

J. Timmis, P. Andrews, N. Owens and
E. Clark

An Interdisciplinary Perspective on Artificial Immune Systems

Received: date / Accepted: date

Abstract This review paper attempts to position the area of Artificial Immune Systems (AIS) in a broader context of interdisciplinary research. We review AIS based on an established conceptual framework that encapsulates mathematical and computational modelling of immunology, abstraction and then development of engineered systems. We argue that Artificial Immune Systems are much more than engineered systems inspired by the immune system and that there is a great deal for both immunology and engineering to learn from each other through working in an interdisciplinary manner.

Keywords artificial immune systems, immunological modelling, mathematical modelling, computational modelling, applications of artificial immune systems, immune inspired computing, immunocomputing, computational immunology

1 Introduction

Artificial Immune Systems (AIS) 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. Over recent years there have been a number of review papers written on AIS with the first being [25] followed by a series of others that either review AIS in

Paul Andrews is supported by EPSRC grant number EP/E053505/1, Nick Owens is supported by EP/E005187/1 and Ed Clark by EP/D501377/1.

Department of Computer Science and Department of Electronics
University of York, Heslington, York. YO10 5DD. UK
Tel.: +44 1904 432348
E-mail: jtimmis@cs.york.ac.uk

general, for example, [29, 30, 68, 43, 103], or more specific aspects of AIS such as data mining [107], network security [71], applications of AIS [58], theoretical aspects [103] and modelling in AIS [39]. In the context of this paper, we feel that it would not be appropriate to attempt to reproduce the effort of those papers. Rather, the aim of this paper is to draw together ideas from the majority of these papers into a single review which forms a current opinion and review of the state of Artificial Immune Systems research. Therefore this paper is not an extensive bibliography of all AIS research, as this has been done by the other review papers, but we hope a good resource for researchers new to the area of AIS, and existing practitioners of AIS. For information, a good resource of the latest developments in AIS is the International Conference on Artificial Immune Systems (ICARIS ¹) conference series dedicated to AIS [105, 104, 85, 63, 9, 31] where there are an extensive number of papers on all aspects of AIS.

AIS has become known as an area of computer science and engineering that uses immune system metaphors for the creation of novel solutions to problems. Whilst this forms the majority view, we argue that the area of AIS is much wider and is not confined to the development of new algorithms. In a recent paper, Cohen [19] concurs with this view and in fact goes onto define three types of AIS scientists. The first are those of the “literal” school that build systems *in silico* to try and do what the actual immune system does (e.g. build computer security systems that discriminate between self and non-self); those of the “metaphorical” school that look for inspiration from the immune system and build computational systems with the immune system in mind (so the application may be far from analogous to what the immune system does) and a third school of people who aim to understand immunity through the development of computer and mathematical models. It is in this vein that our paper is written, and we would like to broaden the understanding of what AIS is all about, thus driving the area into a true interdisciplinary one of genuine interaction between immunology, mathematics and engineering.

Cohen [19] discusses the notion of the immune system using a “*computational strategy*” to carry out it’s functions of protecting and maintaining the body. An interesting analogy is made to the universal Turing machine that transforms input, which is represented as a sequence of information on a tape, to output, again information on a tape, and this machine operates to a set of rules. He raises interesting questions as to “*what does the immune system compute*” and “*what might we gain from thinking about immune computation?*”. Cohen’s main argument is that the immune system computes the state of the organism, based on a myriad of signals, which endows the immune system the ability to maintain and protect the host. Cohen [19] urges the immunological community to embrace working with computational scientists to aid the understanding of the nature of immune computation: this is, in part, the same spirit of this paper. In recent years, the area of AIS has begun to return to the immunology from which the initial inspiration came. For example, works by Stepney *et al.* [96], Twycross and Aickelin [109], Andrews and Timmis [3], Bersini [8] and Timmis [103], all advocate a deeper understanding of the immune system, in part through the use of modelling

¹ <http://www.artificial-immune-systems.org>

techniques, which will lead to the development of richer, more effective immune inspired engineered systems. This theme underpins our review paper, as we have attempted to structure it in such a way as to reflect the nature of AIS research today, that is one that encompasses (or can encompass) a range of activities from modelling immune systems to engineering systems. To this end our paper is structured as follows: section 2 describes the conceptual framework that we will use as a structure for the paper and discusses the process of going from immunology to engineered systems (indeed, the conceptual framework [96] can be seen as a methodology for the development of bio-inspired systems); section 3 provides a brief overview of the immunology that, has to date, provided the inspiration to AIS ranging from clonal selection theory to danger theory and immune cognition; section 4 provides an overview of mathematical and computational modelling approaches employed to understand the immune system covering topics from differential models, to π -calculus, to state charts, discussion is given as to how this can benefit the development of engineered systems; section 5 provides an introduction to the basics of AIS through a simple engineering framework and we provide an overview of the most common immune algorithms in use at the moment; section 6 reflects on the current opinion in the AIS community and finally section 7 reviews various comments on future directions for AIS research.

2 A Framework for thinking about Artificial Immune Systems

As we have outlined in section 1 there has been a gradual shift in AIS towards paying more attention to the underlying biological system that serves as inspiration, and taking time both to develop abstract computational models of the immune system (to help them understand computational properties of the immune system) and work closer with immunologists to better understand the biology behind the system. This does not mean to say that AIS researchers are now only focussed on the biology, but it would be fair to say that AIS is becoming a more interdisciplinary topic where people are working more on the *biological* aspects and others on the more *engineering* aspects. To highlight this, in a recent paper by Stepney *et al.* [97] (extended in [96]) suggest a methodology for the development of AIS was proposed that takes this shift into account. We will discuss that methodology here, however we also propose that this methodology is a good way to describe AIS in its current form, and indeed has formed the general structure for this paper. In addition, concurring with a view of Andrews and Timmis [3], Bersini [8] makes the argument that the AIS practitioner should take more seriously the role of modelling in the understanding and development of immune inspired solutions, and adopt a more “artificial life” approach. Indeed, Bersini makes a compelling argument for undertaking such an “Alife” approach based on pedagogy and the study of emergent phenomena and qualitative predictions, all of which are beneficial to the immunologist and ultimately engineers. Whilst we have a great deal of sympathy with this view, and indeed advocate the approach, we feel this needs to be tempered by the consideration of the engineering aspects, as after all, it is better engineered solutions that are

the driving force behind the vast majority of research being undertaken in AIS. This is to say that we feel both the approach encouraged by Bersini and the problem oriented approach proposed by Freitas and Timmis [41] can sit together, and this can be achieved via the conceptual framework approach [97,96].

In their paper, Stepney *et al.* [97] 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 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.

As noted by Stepney *et al.* [96] 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




Fig. 1 The Conceptual Framework [96]. 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

these steps toward some particular problem giving necessary focus to the modelling work [41].

3 A Guide to the Immunology of AIS

AIS have been inspired by many different aspects of the human immune system. One of the first questions that might be asked is why, as engineers and mathematicians, are we interested in the immune system? The answer is that the immune system exhibits a number of computationally appealing properties such as pattern recognition, learning, memory and self-organisation. In this section we present an overview of much of the immunology that has inspired AIS to give the reader a better appreciation of the discussions on AIS that follow. It is noted that we do not provide a fully comprehensive background on immunology, for that the reader is referred to [46] and [66].

There are a number of competing theories in immunology as to how the immune system actually achieves host protection and maintenance. In this section we review the majority of these theories and focus primarily on the ones that have acted as inspiration to the development of immune-inspired engineered systems. We do not pass judgement here as to which theory is correct or not, these are widely debated issues in immunology, nevertheless, each theory is interesting and useful from a computational and engineering perspective.

3.1 Overview of the Immune System

The immune system is typically described as a defense system that has evolved to protect its host from pathogens (harmful micro-organisms such as bacteria and viruses) [46]. It comprises a variety of specialised cells and molecules along with immune organs that provide a place for the immune cells to mature and function. The interactions between immune cells and

other cells of the body create a rich and complex set of immune behaviours, resulting in the recognition of pathogens and the evocation of a suitable pathogen ridding response.

The vertebrate immune system can be split functionally into two components, the innate immune system and the adaptive (or acquired) immune system. The innate immune system incorporates general pathogen defence mechanisms that have evolved over the *germline* of the organism. These mechanisms remain essentially unchanged during the lifetime of an individual and include the inflammatory response, phagocytosis (ingestion of pathogens by specialised immune cells), and physiologic barriers such as temperature. The mechanisms of the adaptive immune system also develop as the organism evolves, however they also have the ability to change *somatically* (i.e. during the lifetime of an individual). This results in the ability of the adaptive immune system to recognise previously unseen pathogens (learning) and to remember them for future encounters (memory). The innate and adaptive immune systems typically operate over different timescales. The innate operates on a small time scale often initiating a reaction either instantly or within a matter of minutes, whilst the adaptive immune system operates over a longer time period, taking of the order of days to initiate a reaction. It is the combination and interaction of both the innate and adaptive immune mechanisms that provides us with an effective immune system.

