

Swarm intelligence: the state of the art special issue of natural computing

Eric Bonabeau · David Corne · Riccardo Poli

Published online: 5 January 2010
© Springer Science+Business Media B.V. 2010

It has long been recognised that nature provides many stunning and intriguing examples of group behaviour. Different species of bird swarm in varieties of numbers and in varieties of formations; ants collectively hunt food in huge numbers; schools of fish adopt close knit and ever-changing formations as they bewilder predators and bewitch prey. The activities of so-called swarms of organisms in many cases lead to beneficial effects for the organisms in question, and this has led to the concept of *swarm intelligence*, which encapsulates the idea that the behaviour of a swarm can exhibit useful, functional and intelligent outcomes which seem well beyond the capabilities, as far as we understand, of any individual in the swarm.

A common way to define swarm intelligence is as the emergent and collective intelligence of groups of simple and autonomous agents. Implicit in such a definition is the fact that there is assumed to be no central controller—no ‘master’ agent that directs and conducts the activities of its fellows. Each agent is a peer, acting independently, interacting with its brethren, and other features of its environment, according to relatively simple rules. That such a system, with no central controller, can exhibit purposeful and robust behaviour is one of the appealing factors for scientists and engineers, suggesting strategies for building systems that are robust to failures. E.g. when systems rely on a centralised controller, damage to the controller will clearly be disastrous. Swarm intelligence is of great interest to scientists and engineers for two other main reasons. First, scientists need to understand swarms—to understand how the interactions within swarms help social animals achieve their various goals, to understand how swarm behaviour evolved, and so on. Secondly, and of most import to us here, is the fact that the study of natural swarm intelligence leads directly to novel algorithms that have a wide variety of applications. In

E. Bonabeau
Icosystem, 10 Fawcett Street, Cambridge, MA 02138, USA
e-mail: eric@icosystem.com

D. Corne (✉)
Heriot-Watt University, Edinburgh EH14 4AS, UK
e-mail: dwcorne@gmail.com

R. Poli
School of Computer Science and Electronic Engineering, University of Essex, Colchester CO4 3SQ, UK
e-mail: rpoli@essex.ac.uk

particular, swarm intelligence based algorithms provide new ways to solve complex and difficult optimisation problems, new ways to manage and control traffic and communications networks, accomplished and effective ways to generate realistic simulations of swarms, and last, but not least, strategies and ideas for organising goal-oriented groups of robots (so-called ‘swarm-bots’) for potential future applications ranging through agriculture, flexible manufacturing and space exploration.

Swarm intelligence is now a large and wide ranging field of research, however it remains rather dominated by two of the prominent swarm-inspired algorithms whose early successes inspired much of this field. These are *ant colony optimization (ACO)* and *particle swarm optimization (PSO)*. Of these two, PSO has attracted extraordinary levels of attention in recent years, and the current special issue clearly reflects this activity in being dominated by papers that introduce variants of PSO algorithms. Meanwhile, the issue’s papers also include real-world applications of scientific and industrial interest, progress in multi-objective optimisation, progress in algorithms for dealing with dynamic optimisation problems, and, reflecting the still continuing exploration of new natural inspirations for optimisation, a paper describing a swarm-based optimisation method inspired by the laws of gravity. Below we provide a brief summary of the papers in the special issue in order of their appearance.

The paper by Chen, Zhu and Hu, “Discrete and Continuous Optimization based on Multi-swarm Coevolution”, presents a novel variant on PSO which introduces a second level of grouping over individuals. Their algorithm is essentially a swarm of swarms, with the higher level swarm imposing an interaction topology over the individual swarms. A key feature of the new algorithm is the dynamical update equation—essentially, a revision of the standard PSO equations that suitably exploits the multi-swarm environment. This new algorithm is benchmarked against four other successful recent PSO variants, and shown to have impressive relative performance.

