

Immuno-engineering

Jon Timmis¹, Emma Hart³, Andy Hone⁴, Mark Neal⁵ Adrian Robins⁶, Susan Stepney¹, and Andy Tyrrell²

Abstract In this position paper, we outline a vision for a new type of engineering: immuno-engineering, that can be used for the development of biologically grounded and theoretically understood Artificial Immune Systems (AIS). We argue that, like many bio-inspired paradigms, AIS have drifted somewhat away from the source of inspiration. We also argue that through an interdisciplinary approach, it is possible to exploit the underlying biology for computation in a way that, as yet, has not been achieved. Immuno-engineering will not only allow for the potential development of more powerful AIS, but allow for feed back to biology from computation.

1 Introduction

Advances in technology today enable the construction of complex autonomous systems, which can range in size from a robot, perhaps containing tens of simple devices, to mobile, ad-hoc networks containing thousands of such devices. At both extremes, such systems consist of unreliable heterogeneous sensors and actuators which must make decisions across multiple timescales in unpredictable, and potentially hostile, dynamic environments in order to maintain their integrity and achieve their desired functionality. Current technology allows us to hard-wire responses to foreseeable situations; a considerable void is still to be crossed however to achieve systems which adapt continuously and autonomously to their environments and exhibit what is becoming known as self-CHOP characteristics; self-configure, self-heal, self-optimize and self-protect. A paradigm shift in engineering is required to address this; we propose that a new discipline that will allow for the construction

Department of Computer Science, University of York, York, UK e-mail: jtimmis@cs.york.ac.uk · Department of Electronics, University of York, York, UK · School of Computing, Napier University, Edinburgh, UK · Institute of Mathematics, Statistics and Actuarial Science, University of Kent, Canterbury, Kent, UK · Department of Computer Science, Aberystwyth University, Aberystwyth, UK · University of Nottingham Medical School, Nottingham, UK

of engineered artefacts that are fit for purpose in the same way as their biological counterparts needs to be developed.

Our long-term aim is to develop the foundations for a new kind of engineering – *immuno-engineering* – exploiting principles derived from the human immune system to enable the engineering of robust complex artefacts.

In this position paper we outline the concept of Immuno-engineering and discuss its motivation in the context of current Artificial Immune System (AIS) research, and we hint at the way in which such a discipline may be developed.

We propose a bottom-up approach to the engineering of such systems, which will result in a set of immuno-engineering principles; these can be generalised to the future development of a wide range of bio-inspired, autonomic systems. This is achieved via an interdisciplinary approach which cuts across immunology, mathematics, computer science and engineering. In a recent paper [1] we discuss the interdisciplinary nature of AIS research. Our vision is inspired by recent work in immunology which attempts to reposition the immune system away from a pure defence mechanism to a complex, self-organising *computational system*, which computes the state of the body and then responds to it in order to achieve host maintenance and protection [2]. Autonomic systems which operate in a dynamic and information rich environment need to compute their state and then respond in an analogous way if they are to remain operational in order to continuously deliver the services expected of them.

2 Exploiting Immunology for Computation

Artificial Immune Systems (AIS) [3] is a diverse area of research that attempts to bridge the divide between immunology and engineering and are developed through the application of techniques such as mathematical and computational modeling of immunology, abstraction from those models into algorithm (and system) design and implementation in the context of engineering. Many early attempts to apply immunological inspiration to engineering began with efforts to mimic the perceived role of the natural immune system as a mechanism for identifying and then eliminating harmful pathogens from the body in a computer intrusion detection system [4]. However, work previous to this explored the immune system for inspiration in fault diagnosis [5] and control [6]. These investigations sparked a host of attempts to apply aspects of immunology to a wider range of engineering problems, and the reader is referred to the International Conference on Artificial Immune Systems (ICARIS) for a comprehensive collection of papers [7, 8, 9, 10, 11, 12]. Over recent years there have been a number of review papers written on AIS with the first being [13] followed by a series of others that either review AIS in general, for example, [14, 15, 16, 17, 18], or more specific aspects of AIS such as data mining [19], network security [20], applications of AIS [21], theoretical aspects [22] and modelling in AIS [23].

