

Towards a Conceptual Framework for Artificial Immune Systems

Susan Stepney¹, Robert E. Smith², Jonathan Timmis³, and Andy M. Tyrrell⁴

¹ Department of Computer Science, University of York

² The Intelligent Computer Systems Centre, University of the West of England

³ Computing Laboratory, University of Kent

⁴ Department of Electronics, University of York

Abstract. We propose that bio-inspired algorithms are best developed and analysed in the context of a multidisciplinary conceptual framework that provides for sophisticated biological models and well-founded analytical principles, and we outline such a framework here, in the context of AIS network models. We further propose ways to unify several domains into a common meta-framework, in the context of AIS population models. We finally hint at the possibility of a novel instantiation of such a meta-framework, thereby allowing the building of a specific computational framework that is inspired by biology, but not restricted to any one particular biological domain.

1 Introduction

The idea of biological inspiration for computing is as old as computing itself. It is implicit in the writings of von Neumann and Turing, despite the fact that these two fathers of computing are now more associated with the standard, distinctly non-biological computational models.

Computation is rife with bio-inspired models (neural nets, evolutionary algorithms, artificial immune systems, swarm algorithms, ant colony algorithms, L-systems, ...). However, many of these models are naïve with respect to biology. Despite the fact that these models often work extremely well, their naivety often blocks understanding, development, and analysis of the computations, as well as possible feedback into biology.

We propose that bio-inspired algorithms are best developed and analysed in the context of a multidisciplinary conceptual framework that provides for sophisticated biological models and well-founded analytical principles.

2 A Conceptual Framework

The next step in bio-inspired computation should both develop more sophisticated biological models as sources of computational inspiration, and also use a conceptual framework to develop and analyse the computational metaphors and algorithms.

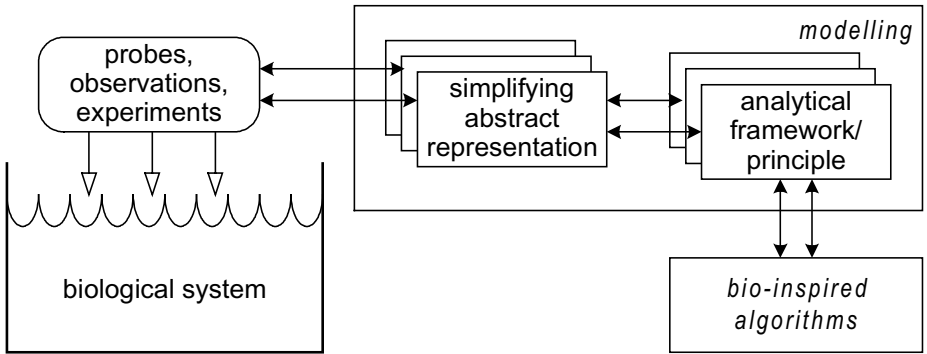


Fig. 1. An outline conceptual framework for a bio-inspired computational domain

Figure 1 illustrates a possible structure for such a conceptual framework. Here *probes* (observations and experiments) are used to provide a (partial and noisy) view of the complex *biological system*. From this limited view, we build and validate simplifying abstract representations, *models*, of the biology. From these biological models we build and validate *analytical computational frameworks*. Validation may use mathematical analysis, benchmark problems, and engineering demonstrators. These frameworks provide principles for designing and analysing *bio-inspired algorithms* applicable to non-biological problems, possibly tailored to a range of problem domains, and contain as much or as little biological realism as appropriate. The concept flow also supports the design of algorithms specifically tailored to modelling the original biological domain, permits influencing and validating the structure of the biological models, and can help suggest ideas of further experiments to probe the biological system. This is necessarily an interdisciplinary process, requiring collaboration between (at least) biologists, mathematicians, and computer scientists to build a complete framework.

An important observation is that none of the representation and modelling steps outlined above is *unbiased*. There are many possible probes, and many possible representations of the same systems even given the same probes, and they all provide different insights. In particular, models derived specifically for the goals of biological simulation may provide insights that are distinct from those that serve computational goals. It is very seldom that the modelling steps used in these distinct activities are examined for common properties, and comparative biases.

