) | 979-8-3503-3984-0/23/$31.00 ©2023 IEEE | DOI: 10.1109/CAI54212.2023.00113

AI for Expensive Optimization Problems in Industry

Niki van Stein
LIACS, Leiden University
Leiden, The Netherlands
b.van.stein@liacs.leidenuniv.nl
0000-0002-0013-7969

Roy de Winter
LIACS, Leiden University
Leiden, The Netherlands
r.de.winter@liacs.leidenuniv.nl
0000-0002-0865-4709

Thomas Bäck
LIACS, Leiden University
Leiden, The Netherlands
t.h.w.baeck@liacs.leidenuniv.nl
0000-0001-6768-1478

Anna V. Kononova
LIACS, Leiden University
Leiden, The Netherlands
a.kononova@liacs.leidenuniv.nl
0000-0002-4138-7024

***Abstract*—The optimization of real-world engineering problems can be a challenging task, owing to the limited understanding of problem characteristics and the high cost of evaluating objectives and constraints in terms of computing time or licenses. This study proposes an AI-assisted optimization pipeline that addresses these challenges by using proxy functions in order to select and optimize the optimization algorithm and its hyper-parameters, thereby significantly accelerating the optimization process on the real (expensive) problem. These proxy functions are inexpensive to evaluate and are selected to exhibit similar landscape characteristics as the original problem. To obtain such proxy functions, we adopt an approach, which involves computing Exploratory Landscape Analysis (ELA) features to characterize the problem's landscape. The ELA features are then used to identify an artificial function that replicates the original problem's properties, which can then be employed as a low-cost proxy function for the hyper-parameter optimization of our pipeline. Several real-world industrial applications are discussed as use-case of our proposed approach.**

***Index Terms*—AI Assisted Optimization, Surrogate Based Optimization, Industrial applications**

I. Introduction

The identification and selection of the most efficient algorithm for a specific black-box multi-objective optimization problem pose a significant challenge. This task, referred to as the algorithm selection problem (ASP) [1], has been addressed in recent evolutionary computation works that focus on landscape-aware automated algorithm selection using machine learning approaches [2]–[7]. These approaches utilize landscape features, which quantify different landscape characteristics of a given problem instance, to predict the performance of an optimization algorithm on such a problem instance.

In essence, once the landscape characteristics of a problem instance have been identified, the performance of an optimization algorithm on an unseen instance can be roughly estimated. This is advantageous, as landscape analysis can provide additional insights into the effectiveness of an algorithm across different problem instances [8]–[11].

However, previous ASP works have primarily focused on academic benchmark functions, such as the black-box optimization benchmarking (BBOB) test set [2], [5], [10]. Conversely, little research has been done to investigate algorithm selection in real-world expensive black-box optimization problems, especially for constrained multi-objective optimization problems where multiple objective and constraint functions can be computationally expensive. Addressing this challenge remains an open research area.

In this work, an automated AI-assisted pipeline for selecting and tuning an optimization algorithm is proposed. The pipeline uses cheap to evaluate proxy functions for the algorithm selection and hyper-parameter tuning. These proxy functions are found using the methodology proposed by Long et al. [12]. In this approach, Exploratory Landscape Analysis (ELA) [13] features are computed and utilized to select appropriate generated artificial functions such that these functions closely resemble the optimization landscape. These cheap to evaluate artificial functions can then be used to emulate the expensive objectives and constraint functions of the real-world engineering problem.