However, despite the many successes of the immune inspired approach we claim that the real potential of the approach has yet to be met [18]. We claim that this results from two limitations in the approach taken. Firstly, all of these applications have cherry picked one (or occasionally a few) features of the vertebrate immune system and attempted to apply them in isolation. Thus, we observe algorithms based on clonal selection e.g. [24], on negative selection [25], on idiotypic networks [26] and dendritic cells [27] with many recent developments in AIS being based around one of these four types of algorithms. Moreover, almost without exception there has been a tendency to exploit only the adaptive component of the vertebrate immune system. It is clear from immunological studies that the innate and adaptive components operate in tandem, and furthermore, regulate each other's effects. Therefore, by selecting only individual components of a complex, interacting system, a huge opportunity to exploit the true potential of the metaphor is being missed. Secondly, the focus on individual applications has followed an approach common to much bio-inspired research: an algorithm is designed and tuned empirically to a particular problem, thereby making it difficult to generalise any principles applicable to other applications. The EPSRC "Danger Theory" project¹, the outlines of which were proposed in [28] was the first attempt to combine current immunological experimentation with computational research. However, even it has focussed on a single application (intrusion detection [20]) and a single aspect of the immune system (danger theory) [27, 29, 30, 31]. Whilst this research has provided significant developments in the area of intrusion detection, we feel that for the area of AIS to progress we need to find more general principles that are applicable in a range of application areas. It should be noted that it is not feasible to capture the whole immune system in a single application, the sheer complexity would be overwhelming, however a focus on higher-level key properties, such as multiple-timescales of response, of the innate and adaptive components, may prove useful in their generic applicability to engineering.

Based on these observations, we feel the time is ripe for a *step change* in the development of AIS, through a principled engineering approach. We now discuss such an approach.

3 Defining Immuno-engineering

We follow Orosz's definitions of 'immuno-ecology' and 'immuno-informatics' [32]:

immuno-ecology : "the study of immunological principles that permit effective immunological function within the context of the immensely complex immunological network . . . the principles serve mainly to provide an infrastructure for the immune system."

¹ <http://www.dangertheory.com/>

immuno-informatics : “the study of the immune system as a cognitive, decision-making device . . . addresses mechanisms by which the immune system converts stimuli into information, how it processes and communicates that information, and how the information is used to promote an effective immuno-ecology . . . how the immune system generates, posts, processes, and stores information about itself and its environment”

and so we now define:

immuno-engineering : the abstraction of immuno-ecological and immuno-informatics principles, and their adaptation and application to engineered artefacts (comprising hardware and software), so as to provide these artefacts with properties analogous to those provided to organisms by their natural immune systems.

Immuno-engineering takes into account the differences between artificial systems and biological systems: for example, the different numbers, kinds, and rates of signals that need to be monitored and processed; the different kinds of decisions that need to be made; the different effectors available to support and implement those decisions; and the different constraints of embodiment, either physically or virtually engineered. For example, Orosz [32] hypothesises that the major design features of the biological immune system that provides speed, flexibility and multiple response options rely on a parallel-processing system which has ‘wasteful’ use of resources, countless back-up systems, and requires the ability to immediately and continuously monitor physical sites. This is of enormous consequence to the engineer, who is constrained by processing speeds, communication overheads, and physical resources, and furthermore hindered by hardware requirements such as transmitting signals from sensors, but who can freely make numerous copies of software agents, subject only to storage constraints.

3.1 A Conceptual Framework for the Development of Immuno-engineering

In their paper, Stepney *et al.* [33] propose that bio-inspired algorithms, such as AIS, are best developed in a more principled way than was currently being undertaken in the literature. To clarify, the authors suggested that many AIS recently developed had drifted away from the immunological inspiration that had fueled their development and that AIS practitioners were failing to capture the complexity and richness that the immune system offers. In order to remedy this, the authors suggest a conceptual framework for developing bio-inspired algorithms within a more principled framework that attempts to capture biological richness and complexity but, at the same time, appreciate the need for sound engineered systems that need to work. This should avoid the “reasoning by metaphor” approach often seen in bio-inspired computing whereby algorithms are just a weak analogy of the process on which they are based, being developed directly from (often naive) biological models and observations. One of the main problems involved in designing bio-inspired

algorithms is deciding which aspects of the biology are necessary to generate the required behaviour and which aspects are surplus to requirements. Thus, the conceptual framework takes an interdisciplinary approach, involving the design of AIS through a series of observational and modelling stages in order to identify the key characteristics of the immunological process on which the AIS will be based. The first stage of the conceptual framework, as outlined in figure 1, aims to probe the biology, utilising biological observations and experiments to provide a partial view of the biological system from which inspiration is being taken. This view is used to build abstract models of the biology. These models can be both mathematical and computational, and are open to validation techniques not available to the actual biological system. From the execution of the models and their validation, insight can be gained into the underlying biological process. It is this insight that leads to the construction of the bio-inspired algorithms. This whole process is iterative, and can also lead to the construction of computational frameworks that provide a suitable structure for specific application-oriented algorithms to be designed from.