In many instances not all of the representational steps outlined above are taken. In particular, bio-inspired computational algorithms usually proceed directly from a (naïve) biological model to an algorithm, with little analytical framing of the representation's properties. Such "reasoning by metaphor" is a troubling aspect of these algorithms. Without the application of available analysis techniques to the simplified representations of biological systems, algorithms derived from these representations rely only on the (often weak) analogy to the biological system to support their use. We feel that it is important to recognize the distinct levels of the modelling process outlined above, to avoid naïve assumptions.

One example that *can* be described in terms of such a framework, at least partially, is John Holland's original adaptive system theories [Holland 1975; Goldberg 1989], founded on a simplified binary-encoded representation of genetics, and analytical

principles of building blocks, k -armed bandit theories, the schema theorem, and implicit parallelism. Evolutionary computation theory has developed and deepened in the wake of this work, and it continues to influence the prescription of genetic algorithms. We propose that other bio-inspired computational domains, including Artificial Immune Systems, should be put on a similarly sound footing.

3 Instantiating for Artificial Immune Systems (AIS)

The natural immune system is a complex biological system essential for survival. It involves a variety of interacting cellular and molecular elements that control either micro- or macro-system dynamics. The effectiveness of the system is due to a set of synergetic, and sometimes competitive, internal strategies to cope with chronic and/or rare pathogenic challenges (antigens). Such strategies remodel over time as the organism develops, matures, and then ages (immuno-senescence). The strategies of the immune system are based on task distribution to obtain distributed solutions to problems (different cells are able to carry out complementary tasks) and solutions to distributed problems (similar cells carrying out the same task in a physically distributed system). Thus, cellular interactions can be envisaged as parallel and distributed processes among cells with different dynamical behaviour, and the resulting immune responses appear to be emergent properties of self-organising processes. Many theories abound in immunology pertaining to how the immune system remembers antigenic encounters (maintenance of memory cells, use of immune networks), and how the immune system differentiates between self and non-self molecules (negative selection, self-assertion, danger theory).

We can explicitly exploit the conceptual framework, in order to develop, analyse and validate sophisticated novel bio-inspired computational schemes, including those inspired by complex processes within the natural immune system. This work needs to be done; here we outline a suggested route.

3.1 A First Step: Interdisciplinary Research

AIS is a relatively new and emerging bio-inspired area and progress has been made from naively exploiting mechanisms of the immune system. Computer security systems have been developed, anti virus software has been created, optimisation and data mining tools have been created that are performing as well as the current state of the art in those areas.

However, as discussed above, in order to push forward the state of the art in this area, a greater interaction between computer scientists, immunologists and mathematicians is required. AIS are poorly understood on a theoretical level, and the metaphors employed have been typically limited: a deeper understanding of these systems is required as to their broader applicability, and a more radical view of their application is needed to break new ground in the research field, rather than reinvent old technology. In order for this emerging area of research to mature, a detailed investigation of the natural immune system properties and their interactions, coupled with sound theoretical development of corresponding AIS, and testing on hard real world prob-

lems, is needed. Such research is essential to ensure success of the area in being adopted as a serious alternative contender in the computational intelligence community.

In recent literature, it seems to be common that naïve approaches to extracting metaphors from the natural immune system have been taken, but this has not always been the case. The original AIS were developed with an interdisciplinary slant.

The work in [Bersini 1991,1992; Bersini & Varela 1994] are such cases. Clear attention to the development of immune network models is given, and these models are then applied to a control problem characterised by a discrete state vector in a state space \mathcal{X}^L . Bersini's proposal relaxes the conventional control strategies, which attempt to drive the process under control to a precise target, to a specific zone of the state space. He instead argues that the metadynamics of the immune network is akin to a meta-control whose aim is to keep the concentration of the antibodies in a certain range of viability so as to continuously preserve the identity of the system.

