Natural Computing (2022) 21:127–130
https://doi.org/10.1007/s11047-021-09859-x

Preface

Special issue on "Understanding of Evolutionary Optimization Behavior", Part 2

Christian Blum[1] · Tome Eftimov[2] · Peter Korošec[2]

Accepted: 21 June 2021 / Published online: 13 July 2021

1 Introduction

The history of Evolutionary Computation has progressed from formal studies to a method-centric and framework-centric period, where many algorithms are described as methods or frameworks and their development is primarily performance-driven. We are still in a phase of optimization research where many researchers are developing "new" optimization algorithms, while not truly understanding them. With this special issue we are trying to show the importance of the move from a performance-driven community to a community in which scientific understanding is more important, in which the design of new algorithms (heuristics) becomes a science instead of an art (the future).

Understanding of optimization algorithm's behavior is a vital part that is needed for quality progress in the field of stochastic optimization algorithms. Too often (new) algorithms are setup and tuned only focusing on achieving the desired optimization goal. While this might be effective and efficient in short term, in long term this is insufficient due to the fact that this needs to be repeated for every new problem that arises. Such approach provides only minor immediate gains, instead of contributing to the progress in research on optimization algorithms. To be able to overcome this deficiency, we need to establish new standards for understanding optimization algorithm behavior, which will provide understanding of the working principles behind the stochastic optimization algorithms. This includes theoretical and empirical research, which would lead to providing insight into answering questions such as (1) why does an algorithm work for some problems but does not work for others, (2) how to explore the problem fitness landscape to gain better understanding of the algorithm's behavior, and (3) how to interconnect stochastic optimization and machine learning to improve the algorithm's behavior on new unseen problem instances.

We hope that this special issue will show why future research on optimization algorithms should not focus on (re)developing "new" algorithms, but understanding current ones and improving them according to the gained knowledge in a smarter/better way. By doing so, the community is expected to develop into one in which scientific understanding is more important. Without a doubt, this will not only lead to the development of even better, more efficient algorithms, but it will also lead to algorithms that are more easily applicable outside of the developer's lab environment. This will help researchers/users to transfer the gained knowledge from theory into the real world, or to find the algorithm that is best suited to the characteristics of a given real-world problem.

The focus of this part of the special issue is on understanding the behavior of optimization algorithms in connection to better understanding of optimization problems and various optimization features. A more detailed overview is provided in the following paragraphs.

Muñoz et al., in *Analyzing randomness effects on the reliability of Landscape Analysis*, examine the reliability of five well-known Exploratory Landscape Analysis (ELA) feature sets across multiple dimensions and sample sizes.

This work was supported by the financial support from the Slovenian Research Agency (research core funding No. P2-0098 and project No. Z2-1867). Christian Blum was funded by project CI-SUSTAIN, Spanish Ministry of Science and Innovation (PID2019-104156GB-I00)

✉ Christian Blum
christian.blum@iiia.csic.es

Tome Eftimov
tome.eftimov@ijs.si

Peter Korošec
peter.korosec@ijs.si

1 Artificial Intelligence Research Institute (IIIA-CSIC), Campus UAB, Bellaterra, Spain

2 Computer Systems Department, Jožef Stefan Institute, Ljubljana, Slovenia

ELA features are used for identifying strengths and weaknesses of algorithms, automatic algorithm selection, per instance algorithm configuration and generation methods, and benchmark construction techniques. However, they have limitations related to the computational costs associated with generating accurate results. Consequently, only approximations are available in practice which may be unreliable, leading to systemic errors. Therefore, Muñoz et al. propose a comprehensive experimental methodology, which uses resampling techniques to minimize the sampling cost, and statistical significance tests to identify strengths and weaknesses of individual features. The experimental data is collected and made available in the LEarning and OPtimization Archive of Research Data (LEOPARD) v1.0. Muñoz et al. demonstrate that instances of the same function can have feature values that are significantly different; hence, non-generalizable across instances, due to the effects produced by the boundary constraints. In addition, some landscape features under evaluation are highly volatile, and strongly susceptible to changes in sample size. Finally, the results show evidence of a curse of modality, meaning that the sample size should increase with the number of local optima.

The paper *Classifying Metaheuristics: Towards a unified multi-level classification system* by Stegherr et al. proposes a basis for the comprehensive classification of metaheuristics. The authors expect that this kind of comprehensive classification will facilitate keeping an overview on metaheuristics, including novel approaches and extensions, by organising them depending on their functioning, features and, with sufficient information incorporated into the system, their problem-solving capabilities. The system takes previous classification approaches into account to define basic criteria classification. It furthermore extends previous approaches by establishing several levels based on these criteria which allows for the conceptual examination of metaheuristics depending on the level of detail required. This enables an individual classification concerning the framework, the components or the specific algorithmic features of the metaheuristic, while in total resulting in a comprehensive classification depicting all differences and commonalities. An example is provided for the utilisation of the classification system and the authors point out those areas where additional criteria are required.