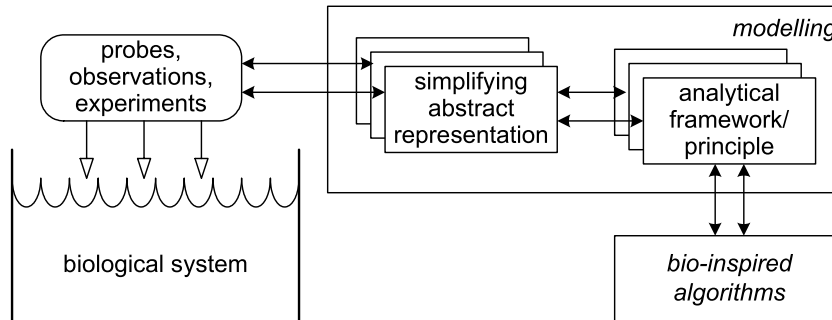


Fig. 1 The Conceptual Framework [33]. This can be seen as a methodology to develop novel AIS allowing true interaction between disciplines where all can benefit, and, a way of thinking about the scope of AIS and how that has broadened over the years once again

As noted by Stepney *et al.* [33] each step in the standard conceptual framework is biased, be it modelling some particular biology mechanism or designing an algorithm for which there is an intended end product or specific concept. The first instantiations of the conceptual framework will produce models specific to certain biological systems and algorithms for solutions to specific problems. One could attempt to produce a computational framework based on some biology without a particular end algorithm/application in mind, that is examining biology and hoping to come across something applicable to a generic computational problem. This, however, would seem to be a very difficult task and one has to ground the development of AIS in some form of application at some point. Therefore, it is far easier to orient these steps toward some particular problem giving necessary focus to the modelling work [34].

4 Towards Immuno-engineering

From an *engineering* perspective, there is a need to design, develop and implement *design libraries* derived from the immuno-engineering principles, tested in diverse practical exemplars ranging from a self-contained and self-maintaining piece of hardware, such as a network switch that reduces downtime and error-propagation on the internet, to a ubiquitous sensing system with reliable message passing that could be embedded into buildings resulting in a system that could rapidly detect and localise survivors following building collapse. From a *biological* perspective, by focussing on the immune system as a computational system we will deliver a framework in which it is possible to reframe experimental immunological data and ask new experimental questions. For example, we might ask how the state of a developing tumor can influence the state of the tissues and therefore how can we induce the immune system to compute this state as abnormal. Recent work in [2] has developed the notion of the “computation of the state” of the immune system: as stated earlier, an immuno-engineering approach should take this into account. Such an interdisciplinary endeavor will thus potentially impact on the understanding of disorders of immune activity such as autoimmune diseases and cancer.

To provide this bridge between immunology and engineering, there is a potential need to utilise state-based modelling techniques which fit well with a computational view of the systems. On the one hand, these will provide a realistic and intuitive environment for immunologists to complement traditional mathematic modelling such as differential equations and, on the other, they can be readily transformed into engineering solutions. One such approach is the π -calculus [35]. This is a formal language used to specify concurrent computational systems. Its defining feature that sets it apart from other process calculi is the possibility of expressing mobility. This allows processes to “move” by dynamically changing their channels of communication with other processes, thus one can model networks that reconfigure themselves. The π -calculus allows composition, choice, and restriction of processes which communicate on potentially private complementary channels. There is a growing similarity between the parallelism and complexity of computer systems today and biological systems. As noted by [36] computational analysis tools such as the π -calculus are just as applicable to biology as they are to computing.

4.1 What do we need to do for Immuno-engineering?

The properties we wish to endow on engineered systems are currently exhibited only by those complex living systems whose immune systems comprise an *innate* component which endows the host with rapid *pre-programmed* responses and an *adaptive* component which is capable of learning through experience. Much of the desired functionality of the system arises from the interplay between these subsystems and the regulatory effect they have on each other. Together, these operate over multiple timescales, from seconds to the entire lifetime of the organism. Therefore,

the modelling of both innate and adaptive components, paying particular attention to the interface between them, enables us to push the boundaries of biologically inspired computing and engineering.