There are other examples of interdisciplinary work, such as the development of immune gene libraries and ultimately a bone marrow algorithm employed in AIS [Hightower *et al* 1995], and the development of the now famous negative selection algorithm and the first application to computer security [Forrest *et al* 1994].

However, in more recent years, work on AIS has drifted away from the more biologically-appealing models and attention to biological detail, with a focus on more engineering-oriented approach. This has led to systems that are examples of the "reasoning by metaphor" outlined above. These include simple models of clonal selection and immune networks [de Castro & von Zuben 2000, 2001; Timmis 2000; Neal 2003], and negative selection algorithms [Bradley & Tyrrell 2002; Gonzalez *et al* 2003; Taylor & Corne 2003]. This is not a criticism of that work, but we are pointing out that these (and many more engineered solutions) may benefit from not only closer interaction with biologists, but also a more principled mechanism for the extraction, articulation, and application of the underlying computational metaphor.

[Freitas & Timmis 2003] outline the need to take into account the application domain when developing AIS. This seems reasonable when one considers the construction of the final solution. However, this stance will not hold completely when developing generic models and frameworks applicable to a variety of domains. The conceptual framework proposal here does not contradict the position taken in [Freitas & Timmis 2003], but rather complements it: once we have a well-developed conceptual framework, we can specialise it for various application domains in a justifiable way.

3.2 Adopting the Conceptual Framework for AIS

[de Castro & Timmis 2002] propose a structure for engineering AIS. The basis is a *representation* to create abstract models of immune organs, cells, and molecules, together with a set of *affinity functions* to quantify the interactions of these "artificial elements", and a set of *general-purpose algorithms* to govern the dynamics of the AIS.

The structure can be modelled as a layered approach (figure 2). To build a system, one typically requires an application domain or target function. From this basis, a *representation* of the system's components is chosen. This representation is domain and problem dependent: the representation of network traffic, say, may well be differ-

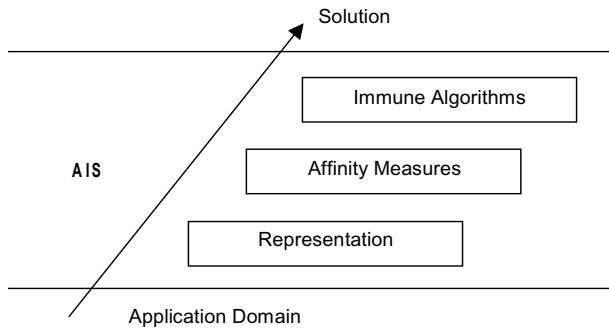


Fig. 2. A structure for AIS, from [de Castro & Timmis 2002]

ent from that of a real time embedded system. Once the representation is chosen, one or more *affinity measures* are used to quantify the interactions of the elements of the system. There are many possible affinity measures (which are also partially dependent upon the representation adopted), such as Hamming or Euclidean distances. The final layer involves the use of *algorithms*, which govern the behaviour (dynamics) of the system, such as negative and positive selection, clonal selection, the bone marrow algorithm, and immune network algorithms. Each algorithm has its own particular range of uses.

However, this layered structure is not complete from the conceptual framework perspective, which we believe is required to allow effective algorithms to be developed.

AIS algorithms in their current form can be classified as *population based* or *network based* [de Castro & Timmis 2002]. In the following sections, we adopt this initial breakdown, and in addition propose that AIS algorithms (in some cases) may benefit from asking questions such as: what is “self”, or “danger”. We propose how one might undergo a development of an AIS algorithms adopting the conceptual framework above.

3.3 Population Based AIS Algorithms

Three common algorithms in AIS, those of positive, negative, and clonal selection, are all based on *populations* of agents trained to recognise certain aspects of interest. (See [de Castro & Timmis 2002] for an overview.) There are similarities between the algorithms: positive and negative selection, for example, are merely two sides of the same coin. There are also differences: positive and negative selection involve essentially random generation of candidate recognisers, whilst clonal selection uses a form of reinforcement based on selection and mutation of the best recognisers.