3.2 Immune Cells, Molecules and Organs

Each of the innate and adaptive immune systems have specialised cells which communicate through direct cell-to-cell interactions and through immune messenger proteins called *cytokines*. This communication is conveyed via special protein molecules called receptors that are expressed on the surface of cells in order to bind extra-cellular molecules. When a sufficiently strong chemical bond occurs between a receptor and another receptor or molecule, a signal is passed into the cell, providing a mechanism for recognition at the molecular level. Receptors of the innate immune system recognise very specific molecules, whilst receptors exist on adaptive immune cells that differ in what molecules can bind in order to recognise previously unseen pathogen related molecules. These receptors are called antigen receptors and any molecule that binds to them is called an *antigen*.

The innate immune systems contains many specialised cells such as macrophages and dendritic cells. Many of these cells are phagocytic, being able to ingest and kill some pathogens such as bacteria. Additionally, many can act as antigen presenting cells (APCs), whereby portions of the ingested pathogen are displayed on their cell membrane as antigens to be recognised by T-cells [46]. T-cells, along with B-cells, are the main actors of the adaptive immune system. They are named after the thymus and bone marrow, which are their respective places of maturation. Collectively B and T-cells are called lymphocytes. Lymphocytes are the only cells that produce antigen receptors, with each lymphocyte producing an antigen receptor that differs from all others. The antigen receptors of B-cells are called antibodies, whereas the antigen receptors of T-cells are called T-cell receptors (TCR). These receptors are

generated via a stochastic process utilising gene libraries to provide a massive potential repertoire of receptors. The specific portions of an antigen to which lymphocyte receptors bind are called epitopes. The ability of a receptor to bind with a single epitope is defined as the receptor's *specificity*. The strength of this binding is measured in terms of an affinity, whereby a high affinity between receptor and epitope results from a tight molecular binding [46].

The antibodies of B-cells can either be attached to the surface of the B-cell, or be secreted from the cell to bind to antigens free in solution, thus marking them for deletion by other immune cells. A B-cell that has not yet encountered any antigen that binds its antibody is called a naive B-cell. Once antigen binding occurs, the naive B-cell is activated, initiating a process of cellular proliferation and differentiation into effector B-cells. This results in the production of a population of clones of the original B-cell, which can secrete large numbers of antibodies. Some B-cells also differentiate into long lived memory cells, providing protection from future infection by the same pathogen whose antigen activated it [46]. Unlike antibodies, TCRs only recognise antigens that are bound to a protein receptor called major histocompatibility complex (MHC). MHC exists in two different forms, the first is expressed on the surface of most cells in the body, whereas the second is only expressed by APCs. T-cells that bind to the MHC of APCs are called T helper (T_H) cells, whereas T cytotoxic (T_C) cells bind to the other MHC type. The job of T_H cells is to activate other immune cells such as B and T_C cells. The role of T_C cells is to kill other cells to which their TCRs bind, these are typically virus infected cells and tumour cells [46]. T-cells follow the same cellular activation, proliferation and differentiation mechanisms described above for B-cells.

As previously mentioned, cytokines are the messenger molecules of the immune system. Different cytokines are produced and received by immune cells at different rates dependent on the cell types and cell states involved. Over 50 different cytokine types are mentioned by Janeway *et al.* [64], however, there is no one-to-one relationship between cytokine and effect with many different cytokines performing similar, or different, functions depending on cellular states. One role of cytokines is to help control an immune reaction by providing amplificatory or suppressive effects.

The immune system incorporates a number of organs distributed around the body. These are classified into two groups, the central (or primary) immune organs, and the peripheral (or secondary) immune organs. The central immune organs include the bone marrow, which is the site of immune cell production, and the thymus whose main purpose is to provide an environment where T-cells can mature and be selected to provide appropriate immune reactivity. The secondary immune organs include the lymph vessels, which provide a transport mechanism for immune cells, and lymph nodes, which provide the immune cells with a place to interact. Many of the immune system's functions such as classification and adaption are carried out by immune cells in the lymph nodes.

3.3 Clonal Selection Theory

It was highlighted above that the antigen receptors of lymphocytes differ. The body contains millions of different lymphocytes, thus millions of different receptors with different specificities are expressed, the combination of which is known as the receptor repertoire of the individual. According to Burnet's 1959 clonal selection theory [12], this repertoire undergoes a selection mechanism during the lifetime of the individual. The theory states that on binding with a suitable antigen, activation of lymphocytes occurs. Once activated, clones of the lymphocyte are produced expressing identical receptors to the original lymphocyte that encountered the antigen. Thus a clonal expansion of the original lymphocyte occurs. This ensures that only lymphocytes specific to an activating antigen are produced in large numbers. The clonal selection theory also stated that any lymphocyte that have antigen receptors specific to molecules of the organism's own body must be deleted during the development of the lymphocyte [66]. This ensures that only antigen from a pathogen might cause a lymphocyte to clonally expand and thus elicit a destructive adaptive immune response. In this sense, the immune system can be viewed as a classifier of antigens into either *self* antigen or *non-self* antigen, with non-self antigen assumed to be from a pathogen and thus needs to be removed from the body.

The process of deleting self-reactive lymphocytes is termed clonal deletion and is carried out via a mechanism called negative selection that operates on lymphocytes during their maturation. For T-cells this mainly occurs in the thymus, which provides an environment rich in APCs presenting self-antigens. Immature T-cells that strongly bind these self-antigens undergo a controlled death (apoptosis). Thus, the T-cells that survive this process should be unreactive to self-antigens. The property of lymphocytes not to react to the self is called immunological tolerance [28].

During the clonal expansion of B-cells (but not T-cells), the average antibody affinity increases for the antigen that triggered the clonal expansion. This phenomenon is called affinity maturation, and is responsible for the fact that upon a subsequent exposure to the antigen, the immune response is more effective due to the antibodies having a higher affinity for the antigen. Affinity maturation is caused by a somatic hypermutation and selection mechanism that occurs during the clonal expansion of B-cells. Somatic hypermutation alters the specificity of antibodies by introducing random changes to the genes that encode for them. This hypermutation mechanism is proportional to the affinity of the antigen-antibody binding, so that the higher the antibody affinity the less mutations it suffers. After the mutations have occurred, the B-cells that produce higher affinity antibodies are preferentially selected to differentiate into effector and memory cells, thus over the course of an immune response, the average population affinity of antibodies increases [46,28].

3.4 Immune Network Theory

In 1974, Jerne [67] proposed an immune network theory to help explain some of the observed emergent properties of the immune system, such as learning and memory. The premise of immune network theory is that any lymphocyte receptor within an organism can be recognised by a subset of the total receptor repertoire. The receptors of this recognising set have their own recognising set and so on, thus an immune network of interactions is formed. To formalise this [67] made use of the following terminology:

- *Epitope*: also known as the antigenic-determinant is the particular portion of an antigen to which the antibody binds.
- *Paratope*: the portion of the immune receptor which binds to a particular epitope complementarily.
- *Idiotypic*: the set of epitopes displayed by a set of antibody molecules.
- *Idiotope*: a single idiotypic epitope.

Following this terminology, immune networks are often referred to as idiotypic networks (figure 2). In the absence of foreign antigen, Jerne concludes that the immune system must display a behaviour or activity resulting from interactions with itself, and from these interactions immunological behaviour such as tolerance and memory emerge. A history of the first twenty-one years of network theory is provided by [21], it documents that the development of network theory has been hampered from a lack of biological evidence and an immunological field sold on clonal selection theory. However [21] states that network theory has undergone a revival and provides perhaps the best explanations for immune self-tolerance, with idiotypic networks providing an internal image of self [7].

3.5 Danger Theory

In [76], Matzinger explains how the clonal selection theory placed the antigen-specific cells of adaptive immunity (most notably the T_H cell) at the centre of the decision of whether or not to initiate an immune response. This decision was achieved through the deletion of the self-reacting lymphocytes, so that responses will only be initiated against non-self. It was discovered, however, that T_H cells themselves require a co-stimulatory signal from non-antigen-specific APCs in order to initiate an effective adaptive immune response. As a consequence, it could not be assured that immunity only be directed against non-self, as APCs express on their surfaces both self and non-self antigens. To address this, Janeway [65] proposed the infectious–non-self model that suggested APCs could discriminate between self and non-self by detecting, via the use of germline encoded receptors, evolutionarily conserved pathogen-associated molecular patterns unique to bacteria.