In order to achieve our aim of laying the foundations for immuno-engineering, a number of key objectives need to be achieved:

1. **derive** mathematical models of the interplay between the innate and adaptive immune systems
2. **develop** and **verify** computational models that capture the interplay of innate and adaptive immunity
3. **implement** an immuno-engineering design and implementation library
4. **develop** and **assess** immuno-engineering insights to inform modulation of the natural immune system
5. **deploy** and **evaluate** the immuno-engineering library in a diverse set of case studies

5 Instantiating Immuno-engineering

In order to develop the Immuno-engineering approach a combination of the conceptual framework [33] and the problem-oriented perspective [34] can be adopted and requires interactions between computation, mathematical analysis, practical implementation, and biological experimentation. These will be rooted in the conceptual framework [33], which formulates the principled abstraction of bio-inspired algorithms through a process of mathematical modelling, computational modelling, and algorithm development for application domains, and makes use of the problem-oriented perspective [34] through the use of case studies to develop and refine the Immuno-engineering libraries.

Work should be based on a combination of mathematical and computational modelling, which leads to the development of an immuno-engineering library. This library should be tested on a number of carefully selected, realistic and diverse case studies that exhibit a broad and diverse spectrum of engineering features, thus allowing for the refinement of the approach. The library should thus be exercised in various different forms (in hardware or software, in open or closed environments), and therefore be tested and evaluated. Such suitable case studies might include web-mining, condition-monitoring and distributed sensing, in line with suggestions for future applications of AIS [21] which emphasises the notion of dynamic, life-long learning and homeostasis. In addition, elements of our immuno-engineering library should be used to computationally model aspects of the human immune system, which will help inform *in vitro* experiments: hence, all disciplines in this endeavor will benefit from the whole approach.

Figure 2 illustrates the interdisciplinary nature of the work. In order to develop the Immuno-engineering approach we would advocate focussing on the development of mathematical models of interactions between the innate and adaptive im-

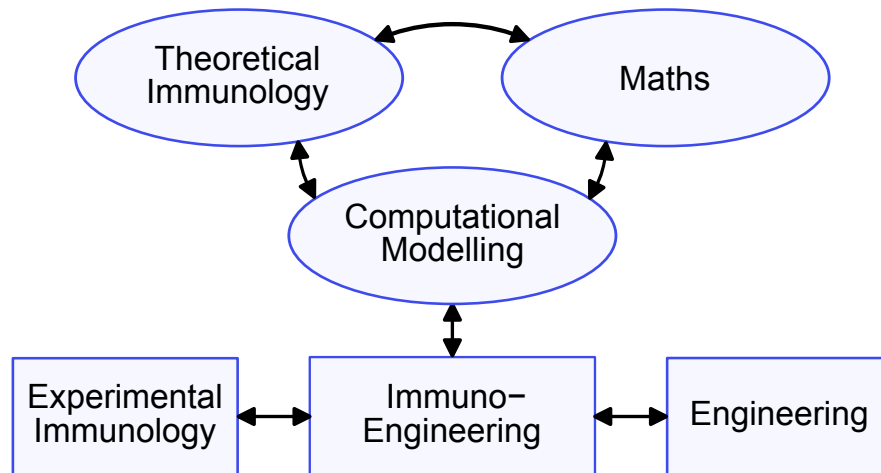


Fig. 2 Interactions between disciplines that leads to the development of immuno-engineering which itself acts as the bridge between experimental immunology and engineering.

immune systems, capturing the essentials of immuno-ecology (the interaction between theoretical immunology and mathematics). Once mathematical models have been developed, we can proceed to develop computational models from them: this will lay the foundation for developing immuno-engineering and develop an *immuno-engineering library*. This library will *bridge the gap* between experimental immunology and engineering, thus breaking the mould of typical biologically inspired computing, which simply jumps from simplistic views of biological systems straight to simplistic engineered solutions, in line with the approach advocated in [33]. With the adoption of this method, it should be possible to generate a set of mathematically sound, biologically grounded techniques applicable to engineering. In addition, the library will drive further investigations into experimental immunology.