Rather than describe how a population based AIS-specific framework might be developed, however, we defer discussion of these models to the meta-frameworks of section 4, and population based models in general.

3.4 Network Based AIS Algorithms

The original immune network theory [Jerne 1974] suggests an immune system with a dynamic behaviour even in the absence of non-self antigens. This differs from the biological basis of the clonal and negative selection algorithms, as it suggests that B-cells are capable of recognizing each other.

Several theoretical immunologists have been interested in creating models of immune networks in order to introduce new ways of explaining how the immune systems works [Perelson 1989; Farmer *et al* 1986]. This has been followed by a number of researchers who have translated some of these ideas into the computing domain, in applications such as optimization and control [Bersini 1991; Bersini & Varela 1994]. This work has also inspired the development of machine learning network models with applications mainly in data analysis [Timmis 2000; de Castro & Von Zuben 2001].

However, as we stated earlier, the later work has somewhat deviated from the biological model, being adapted to a particular problem. In addition, the immune network theory itself is controversial, and not widely accepted by immunologists. This has an impact on the AIS algorithm: if the biology is not correct, then one must re-examine it and understand what is *really* going on; this would hopefully shed light on the more complex nature of the immune systems, and the networks that are clearly present therein.

In order to achieve this understanding, the first step would be to probe the biological system from the perspective of interpreting the system as a *network* of interaction, cooperation and competition amongst molecules, cells, organs and tissues. The discovered properties could then be exploited in order to identify emergent behaviour and formulate a suitable mapping between biological properties and framework components. These components could then be used as the basis for the topology and dynamics of new computational and mathematical models, in addition to re-examining existing models such as [Perelson 2002; Romanyukha & Yashin 2003]. The resulting new models would allow a greater understanding of the operation of such systems to be developed in an artificial context. Within the context of immune networks, one could examine the danger theory [Matzinger 2002], context of response [Janeway & Medzhitov 2002; Kourilsky & Truffa-Bachi 2001], memory mechanisms [Sprent & Surh 2002], general alarm response or stress response [Padgett & Glaser 2003] and self/non-self recognition [Medzhitov & Janeway 2003]. Additionally, the constructive role of noise in biological systems, which is an intrinsic feature of such systems, could also be examined [Gammaitoni 1998].

From these biological models, suitable new computational metaphors could be created. An analytical framework could be created to include appropriate representations for components, methods of assessing interactions between components, and processes to act on components. The framework should also provide features that allow biological models to be represented and manipulated in a number of ways. The framework should permit the analysis and identification of generic properties, and an instantiation of the framework should permit the capture of properties relevant to the application being developed. In an iterative process, the framework algorithms should be implemented and tested in order to test and develop the biological metaphors prior to their implementation and experimental exercises on the intended platform.

Taking this fuller view of immune networks may yield AIS algorithms that truly mimic the qualities of the diversity of immune network memory mechanisms, and may inform us as to the scalability of immune networks, their ability to cope more effectively with noise, their open nature, and the level of interaction both within the network and external to the network. Biology would benefit from the resulting sophisticated models, too.

3.5 Self or Danger?

Immunologists are asking the question, does the immune system distinguish between self and non-self, or is there something more going on? Adopting the framework approach would allow AIS to adopt a principled approach to the review and analysis of such theories as danger theory [Matzinger 2002] and self-assertion [Varela *et al* 1998], and bring forth new and effective bio-inspired algorithms.

A more interdisciplinary slant has already begun in this area. For example, [Aicklen *et al* 2003] describe a large ambitious interdisciplinary project investigating novel ideas from immunology such as danger theory, with the application to computer security. Here the authors propose to observe the biological system by undertaking new experiments to identify key signals involved in cell death, and identify the functions of such signals and how these affect immune cells. The hope is that this research will shed light on how the immune system distinguishes self from non-self, in order to build effective immune-inspired computer security systems that no longer rely on the need to define *a priori* the self of the system. Although those authors make no reference to adopting a framework approach such as outlined above, we believe that it would help ensure not only biologically-plausible algorithms, but effective and general solutions.

