



Guest Editorial

Heuristic optimisation

Journal of the Operational Research Society (2011) **62**, 251–252.
doi:10.1057/jors.2010.160

A heuristic is a method that *‘on the basis of experience or judgement seems likely to yield a good solution to the problem, but cannot be guaranteed to produce an optimum’*. Although over the years there have been some advances in exact optimisation, heuristics continue to be very popular in the Operational Research community for a number of reasons:

1. The computational complexity of many problems means that optimal solutions are unlikely to be found in reasonable time in larger instances.
2. Problems may be ill-defined or data imprecise, so that an optimal solution based on estimated data will almost certainly not be optimal for the actual data. In such a situation it is preferable to obtain a robust solution that will be near-optimal over most scenarios.
3. The user may require several different solutions to make the final choice, particularly in situations where several criteria needed to be balanced using human judgement rather than technical measures.

These perennial concerns have motivated this special issue whose aim is to show some of the most current developments in heuristics from the Operational Research community.

The paper ‘Motif Detection Inspired by Immune Memory’ by Wilson, Birkin and Aickelin presents a novel immune-inspired heuristic called the Motif Tracking Algorithm. It is a new pattern identification tool that is able to identify variable length unknown motifs that repeat within time series data. Motifs are repeating patterns in data and although highly significant for data analysis, relatively little research has been carried out on how to extract them if their length is unknown. This paper shows with the help of some real-life data sets how motifs can be found and what they might mean.

‘Comparing Local Search Metaheuristics for the Maximum Diversity Problem’ by Aringhieri and Cordone compares four local search metaheuristics for the Maximum Diversity Problem, with the aim to identify which additional heuristic elements provide the strongest improvement. Interestingly, the authors identify the best

heuristics are the simplest ones, that is Random Restart and a Variable Neighbourhood Search. The authors suggest that this is probably somewhat related to the structure of the Maximum Diversity Problem, but also conclude that often simpler algorithms might be better.

Wauters, Verbeeck, Vanden Berghe and De Causmaecker present a paper on ‘Learning agents for the multi-mode project scheduling problem’. They use multi-agent reinforcement learning heuristics to build high-quality solutions for the multi-mode resource-constrained project scheduling problem. The contributions of this paper include one of the first applications of multi-agent reinforcement learning to the project scheduling problem, the introduction of new best benchmark results obtained with this method and the introduction of new large benchmark instances.

The paper ‘Exact and Heuristic Methods for Cell Suppression in Multi-Dimensional Linked Tables’ by Roehrig, Padman, Krishnan and Duncan addresses the well-known and challenging problem of how to avoid confidential information disclosure in tabular data. This is an important issue faced by all data-intensive organisations, including national statistical agencies and more recently, healthcare organisations collecting vast amounts of patient data via electronic health records. The methods proposed in this paper mitigate the privacy concerns of patients and other participants about whom data are collected and published in these tables.

‘Revisiting the Big Valley Search Space Structure in the Travelling Salesman Problem (TSP)’ by Hains, Whitley and Howe presents a previously unreported characteristic of the TSP search space under 2-opt based on the big valley hypothesis. The big valley hypothesis states that a strong positive correlation exists between the evaluations of TSP solutions and the number of uncommon edges between them. The authors present an empirical analysis of the search space consisting of tours with evaluations extremely close to that of the global optimum for a number of representative instances. This shows that there is not a single big valley, but rather multiple valleys or funnels composed of tours with similar evaluations but the funnels are separated by a non-trivial distance.

Xu and Qu apply an Evolutionary Simulated Annealing heuristic to multi-objective optimisation problems in the paper ‘Solving Multi-objective Multicast Routing Problems by Evolutionary Multi-objective Simulated Annealing Algorithms with Variable Neighbourhoods’. The author produces experimental results on benchmark

instances and is able to find high-quality non-dominated solutions for multi-objective multicast routing problems. Moreover, the paper shows how variable neighbourhood structures make a significant improvement to the performance of the proposed Variable Neighbourhood algorithm compared with single neighbourhood heuristics.

Naud and Potvin propose a tabu search heuristic with a neighbourhood structure based on ejection chains to solve a variant of the vehicle routing problem, where requests are served by a company's private fleet or an external carrier. Their algorithm outperforms the best approaches in the literature on a set of benchmark instances.

Nance, Roesener and Moore develop a two-dimensional bin-packing algorithm based on tabu search to solve the Mixed Payload Airlift Loading Problem. Its originality lies in being able to accommodate rolling stock cargo, as well as pure pallet cargo loads, while still maintaining feasibility. The load plans require the same or fewer aircraft than AALPS, the load planning software mandated by the US Department of Defense.

Yapicioglu uses a statistical interpolation to expand optimal solutions obtained by multi-criteria heuristic search in order to increase the number of Pareto optimal solutions while limiting computational effort. Tests on non-linear bi-objective problems illustrate how their approach expands the Pareto frontier.

Burke, Curtois, van Draat, van Ommeren and Post investigate a heuristic approach in which local search is alternated with a process that 'jumps' to another point in the search space. The authors propose a model for estimating the quality of new local optima to avoiding

potentially bad neighbourhoods, and test it using five challenging real nurse-rostering problems.

Mansour, Haouari, Sherali and Aissaoui develop tailored optimization-based heuristics to solve a Flexible Aircraft Fleeting and Routing Problem at an airline. Tests with actual data demonstrate the heuristics' effectiveness and robustness in much improving the airline's scheduling and assignment of aircraft.

Vásquez-Rodríguez and Ochoa develop a genetic programming approach to find variants of the NEH heuristic for permutation flowshop problems. Ranking functions are used to prioritise operations during the construction of flowshop schedules. Tests show an improvement on the original NEH heuristic and its stochastic version.

Montoya-Torres, Aponte and Rosas adapt Greedy Randomized Adaptive Search (GRASP) to solve the multi-item version of the NP-hard three-echelon uncapacitated facility location problem. Test show that it compares well to MIP methods in producing fast near-or-actually optimal solutions.

We received nearly 60 submissions for this special issue. We hope the papers selected will be a source of useful results on heuristic optimisation, and provide a direction for future research. We are very grateful to all referees who have provided their constructive comments in order to improve the quality of the papers. We would like to thank the authors who have contributed their work in this special issue.

The University of Nottingham, UK

University of the West of England, UK

Uwe Aickelin

Alistair Clark