5.1 Modelling and Immuno-Ecology

Modelling provides us with a fundamental insight into the workings of the immune system. Inherent in our proposal is the desire to gain a detailed understanding of immunological principles which will ultimately lead to the development of the immuno-engineering library. Modelling affords us the opportunity to investigate a complex system from different perspectives: from the level of individual components (molecules and cells), to the level of populations of cells, to an overall systems level. For the modelling, a wide variety of options are open, all with their own advantages and disadvantages [23] such as dynamical systems, optimal control theory, information and coding, probability, stochastic π -calculus and complex network theory.

5.1.1 Modelling of Information Processing

A central part of the interaction between the innate and adaptive immune system is the process by which antigen is recognised. In order to derive mathematical models of *antigen processing*, one could adopt an information-theoretic approach to the identification of antigens, regarding them as salient chunks of data to be processed by the immune system. Antigens usually take the form of proteins which are recognised as being “foreign” by the immune system, and, for this recognition to occur, suitable features of the structure and composition of the protein must first be isolated. The latter is achieved by immune cell receptors which bind to specific chunks of the antigen. Regarding proteins as streams of data will enable the use of information theory to model the processing of such data.

Having obtained an abstract representation of antigen processing, it should then be possible to produce dynamic models of the *regulation* of receptor-bearing agents, based on, for example, optimal control theory. These would involve both continuous and discrete dynamics (differential and difference equations respectively). By performing analytical studies of the mathematical models, their asymptotic behaviour can be determined. These dynamical systems will also be converted into numerical simulations suitable for developing computational models, which will provide numerical predictions that can be compared with the results of the development of the algorithms.

5.1.2 Modelling of Network Topologies

It is possible to examine the way that the overall response is mediated by the network of signals created by cytokines. A sensible approach would be to extract some generic topological features by focusing on particular small subsets of the immune network, chosen for the relevance to the desired engineering properties. In order to do this one could look towards complex network theory, using as a starting point the models of Barabasi, Watts & Strogatz, and extending this with ideas of Alon on network motifs for biological systems [37]. Specific subsets of immunological networks to be examined include: self regulation and switching between immune cell responses such as T-helper cells e.g. Th1/Th2; innate modulation of adaptive immune response via the complement system (which helps clear pathogens from an organism); mutually inhibitory effects of cytokines such as IL4 and IF γ upon one another.

Network models of specific subsystems such as these will be developed, and can then be tested against experiments, leading to further refinement of these models. Through the biological experiments it should be possible to verify features of the local topology of subsystems within immune networks. It would be possible to investigate how the particular architecture of these networks affects the robustness of their response, allowing the selection of suitable network motifs to be incorporated into the development of the Immuno-engineering library.

5.2 *Immuno-engineering Library*

Adopting a problem-oriented approach [34] it should be possible to determine the immuno-requirements of the case studies helping to drive the research and determine the relevant Immuno-Ecology and Immuno-Informatics principles. From these bio-specific principles, it will be possible to develop abstract descriptions in the form of UML (statecharts, sequence diagrams) and develop design patterns [38] that capture salient properties and encapsulate constraints. Their properties could be analysed to ensure that they have not lost the desired immuno-properties during abstraction. This should be iterated as appropriate, incorporating extra components as discovered through the modelling work and the case studies. The output of this task would comprise generative pattern languages of immuno-engineering analysis and design: tools that aid a system designer in analysing a specific application, and in applying immuno-engineering principles and concepts to its particular sensors, tasks, and embodiment that respects the biological underpinnings.

The abstractions and pattern languages that have now been developed will provide the foundation for building a software library for developing immuno-engineered systems. This would result in an implementation library in the form of architectures and algorithms that can be instantiated in the case studies. One of the fundamental properties of the immune system is that it is a highly parallel system of communicating agents. Therefore, consideration should be given to potential parallel processing architectures, platforms and programming languages and determine which will allow the immuno-engineering properties to be implemented effectively, and what physical constraints they will impose. For example, technologies such as VHDL, JSCP and *occam- π* allow for the development of truly parallel software systems and allow for a natural mapping from the stochastic- π calculus.

5.3 *Experimental Immunology*

Central to the concept of immuno-engineering is the ability to employ immuno-engineering principles in the context of actual experimental immunology. Immuno-engineering principles and mathematical models developed through this process should direct experimental work. Indeed, one of the major motivations for the use of modelling in experimental work is to make use of the predictive nature of the models and to tie them closely to experimental work: otherwise models are developed in a vacuum. We do not expect that work from experimental immunology will feed directly into the development of algorithms, but rather assist in the validation of immuno-engineering principles, which will then feedback into the development of models which then influences the library development. We now discuss possible avenues for experimental work in the context of immuno-engineering.