4 Meta-frameworks for Bio-inspired Computation

The earlier figure 1 shows potentially many representations of the same systems under the same observations, each of which may provide different insights. Such distinct representations, although common, are seldom examined for unifying properties. Once we have a conceptual framework, we can not only make such comparisons, we can go a step further: to examine and compare the separate conceptual, mathematical and computational frameworks, to develop more *integrated and generic frameworks*, and to expose essential differences.

To achieve this, we can apply the same conceptual model, at a higher level (figure 3). The key probes here are meta-questions. Just as the questions at the biological level influence the kinds of models developed, so the meta-questions influence the kinds of meta-models developed.

4.1 Meta-probes for Complex System Frameworks

What kind of meta-questions might we ask? Clearly, the questions asked influence the resulting framework. We have identified some initial areas thought to affect complex

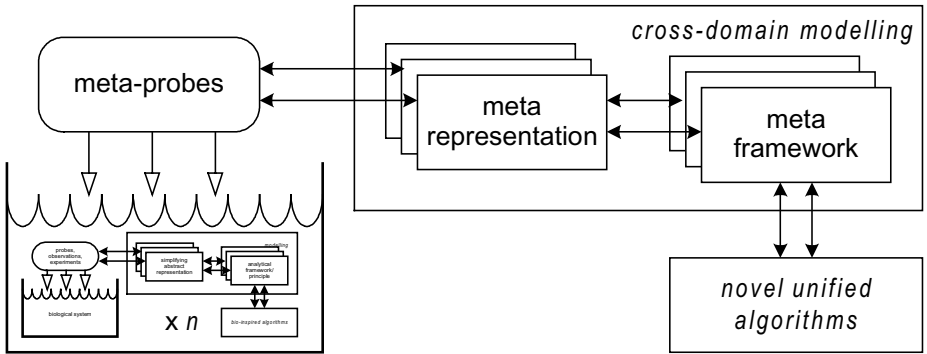


Fig. 3. An outline conceptual framework for integrating bio-inspired computational domains

behaviour in general; questions that address notions such as openness, diversity, interaction, structure, and scale might lead to models of complex adaptive systems. The idea is to ask each question (suitably refined and quantified) across the range of frameworks being incorporated, and to use the answers as part of the input to build the meta-framework.

Openness: We do not want our computations to halt; we want continual evolution, continual growth, continual addition of resources: that is, *open, far-from-equilibrium* systems. How much openness is necessary? How is openness controlled by structure and interaction? How is system unity maintained in the presence of openness?

Diversity (heterogeneity) is present in all complex biological systems, and occurs in structure, behaviour, and interactions. When can we talk of an *average* agent? How much diversity is necessary within a level of a structure? between levels? What does it cost? How does it combat fragility?

Interaction: Agents interact with their environment and with each other. What are the features of interaction within structural levels? between levels? What is the balance between computation and communication?

Structure: Biological systems have structure on a variety of levels, yet the levels are not crisply delineated. Are the levels we discern artefacts of our modelling framework? How can we recognise levels? When is a hierarchy an adequate structural model? How does structure affect interaction? What are the relationships between physical structures and information structures? What is the relationship with specialisation of function? with localisation of function?

Scale: Biological systems have a vast scale, a vast number of components. When and how does ‘more’ become ‘different’? What are the critical points/phase transitions? How small can a system be, and still be emergent? When is a system too big? How important is multi-scale modelling? What are the relationships between scale and diversity?

Generic questions additionally apply to each meta-probe question area X : What is the role of X within a system? What is the balance between X and not- X at the peak of complexity? How and when does X emerge? How does X evolve? How does physical embodiment affect X ? How can we exploit X ?

4.2 A Meta-framework for Population Models

Many bio-inspired algorithms are based on *populations* of agents trained to perform some task, or optimise some function. The most obvious one is the area of evolutionary algorithms, based on analogy to populations of organisms breeding and selecting to become “fitter” [Mitchell 1996]. In AIS, there are the positive and negative selection, and clonal selection algorithms. Swarm algorithms and social insect algorithms [Bonabeau 1999] are based on populations of agents whose co-operations result in problem solving. Even a neural network could be viewed as a population of neurons cooperating to perform a recognition task.