II. Proposed pipeline

The proposed automated optimization pipeline for real-world expensive black-box optimization problems is visualised in Figure 1. The pipeline starts with an initial Design of Experiments (DoE) sample $\mathcal{X} = \{\mathbf{x}_1, \ldots, \mathbf{x}_n\}$, $\mathbf{x}_i \in \mathcal{S} \subset \mathbb{R}^d$ and the corresponding objective function values $\mathcal{F} = \{f(\mathbf{x}_1), \ldots, f(\mathbf{x}_n)\}$ as obtained by (time-consuming) simulation of the real-world problem (Step 1). Based on the landscape characteristics (ELA features) of the given problem instance (Step 2), which are then computed on the sample $(\mathcal{X}, \mathcal{F})$, the objective of the pipeline is to identify representative, fast-to-evaluate functions with similar characteristics (Step 3). At the moment, this is done by using a random function generator by which we generate thousands of random functions from which we select one that minimizes the Euclidean distance between its ELA feature vector and the target feature vector [12]. These fast-to-evaluate functions are used to select and configure the optimal optimization algorithm

979-8-3503-3984-0/23/$31.00 ©2023 IEEE
DOI 10.1109/CAI54212.2023.00113

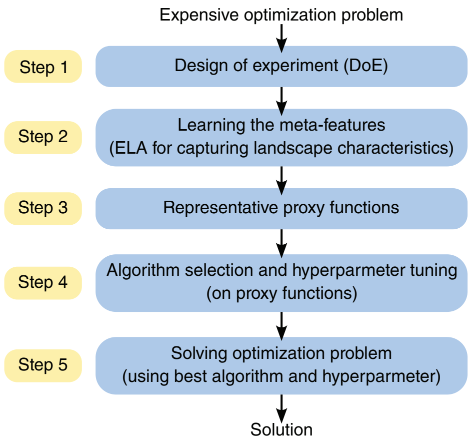

Fig. 1. Proposed pipeline using ELA features to find similar proxy functions and to select and hyper-parameter optimize the best optimization algorithm for solving a given expensive black-box optimization problem.

$\mathcal{A}$ for this problem (Step 4). The final step involves applying the optimally fine-tuned optimization algorithm $\mathcal{A}$ on the real-world problem instance (Step 5).

III. Industrial Applications

In this section, we discuss three examples from the automotive and maritime industries which are representative of costly real-world optimization applications, with objective function evaluation requiring up to several hours of computation time.

A. Crashworthiness Optimization

The design of a car must meet several safety requirements, including providing sufficient protection to passengers in case of a crash, while simultaneously fulfilling other necessary design parameters such as durability and weight [14]. With car designs becoming increasingly sophisticated, crashworthiness optimization has become a challenging and tedious task. Moreover, simulating a highly nonlinear crash through Finite Element analysis is computationally expensive, thereby severely limiting the number of affordable function evaluations for optimization. Despite recent advancements in developing surrogate-based optimization methods for automotive crashworthiness problems [15]–[22], there remains a dearth of understanding of problem characteristics.

To address this, a more comprehensive and thorough understanding of the problem characteristics is necessary. The existing literature primarily focuses on surrogate-based optimization methods; however, a holistic approach to the problem is lacking. With the proposed pipeline, the optimization problem can be emulated with artificial functions to fine-tune and select a well-performing optimization algorithm using a limited evaluation budget before applying the optimized algorithm to the expensive objective function. The advantage of using artificial functions instead of surrogate models is that we *do not* impose any assumptions on the optimization landscape, while surrogate models are heavily subject to the assumptions by the type of surrogate model used. The proposed pipeline would allow for more efficient and effective crashworthiness optimization.

We have illustrated the value of using machine learning approaches for training a response surface and then optimizing on the response surface in recent work on multidisciplinary optimization (MDO) for crashworthiness [23]. The resulting algorithm combines several techniques such as a global sensitivity matrix (showing the relevance of each design parameter for each of the disciplines [24]), optimization algorithms using a local prediction uncertainty measure [25], methods for reducing the finite element computation effort, and an adaptive model complexity control approach. On a full vehicle body design optimization example, the approach is able to reduce car body weight by 17.75 kg, finding a practically feasible solution with a reduction of the computational effort by 64.55%, in comparison to the state-of-the-art MDO process [23].