5.3.1 Cytokine Interactions

As a starting point for part experimental work, we would propose the testing of cytokine interaction models which can be done using in-vitro systems. In such experiments, control of the T-helper subset balance by innate signals derived from allergen can be demonstrated: proteolytically active allergen favoured the development of T-cells producing of IL4 (T-helper 2, Th2), and reduced numbers of interferon gamma producing T- cells (T-helper 1, Th1).

5.3.2 Regulatory Interactions

From insights gained from the modeling and immuno-engineering phases, these models can be developed to study regulatory interactions between cytokine producing subsets of cells. Multi-parameter flow cytometry and intracellular staining can be used (which allows visualisation of actual immune cells in a test-tube) to analyse cytokine profiles at the single cell level. Such techniques are able to analyse 6 colours of fluorescence simultaneously, allowing detailed characterisation of large populations (at least 10^6 individual cells) of responding lymphocytes. This approach has been used to model signalling molecule relationships [39]. The complexity of cytokine mediated control mechanisms has recently been emphasised by the description of a new helper T-cell subset producing interleukin 17 (Th17), which may have a critical role in immune mediated tissue damage and cancer [40]. Modelling of these complex interactions may provide a basis for intervention in this important group of diseases

5.4 Case Studies and a Comment on Applications of AIS

The case studies should be used to utilise and evaluate the architectures and algorithms developed as part of the above process. In combination, such case studies should display a range of immuno-informatics characteristics e.g. virtual, physical, open, dynamic and a range of space and time scales. The diversity of these characteristics should ensure that after case study development and feedback, the applicable scope of the immuno-engineering library will be sufficiently broad to support a wide range of applications.

Hart and Timmis [41] state that considering the application areas to date, AIS have been reasonably successful but, as yet, do not offer sufficient advantage over other paradigms available to the engineer. To address this and therefore tap the unexploited potential of AIS, one of the suggestions they make is that life-long learning is a key property of the immune system but true life-long learning, whereby a system is required to improve its performance as a consequence of its lifetime's experience, has not been utilised in AIS. Hart and Timmis propose a list of features they believe

AIS will be required to possess a combination of, if the field of AIS is to carve out a computational niche. These future AIS features, quoted verbatim from [21], are:

1. They will exhibit *homeostasis*
2. They will benefit from interactions between *innate* and *adaptive* immune models
3. They will consist of *multiple, interacting components*
4. Components can be easily and naturally *distributed*
5. They will be required to perform *life-long learning*

As for the future roles of AIS, Garrett [17] states that the biggest difficulty facing AIS is the lack of application areas to which it is clearly the most effective method. It is suggested that hybrid AIS may help to provide more powerful methods to solve certain problems. The current types of AIS used are also classified by Garrett into those that detect antigens (negative selection and danger theory models), and those that focus on destroying them (clonal selection and immune network models). It is pointed out, however, by Garrett [17] that the immune system has more to offer than this, with the mechanisms of the innate immune system and the view that the immune system is a homeostatic control system, being highlighted as future areas for AIS to exploit, and indeed this call seems to be taken up in a small part in recent times [42, 43, 44]. Concurring with the above, Bersini [45] argues that the immune system is much more than a simple classifier and performs much more than “pattern matching” and urges people in AIS to think about applications that are far removed from such applications which are the dominant force in AIS [23]. This challenges the community to find that *niche* application that AIS alone can tackle. This may come in the form of certain engineering type applications such as robotics and real-time systems where the system is embodied in the world and needs to be able to cope with extreme challenges that are constantly changing. Adopting an Immuno-engineering approach should enable us to begin to tackle this challenges.

6 Conclusions

In this paper, we have presented a new way of thinking about the development of immune-inspired systems and we have proposed the Immuno-engineering approach. Currently, like many bio-inspired paradigms, Artificial Immune Systems (AIS) are pale counterparts to their natural system. This is not to say that AIS should be like an immune system or copy exactly what the immune systems does, this would lead to not only complexity issues when instantiating such AIS, but also conceptual issues: just because the immune system does something it does not mean that we should do the same in an engineering context. Rather, we advocate the interdisciplinary interaction to develop biologically grounded, theoretically understood and well tested Immuno-engineering principles that can be deployed in a wide variety of application domains. Should such an Immuno-engineering approach be developed, we believe that AIS will then begin to not only capture the computationally interesting prop-

erties of the immune system, but be able to make a significant contribution to the immunology that serves as its inspiration.