Given the number of underlying commonalities, it seems sensible to abstract a meta-framework from individual population based models. (Since these individual frameworks themselves do not yet exist, this section is somewhat meta-speculative!)

What are the key properties of population models, and how are they realised in the various individual models? Here we outline just a few similarities and differences of these models, which could be used in constructing a population based meta-framework.

All these models contain a population of individual agents. Members of the population usually exhibit a range of *fitnesses*, where these *fitnesses* are used when calculating a new population: fitter individuals have a greater effect on the composition of the next generation than do the less-fit individuals. The aim is to find a population that is sufficiently fit for the task at hand.

In evolutionary algorithms (EAs), a population of *chromosomes* reproduces in a *fitness landscape*. Fitter individuals are *selected* more frequently, to breed the next generation. When described in these terms, the clonal selection algorithm looks very similar: the population comprises a collection of *antibodies*, which proliferate in an *affinity landscape*. The higher affinity individuals are cloned more, and mutated less, when producing the next generation. Additionally, the lowest affinity cells are replaced by random cells (providing automatic diversity maintenance.) In swarm algorithms, a population of *particles* exists and adapts in a *fitness landscape*. Fitter individuals’ properties are *copied* more by the next generation. In ant colony algorithms, a population of paths exist in a local fitness (path component length) landscape. The use of components from fitter (shorter) paths are reinforced (by “pheromones”) in the next generation, which is then constructed by “ants” following pheromone trails.

In EAs, clonal AIS, and ant algorithms, the fitness of the entire population is evaluated and used for selection and construction of next generation. Swarm algorithm evaluate the fitness of each individual relative to the others in its *local neighbourhood*. (Some EA variants incorporate *nicheing*, which provides a degree of locality.)

In EAs, swarm and ants algorithms, the result is the fittest member of the final population. In clonal AIS, however, the result is the entire final population of detectors; the individual detectors are each partial and unreliable, yet their combined cooperative effect makes the full robust detector.

Such commonalities and differences as outline above, once exposed and analysed, can be used to suggest more general algorithms. For example, the natural diversity maintenance of clonal AIS suggests ways for similar mechanisms to be added to other population algorithms, in a less *ad hoc* manner than currently. Also, many population algorithms find themselves forced to add some form of *elitism* to preserve the best solution so far: clonal AIS is naturally elitist. One potentially interesting feature to

explore is the relationship between the natural locality of swarm algorithms, and the locality inherent in AIS Danger Theory.

Such a combination of models permits many of the meta-probe question outlined above to be asked. *Diversity* is a key question: how to maintain diversity within a single population, but additionally, should there be different “species”, too? *Interaction* with the environment (laying and sensing pheromones) is crucial in the ant algorithms, and with other agents (at least at the level of copying their behaviour) in swarm algorithms. Co-evolution, with its effect on mutual fitness landscapes, can be regarded as a form of interaction. What *scale*, that is, what population size, is appropriate? The probes also force us to think of new issues: is *openness* a relevant aspect? Should we be concerned with flows of agents into and out of the population (other than by internal mechanisms of generational breeding)? And is there any way to exploit *structure*, given the homogeneity of most population algorithms?

This somewhat simplistic meta-framework sketch is built on the correspondingly simplistic population models. More sophisticated population models developed in terms of full conceptual frameworks would doubtless lead to much richer and more powerful meta-frameworks.

4.3 A Meta-framework for Network Models

AIS networks, metabolic networks, auto-catalytic chemical reaction networks, intracellular protein interaction networks, inter-cellular cytokine, hormone and growth factor signalling networks, ecological food webs, are all examples of biological networks. Indeed, most biological processes operate through a complicated network of interactions, with positive and negative feedback control by factors that are themselves subject to similar controls. These networks function in a distributed fashion: most components have a variety of roles, and most functions depend on more than one component. This presumably underpins their robustness, whilst keeping the malleability required for adaptability and evolution. How this is achieved in practice is poorly understood.