B. Vehicle Dynamics Optimization

The dynamics of modern vehicles is largely dictated by control systems dynamically regulating aspects such as engine control, active suspension, rollover prevention, etc. Such control systems are built based on complex general models implemented by vehicle manufacturers which capture a multitude of highly interdependent parameters. During the design stage, such models need to be optimised (in a single- or multi-objective fashion) to meet various performance, safety, cost and environmental requirements given the constraints for a vehicle model. However, due to the computational complexity of these models, their optimisation is far beyond the capabilities of traditional methods. This makes such problems highly suitable for the approach discussed in Section II.

As an example, anti-lock braking and variable damper control systems (ABS and VDS, respectively) which are responsible for driving safety depend on hundreds of parameters of a vehicle that cover different driving situations [26]. While the former controls the relative motion between a tire and a road during braking by adjusting the brake pressure, the latter manipulates braking force by adjusting damper constants of the shock absorbers. The approach discussed in Section II has been successfully used to find a set of ABS and VDC parameters that deliver minimal (simulated) braking distance [26].

Based on ELA features computed for a DoE, a representative function has been identified from a set of benchmarking functions. On this representative function, an investigation has been carried out to identify the best method to construct a surrogate model (Random Forest, Support Vector Machine or Multi-layer Perceptron) which approximates the objective function, i.e., we investigated on a fast-to-evaluate proxy function which machine learning algorithm was able to represent the proxy function - and therefore the original real-world problem - well and could serve as a good predictor of the optimum. The validity of such a choice of surrogate has then been verified for the expensive objective function of the aforementioned ABS and VDS optimisation problem. Results shown in Figure 2 demonstrate the validity of such an

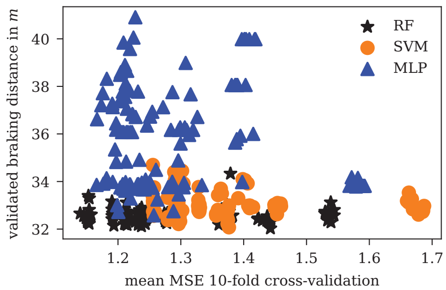

Fig. 2. Ninety different response surface models (three machine learning algorithms - random forest RF, support vector machines SVM, multi-layer perceptron MLP, ten runs each, three different DoE sizes) were trained as proxy functions and optimization using a CMA-ES was run on the response surfaces. The plot shows the quality of verified (using the simulator) final optimization results (y-axis) over the quality of the response surface model, measured as mean squared error (x-axis).

approach. The x-axis shows the mean squared error obtained through 10-fold cross-validation of the machine learning algorithm for ten different runs and three different DoE sizes, and the y-axis show the validated objective function value associated to the optimum that was obtained by optimizing on the response surface with a Covariance Matrix Adaptation Evolution Strategy (CMA-ES) [27].

Results show that MSE is not a good indicator of the suitability of the response surface model for optimization. However, despite relatively large variation in MSE values, good braking distance values close to 32 m can be obtained based on SVM and RF. The value of 32 m is a *new best* (simulator-based) solution for the corresponding vehicle dynamics problem; this solution is still to be evaluated in the field on a real physical vehicle. It should be noted here that thanks to the approach described here, we can operate at four levels of abstraction: (i) the real car, (ii) its simulation, (iii) the proxy function for the simulation, (iv) a response surface model that was trained to approximate the proxy function (to see whether a response surface model can, in principle, approximate the simulated problem well enough).

A similar approach has been subsequently taken in a follow-up paper [28], where the best parameter settings of a single-objective optimiser have been investigated in terms of a compromise between the attained breaking distance and available computing budgets, as two conflicting objectives in a multi-objective setting.

C. Ship Design

When designing a ship the aim is to find the ideal size and hull shape that fit all required cargo and technical spaces [29]. When searching for the ideal size, there are often two conflicting objectives and many constraints. Objectives that are optimized are often the steel weight (a strong indicator for required capital investments) and resistance at operational speed (a strong indicator for operational expenses). The steel weight is usually computed from a design that complies to the international rules that determine the height, thickness, and distance between the steel beams. The resistance of the vessel at operational speed however is much more computationally expensive since this is typically evaluated by performing Reynolds-averaged Navier–Stokes calculations which can take hours on high-performance clusters. The constraints in ship design are often volumetric constraints on the cargo hold, engine room, fuel tanks ect. and calculations for intact- and the computationally expensive damage-stability criteria.