As an alternative explanation, Matzinger [75] proposed the danger theory in 1994, which has gained much popularity amongst immunologists in recent years as an explanation for the development of peripheral tolerance (tolerance to agents outside of the host). The danger theory states that APCs are themselves activated via an alarm: danger signals. These activated APCs will

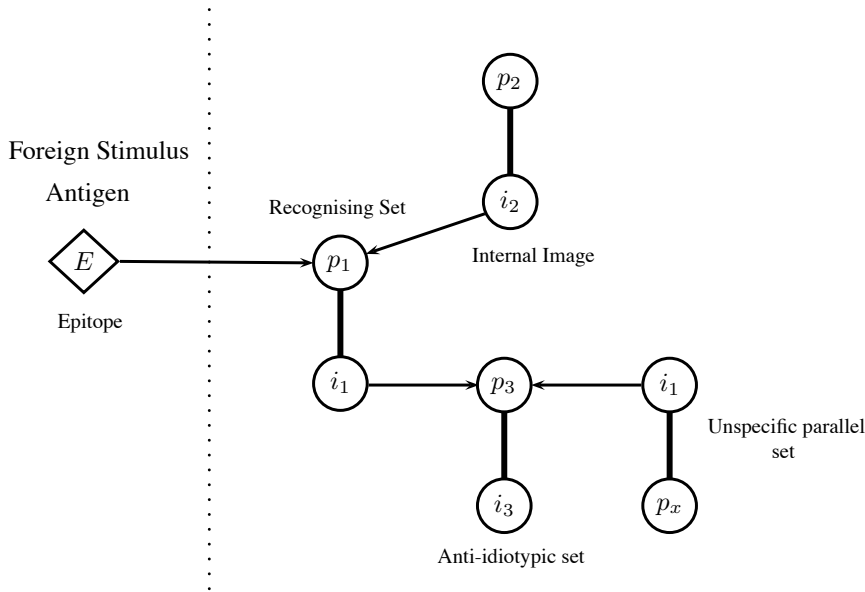


Fig. 2 Eigen-behaviour of the immune system [67]. Epitope E of an antigen is presented to the immune system, there is a set p_1 of paratopes recognising E which has an associated set i_1 of idiotopes. Set p_1 also recognises a set i_2 of idiotopes, this can be thought of as an internal image from epitope E . Set i_1 is recognised by internal paratope set p_3 , which represent an anti-idiotypic set. There is also a set of paratopes p_x which display idiotopes i_1 but do not fit the foreign epitope. The internal image has a stimulatory effect on the recognising set, whereas the anti-idiotypic has an inhibitory effect.

then be able to provide the necessary co-stimulatory signal to the T_H cells that subsequently control the adaptive immune response. The danger signals are emitted by ordinary cells of the body that have been injured due to attack by pathogen. For example, the intra-cellular contents released due to uncontrolled (necrotic) cell death could provide such signals. These signals are detected by specialised innate immune cells called dendritic cells that seem to have three modes of operation: immature, semi-mature and mature [74]. In the dendritic cell's immature state it collects antigen along with safe and danger signals from its local environment such as: *pathogen-associated molecular patterns (PAMPs)* and inflammatory cytokines. The dendritic cell is able to integrate these signals [82] to decide whether the environment is safe or dangerous. If safe the dendritic cell becomes semi-mature and upon presenting antigen to T-cells the dendritic cell will cause T-cell tolerance. If dangerous the dendritic cell becomes mature and causes the T-cell to become reactive on antigen-presentation. Through this process the dendritic cell contributes to immune homeostasis [74]. The transition from immature to semi-mature is illustrated in figure 3.

The danger theory presents a number of important consequences for the field of immunology. Firstly, with the danger signals arising from normal cells of the body, an immune response is no-longer initiated by the specialised

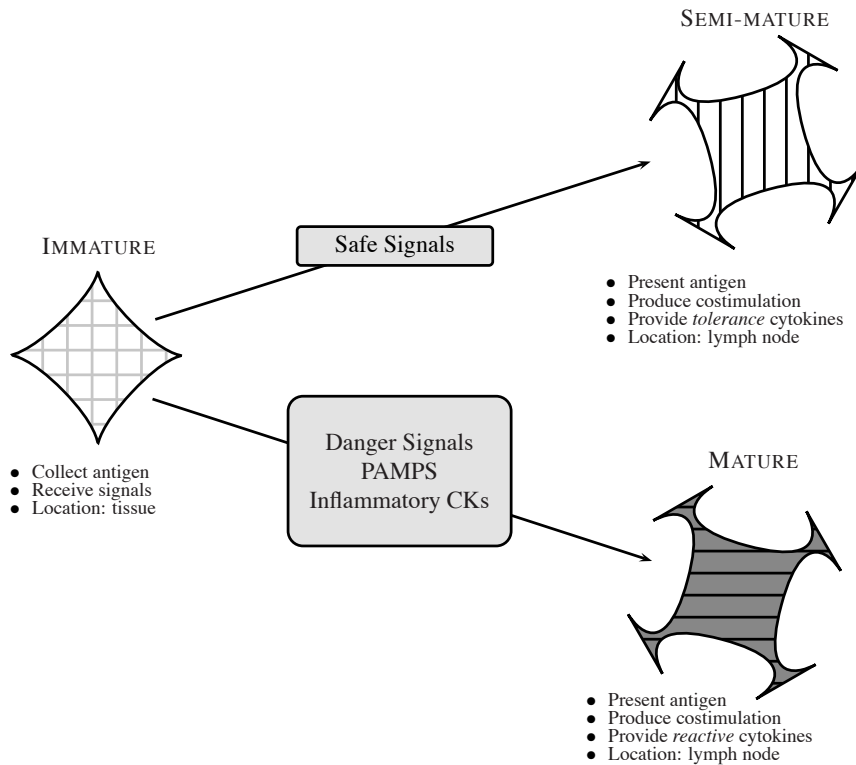


Fig. 3 Dendritic cell differentiation from immature to semi-mature, mature. Diagram reproduced from [49]

cells of the immune system. Secondly, the adaptive immune response is itself controlled by the action of innate immune cells, thus blurring the distinction between the adaptive and innate arms of the immune system. Lastly, the notion of self–non-self discrimination is replaced with a danger–non-danger metaphor, whereby foreign non-self no longer necessarily initiates an immune response.

3.6 Competing Immune Theories

By the 1970s the immune self had become the defining idea in immunology with the field itself being referred to as the science of self–non-self discrimination [102]. It is evident, however, that in recent years this view has changed. In the editorial summary to an issue of the journal *Seminars in Immunology* that examined competing theories of the immune system, Langman and Cohn [72] state:

“There is an obvious and dangerous potential for the immune system to kill its host; but it is equally obvious that the best minds in im-

munology are far from agreement on how the immune system manages to avoid this problem.”

On examination of the articles presented in the journal issue [2,11,17,51,73,79,102,95], the extent to which many immunologists differ in their views becomes clear. In his commentary on the models proposed by the other immunologists in the journal issue, Tauber [102] believes that they all fall to various degrees between the ideas of Burnet (clonal selection theory) and Jerne (immune network theory), and are thus a continuation of the arguments between these two points of view.

As an example of a competing immune theory, we highlight details of the ideas presented by Cohen [17]. In a major departure from the classical immune system view, Cohen’s model removes the requirement of self–non-self discrimination entirely. Instead, all immune cells recognise both self and non-self antigens and form an immune dialogue with the body’s tissues in order to fulfil the role of body maintenance. In order to achieve body maintenance, the immune system must select and regulate the inflammatory response according to the current condition of the body. This condition is assessed by both the adaptive and innate immune agents, which are required to recognise both the presence of pathogens (non-self antigens) and the state of the body’s own tissues (self antigens). The specificity of the immune response, therefore, is not just the discrimination of danger, or the distinction of self–non-self, but the diagnosis of varied situations, and the evocation of a suitable response. In summary, Cohen’s maintenance role of the immune system requires it to provide three properties [18]:

- *Recognition*: to determine what is right and wrong
- *Cognition*: to interpret the input signals, evaluate them, and make decisions
- *Action*: to carry out the decisions

These properties are provided via a cognitive strategy in which self-organisation of the immune system is used make deterministic decisions. It is the outcome of these decisions that are proposed to perform body maintenance. Self-organisation of the adaptive immune system occurs in the construction of the T-cell and B-cell antigen receptor repertoires. In the innate immune system self-organisation takes place in the generation of the actual response repertoire of the immune system via the fine tuning of the set of innate immune responses. This determines the types of response that will be connected to the signals perceived by the lymphocyte receptor repertoire. As we discussed in section 1, Cohen’s new ideas extend this notion of cognition into a computational paradigm where the immune system computes the state of the body in order to achieve maintenance [19].

4 Modelling the Immune System

Within the context of the conceptual framework (section 2) modelling plays an important role in the understanding of the computational aspects of the immune system. There is a vast range of modelling approaches available, each

with their own advantages and disadvantages operating at different levels of abstraction [39]. What we present in this section is an overview of some of the techniques that are common place in the immunological world and help us, from a computational and engineering background, understand how the immune system computes.

A recent paper by Forrest and Beauchemin [39] provides an excellent review of modelling approaches in immunology (and further discussions on engineering immune systems for computer security). The authors highlight that there are a number of ways in which one can model the immune system, with each approach offering different perspectives to the modeller. Within the paper, the authors focus more on Agent Based Modelling (ABM) as a tool where cells might be represented as individual *agents*, rather (as in the more traditional differential equations) than as a population of cell types. An agent in the system may be a certain type of cell that is encoded with simple rules that govern its behaviours and interactions. Within ABM it is possible to observe quite easily the dynamics of the agent population that arise as a result of the interactions between the agents. One difficult aspect of ABM is defining the right level of abstraction for each agent in the model, as this will clearly affect how the simulation operates. Forrest and Beauchemin [39] argue that ABM might be a more appropriate tool for modelling immunology due to the ease of which one can incorporate knowledge into the model that might not be able to be expressed mathematically and that multiple tests (or experiments) can be run with great ease, thus allowing the experimental immunologist a chance to perform experiments (albeit ones at a certain level of abstraction) *in silico*. This concurs with the view of Bersini [8] who advocates the use of object oriented (OO) technologies, and indeed ABM is a natural implementation of the OO paradigm. Another modelling approach is one of state charts first proposed by Harel [53] as a mechanism for representing computational processes by means of states and events that cause a transition between states. Such state charts can be developed to model complex interactions between elements and have demonstrated themselves to be useful in the context of immunological modelling [8,39].