The initial population in genetic programming (GP) should form a representative sample of all possible solutions (the search space). While large populations accurately approximate the distribution of possible solutions, small populations tend to incorporate a sampling error. The article *On sampling error in genetic programming* by Schweim et al. analyzes how the size of a GP population affects the sampling error and contributes to answering the question of how to size initial GP populations. First, the article presents a probabilistic model of the expected number of subtrees for GP populations initialized with full, grow, or ramped half-and-half. Second, based on the frequency model, a model that estimates the sampling error for a given GP population size is presented. The models are validated empirically and it is shown that, compared to smaller population sizes, the recommended population sizes largely reduce the sampling error of measured fitness values. Increasing the population sizes even more, however, does not considerably reduce the sampling error of fitness values. Last, the article recommends population sizes for some widely used benchmark problem instances that result in a low sampling error. A low sampling error at initialization is necessary (but not sufficient) for a reliable search since lowering the sampling error means that the overall random variations in a random sample are reduced. The proposed model allows practitioners of GP to determine a minimum initial population size so that the sampling error is lower than a threshold, given a confidence level.

The paper *Comparison of synchronous and asynchronous parallelization of extreme surrogate-assisted multi-objective evolutionary algorithm* by Harada et al. investigates the integration of a surrogate-assisted multi-objective evolutionary algorithm (MOEA) and a parallel computation scheme to reduce the computing time until obtaining the optimal solutions in evolutionary algorithms (EAs). A surrogate-assisted MOEA solves multi-objective optimization problems while estimating the evaluation of solutions with a surrogate function. A surrogate function is produced by a machine learning model. This paper uses an extreme learning surrogate-assisted MOEA/D (ELMOEA/D), which utilizes one of the well-known MOEA algorithms, MOEA/D, and a machine learning technique, extreme learning machine (ELM). A parallelization of MOEA, on the other hand, evaluates solutions in parallel on multiple computing nodes to accelerate the optimization process. The authors consider a synchronous and an asynchronous parallel MOEA as a master-slave parallelization scheme for ELMOEA/D. Moreover, they carry out an experiment with multi-objective optimization problems to compare the synchronous parallel ELMOEA/D with the asynchronous parallel ELMOEA/D. In the experiment, they simulate two settings of the evaluation time of solutions. One determines the evaluation time of solutions by the normal distribution with different variances. On the other hand, another evaluation time correlates to the objective function value. The authors compare the quality of solutions obtained by the parallel ELMOEA/D variants within a particular computing time. The experimental results show that the parallelization of ELMOEA/D significantly reduces the computational time. In addition, the integration of ELMOEA/D with the asynchronous parallelization scheme obtains higher quality of solutions quicker than the synchronous parallel ELMOEA/D.

Evolutionary Algorithms have become state of the art in solving complex optimization problems. However, it is often difficult for practitioners to classify their search behavior in the greater context of global optimization. The article *A new taxonomy of global optimization algorithms* by Stork et al. presents a taxonomy of the field of modern optimization algorithms, which explores, matches, and describes their behaviors by extracting similarities and differences in their search strategies. A particular focus lies on algorithms using evolutionary designs, surrogates, and those created by automatic algorithm generation. The extracted features of algorithms, their main concepts, and search operators, are used to create a set of classification indicators to distinguish between a representative set of algorithm behaviors. The features allow a deeper understanding of components of the search strategies and further indicate the close connections between the different optimization algorithms. The new taxonomy differentiates between algorithm classes by presenting intuitive analogies of their search behaviors, eligible to understand their core concepts. It focuses on simplicity, enabling an easy understanding of global optimization algorithms by novices in this research field. Moreover, it provides further knowledge and references for experts and can support the selection of suitable algorithms for a new optimization problem.

The ability for the results of data-mining to be human-readable has been emphasised with increased research into explainable artificial intelligence (XAI). Evolutionary machine learning enables human-readable symbols to be evolved in the form of if-then rules. However, when the mined knowledge contains thousands of cooperating rules the discovered data patterns become opaque to humans. The paper *Visualizations for rule-based machine learning* by Liu et al. contributes three graphical based methods to enable the discovered rich pattens to be visualised: (1) Feature Importance Map (FIM), shows the importance of each feature to the problem - identifying redundancy, highlighting epistatic relationships, and revealing levels of generality. (2) Action-based Feature Importance Map (AFIM), whereas FIMs show macro-level interactions between features, AFIMs use the micro-level to show the importance of the actual value of the feature. Further, this shows if different classes are formed by different patterns (rather than just varying values in the same pattern). (3) Action-based Feature's Average value Map (AFVM), also works at the micro-level to show the importance of individual feature values to the target action value, which identifies any correlations. These techniques successfully produce interpretable results for tens to thousands of rules, highlighting previously obscure patterns, such as varying class imbalance as a problem scales.