To limit the number of expensive function evaluations and this way save time, an algorithm is proposed which exploits machine learning models as surrogates for the expensive constraints and objectives [30], [31]. Since not many results are available from ship design optimization problems, and since every ship design is different, the characteristics of the problem landscape are learned on the fly. The self-adaptive algorithm optimizes the hyper-parameters of the machine learning models, and selects the best fitting surrogate models based on mean squared error, while it simultaneously tries to find feasible and optimal solutions.

The self-adaptive algorithm has proven to be very effective at finding Pareto-efficient solutions which are also feasible. The algorithm is able to find better Pareto frontier approximations compared to other algorithms which only exploit one type of machine learning method to learn the constraints and objectives [31]. The algorithm has effectively been used to find an optimal offshore service vessel. Compared to the original service vessel design, the algorithm found a design with 19% less steel weight, and 10% less resistance [32].

IV. Conclusions and Outlook

We propose to tune optimization algorithms for expensive real-world optimization tasks by using proxy functions that exhibit the properties of the real-world task but are fast to evaluate, such that many function evaluations can be performed for optimizer tuning and the resulting tuned optimization algorithm can then be applied to the real-world task (and similar instances from the same problem class). As representatives, we picked two examples from the automotive industry, namely crashworthiness and vehicle dynamics optimization. In both cases, we showed that: (i) it is possible to find such proxy functions, (ii) optimization on the proxy functions or by using response surface models can be used to find better solutions faster. The latter holds for the full vehicle body optimization (reducing weight by 17.75 kg and CPU time by 64.55%) and the braking distance example (reducing braking distance to 32 m, which is an unexpected improvement over the current best value). In addition to the approach with ELA features and learning from proxy functions, an alternative self-adaptive algorithm has been used in ship design. The self-adaptive algorithm learns hyper-parameters on the fly and chooses the best surrogate model adaptively. It has been used to optimize an offshore service vessel by reducing the steel-weight by 19% and resistance by 10%.

As a next step, the automated machine learning pipeline for generating the closest proxy function for an arbitrary real-world problem will be finished. In parallel, we are developing an automated configuration and hyper-parameter optimization framework for efficient global optimization frameworks, allowing to tune such optimization frameworks for the derived proxy function. The resulting tuned optimization algorithm would then be a solver that is tuned for the requirements of the corresponding real-world problem class.

References

[1] J. R. Rice, "The algorithm selection problem," ser. Advances in Computers, M. Rubinoff and M. C. Yovits, Eds. Elsevier, 1976, vol. 15, pp. 65–118.

[2] B. Bischl, O. Mersmann, H. Trautmann, and M. Preuß, "Algorithm selection based on exploratory landscape analysis and cost-sensitive learning," in *Proceedings of the 14th Annual Conference on Genetic and Evolutionary Computation*, ser. GECCO '12. New York, NY, USA: Association for Computing Machinery, 2012, p. 313–320.

[3] P. Kerschke and H. Trautmann, "Automated algorithm selection on continuous black-box problems by combining exploratory landscape analysis and machine learning," *Evolutionary Computation*, vol. 27, no. 1, p. 99–127, 2019.

[4] J. Dréo, C. Doerr, and Y. Semet, "Coupling the design of benchmark with algorithm in landscape-aware solver design," in *Proceedings of the Genetic and Evolutionary Computation Conference Companion*, ser. GECCO '19. New York, NY, USA: Association for Computing Machinery, 2019, p. 1419–1420.

[5] A. Jankovic and C. Doerr, "Landscape-aware fixed-budget performance regression and algorithm selection for modular cma-es variants," in *Proceedings of the 2020 Genetic and Evolutionary Computation Conference*, ser. GECCO '20. New York, NY, USA: Association for Computing Machinery, 2020, p. 841–849.