It seems clear that there is a great deal to be learnt from the examination of the immune system in more computational terms. Indeed, our position is to concur with Forrest and Beauchemin [39], Andrews and Timmis [3], Stepney *et al.* [96], Bersini [8], Timmis [103] and Cohen [19] that there is a great benefit from the AIS practitioner engaging with the immunological modelling community to help not only the engineers but also the immunologists. Having now motivated the study of immunological modelling, and the role it can play in not only understanding the immune system, but also its potential role in the development of AIS, we briefly review immunological modeling in terms of mathematical models (the more traditional differential equation type models), and computational models, with a focus on π -calculus and state chart modelling. We have chosen these topics because from the mathematical side, the differential models are still very prevalent in the literature and have held to shape, to some degree, the nature of AIS to date. From a computational modelling perspective, we focus on π -calculus as this is potentially a very powerful tool from computer science that can be used to

model highly complex and parallel systems with large scale (something that simple OO methods would struggle at) and state charts, as these are simple tools that can be used to capture quite complex interactions between components in a visual manner (thus making them more understandable to the non-mathematician). However, for a full review of this topic, the reader is directed to Forrest and Beauchemin [39].

4.1 Mathematical Modelling of Different Immune Networks

In this section, we have decided to focus on two different aspects of immunology that have been modelled using differential equations. We begin our discussions by examining the seminal work of Farmer *et al.* [37] on idiotypic immune networks (section 3.4). We have selected this area to focus on due to the impact that it had on the early development of AIS and it still continues to influence the thinking of AIS practitioners. We then move to modelling a different type of immune network: cytokine networks. This work by Hone and van den Berg [61] is an interesting recent addition to the modelling literature and presents a network that exhibits decision making properties, that could be useful in an engineering context.

4.1.1 Idiotypic Networks

The seminal work of Farmer *et al.* [37] investigated Jernes' work [67] in great depth and provided insights into some of the mechanisms involved in the production and dynamics of the immune network, that up to that point had been only discussed. The authors created a simplistic model to simulate the immune system which ignored the effect of T-cells and of macrophages in an attempt to capture the essential characteristics of the immune network. Central to their work was the calculation of the dynamics of B-cell population related to a B-cell's stimulation level. The authors proposed a simple equation that they consider takes into account the three main contributing factors to B-cell stimulation level, these are: (i) the contribution of the antigen binding (ii) the contribution of neighbouring B-cells and (iii) the suppression of neighbouring B-cells. The rate of change of antibody concentration is given by :

$$\dot{x}_i = c \left[\sum_{j=1}^N m_{j,i} x_i x_j - k_1 \sum_{j=1}^N m_{i,j} x_i x_j + \sum_{j=1}^M m_{i,j} x_i y_j \right] - k_2 x_i \quad (1)$$

where the first term represents the stimulation of the paratope of an antibody type i by the epitope of an antibody j . The second term represents the suppression of antibody of type i when its epitope is recognized by the paratope of type j . The parameter c is a rate constant that depends on the number of collisions per unit time and the rate of antibody production stimulated by a collision. Constant k_1 represents a possible inequality between stimulation and suppression. The stimulation of a B-cell cloning and mutation were included in the model to create a diverse set of B-cells. The amount

by which any one B-cell cloned was in relation to how stimulated the B-cell was. The more stimulated a B-cell, the more clones it produces. Three mutation mechanisms were introduced on the strings: crossover, inversion and point mutation. Crossover is the interchanging of two points on two different strings, inversion is the simple inverting of the value of the bit in a string, a 0 to a 1 and vice versa and point mutation is the random changing of a bit in a given string.

4.1.2 Cytokine Networks

Moving away from idiotypic networks, Hone and van den Berg [61] define a generic dynamical model for the interaction of cytokines and immune cells in which they determine an Artificial Cytokine Network (ACN). The model describes a network of n cytokines with concentrations u_1, \dots, u_n which are produced by m cell types with densities v_1, \dots, v_m within a medium, all cells and cytokines are considered to be diffuse and well mixed (as is the case with all differential models which can be seen as a limitation of the approach as biological systems are not really well mixed at all). The cell types are acted upon by r external stimuli s_1, \dots, s_r . The amount of any cytokine produced by a given cell type is a function of the concentrations of cytokines, external stimuli and the density of that T-cell type. A change in the cell types density is a function of cytokine concentrations, external stimuli and the cell type's current density. It is described by the following:

$$\dot{u}_k = \sum_{\ell=1}^m \psi_{\ell k}(u_1, \dots, u_n, s_1, \dots, s_r) v_{\ell} - \nu_k \mu_k \quad (2)$$

$$\dot{v}_{\ell} = (\varphi_{\ell}(u_1, \dots, u_n, s_1, \dots, s_r) - \mu_{\ell}) v_{\ell} \quad (3)$$

where $k = 1, \dots, n$; $\ell = 1, \dots, m$; the function $\psi_{\ell k} > 0$ expresses the effect of the cytokines and external stimuli on the production of cytokine k by cell type ℓ ; ν_k is the rate of degradation of the k th cytokine; the function $\varphi_{\ell} > 0$ expresses the effect of the cytokines and external stimuli on the proliferation rate of a cell of type ℓ .

There is a difference in time scales between the degradation of cytokines and cells, the cytokines should degrade at a much higher rate than cells. As a consequence [61] are able to discuss the stability of potential specifications of $\psi_{\ell k}$ with the assumption that $v_{\ell}(t) \equiv \bar{v}$, that is the cell density is constant from the point of view of the cytokines.

Hone and van den Berg [61] present numerical results of an ACN, performed with two cytokines $n = 2$; one cell type $m = 1$; and a single external stimulus $s(t)$. It is assumed that $\psi_j, j = 1, 2$ are sigmoid functions $S(x) \in [0, 1]$ for all $x \in \mathbb{R}$. The model is constructed such that the proliferation of cells is encouraged by cytokine u_2 and the external stimulus s , whereas the cytokine measured by u_1 will tend to decrease cell proliferation, ψ_j and φ are defined as:

$$\psi_j(u_1, u_2, s) = \bar{\psi}_j S \left(\sum_{k=1,2} W_{jk} u_k - \tilde{\theta}_j \right), \quad \varphi(u_1, u_2, s) = s u_2 e^{-\gamma u_1} \quad (4)$$




Fig. 4 Behaviour of the ACN defined by (2), (3) and (4) with parameters : secretion rates $\overline{\psi}_1, \overline{\psi}_2 = 1$; decay rates $\nu_1 = 6, \nu_2 = 5, \mu = 2$; threshold values $\tilde{\theta}_1 = 3, \tilde{\theta}_2 = 6$; interactions $W_{11} = -1, W_{12} = 1, W_{21} = 1, W_{22} = 0$. The stimulus is set such that there is a weak stimulus $s(t) = 3$ for $0 \leq t < 0.2$ and a short strong stimulus $s(t) = 70$ for $1.0 \leq t < 1.01$. Starting from initial values $u_1 = 6.5, u_2 = 12.5$ and $v = 12.5$. Two stimuli are applied, a long weak stimulus $s(t) = 3$ for $0 \leq t < 0.2$ and a short strong stimulus $s(t) = 70$ for $1.0 \leq t < 1.01$. The initial long weak stimulus primes the system, the second short stimulus induces a large proliferation of cells. It is important to note that it is not just the larger size of second stimulus which produces the result, the size of the second peak of v is reduced if the time interval between the stimuli is reduced.

for $j = 1, 2$; $\gamma > 0$; $S(x) = 1/(1 + e^{-x})$; $\overline{\psi}_j$ is the maximum secretion rate of cytokine j ; W is an interaction matrix that determines how cytokine j affects the production of cytokine k ; $\tilde{\theta}_j$ represents a stimulation threshold of the cell by the cytokine j . In order to better understand the model, we have plotted a numerical run of the equations (Figure 4). With just 2 cytokines and one cell type the ACN demonstrates a complex variety of behaviours, as demonstrated by other parameter choices given in [61].

4.2 Computational Modelling

In the introduction to section 4 we discussed that there is a large literature on the computational modelling of the immune system. Here we focus mainly on one technique the π -calculus [81]. We focus on this technique primarily as this one has not been covered in previous review papers on immunological modelling, specifically it is not covered by the review by Forrest and Beauchemin [39] and the technique has a great deal to offer both the immunologist and the engineer and can be seen as an excellent bridge between the two areas.

We also briefly discuss the role of UML in modelling immune systems, the use of design patterns and alternative ways of modelling immune networks as these are now beginning to find favour in the AIS community once again.

4.2.1 π -calculus

π -calculus 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 [88] computational analysis tools such as the π -calculus are just as applicable to biology as they are to computing.

