Guest Editorial Special Issue on Theoretical Foundations of Evolutionary Computation

IT IS our pleasure to introduce this special issue on the recent advances in the theoretical foundations of evolutionary computation (EC). While in the early days of this field, theoretical analyses inevitably focused on simplified models of evolutionary algorithms (EAs), the continuous progress made in the development of suitable mathematical techniques for the analysis now allows to derive proven statements regarding the performance of off-the-shelf metaheuristics, such as standard generational and steady-state genetic algorithms with no algorithmic simplifications. Comparisons between the performance of a given EA with the best possible one can also be made nowadays, allowing to assess whether a given algorithm may be improved upon or whether its performance is optimal for a given class of problems. Such understanding often provides insights for the design of new EAs which provably have better performance for given problems. We are glad that examples of results of this kind are present within this special issue. A total of 27 papers were submitted which were the subject of at least three independent reviews, and six manuscripts of the highest quality were selected for publication in the special issue. In the following, we provide a brief summary of these manuscripts.

The paper "A General Dichotomy of Evolutionary Algorithms on Monotone Functions" by Johannes Lengler provides new exciting insights into the benefits of the recombination operator in EC. It is well understood that choosing an appropriate mutation rate for an EA is a delicate matter. Even for the simple problem of optimizing a monotone function, mutation-only EAs may be extremely inefficient if the mutation rate is set too high. The paper provides rigorous evidence that the use of a crossover operator makes the EA considerably more robust to the mutation rate making the difference between quasilinear and exponential expected runtimes for arbitrarily high mutation rates c/n where c is an arbitrary constant.

Per Kristian Lehre and Dirk Sudholt explore the performance limits of EAs in the paper "Parallel Black-Box Complexity With Tail Bounds." Black-box complexity provides a baseline for performance comparisons of different randomized search heuristics by stating the minimum expected number of fitness function evaluations required by *all* black-box algorithms to solve any instance of a given problem class. The authors propose a parallel black-box complexity model that allows to capture an inherent property of EAs: in each generation a population of solutions is sampled based on the information gained in the previous ones. Previous black-box models could not capture this inertia because they assume that new solutions are created in sequence rather than in parallel. The design of an EA that matches the parallel black-box complexity of a given problem class shows the power of the approach: no unary unbiased population-based EA may have better asymptotic performance for the problem.

Estimation of distribution algorithms (EDAs) are the focus of the paper "Significance-Based Estimation-of-Distribution Algorithms" by Benjamin Doerr and Martin S. Krejca. EDAs evolve a probabilistic model of the solution space rather than maintaining an explicit population of candidate solutions. The algorithms iteratively sample solutions from the model and use the gained information to update it. Rather than updating the model in each generation as in traditional EDAs, the authors propose to take into account a longer history of samples and to only update the model when a statistically significant change is detected within the history. The validity of the approach is shown by proving that the proposed EDA, called sig-cGA, can optimize the standard benchmark functions OneMax, LeadingOnes, and BinVal in quasilinear expected time. No other EA has been shown to be capable of achieving such performance for all the three problems.

In "Finite-Sample Analysis of Information Geometric Optimization With Isotropic Gaussian Distribution on Convex Quadratic Functions," Kento Uchida, Shinichi Shirakawa, and Youhei Akimoto analyze the influence of two parameters of the information geometric optimization (IGO) algorithm on convex-quadratic functions. By using a fitness utility function, they show that the expected distance to the optimum and the expected standard deviation must decrease at the same rate, a property observed in practice that is not true with the unrealistic assumption that the algorithm has an infinite population size. They derive the stable value of the ratio of the square distance to the optimum over the distribution variance, as well as the conditions required for the stable value to exist. These properties can be directly connected to the practical behavior of the algorithm.

Efficiently controlling the population diversity with the covariance matrix has been the key insight behind probably the most successful EA for continuous domains, namely the CMA-ES. In "Population Diversity of Non-Elitist Evolutionary Algorithms in the Exploration Phase," Jaroslaw Arabas and

Digital Object Identifier 10.1109/TEVC.2020.3035225

1089-778X © 2020 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission.
See https://www.ieee.org/publications/rights/index.html for more information.

Karol Opara analyze the diversity of a real-coded EA. For this purpose, they compute the expected covariance matrix of the population under different operators and selection schemes.

In "Landscape-Aware Performance Prediction for Evolutionary Multi-Objective Optimization," Arnaud Liefooghe, Fabio Daolio, Sebastien Verel, Bilel Derbel, Hernan Aguirre, and Kiyoshi Tanaka investigate how the performance of evolutionary multiobjective optimization (EMO) algorithms relates to the fitness landscape of the problem instances to be solved. New general-purpose features that are affordable to compute even on high-dimensional problem instances are developed, leading to predictive models that fit well to experimental performance data on two broad classes of multiobjective combinatorial optimization problems. The features and models developed pave the ground for the construction of full-fledged feature-based performance prediction and algorithm selection frameworks for multiobjective combinatorial optimization and have the potential to contribute to the improvement of the design of the EMO algorithms.

We would like to thank the Editor-in-Chief, Prof. Kay Chen Tan, for his support during the organization of this special issue. We are also grateful to the authors for submitting their work and to the numerous referees for their careful reviewing and constructive comments. We hope that the readers of IEEE TRANSACTIONS ON EVOLUTIONARY COMPUTATION will enjoy reading the papers and gain important insights and inspiration to further strengthen our theoretical understanding of the behavior and performance of EAs.

PIETRO S. OLIVETO
Department of Computer Science
University of Sheffield
Sheffield S1 4DP, U.K.
e-mail: p.oliveto@sheffield.ac.uk

ANNE AUGER
INRIA
École Polytechnique
IP Paris
91128 Palaiseau, France
e-mail: anne.auger@inria.fr

FRANCISCO CHICANO
Department of Languages and Computing Sciences
University of Malaga
29071 Málaga, Spain
e-mail: chicano@lcc.uma.es

CARLOS M. FONSECA
CISUC
Department of Informatics Engineering
University of Coimbra
3030-290 Coimbra, Portugal
e-mail: cmfonsec@dei.uc.pt