[6] M. Pikalov and V. Mironovich, *Automated Parameter Choice with Exploratory Landscape Analysis and Machine Learning*, ser. GECCO '21. New York, NY, USA: Association for Computing Machinery, 2021, p. 1982–1985.

[7] A. Jankovic, G. Popovski, T. Eftimov, and C. Doerr, *The Impact of Hyper-Parameter Tuning for Landscape-Aware Performance Regression and Algorithm Selection*, ser. GECCO '21. New York, NY, USA: Association for Computing Machinery, 2021, p. 687–696.

[8] D. Simoncini, S. Barbe, T. Schiex, and S. Verel, "Fitness landscape analysis around the optimum in computational protein design," in *Proceedings of the Genetic and Evolutionary Computation Conference*, ser. GECCO '18. New York, NY, USA: Association for Computing Machinery, 2018, p. 355–362.

[9] M. A. Muñoz, Y. Sun, M. Kirley, and S. K. Halgamuge, "Algorithm selection for black-box continuous optimization problems: A survey on methods and challenges," *Information Sciences*, vol. 317, pp. 224–245, 2015.

[10] P. Kerschke, H. Hoos, F. Neumann, and H. Trautmann, "Automated algorithm selection: Survey and perspectives," *Evolutionary Computation*, vol. 27, no. 1, pp. 3–45, 2019.

[11] K. M. Malan, "A survey of advances in landscape analysis for optimisation," *Algorithms*, vol. 14, no. 2, p. 40, 2021.

[12] F. X. Long, B. van Stein, M. Frenzel, P. Krause, M. Gitterle, and T. Bäck, "Learning the characteristics of engineering optimization problems with applications in automotive crash," in *Proceedings of the Genetic and Evolutionary Computation Conference*, 2022, pp. 1227–1236.

[13] O. Mersmann, M. Preuss, and H. Trautmann, "Benchmarking evolutionary algorithms: Towards exploratory landscape analysis," in *Parallel Problem Solving from Nature, PPSN XI*, R. Schaefer, C. Cotta, J. Kołodziej, and G. Rudolph, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2010, pp. 73–82.

[14] F. Duddeck, "Multidisciplinary optimization of car bodies," *Structural and Multidisciplinary Optimization*, vol. 35, no. 4, pp. 375–389, 2008.

[15] H. Kurtaran, A. Eskandarian, D. Marzougui, and N. Bedewi, "Crashworthiness design optimization using successive response surface approximations," *Computational Mechanics*, vol. 29, no. 4, pp. 409–421, 2002.

[16] T. Jansson, L. Nilsson, and M. Redhe, "Using surrogate models and response surfaces in structural optimization – with application to crashworthiness design and sheet metal forming," *Structural and Multidisciplinary Optimization*, vol. 25, no. 2, pp. 129–140, 2003.

[17] L. Shi, R.-J. Yang, and P. Zhu, "A method for selecting surrogate models in crashworthiness optimization," *Structural and Multidisciplinary Optimization*, vol. 46, no. 2, pp. 159–170, 2012.

[18] X. Liao, Q. Li, X. Yang, W. Li, and W. Zhang, "A two-stage multi-objective optimisation of vehicle crashworthiness under frontal impact," *International Journal of Crashworthiness*, vol. 13, no. 3, pp. 279–288, 2008.

[19] M. Kiani, I. Gandikota, A. Parrish, K. Motoyama, and M. Rais-Rohani, "Surrogate-based optimisation of automotive structures under multiple crash and vibration design criteria," *International Journal of Crashworthiness*, vol. 18, no. 5, pp. 473–482, 2013.

[20] K. Hamza and M. Shalaby, "A framework for parallelized efficient global optimization with application to vehicle crashworthiness optimization," *Engineering Optimization*, vol. 46, no. 9, pp. 1200–1221, 2014.