Acknowledgements

The authors would like to thank Paul Andrews, Alex Freitas, Andy Greensted, Colin Johnson, Nick Owens, Jamie Tycross and Joanne White for input into some of the ideas behind Immuno-engineering. Jon Timmis would like to acknowledge the EPSRC grant EP/E005187/1² during which some ideas in the paper have been developed.

References

1. Timmis, J., Andrews, P.S., Owens, N., Clark, E.: An interdisciplinary perspective on artificial immune systems. *Evolutionary Intelligence* **1**(1) (2008) 5–26
2. Cohen, I.R.: Real and artificial immune systems: Computing the state of the body. *Imm. Rev.* **7** (July 2007) 569–574
3. de Castro, L.N., Timmis, J.: *Artificial Immune Systems: A New Computational Intelligence Approach*. Springer (2002)
4. Forrest, S., Perelson, A., Allen, L., R.Chelukuri: Self-nonsel self discrimination in a computer. In: *IEEE Symposium on Research in Security and Privacy*, Los Alamos, CA, IEEE Computer Society Press (1994)
5. Ishida, Y.: Fully distributed diagnosis by pdp learning algorithm: Towards immune network pdp model. In: *Proc. of the Int. Joint Conf. on Neural Networks*. (1990) 777–782
6. Bersini, H., Varela, F.J.: Hints for adaptive problem solving gleaned from immune networks. In Schwefel, H., Manner, R., eds.: *Proc. of the First Conference on Parallel Problem Solving from Nature*. Springer-Verlag, Berlin, Germany (1991)
7. Timmis, J., Bentley, P.J., eds.: *Proceedings of the 1st International Conference on Artificial Immune Systems (ICARIS 2002)*, University of Kent Printing Unit (2002)
8. Timmis, J., Bentley, P., Hart, E., eds.: *Proceedings of the 2nd International Conference on Artificial Immune Systems (ICARIS 2003)*, LNCS 2787, Springer (2003)
9. Nicosia, G., Cutello, V., Bentley, P.J., Timmis, J., eds.: *Proceedings of the 3rd International Conference on Artificial Immune Systems (ICARIS 2004)*. LNCS 3239, Springer (2004)
10. Jacob, C., Pilat, M., Bentley, P., Timmis, J., eds.: *Proc. of the 4th International Conference on Artificial Immune Systems (ICARIS)*. Volume 3627 of *Lecture Notes in Computer Science*., Springer (2005)
11. Bersini, H., Carneiro, J., eds.: *Proc. of 5th International Conference on Artificial Immune Systems*. *Lecture Notes in Computer Science*, Springer (2006)
12. de Castro, L.N., Von Zuben, F.J., Knidel, H., eds.: *Proceedings of the 6th International Conference on Artificial Immune Systems*. Volume 4628 of *Lecture Notes in Computer Science*. Springer (2007)
13. Dasgupta, D., ed.: *Artificial Immune Systems and their Applications*. Springer (1999)
14. de Castro, L.N., Von Zuben, F.J.: *Artificial immune systems: Part I—basic theory and applications*. Technical Report DCA-RT 01/99, School of Computing and Electrical Engineering, State University of Campinas, Brazil (1999)