Regev et al. [91] apply the π -calculus to model a signal transduction pathway, the authors note that the π -calculus allows the model to be mathematically well-defined, and remain biologically faithful and transparent. The authors also note that the π -calculus only allows qualitative analysis of a biological system. For quantitative analysis Stochastic π -calculus ($S\pi$) [89] is needed. $S\pi$ extends the π -calculus by adding a rate parameter r to interactions, this defines an exponential distribution, such that the probability of an interaction occurring within time t is $F(r, t) = 1 - e^{-rt}$. Thus the average duration of an interaction is the mean $1/r$. The purpose of the conception of $S\pi$ is to allow performance analysis of concurrent computational systems, as a consequence [89] demonstrates how it is possible to turn a system described in $S\pi$ to a continuous time Markov chain. Priami et al. [90] follows the work [91] and applies $S\pi$ to quantitatively examine biological pathways.

The mathematical nature of π -calculus, stochastic or otherwise, can render it inaccessible to non-computer scientists and potentially unwieldy or at least non-intuitive when modelling biological systems. To address this issue Philips et al. [88] define a Graphical Stochastic π -calculus ($GS\pi$) which represents a $S\pi$ specification as a graph of typed nodes with labelled edges. The authors prove this to be reduction equivalent to $S\pi$ to ensure that they both have the same expressive power. A number of advantages to $GS\pi$ are discussed in [88], including the ease with which one can discover cycles in a system (cycles are an important facet found at all levels of biological systems) and the ability to provide a frontend to an $S\pi$ simulator. Clearly the ability to simulate systems specified in $S\pi$ is essential to understanding its behaviour, to this end there are number of stochastic simulators, for example, BioSpi [90] and SPiM [88]. Both make use of the Gillespie Algorithm [45] to simulate bio-molecular interactions. However SPiM would seem to have some advantages over BioSpi, first, it is proved to be a correct simulation of $S\pi$. Second, it is optimised for the simulation of biology, it does this by noting that most biological simulations contain many thousands of the identical processes, i.e. many thousands of the same protein. Third, it provides visualisation through $GS\pi$ and through an animated 3D visualisation of $S\pi$.

A final point raised in [91] highlights that the tools of π -calculus can aid the understanding of biology. For example bisimulation allows formal comparison between two π -calculus programs by an equivalence relation in terms of their behaviour. This may allow abstraction of concepts common to many biological systems. Such ideas have an interesting instantiation here, it may be possible to use them to pin down what in a biological system is necessary for its behaviour and what is superfluous, and thus be a great benefit to the AIS practitioner as they will better understand why the biology behaves as it does. This will give more considered steps on the route through the conceptual framework towards bio-inspired algorithms.

4.2.2 Higher Levels of Abstraction

Work already briefly discussed by Bersini [8] advocates an object oriented approach incorporating tools such as UML [38], design patterns [42] and state charts [53]. Indeed, with regards to the use of state charts a small amount of work had been already undertaken by Cohen and Harel, see [20] for a good review of their work. Let us return to the work of Bersini [8] in the use of UML and design patterns. UML is a standardised specification language that can be used for general purpose modelling allowing for the creation of an *abstract* model of the system under study. Contained within UML is a wide variety of notions and figures that allow for the construction of the model. Design patterns are designs that are reused (or can easily be repeated) to a certain problem, and are widely used in the software engineering domain. These design patterns can be thought of as a template of a solution that are reused over and over again in different circumstances.

How do these two tools relate to modelling immunology? By adopting such approaches it makes the modelling of a complex system that much more comprehensible, adaptable and effective [8] and the fact that the modelling language is visual can act as an aid to interaction between the biologist and computational scientist (as advocated by Cohen [19]). Bersini [8] develops a simple UML class and state diagrams of a clonal selection model. Here he illustrates the use of the “state” design pattern and the “template” design pattern which illustrate how it is possible to encapsulate certain design information easily, but also how it is possible to isolate functionality within the model. Such techniques allow us to understand the operation of such complex systems and the very act of designing such a model prompts us to ask questions of the system under study that we might have never asked. Indeed, other work such as Mendao *et al.* [80] also advocate a UML type approach and in their paper take ideas from Cohen [18], with a focus on degeneracy, to create a simple clonal selection model based on degenerate receptors (receptors that are structurally similar, but can be functionally different). This was based on earlier work by Andrews and Timmis [5] where a simple model of a lymph node was developed incorporating the notion of degenerate receptors, and recognition of antigenic patterns was demonstrated to occur. Both these works can be seen as initial stages in the development of alternative immune inspired algorithms with UML acting as a natural bridge between the im-

munology (as diagrams are an easy way to communicate complex ideas) and the engineering (as UML naturally lends itself to computational settings).

Following on from our discussions on modelling immune networks (section 4.1.1), recent work by Hart and collaborators [54,55] has investigated, in depth, immune networks from a more *computational* perspective. The work of Hart *et al.* attempts to understand the evolution of immune networks and how the topology emerges from the interactions within the network. Their initial studies would indicate the main effect on immune network algorithms may well be the way in which interaction is defined, i.e through the definition of affinity. Through the development of a simple model Hart and Ross demonstrate the evolution of various immune network structures which are considerably affected by the choice of affinity measures between two B-cells, which in turn affects how B-cells interact with each other and how the boundary of tolerance of self/non-self emerges through the interactions of the network, thus producing an internal image of the environment. Such measures considered were single and multiple-site binding, complementarily binding and non-symmetrical binding and they were found to affect such things as network stability and tolerance of non-self agents [56]. The role of such modelling is now starting to be useful in the understanding of the dynamics of immune networks thus enabling the AIS practitioner to more readily make use of the metaphor.

4.3 On the use of Metaphors

Now we have reviewed the immune system and techniques that can be used to aid our understanding of that system, we feel it necessary to address a fundamental issue within AIS. In a recent paper, Neal and Trapnel [83] discussed the role of drawing separation between subsystems within biological systems. Of most interest to us here is their arguments concerning the separation of innate and adaptive immunity. They argue that drawing separations between these two systems (as is prevalent in the AIS literature, and indeed to some degree even this review paper) might not be that helpful, as the two systems are so intertwined, separation is not really possible (or sensible). Neal and Trapnel go onto argue that such a separation, and indeed the separation of those systems into further systems could be problematic. They argue that types of systems identified by researchers can often be biased by the actual research interests of the people undertaking the decomposition e.g. the bias may often be medical and biased towards pathological cases [83] and the break up of systems that arises from such a bias may not necessarily be the best way of splitting up the system. The authors state a set of assumptions that they argue, “afflicts most computer scientists”:

1. The sub-systems and components as described by biologists are circumscribed by “natural” boundaries and by implication have some degree of functional separability over and above that likely to be found for similar sized arbitrary sets of linked components selected at random.
2. That sub-systems and components *as identified by the biologists* will be of interest and utility to computer scientists seeking inspiration.

3. That biologists really believe that they have identified the important features of the immune system when describing these sub components.
4. That it is valid to make computational simplifications and caricatures by using “off the shelf” components when building computational analogues of the sub-systems and their components.

Neal and Trapnel [83] argue that if none of these assumptions were ever valid undertaking the development of bio-inspired systems would be “essentially meaningless”. However, they counter that if one is prepared to accept (1) and (2) most of the time, to talk to biologists in order to ascertain when (3) holds and to guard against (4) by maintaining sufficient complexity and constantly re-examine the biology to avoid over-simplification then both activities (the biology and the engineering) maintain value. They go onto argue that (1) current technology does not afford us the chance to prove beyond doubt for many systems, so acting in good faith seems a reasonable thing to do, however the “blind acceptance” of (2) is problematic. This, as argued by Neal and Trapnel, is because that what appear to be “clever” techniques in the biology might not necessarily translate to a “clever” computer solution: great care is needed to ensure that the analogy is suitable and not merely manufactured. Finally, they argue that coping with (3) and (4) is more manageable and to some degree, work by Stepney *et al.* [96] guards against that.

5 Artificial Immune Systems

In this section AIS are reviewed in detail, providing a description of the main types of AIS and the applications to which they have been applied. It does not provide a comprehensive summary of all AIS, but provides a simple overview of the main aspects of AIS.

5.1 The Structure of AIS

In [28], de Castro and Timmis propose a layered framework for engineering AIS. This framework succinctly demonstrates the general structure of most AIS, and so is used here as a template for this description of the main AIS types. The layered framework takes the application domain of the AIS as its starting point, followed by three layers to be considered before the required AIS is engineered. These layers are:

- *Component Representations*: how the components of the system are to be represented
- *Affinity Measures*: how the interactions between the components of the system are to be quantified
- *Immune Algorithms*: how the components of the systems are going to interact to determine the system dynamics

The immune algorithms are typically inspired by the immune processes covered in section 3.1, and fall into one of four groups: negative selection, clonal selection, immune networks and danger theory. Each of these will be discussed in detail below, along with examples of specific AIS implementations.

The representations and affinity measures used in AIS tend to be applicable to any type of immune algorithm, and so will be discussed in general terms here.

