Guest Editorial: special issue on evolutionary optimisation and learning

Mengjie Zhang · Michael Kirley · Xiaodong Li

Published online: 5 June 2010
© Springer-Verlag 2010

Problem-solving techniques based on simulated evolution, particularly those inspired by the Darwinian theory of natural selection and biological evolution, have attracted great attention. These evolutionary optimisation and learning algorithms differ from the traditional optimisation and learning methods in that they involve a search from a population of solutions. Over the past decades, evolutionary optimisation and learning algorithms have been successfully applied to a wide variety of real-world problems.

This special issue brings together some recent works from a wide range of topics concerning evolutionary optimisation and learning, including Differential Evolution, Cellular Genetic Algorithms, Particle Swarm Optimisation, Artificial Immune System, Coevolution, Crossover Analysis, Dynamic Combinatorial Optimisation, and Multimodal Optimisation. The eight papers included in this special issue originated from SEAL'08 (7th International Conference on Simulated Evolution and Learning), but have been substantially extended and revised from the conference version. These further extended papers were again rigorously

reviewed in two rounds by at least three anonymous reviewers.

The paper by Lehre and Yao investigates the impact of the parameter settings particularly the acceptance criterion in evolutionary algorithms (EAs) and the crossover operator SSGA when computing unique input output sequences (UIOs) from finite state machines. The objective is to identify simple, archetypical cases where these EA parameter settings have a particularly strong effect on the runtime of the algorithm. This result shows that minor modification in evolutionary algorithms can have an exponentially large runtime impact for computing UIOs and that when computing UIOs, the crossover operator can be essential, and simple EAs can be inefficient.

The paper by Ronknen et al. describes a software framework for generating multimodal test functions. The framework provides an easy way to construct parameterisable functions and offers an environment for testing multimodal optimisation algorithms. A number of function families with different characteristics are included. The framework implements new parameterisable function families for generating desired landscapes. In addition, the framework implements a selection of well-known test functions from the literature, which can be rotated and stretched. The software module can be easily imported into any optimisation algorithm implementation compatible with the C-programming language. As an example, eight optimisation approaches are compared by their ability to locate several global optima over a set of 16 functions with different properties generated by the proposed module.

The paper by Dick and Whigham introduces two modifications to the local sharing method for multimodal optimisation. The first alters local sharing so that parent selection and fitness sharing operate at two different spatial levels: parent selection is performed within small demes,

M. Zhang (✉)
School of Engineering and Computer Science,
Victoria University of Wellington, PO Box 600,
Wellington, New Zealand
e-mail: mengjie.zhang@ecs.vuw.ac.nz

M. Kirley
Department of Computer Science and Software Engineering,
Melbourne School of Engineering, The University of Melbourne,
Melbourne, VIC 3001, Australia
e-mail: mkirley@cs.mu.OZ.AU

X. Li
School of Computer Science and Information Technology,
RMIT University, Melbourne, VIC 3001, Australia
e-mail: xiaodong.li@rmit.edu.au

while the effect of fitness sharing is weighted according to the distance between individuals in the entire population structure. The second modification replaces fitness sharing within demes with clearing to produce a method that we call local clearing. The proposed modifications are examined and compared with the traditional fitness sharing and standard local sharing methods and the results show that the new methods are more efficient. The work also offers a level of parameter robustness that surpasses other elitist niching methods, such as clearing. The analyses of the local clearing method show that the parameter robustness is a result of the isolated nature of the demes in a spatially structured population being able to independently concentrate on subsets of the desired optima in a fitness landscape.

The paper by Rohlfshagen and Yao analyses the properties of the NP-hard dynamic subset sum problem to gain a better understanding of how the dynamics, affecting the parameters of the problem, are materialised in the fitness landscape. The paper highlights the correlation between the dynamic parameters of the problem and the resulting movement of the global optimum. The results show that the degree to which the global optimum moves in response to the underlying dynamics is correlated only in specific cases. The role of the representation used to encode the problem and the impact of the formulation of the objective function on the dynamics are also discussed in the paper.

The paper by Barbosa et al. presents an approach to constructing an ensemble of neural networks using coevolution and the artificial immune system (AIS). A diversity promotion technique is added to CLONENS (CLONal Selection Algorithm for building ENSembles). A coevolutionary approach to build neural ensembles is introduced, where two populations representing the gates and the individual neural networks are coevolved. The first population is responsible for defining the ensemble size and selecting the members of the ensemble, which is evolved using differential evolution. The second population supplies the best individuals for building the ensemble, which is evolved by AIS. The results show that it is possible to automatically define the ensemble size and find smaller ensembles with good generalisation performance on the tested benchmark regression problems. The use of the diversity measure during the evolutionary process does not necessarily improve generalisation.

The paper by Ishibuchi et al. describes several implementations of cellular genetic algorithms with two neighbourhood structures for single-objective and multi-objective optimisation problems. As local selection has already been utilized in cellular genetic algorithms in the literature, this paper is focused on local competition. This paper presents three ideas: local elitism, local ranking, and

local replacement. Local elitism and local ranking are used for single-objective optimisation to increase the diversity of solutions, and local replacement is used for multi-objective optimisation to improve the convergence of solutions to the Pareto front. In this paper, the two neighbourhood structures are specified independently of each other and the effect of each neighbourhood structure on the behaviour of cellular algorithms is examined separately.

The paper by Iorio and Li investigates an enhancement to differential evolution that introduces greater diversity while also directing the search to more promising regions of the search space. The combinatorial sampling differential evolution (CSDE) is introduced which can sample vectors in two ways: highly correlated with the search space or around a ‘better’ individual. This approach to sampling vectors is capable of optimising problems with extensive parameter interactions. It also demonstrates fast convergence towards the global optimum and is highly scalable in the decision space on a variety of single and multi-objective problems due to the balance between sampling highly directed correlated vectors and non-correlated vectors which contribute to sampling diversity.

The paper by Mohammed et al. describes an approach to the use of particle swarm optimisation (PSO) within an AdaBoost framework for object detection. Instead of using exhaustive search for finding good features for constructing weak classifiers in AdaBoost, this paper proposes two new methods based on PSO. The first uses PSO to evolve and select good features only and the weak classifiers use a simple decision stump. The second uses PSO for both selecting good features and evolving weak classifiers in parallel. These two methods are examined and compared on two challenging object detection tasks in images: detection of individual pasta pieces and detection of a face. The experimental results suggest that both approaches can successfully detect object positions, and that using PSO for selecting good individual features and evolving associated weak classifiers in AdaBoost is more effective than for selecting features only. The paper also shows that PSO can evolve and select meaningful features in the face detection task.

We would like to express our gratitude to the reviewers for their expertise and untiring efforts in ensuring the quality of the papers for this special section. We thank especially the following reviewers (in alphabetical order): Peter Andreae, Ying-ping Chen, Raymond Chiong, Grant Dick, Muddassar Farooq, Samuelson W. Hong, Mark Johnston, Krzysztof Krawiec, Takio Kurita, Per Kristian Lehre, Jing Liu, Wenjian Luo, Kourosh Neshatian, Yew Soon Ong, Kai Qin, Ramon Sagarna, Sancho Salcedo-Sanz, Hiroshi Someya, Andy Song, Ke Tang, Yu-Xuan

Wang, Russuel Y. Webb, Peter Whigham, Upali K Wickramasinghe, Shengxiang Yang, Lean Yu. Last but not least, we would like to thank the Area Editor Professor Francisco Herrera for his support and assistance during the editing process of this special issue.

Mengjie Zhang
Michael Kirley
Xiaodong Li
Guest editors, October 2009