The paper *Similarity in metaheuristics: a gentle step towards a comparison methodology* by de Armas et al. proposes a classification methodology that supports the identification, description, and analysis of metaheuristics considering their defining components. This procedure constitutes a way for providing insights about their contribution while increasing the transparency regarding algorithmic design. Moreover, bearing the proliferation of nature-inspired algorithms in mind, the authors apply this methodology to investigate the novelty of some of them in terms of their algorithmic components. In doing so, they compare diverse nature-inspired metaheuristics for continuous optimization and propose categorizations and measures to identify their similarities and quantify their novelty. As a result, the authors indicate that while some methods' components present a large similarity, others' novelty comes from the combination of already defined components in other algorithms but not properly connected to the body of literature. Thus, the authors highlight the importance of analysing whether there is enough novelty in a so-called "new" (nature-inspired or other) algorithm.

The paper *Understanding measure-driven algorithms solving irreversibly ill-conditioned problems* by Sawicki et al. provides a framework for understanding and analyzing the behavior of stochastic population-based algorithms in solving ill-conditioned global optimization problems. In particular, it targets optimization problems' whose graphs contain lowlands, i.e. connected subsets of minimizers of positive measure and are irreparable in that manner. In such cases, the goal of the problem shifts from finding isolated minimizers to identifying the shapes of lowlands. For a class of optimization strategies with a focusing heuristic, it draws some informative conclusions, such as, that a fixed point of a heuristic will contain maximum information about the problem. Moreover, the authors present several ways to approach stopping conditions in such ill-conditioned cases. An algorithm dealing with such problems, i.e. Hierarchic Memetic Strategy coupled with Multi-Winner Evolutionary Algorithm (HMS/MWEA), illustrates the analytical findings. Its Markov chain is ergodic, and the proof of its asymptotic guarantee of success is shown. The authors conclude by showing its performance on actual ill-conditioned problems: two benchmarks and a real-life case. Putting it all into perspective, the paper presents valuable insights into the inner workings of measure-driven algorithms.

The topic of the paper *The fractal geometry of fitness landscapes at the local optima level* by Thomson et al. is a fitness landscape analysis of the classical combinatorial optimisation domain, the Quadratic Assignment Problem (QAP). Landscape features are studied for their relationship to metaheuristic algorithm performance on the instances. Local optima networks (LONs) are studied for their fractal complexity. LONs are a compact representation of a fitness landscape; the network nodes are local optima, while the edges denote search operator connectivity between those optima. In recent literature, LONs have been subject to fractal dimension analysis (the fractal dimension of an object is an index of spatial complexity). The contributions of this paper are as follows. It brings new insight into how multifractal geometry at the local optima level can help explain and predict algorithm performance. There is also a significant expansion of the data-set used for fractal analysis in LONs (using more than 3x the previous number of QAPLIB instances and raising $N \leq 28$ to $N \leq 50$, as well as deploying a recent refined and tested sampling algorithm for constructing the LONs); and enhanced statistical techniques for properly validating the use of LON fractal analysis for algorithm explanation and prediction (random forest to model non-linearities; random repeated subsampling cross-validation; using intelligible predictors such as the extent of multifractality and the median fractal dimension).

In the growing field of swarm-based metaheuristics, it is widely agreed that the behaviour of an algorithm, in terms of a good balance of exploration and exploitation, plays an important part in its success. Despite this, the influence that the characteristics of an optimisation problem may have on the behaviour of an algorithm is largely ignored. The characteristics of an optimisation problem can be intuitively understood and quantified in terms of fitness landscapes characteristics. Similarly, the behaviour of a swarm-based algorithm can be quantified in terms of its diversity rate-of-change (DRoC). The paper *The influence of fitness landscape characteristics on particle swarm optimisers* by Engelbrecht et al. investigates correlations between the landscape features of optimisation problems and the behaviour, measured via DRoC, of particle swarm optimisation (PSO) algorithms. The authors show that for problems with a simpler global and local structure (characterised by higher fitness distance correlation values and less variation in gradients) the PSO algorithms transitioned to exploitation at a faster rate. They also show that multi-funnelled landscapes deceived the algorithms into faster convergence, pointing out a potential for improvement in PSO algorithms. The approach followed in this paper can be used as a template for analysing the behaviour of other algorithms to lead to a deeper understanding of the strengths and weaknesses of algorithms on problems with different characteristics.

The paper *Gamesourcing: an unconventional tool to assist the solution of the traveling salesman problem* by Zelinka and Das presents an approach to solve a variant of the well-known Travelling Salesman Problem (TSP) by using a gamesourcing approach. The TSP is solved in contemporary literature by a wide spectrum of modern as well as classical computational methods. In contrast, gamesourcing can be understood as a version of game-driven crowdsourcing. The main part and contribution of this paper is a demonstration of the use of gamesourcing in the game called Labyrinth that resembles the TSP structure. The game is based on a maze that enables players to move through it, similar to the way in which artificial ants move in the metaheuristic ant colony optimization. The problem is then solved by playing. The performance of this "human ant-like system" is then evaluated and compared against some well-known versions of ACO. The presented experiments suggest that this approach can serve to assist combinatorial optimizers in achieving better results on well-known NP-hard optimization problems.