Theoretical Computer Science 561 (2015) 1–2

Contents lists available at ScienceDirect

Theoretical Computer Science

www.elsevier.com/locate/tcs

Preface

Genetic and Evolutionary Computation

Evolutionary computation and other nature-inspired search heuristics are known and applied on a daily basis for 50 years. They remain an important tool in situations where difficult problems need to be solved and no good problem-specific solution is available. Thanks to continuous efforts directed at a theoretical foundation of this broad and complex set of heuristics we have a much improved understanding of their properties, strengths and limitations.

This special issue contains a collection of theoretical analyses of quite different nature-inspired heuristics, all presented in preliminary form either at the field's largest conference, the Genetic and Evolutionary Computation Conference (GECCO 2013), or at a small and specialized workshop on Theory of Randomized Search Heuristics (ThRaSH 2013) that provides an opportunity for theoreticians in the field to meet and discuss their work since 2007. All articles presented here have been reworked and significantly expanded beyond their initial presentation and they all witness that theory is now developed well beyond the understanding of very simple evolutionary algorithms on simple example functions.

Doerr and Künnemann consider the effects of the offspring population size in a mutation-based evolutionary algorithm on the class of linear pseudo-Boolean functions. It is a classical result in the field that with a single offspring the run time is $\mathcal{O}(n \log n)$ for all such functions and $\Theta(n \log n)$ for all such functions that depend on each bit, regardless of the particular function. It is therefore very surprising to see that introducing a larger offspring population suddenly not all linear pseudo-Boolean functions are equal, that the expected run time is asymptotically smaller for some than for others. Beyond this surprising result the article is a shining example of how modern analytical tools can be employed to elegantly derive insights.

Multi-objective optimization is an important area of applications where evolutionary algorithms are applied with great success. Nguyen, Sutton and Neumann analyze evolutionary algorithms and tree-based genetic programming using a popular modern approach to multi-objective evolutionary computation, the hypervolume indicator. Considering classical example functions they prove that performance of more involved evolutionary algorithms differs fundamentally from simpler counterparts. By stepping into the field of multi-objective tree-based genetic programming they demonstrate how the available analytical methods carry much further than the areas where they have been originally developed.

Another broad field where evolutionary algorithms are successfully applied is dynamic optimization where the objective function changes over time. Significantly extending known results about very simple evolutionary algorithms Oliveto and Zarges investigate the benefits a population can provide for an evolutionary algorithm. They prove that a population alone can be insufficient to guarantee crucial improvements and that employing diversity mechanisms together with the population makes the decisive difference. In addition to theorems with asymptotic results they present thorough statistical analyses of carefully conducted experiments and give an exemplarily demonstration how experiments can yield valuable additional insights if carried out and analyzed properly.

Schmitt and Wanka take us outside the realm of evolutionary computation and present a thorough analysis of particle swarm optimization, an important nature-inspired heuristic for the optimization of continuous optimization problems. The theory of particle swarm optimization is much less developed than that for evolutionary algorithms and it is a very significant achievement to prove for a broad class of objective functions that particle swarm optimization finds a local optimum almost surely. In an insightful and accessible analysis the authors lead us to a slightly modified particle swarm optimization that achieves this.

Ant colony optimization is another example of a nature-inspired heuristic that is successfully applied in practice but where our theoretical understanding is not yet as far developed as it is for evolutionary algorithms. Lissovoi and Witt extend known results for a very simple ant colony optimizer using only a single ant to a system with multiple ants and investigate what such a system can achieve when confronted with the problem of finding shortest paths in a dynamic setting where the lengths of paths changes over time. Considering a wide range of scenarios covering single changes, periodic local changes as well as periodic global changes they enhance our understanding of the potential and limitations of ant colony optimization. They also include thorough statistical analyses of carefully conducted experiments to make the asymptotic results more tangible.

http://dx.doi.org/10.1016/j.tcs.2014.11.022
0304-3975/

We feel that the five articles in this special issue are an excellent example of what modern theory of evolutionary computation and other nature-inspired heuristics has to offer to further our understanding of these important and useful tools. They demonstrate that the field is productive and evolving and will remain an exciting area of research.

Tobias Friedrich
Jun He
Thomas Jansen
Alberto Moraglio