[21] G. Sun, Y. Tian, R. Wang, J. Fang, and Q. Li, "Parallelized multi-objective efficient global optimization algorithm and its applications," *Structural and Multidisciplinary Optimization*, vol. 61, no. 2, pp. 763–786, 2020.

[22] J. Fang, G. Sun, N. Qiu, N.-H. Kim, and Q. Li, "On design optimization for structural crashworthiness and its state of the art," *Structural and Multidisciplinary Optimization*, vol. 55, no. 3, pp. 1091–1119, 2017.

[23] J. Büttner, A. Schumacher, T. Bäck, S. Schwarz, and P. Krause, "Making multidisciplinary optimization fit for practical usage in car body development," *Structural and Multidisciplinary Optimization*, vol. 66, no. 3, p. 62, 2023.

[24] J. Büttner, S. Schwarz, A. Schumacher, and T. Bäck, "Global sensitivity matrix for vehicle development," *ATZ worldwide*, vol. 123, no. 3, pp. 26–31, 2021.

[25] B. van Stein, H. Wang, W. Kowalczyk, and T. Bäck, "A novel uncertainty quantification method for efficient global optimization," in *Information Processing and Management of Uncertainty in Knowledge-Based Systems. Applications*, J. Medina, M. Ojeda-Aciego, J. L. Verdegay, I. Perfilieva, B. Bouchon-Meunier, and R. R. Yager, Eds. Cham: Springer International Publishing, 2018, pp. 480–491.

[26] A. Thomaser, A. V. Kononova, M.-E. Vogt, and T. Bäck, "One-Shot Optimization for Vehicle Dynamics Control Systems: Towards Benchmarking and Exploratory Landscape Analysis," in *Proceedings of the Genetic and Evolutionary Computation Conference Companion*. New York, NY, USA: Association for Computing Machinery, 2022, p. 2036–2045.

[27] N. Hansen and A. Ostermeier, "Adapting Arbitrary Normal Mutation Distributions in Evolution Strategies: The Covariance Matrix Adaptation," in *Proceedings of the IEEE International Conference on Evolutionary Computation*, 1996, pp. 312–317.

[28] A. Thomaser, M.-E. Vogt, A. V. Kononova, and T. Bäck, "Transfer of multi-objectively tuned cma-es parameters to a vehicle dynamics problem," in *Evolutionary Multi-Criterion Optimization*, M. Emmerich, A. Deutz, H. Wang, A. V. Kononova, B. Naujoks, K. Li, K. Miettinen, and I. Yevseyeva, Eds. Cham: Springer Nature Switzerland, 2023, pp. 546–560.

[29] R. de Winter, J. Furustam, T. Bäck, and T. Muller, "Optimizing ships using the holistic accelerated concept design methodology," in *Practical Design of Ships and Other Floating Structures*, T. Okada, K. Suzuki, and Y. Kawamura, Eds. Singapore: Springer, 2021, pp. 38–50.

[30] R. d. Winter, B. v. Stein, and T. Bäck, "SAMO-COBRA: A fast surrogate assisted constrained multi-objective optimization algorithm," in *International Conference on Evolutionary Multi-Criterion Optimization (EMO)*, H. Ishibuchi, Q. Zhang, R. Cheng, K. Li, H. Li, H. Wang, and A. Zhou, Eds. Springer, 2021, pp. 270–282.

[31] R. de Winter, P. Bronkhorst, B. van Stein, and T. Bäck, "Constrained multi-objective optimization with a limited budget of function evaluations," *Memetic Computing*, vol. 14, pp. 151–164, 2022.

[32] P. Bronkhorst, R. d. Winter, T. Velner, and A. Kana, "Enhancing offshore service vessel concept design by involving seakeeping: developing a framework to efficiently design high performance offshore service vessel concepts," in *International Conference on Computer and IT Applications in the Maritime Industries*, V. Bertram, Ed. Technische Universität Hamburg-Harburg, 2022, pp. 273–287.