The most influential concept to affect the representation of components in AIS was introduced by Perelson and Oster [87], who defined the notion of shape space. Their theoretical investigation into how large an antibody repertoire must be in order to provide effective defence, viewed the immune system as a molecular recognition device designed to identify foreign shapes. For an antibody, they define the relevant antigen combining region in terms of N shape parameters (the authors do not list these parameters, but suggest that they may include geometric quantities specifying molecular sizes and shapes, or molecular charges). Likewise, an antigen that binds this antibody combining region can be described by the same N parameters. By combining these N parameters into a vector, both the antibody and antigen can be represented as points \mathbf{Ab} and \mathbf{Ag} respectively in an N -dimensional vector space called shape space, S . Following on from this, Perelson and Oster [87] state that the antibody and antigen bind perfectly when $\mathbf{Ab} = \mathbf{Ag}$ if the fact that the shapes of the antibody and antigen are complementary is ignored. The antibody-antigen complementarity can, therefore, be measured as the distance between \mathbf{Ab} and \mathbf{Ag} using a suitable metric in S . If all of the N parameters contribute equally to the antibody specificity, then a metric such as Euclidean distance can be used. A volume, V , can then be defined in S as the volume in which the possible repertoire of antibodies and antigens can fall. For each individual antibody, a small region V_ε can also be defined, where ε is a distance threshold in S . Any antigen that falls within the V_ε of an antibody is said to be bound by that antibody. Perelson and Oster [87] note that this strict threshold could be replaced with a probability of binding that decreases with distance from the antibody, however, this was not investigated. They also note that previous work by Edelstein and Rosen [33] showed that the shape of an antibody could be defined by a continuous function instead of a point vector. Even though the work of Perelson and Oster [87] used antibodies and antigens, de Castro and Timmis [28] point out that this shape space representation can be applied to any type of receptor and molecule that binds it.

Within the majority of AIS that utilise immune receptor and antigen components, the notion of shape space is employed. Thus, the AIS components will typically be represented as attribute strings of N parameters, for example an antibody component would be represented as the vector $\mathbf{Ab} = \langle Ab_1, Ab_2, \dots, Ab_N \rangle$. The choice of these vector attributes will determine the shape space type. de Castro and Timmis [28] highlight three main types that can be used in AIS (but this list is not exhaustive):

- *Real-Valued Shape Space*: attributes of all the components are real numbers
- *Hamming Shape Space*: attributes of all the components are from a finite alphabet
- *Symbolic Shape Space*: attributes can be of any type, including symbols such as age and name.

The choice of attribute types used in the AIS components will naturally determine the metric in S that can be used to measure the similarity (affinity) between these components. It is normal in AIS that each of the N attributes in the component contribute equally to the affinity measure. For real valued shape spaces, this affinity measure is typically Euclidean distance, although similar measures such as Manhattan distance are equally applicable. The affinity measures used for the Hamming shape spaces are determined by the alphabet used. Two of the most popular affinity measures for the binary alphabet are the Hamming distance and the r -contiguous bit rule. The Hamming distance is simply calculated by applying the XOR operator to the two components that are being measured. The r -contiguous bit rule calculates the affinity as the length of the largest contiguous region between the two components. The affinity measures used for symbolic shape spaces are very much dependent on the nature of the symbols that make up the attribute strings, and thus defining this measure can be a non-trivial task. For all affinity measures, a recognition threshold akin to the shape space parameter ε described above, can be set to determine whether recognition between components has occurred. As a small aside, it should be noted that affinity measures such as r -contiguous do not satisfy the triangle inequality and therefore can not be strictly considered to be a distance metric in the mathematical sense.

5.2 Negative Selection Algorithms

Negative selection algorithms are inspired by the main mechanism in the thymus that produces a set of mature T-cells capable of binding only non-self antigens. The first negative selection algorithm was proposed by Forrest *et al.* [40] to detect data manipulation caused by a virus in a computer system. The starting point of this algorithm is to produce a set of self strings, S , that define the normal state of the system. The task then is to generate a set of detectors, D , that only bind/recognise the complement of S . These detectors can then be applied to new data in order to classify them as being self or non-self, thus in the case of the original work by Forrest *et al.* [40], highlighting the fact that data has been manipulated. The algorithm of Forrest *et. al.* [40] produces the set of detectors via the process outlined in algorithm 1.

```

input :  $S_{seen}$  = set of seen known self elements
output:  $D$  = set of generated detectors
begin
  repeat
    Randomly generate potential detectors and place them in a set  $P$ 
    Determine the affinity of each member of  $P$  with each member of
    the self set  $S_{seen}$ 
    If at least one element in  $S$  recognises a detector in  $P$  according to a
    recognition threshold, then the detector is rejected, otherwise it is
    added to the set of available detectors  $D$ 
  until Until stopping criteria has been met
end

```

Algorithm 1: Generic Negative Selection Algorithm

The negative selection algorithm formed the foundation of a significant development within AIS, that being the work of Hofmeyer and Forrest [60] where they proposed an architecture for an AIS as applied to network security. This remains one of the most influential works in the area to date, spawning a whole avenue of AIS research with examples including [6, 48, 47, 26, 69]. On more theoretical aspects, work by Esponda *et al.* [35, 34] showed a connection between the boolean satisfiability problem (SAT) and a negative database. In work by Stibor *et al.* [101] and further in Stibor's thesis [98] the authors specialised the approach presented in by Esponda *et al.* in that they showed that the problem of generating r -contiguous detectors can be transformed in a k -CNF satisfiability problem. This is further extended by Stibor [99] where he showed that generating such detectors is hardest in the "phase transition" region depending on the parameter settings of size of self and of r . It was concluded that it is either very easy or very hard to generate detectors, with most interesting problems being in the very hard domain, thus reducing the applicability of such an approach. Through such theoretical approaches we can begin to understand in a deeper way, issues surrounding the generation of detectors using the r -contiguous matching rule and the problems inherent in that process.

5.3 Clonal Selection Algorithms

The clonal selection theory has been used as inspiration for the development of AIS that perform computational optimisation and pattern recognition tasks. In particular, inspiration has been taken from the antigen driven affinity maturation process of B-cells, with its associated hypermutation mechanism. These AIS also often utilise the idea of memory cells to retain good solutions to the problem being solved. In [28], de Castro and Timmis highlight two important features of affinity maturation in B-cells that can be exploited from the computational viewpoint. The first of these is that the proliferation of B-cells is proportional to the affinity of the antigen that binds it, thus the higher the affinity, the more clones are produced. Secondly, the mutations suffered by the antibody of a B-cell are inversely proportional to the affinity of the antigen it binds. Utilising these two features, de Castro and Von Zuben [14] developed one of the most popular and widely used clonal selection inspired AIS called CLONALG, which has been used to performed the tasks of pattern matching and multi-modal function optimisation.

When applied to pattern matching, a set of patterns, S , to be matched are considered to be antigens. The task of CLONALG is to then produce a set of memory antibodies, M , that match the members in S . This is achieved via the algorithm outlined in algorithm 2.

The algorithm shares many similarities with evolutionary algorithms [84], although importantly the selection and mutation mechanisms are influenced by the affinities of antibody-antigen matching [28]. Other AIS that have been inspired by the adaptive immune mechanisms of B-cells are AIRS [113], a supervised learning algorithm (or more recently described as an *instance creation algorithm* [93], and the B-cell algorithm [70], an optimisation algorithm with a unique contiguous hypermutation operator. Like negative selection,

```

input :  $S$  = set of patterns to be recognised,  $n$  the number of worst
         elements to select for removal
output:  $M$  = set of memory detectors capable of classifying unseen patterns
begin
  Create an initial random set of antibodies,  $A$ 
  forall patterns in  $S$  do
    Determine the affinity with each antibody in  $A$ 
    Generate clones of a subset of the antibodies in  $A$  with the highest
    affinity. The number of clones for an antibody is proportional to its
    affinity
    Mutate attributes of these clones inversely proportional to its
    affinity. Add these clones to the set  $A$ , and place a copy of the
    highest affinity antibodies in  $A$  into the memory set,  $M$ 
    Replace the  $n$  lowest affinity antibodies in  $A$  with new randomly
    generated antibodies
  end
end

```

Algorithm 2: Generic Clonal Selection Algorithm

clonal selection has proven to be very popular in the AIS community spawning a great deal of research with recent examples including [23, 77, 62].

From a theoretical perspective, given the stochastic nature of the algorithms, it should be possible to consider, in particular with clonal selection algorithms, the evolution of a population belonging to a discrete state space and changing according to probabilistic rules [106]. As long as the probabilities for transitions to a new state depend only on the current state of the system (and not on the previous history), all the properties of a Markov chain are satisfied, so it is natural to describe AIS clonal selection algorithms in these terms [106]. Indeed, there has been limited work by Villalobos-Arias *et al.* [111], which used Markov chain theory to prove convergence of MISA [22] (a clonal selection based algorithm), with the proviso that an elitist memory set must be maintained. Other work by Clark *et al.* [15], who after modelling the hypermutation operator associated with the B-cell algorithm [70], adopted the same method of proof and simplified it in order to prove convergence, in the sense that it finds (at least one) global optimum solution with probability one, in the limit as $t \rightarrow \infty$, for the case of the B-cell algorithm [70].

Work by Cutello *et al.* [24] adapts the criteria of Rudolph [92] to the setting of a generic clonal selection-type algorithm, called the Immune Algorithm (IA). This IA includes the possibility of a variety of different schemes for hypermutation, aging and so on, and two sufficient conditions for convergence are given in [24]. In particular, it appears that the convergence of the B-cell Algorithm can also be proved by minor modifications of the approach in that paper [106].

