

## Advances in artificial immune systems

Emma Hart · Chris McEwan · Jon Timmis ·  
Andrew Hone

Published online: 19 May 2011  
© Springer-Verlag 2011

The field of Artificial Immune Systems (AIS) derives inspiration from processes and mechanisms apparent in the biological immune system. After early applications of this paradigm to problems in anomaly detection and classification, the field rapidly expanded, leading to a number of successful optimization algorithms, and more recently, applications in fields as diverse as swarm robotics and wireless sensor networks.

In contrast to the situation in most other areas of biologically-inspired computing, practitioners of AIS maintain close links with the field that inspires their work. The field of immunology itself is continually moving, with new discoveries constantly refining the knowledge and understanding we have of the natural immune system. New immunological knowledge is translated to AIS in the form of new algorithms and architectures, but in addition, research driven by computational modelling also informs immunology. This two-way flow of knowledge between

engineering and immunology distinguishes the field from other biological paradigms in computing, where the underlying biological phenomena are perhaps more clearly understood. In parallel to this, a body of literature has now emerged which provides a theoretical underpinning for the field, employing many of the same techniques that are used to analyze other randomized search algorithms. This facilitates a more rigorous comparison between AIS algorithms and other search algorithms, and further enables crossover of ideas between AIS and other paradigms.

The papers in this special issue—selected from papers submitted following the 9th International Conference on Artificial Immune Systems (ICARIS) held in Edinburgh in July 2010—capture the essence of the current state of the art in the field, illustrating the aforementioned trends.

In their paper *Collective Classification of Textual Documents by Guided Self-Organization in T-Cell Cross-regulation Dynamics*, Abi-Haidar and Rocha analyse a computational extension of the model of T-Cell cross-regulation initially proposed by Carneiro et al. This work provides several interesting contributions to the AIS literature. On the one hand, it elaborates the original, minimalist analytical model to incorporate heterogeneous surface presentation and T-Cell specificities, which in turn, provides insight into how the hypothesised mechanism might unfold in a more realistic setting. On the other hand, this demonstration is provided in the context of document classification, where amino-acid sequences are replaced with words and tolerance and immunity replaced with skewed, two-class classification. The subsequent empirical demonstration of the model's performance in comparison to established statistical algorithms provides a compelling insight into the robustness of the biological model that would be very difficult to assert using traditional means, due to the lack of sufficient experimental data.

E. Hart (✉)  
Institute for Informatics and Digital Innovation,  
Edinburgh Napier University, Edinburgh, UK  
e-mail: E.Hart@napier.ac.uk

C. McEwan  
I2D3 Integrative Immunology: Differentiation,  
Diversity, Dynamics, UPMC Univ Paris 06, Paris, France  
e-mail: chris.mcewan@upmc.fr

J. Timmis  
Department of Computer Science and Department  
of Electronics, University of York, York, UK  
e-mail: jtimmis@cs.york.ac.uk

A. Hone  
School of Mathematics, Statistics & Actuarial Science,  
University of Kent, Kent, UK  
e-mail: A.N.W.Hone@kent.ac.uk

In *Surrogate-assisted Clonal Selection Algorithms for Expensive Optimization Problems*, Bernardino et al. provide a thorough empirical assessment of using surrogate functions in the CLONALG optimization algorithm. Stochastic search algorithms, of which CLONALG is an immune-inspired variant, often require a large amount of function evaluations. Surrogate functions provide a light-weight proxy objective function that can be used in place of the true underlying objective function, which can be exploited in several ways. Bernardino et al. show how surrogate functions can be used to improve solutions given the constraint of a fixed number of true objective evaluations. Such a constraint is typical of inverse problems in scientific computing, where the objective function may be an expensive simulation. Their paper demonstrates the applicability of these techniques to scale immune-inspired optimisation algorithms to this important class of problems, by comparison with a default CLONALG-based approach. They also provide some empirical insight into the best method of approximating the underlying objective function by comparing nearest-neighbour and locally linear variants.

Finally, in *On the Role of Age Diversity for Effective Aging Operators*, Jansen and Zarges perform a theoretical study of the aging mechanisms in randomized search heuristics. Aging is a procedure employed in both evolutionary computation and AIS algorithms, whereby each search point is assigned an age which increases with each generation, and when a search point becomes older than some pre-defined maximal age, it is deleted from the population. New search points, or offspring, are then introduced, with the overall effect being to increase the diversity of the population, but there are various different strategies for assigning an age to the offspring. This paper considers different aging operators and provides asymptotic bounds on their performance when applied to optimization of a particular example objective function. These theoretical bounds are also compared with experimental results for small problem sizes, revealing a gap between theory and practice. The overall conclusion is that not only diversity of search points, but also diversity with respect to age plays an important role in determining the performance of randomized search algorithms that employ aging strategies.