Currently mathematical and computational descriptions of the structure of biological networks tend to be static (there is no time component to the architecture), closed (no inputs from the environment), and homogeneous (the types of nodes and connections are uniform, and new instances, and new kinds, of connections and nodes, are not supported). It will be necessary to develop novel mathematical approaches to model real complex biological networks. Developing these new mathematical models in the context of the proposed conceptual framework will provide mechanisms for evaluating their appropriateness and power.

5 Discussion and Conclusion

We have argued that bio-inspired algorithms would benefit from exploiting more sophisticated biological models, and from being based on sound analytical principles; we believe that biology could benefit from the resulting sophisticated models, too. We have outlined what we believe to be a suitable conceptual framework including these various components. We have suggested how AIS network models might fit into this framework.

We have additionally sketched how meta-frameworks, based on the same underlying structure, might be applied at higher levels to unify various kinds of bio-inspired architectures, and we have suggested how population based models, including AIS models, might form one such meta-framework. We do not expect that every individual model will fit perfectly into an integrated model: part of the development process will be to expose essential differences as well as to integrate common abstractions.

One exciting prospect of a unified meta-framework is the possibility of a *novel instantiation*, possibly using concepts from across a range of biological domains, and possibly using concepts from outside biology (since words like “Lamarck” and “teleology” need not be so necessarily dismissed in the artificial domain). This would allow the building of a chimerical computational framework that is inspired by biology, but not restricted to any one particular biological domain.

Acknowledgements. The conceptual framework described in this paper was developed during the EPSRC-funded EIVIS project. We would like to thank the other members of the project for their invaluable contributions to some of the ideas in this paper: Andrew Anderson, Jim Austin, Brian Bell, Peter Bentley, David Broomhead, Robin Callard, Steve Furber, David Halliday, Andy Hone, Douglas Kell, Alan Murray, Jaroslav Stark, Stefan Wermter, David Willshaw, Xin Yao, Peter Young.

References

- [1] U. Aicklen, P. Bentley, S. Cayzer, J. Kim, J. McLeod. Danger Theory: The Link Between AIS and IDS? In J. Timmis, P. Bentley, E. Hart (eds) *LNCS 2787*, 156-167, Springer, 2003
- [2] H. Bersini. Immune Network and Adaptive Control. *Proc. First European Conference on Artificial Life*, 217-226. MIT Press, 1991
- [3] H. Bersini. Reinforcement and Recruitment Learning for Adaptive Process Control. *Proc. Int. Fuzzy Association Conference (IFAC/IFIP/IMACS) on Artificial Intelligence in Real Time Control*, 331-337, 1992
- [4] H. Bersini, F. J. Varela. The Immune Learning Mechanisms: Reinforcement, Recruitment and Their Applications. In R. Paton (ed) *Computing with Biological Metaphors*, 166-192. Chapman & Hall, 1994
- [5] E. W. Bonabeau, M. Dorigo, G. Theraulaz. *Swarm Intelligence: from natural to artificial systems*. Addison Wesley, 1999
- [6] D. W. Bradley, A. M. Tyrrell. Immunotronics: Novel Finite State Machine Architectures with Built in Self Test using Self-Nonself Differentiation. *IEEE Transactions on Evolutionary Computation*, 6(3) 227-238, June 2002
- [7] L. N. de Castro, J. Timmis. *Artificial Immune Systems: A New Computational Intelligence Approach*. Springer, 2002
- [8] L. N. de Castro, F. J. Von Zuben. The Clonal Selection Algorithm with Engineering Applications. *Workshop on Artificial Immune Systems and Their Applications, Genetic and Evolutionary Computation Conference*, 36-37, 2000
- [9] L. N. de Castro, F. J. Von Zuben. aiNet: An Artificial Immune Network for Data Analysis. In H. A. Abbass, R. A. Sarker, C. S. Newton (eds) *Data Mining: A Heuristic Approach*, Chapter XII. Idea Group Publishing, 2001
- [10] J. D. Farmer, N. H. Packard, A. S. Perelson. The Immune System, Adaptation, and Machine Learning. *Physica D* 22 187-204, 1986