5.4 Immune Network Algorithms

Immune network models fall into two distinct categories, the continuous models and the discrete models. The continuous models are based on ordinary differential equations and are typically used by theoretical immunologists to

explore the perceived behaviour of real immune networks. Examples include the models by Farmer *et al.* [37] and Varela and Coutinho [110]. These continuous models have been subsequently used as inspiration for the discrete immune network models found in AIS. The behaviour of these discrete models are based on iterative procedures of adaptation and they are applied to problem solving rather than understanding immune concepts. This produces one of the major differences between the discrete and continuous models, as discrete models interact with their environment (i.e. antigens), whereas the continuous models typically do not [28]. Farmer [36] outlined the link between immune network models, neural networks, autocatalytic networks and classifier systems in terms of connectionism. This was the first attempt at positioning the idea of immune networks as an adaptive learning system, and these ideas have been pursued by AIS researchers.

The main difference between immune network algorithms and other immune algorithms is that the components of the system not only interact with antigenic components, but with the other components in the AIS. Two examples of immune network algorithms are AINE [108] and aiNet [13], which attempt to utilise the basic concepts of immune network theory to solve problems such as pattern recognition and data clustering. For example, aiNet was originally designed for the task of data clustering [13], and consists of a network of antibody components that adapt to match a population of input components (antigens) to be clustered. aiNet is a modified version of CLON-ALG (described above) with an added mechanism of suppressive interactions between the antibody components, and is outlined in algorithm 3.

The resulting set of network antibodies, N , represents an internal image of the antigens to which they have been exposed [28] as a reduced number of elements. There is less theoretical work on the analysis of immune network algorithms, but one recent paper by Stibor and Timmis [100] has shown that aiNet suffers problems when clustering non-uniformly distributed data. This is primarily because of the way the suppression mechanism operates in the algorithm (via a distance metric) and it leads to either insufficient retention of information in the clusters from when it started (so the data has been compressed, but significant information has been lost), or the algorithm actually produces more elements to represent the input space than it began with. However, aiNET has found itself widely used, particularly in the area of optimisation, where it performs very well with such examples include [10, 4, 16].

5.5 Danger Theory Algorithms

Danger theory is a relatively new addition to the field of immunology, and thus danger theory inspired algorithms are still in their infancy. In a summary of danger theory inspired approaches, Garrett [43] states that these algorithms should provide an alternative to the negative selection approach, as danger theory has a number of appealing properties from a computational perspective. For example, ideas from danger theory should focus events on what is harmful instead of just non-self, reporting only the detection of something dangerous. A number of proposed danger theory inspired AIS have been

```

input :  $S$  = set of patterns to be recognised,  $nt$  network affinity threshold,
          $ct$  clonal pool threshold,  $h$  number of highest affinity clones,  $a$ 
         number of new antibodies to introduce
output:  $N$  = set of memory detectors capable of classifying unseen patterns
begin
  Create an initial random set of network antibodies,  $N$ 
  repeat
    forall patterns in S do
      Determine the affinity with each antibody in  $N$ 
      Generate clones of a subset of the antibodies in  $N$  with the
      highest affinity. The number of clones for an antibody is
      proportional to its affinity
      Mutate attributes of these clones inversely proportional to its
      affinity, and place the  $h$  number of highest affinity clones into a
      clonal memory set,  $C$ 
      Eliminate all members of  $C$  whose affinity with the antigen is
      less than a pre-defined threshold ( $ct$ )
      Determine the affinity amongst all the antibodies in  $C$  and
      eliminate those antibodies whose affinity with each other is less
      than a pre-specified threshold ( $ct$ )
      Incorporate the remaining clones in  $C$  into  $N$ 
    end
    Determine the affinity between each pair of antibodies in  $N$  and
    eliminate all antibodies whose affinity is less than a pre-specified
    threshold  $nt$ 
    Introduce a number ( $a$ ) of new randomly generated antibodies into
     $N$ 
  until until a stopping condition has been met
end

```

Algorithm 3: Generic Immune Network Algorithm

presented by Secker *et al.* [94], Aickelin *et al.* [1] and Greensmith *et al.* [49, 50]. Secker *et al.* [94] explores the relevance of danger theory to web mining by investigating the use of danger signals to provide context for searches. The second looks at the possibility of an intrusion detection system based on danger theory by investigating how various intrusion scenarios can be detected via the use of different computer danger signals. This work was later extended by Greensmith *et al.* [49, 50] with the dendritic cell algorithm (DCA) (algorithm 4), which introduced the notion of danger signals, safe signals and PAMP signals which all contribute to the context of a data signal at any given time. This context is integrated via a process inspired by the role of dendritic cells (a specialised APC of the innate immune system). This removes the need to define what *self* is, but adds the necessity to define the danger, safe and PAMP signals, guidance on this is given in [50].

5.6 Applications of AIS

Although not as widely used as other bio-inspired paradigms, such as evolutionary algorithms and neural networks, the body of work describing the practical applications of AIS has become substantial over the last decade. Accordingly, Hart and Timmis [59] have investigated the application areas of AIS, and considered the contribution AIS have made to these areas. Their

```

input :  $S$  = set of data items to be labelled safe or dangerous
output:  $L$  = set of data items labelled safe or dangerous
begin
  Create an initial population of dendritic cells (DCs),  $D$ 
  Create a set to contain migrated DCs,  $M$ 
  forall data items in S do
    Create a set of DCs randomly sampled from  $D$ ,  $P$ 
    forall DCs in P do
      Add data item to DCs' collected list
      Update danger, PAMP and safe signal concentrations
      Update concentrations of output cytokines
      Migrate dendritic cell from  $D$  to  $M$  and create a new DC in  $D$  if
      concentration of costimulatory molecules is above a threshold
    end
  end
  forall DCs in M do
    Set DC to be semi-mature if output concentration of semi-mature
    cytokines is greater than mature cytokines otherwise set as mature
  end
  forall data items in S do
    Calculate number of times data item is presented by a mature DC
    and a semi-mature DC
    Label data item as safe if presented by more semi-mature DCs than
    mature DCs otherwise label as dangerous
    Add data item to labelled set  $M$ 
  end
end

```

Algorithm 4: Dendritic Cell Algorithm

survey of AIS is not exhaustive, but attempts to produce a picture of the general areas to which they have been applied. Over 100 papers were classified into 12 categories that were chosen to reflect the natural groupings of the papers. Some of these categories are broad, whereas some are narrow, as a new category was created when there was more than one paper reporting a particular application area. The 12 identified categories, in the order of most papers first, were: clustering/classification, anomaly detection (e.g. detecting faults in engineering systems), computer security, numerical function optimisation, combinatoric optimisation (e.g. scheduling), learning, bio-informatics, image processing, robotics (e.g. control and navigation), adaptive control systems, virus detection and web mining. The authors go on to note that these categories can be summarised into three general application areas of learning, anomaly detection and optimisation.

6 Reflections on Artificial Immune Systems

After reviewing the more technical side of AIS, we now try and bring together ideas from other review papers that have taken a more reflective standpoint. This will allow us to reflect on the progress made so far within AIS, identify common themes that are emerging within the literature and identify areas where things can be improved and maybe where future focus should be concentrated.

A review paper by Garrett [43] aimed to assess the usefulness of different types of AIS requires the definition of the term usefulness with respect to a computational method, followed by a set of criteria to enable its measurement. It is argued that a ‘useful’ algorithm is both ‘distinct’ from other algorithms, and ‘effective’ at performing its required function. Thus, ‘distinctiveness’ and ‘effectiveness’ are the metrics chosen to measure the ‘usefulness’ of different AIS types. In order to determine the distinctiveness and effectiveness of these AIS, two sets of questions are asked. The distinctiveness questions, quoted verbatim from [43], are:

1. Does the new method contain unique *symbols*, or can the features of this method be transformed into the features of another method, without affecting the dynamics of the new method?
2. Are the new method’s symbols organized in novel *expressions*, or can its expressions be transformed to become the same as some other method, without affecting its dynamics?
3. Does the new method contain unique *processes* that are applied to its expressions, or can its processes be transformed to become identical to some other method, without affecting its dynamics?

A method is deemed distinctive if the answer to at least one of these questions is ‘yes’. Likewise, a method is deemed effective if the answer to at least one of the following questions, quoted verbatim from [43], is ‘yes’:

1. Does the method provide a *unique* means of obtaining a set of results?
2. Does the method provide *better* results than existing methods, when applied to a shared benchmark task?
3. Does the method allow a set of results to be obtained *more quickly* than another method, on a benchmark test?

A truly useful AIS is, therefore, considered by Garrett [43] to be one which has been classified as being both distinctive and effective.

After outlining the above criteria, Garrett [43] proceeds to assess the usefulness of AIS via an empirical process based on observations in the body of AIS literature. This process is carried out separately for the three main types of AIS described in the last chapter: negative selection, clonal selection and immune network models. For each of these, a description is given of the major developments within the area along with an identification of the state of the art and some typical applications. Based on this analysis, the distinctiveness and effectiveness questions described above are applied. For the distinctiveness questions, all three AIS types Garret answers ‘no’ to the first two question and ‘yes’ to the last, thus all are identified as being distinctive as they contain unique algorithmic processes. For the effectiveness questions, negative selection AIS answer ‘yes’ to only the first question having been identified as providing unique results, and so are classified as effective. Clonal selection AIS answer ‘no’, ‘sometimes’ and ‘not yet clear’ respectively to these questions, and are therefore deemed to be effective only sometimes. Immune network AIS answer ‘yes’ to the first two questions, so are also deemed effective.