In Schoeman and Engelbrecht’s contribution, “A Novel Particle Swarm Nicheing Technique Based on Extensive Vector Operations”, a related theme is considered: again, the development of a new variant of PSO that attempts to improve performance on larger scale and more difficult problems by somehow organising a larger population into multiple swarms. In the vector-based PSO of this paper, however, individual swarms (niches) within the larger population are inferred automatically and dynamically, by using specific vector operations to estimate the boundaries of groups of individuals that are exploring a particular niche in the search space. This approach has advantages over other attempts to group populations into subpopulations associated with niches; usually such approaches need to assume a pre-defined number of niches, and sizes (radii) for these niches, however these are both determined automatically within the vector-based PSO algorithm.

The next paper, by Wang, Yang, Ip and Wang, is “A Particle Swarm Optimization Based Memetic Algorithm for Dynamic Optimization Problems”. This provides a state of the art snapshot of at least two strong directions in swarm intelligence research which converge in this paper: the hybridisation of swarm intelligence algorithms with other algorithms (typically local search), and the exploration of their performance on important and different challenges beyond pure ‘static’ optimisation. In dynamic optimisation problems, of which there are countless real-world examples, the search landscape itself, and therefore the location of optima, is ever changing. This provides various challenges for optimisation algorithms, which need, as ever, to be effective in finding optima in the first place, but also need to be able to adapt continually to the shifting landscape, without getting entrenched in previously attractive, but now relatively unfit positions. The

algorithm described here is shown to be highly effective on a commonly used benchmark problem in the field of dynamic optimisation.

The paper by Rashedi, Nezamabadi-pour and Saryazdi entitled “BGSA: Binary Gravitational Search Algorithm”, provides an example of an ever-continuing aspect of swarm intelligence research (and natural computing research in general), which is the quest for new and novel algorithms from alternative natural inspirations. In this case, the inspiration is the action of Newtonian gravity, and both local and global interactions among a swarm of particles are organised in a way that echoes the nature of Newton’s laws. The authors’ previous “Gravitational Search Algorithm” has been published elsewhere; the contribution in the current article is a version of this algorithm that is able to deal with binary encodings—i.e. either binary encodings of real-valued parameter problems, or combinatorial problems with a natural binary representation. On a range of test functions, it is shown the BGSA has promising performance characteristics.

The paper by Abido, entitled “multiobjective particle swarm optimization with non-dominated local and global sets”, exemplifies a further two aspects of how swarm intelligence methods are reaching out in a big way across the frontiers of research and practice. The paper concerns multi-objective optimisation, now a huge and still growing area of optimization, in which it is recognised that most real-world problems have several conflicting objectives, and specialised techniques are required to address them, rather than simplifying the problems so that ‘single-objective’ optimisation can be applied. The second notable aspect of this paper is the concentration on a real-world application, namely the electrical power dispatch problem, in which the problem is to minimize overall costs and emissions for operating a fossil-fuel power-station subject to demand constraints. On this problem, as well as some standard test function optimisation problems, the method in the paper has strong relative performance.

Finally, in “A flocking-based approach for data clustering”, by Ahmadi, Karray and Kamel, we see a swarm-inspired algorithm for solving a ubiquitous problem in data mining, as well as science in general—the problem of finding structure in multidimensional data via seeking a good clustering of the data. PSO-based clustering, often involving hybridisation with data mining strategies that are associated with clustering, such as self-organising maps, or k -means, is an active area of research. The paper reviews this area, and then describes the new approach, which (echoing developments elsewhere) relies on suitably organising multiple swarms. Extensive tests and comparisons are done, involving benchmarking against several state of the art approaches to clustering, and it seems clear that the method of the paper provides a highly competitive solution.

This special issue presents a snapshot of ongoing research in swarm intelligence with a practical orientation. As with all small samples, the distribution of activities and areas present in these papers is not a particularly good estimate of the real distribution of current research in the landscape of swarm intelligence. Nevertheless, there are clear elements of that landscape that are faithfully retained in this sample. Namely, we see a strong focus on optimisation (and that mainly, these days, explored via PSO variants), we see the arms of swarm intelligence reaching out to important real-world problems, we see methods from swarm intelligence adapting to up-to-the-mark challenges in optimisations such as dynamic and multi-objective problems, and we see an ever-active backbone of research that seeks new algorithms from new inspirations.

In conclusion, we would like to thank the several academics who supported us by providing reviews of papers in this special issue. Their time, expertise and energy has been immensely appreciated.