- [11] S. Forrest, A. Perelson, L. Allen, R. Cherukuri. Self-Nonself Discrimination in a Computer. *Proc. IEEE Symp. on Research in Security and Privacy*, 202-212, 1994
- [12] A. Freitas, J. Timmis. Revisiting the Foundations of Artificial Immune Systems. In J. Timmis, P. Bentley, E. Hart (eds) *ICARIS 2003*, LNCS 2787, 229-241. Springer, 2003.
- [13] L. Gammaitoni, P. Hanggi, P. Jung, F. Marchesini. Stochastic Resonance. *Rev. Mod. Phys.* **70**(1) 223-287, 1998
- [14] D. E. Goldberg. *Genetic Algorithms in Search, Optimization and Machine Learning*. Addison Wesley, 1989
- [15] R. R. Hightower, S. A. Forrest, A. S. Perelson. The Evolution of Emergent Organization in Immune System Gene Libraries. In L. J. Eshelman (ed) *Proc. 6th Int. Conf. on Genetic Algorithms*, 344-350. Morgan Kaufmann, 1995
- [16] J. H. Holland. *Adaptation in Natural and Artificial Systems*. University of Michigan Press, 1975
- [17] C. A. Janeway Jr, R. Medzhitov. Innate immune recognition. *Ann. Rev. Immunol.* **20** 197-216, 2002
- [18] N. K. Jerne. Towards a Network Theory of the Immune System. *Ann. Immunol. (Inst. Pasteur)* **125C** 373-389, 1974
- [19] P. Kourilsky, P. Truffa-Bachi. Cytokine fields and the polarization of the immune response. *Trends Immunol.* **22** 502-509, 2001
- [20] M. Neal. Meta-stable Memory in an Artificial Immune Network. In J. Timmis, P. Bentley, E. Hart (eds) *ICARIS 2003*, LNCS 2787, 168-180. Springer, 2003.
- [21] P. Matzinger. The danger model: a renewed sense of self. *Science* **296** 301-305, 2002
- [22] R. Medzhitov, C. A. Janeway Jr. Decoding the patterns of self and nonself by the innate immune system. *Science* **296** 298-300, 2002
- [23] M. Mitchell. *An Introduction to Genetic Algorithms*. MIT Press, 1996
- [24] D. A. Padgett, R. Glaser. How stress influences the immune response. *Trends Immunol.* **24** 444-448, 2003
- [25] S. Perelson. Immune Network Theory. *Imm. Rev.* **110** 5-36, 1989
- [26] S. Perelson. Modelling viral and immune system dynamics. *Nat Rev Immunol.* **2** 28-36, 2002
- [27] A. Romanyukha, A. I. Yashin. Age related changes in population of peripheral T cells: towards a model of immunosenescence. *Mech Ageing Dev.* **124** 433-443, 2003
- [28] J. Sprent, C.D. Surh. T cell memory. *Ann. Rev. Immunol.* **20** 551-579, 2002
- [29] D. Taylor, D. Corne. An Investigation of the Negative Selection Algorithm for Fault Detection in Refrigeration Systems. In J. Timmis, P. Bentley, E. Hart (eds) *ICARIS 2003*, LNCS 2787, 34-45. Springer, 2003.
- [30] J. Timmis. *Artificial Immune Systems: A Novel Data Analysis Technique Inspired by the Immune Network Theory*, Ph.D. Dissertation, Department of Computer Science, University of Wales, September. 2000
- [31] F. Varela, A. Coutinho, B. Dupire, N. N. Vaz. Cognitive Networks: Immune, Neural and Otherwise. In A. S. Perelson (ed) *Theoretical Immunology, part 2*, 359-375. Addison-Wesley, 1988