Intl. Trans. in Op. Res. xx (2021) 1–2

Special section for Stochastic Local Search: Recent Developments and Trends

Holger Hoos[a,b], Laetitia Jourdan[c,*], Marie-Eléonore Kessaci[c], Thomas Stützle[d] and Nadarajen Veerapen[c]

[a]*University Leiden, Leiden, The Netherlands*
[b]*University British Columbia, Vancouver, Canada*
[c]*University of Lille, Lille, France*
[d]*Université libre de Bruxelles, Bruxelles, Belgium*
E-mail: hh@liacs.nl, laetitia.jourdan@univ-lille.fr, marie-eleonore.kessaci@univ-lille.fr, stuetzle@ulb.ac.be, nadarajen.veerapen@univ-lille.fr

Received DD MMMM YYYY; received in revised form DD MMMM YYYY; accepted DD MMMM YYYY

Stochastic local search (SLS) algorithms are among the most powerful techniques for tackling computationally hard problems in various areas of computing science, operations research, business administration and engineering. SLS techniques range from simple constructive and iterative improvement algorithms to general-purpose SLS methods, also known as metaheuristics.

In recent years, it has become evident that the development of SLS algorithms is complex, and that it typically combines aspects of algorithm design and implementation with empirical analysis and problem-specific background knowledge. The difficulty of this process is in part due to the complexity of the problems being tackled, and in part due to the large number of degrees of freedom researchers and practitioners face when developing SLS algorithms. This development process needs to be supported by a sound methodology that addresses the issues arising in the phases of algorithm design, implementation, tuning and experimental evaluation. In addition, more research is required to understand which SLS techniques are best suited for particular problem types, and to gain better insights into the relationships between algorithm components, parameter settings, problem characteristics and performance.

The SLS 2019 workshop has given the opportunity to researchers interested in the principles and practice of the design, implementation and analysis of stochastic local search algorithms to meet, to present their latest research, and to discuss current developments and applications. The SLS special cluster of papers has attracted submissions by numerous researchers who attended the workshop, but also a number of contributions by other researchers interested in recent advances in the area of SLS algorithms; among these latter submissions we selected six.

The first article, written by Vincent Henaux, Adrien Goeffon and Frédéric Saubion, is titled ”from fitness landscapes evolution to automatic local search algorithm generation”. It proposes an approach based on evolving the fitness function of a local search algorithm to automatically generate search algo-

*Author to whom all correspondence should be addressed (e-mail: laetitia.jourdan@univ-lille.fr).

rithms and tackles discrete optimization problem by exploring an alternative search space. Experiments are conducted on NK-landscapes; these demonstrate the ability of the approach to find alternative fitness landscapes, which are more conducive to effective search than the reference landscape.

Sara Tari, Matthieu Basseur and Adrien Goëffon study "partial neighbourhood local searches" and define the partial neighborhood local search (PNLS) techniques, where moves are based on random subsets of the neighbourhood of the current solution. In PLNS techniques, the balance between intensification and diversification is essentially defined by a parameter λ, which defines the subset size. They analyze three PNLS variants, applying them to four combinatorial optimization problems, and perform a parameter study for two computational budgets. The experiments provide a way to determine good parameter settings for the algorithms, which on some problems perform better than classical variants of iterated local search and tabu search.

Federico Pagnozzi and Thomas Stützle perform an evaluation of the impact of grammar complexity in automatic algorithm design. In such automatic algorithm design systems, the algorithms are composed of components as well as a grammar that is used to describe how these components may be composed to create an algorithm. This approach is used to instantiate and hybridize different metaheuristics templates, but it leads to a large configuration space that is challenging to explore effectively. Here, they investigate what happens if they permit hybridization at most two, one or zero times. They show on the three widely studied permutation flow-shop problems that their system only uses hybridization when doing so provides performance improvements.

"Pattern mining-based pruning strategies in stochastic local searches for scheduling problems" by Arnaud Laurent, Damien Lamy, Benjamin Dalmas and Vincent Clerc proposes a module for knowledge extraction in stochastic local search. The knowledge extension is independent and can be plugged into optimisation methods that rely on various runs of the stochastic local search algorithm. The aim of the knowledge extension is to prune some elements of the search space that do not result in high-quality solutions. Example results on two scheduling problems show that this can speed up convergence and also result in better local optima than the same method applied without these techniques.

In the article "two heuristics for the label printing problem", Federico Alonso-Pecina, Javier Arellano-Verdejo and Rocio Diego-Celis propose to solve the Label Printing Problem using two metaheuristics: threshold accepting and tabu search, each with two neighborhoods. Experiments show that the proposed heuristics can improve the results obtained for most instances from the literature. The metaheuristics are also applied successfully to real-world instances from an industrial setting.

Salim Haddadi discusses "exponential neigborhood search for consecutive block minimization". A simple local search method where the size of the neighborhood is exponential in the instance size, which is called the exponential neighborhood search, is proposed to solve the consecutive block minimization problem. Results on benchmarks from the literature show that the proposed approach is faster and gives better results than methods known from the literature.

Last but not least, we would like to thank the Editor for supporting this special issue, the many people at Wiley who helped us in the production phase, the referees who helped us to select these articles among the best submissions, and finally the authors, who submitted some of their best work to ITOR.