² <http://www.bioinspired.com/research/xArch/index.shtml>

15. de Castro, L.N., Von Zuben, F.J.: Artificial immune systems: Part II—a survey of applications. Technical Report DCA-RT 02/00, School of Computing and Electrical Engineering, State University of Campinas, Brazil (2000)
16. Ji, Z., Dasgupta, D.: Artificial immune system (AIS) research in the last five years. In: Congress on Evolutionary Computation. Volume 1., Canberra, Australia, IEEE (December 8–12 2003) 123–130
17. Garrett, S.: How do we evaluate artificial immune systems? *Evolutionary Computation* **13**(2) (2005) 145–177
18. Timmis, J.: Artificial immune systems: Today and tomorrow. *Natural Computing* **6**(1) (Feb. 2007) 1–18
19. Timmis, J., Knight, T.: Artificial immune systems: Using the immune system as inspiration for data mining. In: *Data Mining: A Heuristic Approach*. Idea Group (2001) 209–230
20. Kim, J., Bentley, P., Aickelin, U., Greensmith, J., Tedesco, G., Twycross, J.: Immune system approaches to intrusion detection - a review. *Natural Computing* **in print** (2007)
21. Hart, E., Timmis, J.: Application areas of AIS: The past, the present and the future. *Applied Soft Computing* **8**(1) (2008) 191–201 In Press, Corrected Proof, Available online 12 February 2007.
22. Timmis, J., Hone, A., Stibor, T., Clark, E.: Theoretical advances in artificial immune systems. *Journal of Theoretical Computer Science* **In press**(doi:10.1016/j.tcs.2008.02.011) (2008)
23. Forrest, S., Beauchemin, C.: Computer Immunology. *Immunol. Rev.* **216**(1) (2007) 176–197
24. de Castro, L.N., Von Zuben, F.J.: Learning and optimization using the clonal selection principle. *IEEE Transactions on Evolutionary Computation* **6**(3) (2002) 239–251
25. Gonzalez, F.A., Dasgupta, D.: Anomaly detection using real-valued negative selection. *Genetic Programming and Evolvable Machines* **4**(4) (2003) 383–403
26. Neal, M.: Meta-stable memory in an artificial immune network. [8] 168–180
27. Greensmith, J., Aickelin, U., Cayzer, S.: Introducing dendritic cells as a novel immune-inspired algorithm for anomaly detection. [10]
28. Aickelin, U., Bentley, P., Cayzer, S., Kim, J., McLeod, J.: Danger theory: The link between AIS and IDS? [8] 147–155
29. Bentley, P.J., Greensmith, J., Ujjin, S.: Two ways to grow tissue for Artificial Immune Systems. [10] 139–152
30. Twycross, J., Aickelin, U.: Towards a conceptual framework for innate immunity. [10] 112–125
31. Greensmith, J., Aickelin, U., Twycross, J.: Articulation and clarification of the dendritic cell algorithm. [46] 404–417
32. Orosz, M.: An Introduction to Immuno-Ecology and Immuno-Informatics. In: *Design Principles from the Immune System*. Sante Fe (2001) 125–150
33. Stepney, S., Smith, R., Timmis, J., Tyrrell, A., Neal, M., Hone, A.: Conceptual frameworks for artificial immune systems. *Int. J. Unconventional Computing* **1**(3) (2006) 315–338
34. Freitas, A., Timmis, J.: Revisiting the foundations of artificial immune systems for data mining. *IEEE Trans. Evol. Comp.* **11**(4) (2007) 521–540
35. Milner, R.: *Communicating and Mobile Systems: the π -Calculus*. Cambridge University Press (1999)
36. Phillips, A., Cardelli, L.: Efficient, correct simulation of biological processes in the stochastic pi-calculus. In: *Proceedings of Computational Methods in Systems Biology (CMSB'07)*. Volume 4695. (2007) 184–199
37. Alon, U.: Uri alon, network motifs: theory and experimental approaches. *Nature Reviews Genetics* **8** (2007) 450–461
38. Gamma, E., Helm, R., Johnson, R., Vlissides, J.: *Design Patterns*. Addison Wesley (1995)
39. Sachs, K., Perez, O., Pe'er, D., Lauffenburger, D., Nolan, G.: Causal protein-signaling networks derived from multiparameter single-cell data. *Science* **308** (2005) 523–529
40. Steinman, L.: A brief history of t(h)17, the first major revision in the t(h)1/t(h)2 hypothesis of t cell-mediated tissue damage. *Nature Medicine* (2007) 139–145
41. Hart, E., Timmis, J.: Application areas of AIS: The past, the present and the future. [10] 483–497

42. Owens, N., Timmis, J., Greensted, A., Tyrrell, A.: On immune inspired homeostasis for electronic systems. [12] 216–227
43. Davoudani, D., Hart, E., Paechter, B.: An immune-inspired approach to speckled computing. [12] 288–299
44. Guzella, T., Mota-Santos, T., Caminhas, W.: Towards a novel immune inspired approach to temporal anomaly detection. [12] 119–130
45. Bersini, H.: Immune system modeling: The OO way. [46] 150–163
46. Bersini, H., Carneiro, J., eds.: Proceedings of the 5th International Conference on Artificial Immune Systems. Volume 4163 of LNCS. Springer (2006)