In conclusion, Garrett [43] states that the AIS paradigm has provided three distinct types of method that can in some cases produce effective results. However, these two properties have not been combined in a reliable

way. Although this process of reflection is vital part to the scientific process, a number of flaws seem apparent in Garrett's approach to assessing the usefulness of AIS. Firstly, all AIS approaches have been grouped into one of the identified general types of AIS to be assessed for their usefulness. This does not seem to be helpful, as too many generalisations and assumptions have to be made about specific AIS approaches. For example, two AIS inspired by the clonal selection algorithm, such as CLONALG [14] and the B-cell Algorithm [70], may possess very different, but equally useful properties. By considering them as the same type of algorithm, these properties may be lost for the purposes of evaluation. Secondly, as the distinctiveness and effectiveness questions are answered via an empirical investigation, it can be argued that the answers to these questions are entirely subjective, depending on the level of abstraction at which the AIS are examined. For example, the third of the distinctiveness questions looks at the distinctiveness of the processes in the algorithm. On one level, it could be argued that evolutionary algorithms and clonal selection algorithms do not contain distinct processes as both are population based algorithms with mutation and selection mechanisms. Garrett [43] argues differently, stating that they are distinct as these mechanisms in clonal selection algorithms are unique, being related to fitness of the solution. However, by this argument two evolutionary algorithms with slightly different mutation and selection mechanism to each other could also be defined as distinct. In the end, this becomes simply a process of arbitrary classification, and it is not clear how this can benefit the future development of AIS.

Hart and Timmis [59] state that unlike the field of evolutionary algorithms, there are few exemplars that stand out in the AIS literature where AIS have been successfully applied to hard real-world problems, or used in industry (one exception with use in industry is work by De Lemos *et al.* on automated teller machines [32]). It is important, therefore, that effort be made to establish a distinctive niche for AIS. Thus, in a similar vein to Garrett [43], Hart and Timmis [59] argue that for a paradigm such as AIS to be considered useful, it is not sufficient for it to simply outperform other algorithms, but it should contain features not contained within other paradigms. It is these features that make a paradigm distinctive and that Hart and Timmis attempt to identify for AIS in [59]. The way forward for AIS should therefore involve the selection of application areas that map the problem features to mechanisms expressed by the immune system. In addition, they argue that this needs to be backed up by more theoretical research into the workings of AIS.

For each of the three identified AIS application areas highlighted in section 5.6, Hart and Timmis [59] review whether or not applying AIS to these areas brings any benefits that could not have been gained via alternative methods. As a guide, their review considers the list of immune system features computationally relevant to AIS that were identified by Dasgupta [25]: recognition, feature extraction, diversity, learning, memory, distributed detection, self recognition, thresholds, co-stimulation, dynamic protection and probabilistic detection. The purpose of anomaly detection techniques is to decide whether an unknown data item is produced by the same probability

distribution as a training set of data, so as to classify this unknown data item. Thus, immune inspired anomaly detection techniques (typically negative selection AIS) attempt to generate detectors capable of detecting when the normal state of a system has changed. Theoretical and empirical investigations by Stibor [98] have highlighted, however, that both real-valued and hamming shape space versions of negative selection AIS are problematic for real-world anomaly detection problems. In addition, they compare negative selection to a support vector machine approach and conclude that results for support vector machines are as good, if not better than the immune-inspired approach of negative selection. Thus, for anomaly detection, Hart and Timmis [59] conclude that, at that point, immune inspired approaches do not seem to offer anything over alternative approaches, this concurs to some degree with Garret [43] but as Garret points out the fact that the subsets of the detectors generated via this approach can be easily distributed might offer an argument for being unique (according to Garret's criteria). Work by Esponda *et al.* extended the ideas of negative selection to that of negative information [34] that can be applied to a security context. Recent work by Greensmith *et al.* [49,50] offer an alternative to the negative selection approach for anomaly detection. Their basic system is outlined in section 5.5 and from initial investigations seems to be performing well in terms of identification of anomalies.

Most AIS designed for optimisation, Hart and Timmis [59] note, are based on the clonal selection principle, and are applied to static optimisation problems. These algorithms require only two of the features highlighted by Dasgupta [25], a diversity mechanism and a memory mechanism. These mechanisms are common to many other optimisation techniques such as evolutionary algorithms. In fact, clonal selection algorithms are considered by many to be a form of evolutionary algorithms. Thus, for static optimisation functions, Hart and Timmis [59] conclude that AIS provide no added value. They state, however, that immune approaches may be more applicable to dynamic optimisation, whereby a solution must be found and tracked in a continuously moving environment. Despite this, they conclude that the immune system is not a suitable model for the inspiration of optimisation methods. Like optimisation AIS, the majority of clustering and classification AIS have been applied to problems with static data sets. Of Dasgupta's list [25], these AIS require feature extraction, recognition and learning. According to Hart and Timmis [59], these features are also key to all machine-learning algorithms, and so again there are no unique features in AIS indicating they offer anything over other machine-learning techniques. Again, however, AIS may be more applicable to dynamic clustering or classification in which patterns and trends are tracked in data over time. To be effective, this should require a form of memory, one of the highlighted properties of AIS. Thus, AIS may be able to outperform machine-learning methods that do not possess a memory mechanism for this task of dynamic clustering/classification. This view of applying AIS to such dynamic problems was echoed in discussion sessions at the recent (at the time of writing this paper) ICARIS conference [31] where it was felt that in order for AIS to compete, or maybe more importantly to

offer something unique, AIS would be better suited to dynamic problems as inernet in many AIS is the ability to continually learn, forget and adapt.

In a recent position paper, Timmis [103] states that the area of AIS has reached “an impasse” and is being hampered by the lack of attention being payed to the underlying biological system (both in terms of immunology and interactions with other systems), the lack of theoretical foundations being laid and the lack of challenging application areas to drive forward the engineering aspect to AIS. This paper takes a slightly different perspective to that of Garrett [43] in so much that Timmis argues there are a number of factors, which when combined, are affecting the progression of AIS from yet another evolutionary technique to something that is, to use Garret’s terms, useful and distinctive. Garrett attempts to assign some meaning to the usefulness and distinctive criteria, but this, as we have discussed, is potentially problematic and by it’s very nature, subjective. To address some of the concerns of Timmis [103], we can look at the papers of Bersini [8], Forrest and Beauchemin [39] and Cohen [19] and conclude that modelling and greater interaction with immunologists can help the development of AIS in bringing greater understanding of the immune system. Through this interaction it may well be possible to begin the development of new, useful and distinctive algorithms and systems, that go way beyond what engineering has to offer to date. Indeed, at the recent ICARIS conference a number of papers were dedicated to this and explore the usefulness of tunable activation thresholds [86, 52], Cohen’s cognitive model [27, 112] and immune networks [57, 78]. However, there is one word of caution in the excitement of modelling, and we echo the concerns of Neal and Trapnel [83] (section 4.3) in that just because the immune system does a certain task in a certain way, it does not mean that an AIS can do the same task in the same way: immune and engineered systems are fundamentally different things. What is key, is to abstract out the computational properties of the immune system, and by seeing the immune system as a computational device [19], this may be the key to future development. It would be primarily *logical* properties that would be extracted, but in contrast to [41] who advocate only logical principles, it is possible that there are *physical* properties that can be used as inspiration (such as the physical structure of lymph nodes), but being mindful that physical properties are difficult to translate from natural to artificial systems. A greater collaboration with immunologists should help us understand in a better way the intricate interaction both within and outside of the immune system: as outlined in another challenge by Timmis [103]. Neal and Trapnel [83] outline such interactions within the immune system and it is clear from this simple view, that the interactions are complex and effective tools are going to be needed for us to even begin to understand such interactions, let alone abstract useful and distinctive computational properties for our artificial systems.

7 Comments on the Future

It is always dangerous to pass comments on what you think will happen in 5 or 10 years time in any area, however a certain amount of speculation is healthy from the point of view of stimulating discussion and debate:

as not everyone will agree with your opinion about the future direction of the research area. In suggesting the way forward for AIS, Hart and Timmis [58] 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, they identify three key ideas mostly missing in the AIS domain. Firstly, the innate immune system has been mainly ignored by AIS practitioners so far (with the recent exceptions of danger theory inspired AIS). This last decade, however, has seen a resurgence of interest in the mechanisms of the innate immune system, such as signalling, by immunologists [44], and its role in controlling the adaptive response. Secondly, the immune system does not operate in isolation. Through the interactions of the immune, neural and endocrine systems, organisms achieve a steady internal state in varying environments, a process called homeostasis. This type of property would be useful in many applications such as anomaly detection and maintenance of engineered systems as outlined by Owens *et al.* [86]. Third and finally, 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. In summary, 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 [59], 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 [43] 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 [43] 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 [86, 27, 52]. Concurring with the above, Bersini [8] 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 [39]. 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.

However, it is not only about the application area, serious developments in theory are required to fully understand how and why the algorithms work the way they do and there are many advances that can be made with respect to modelling the immune system. As a final point we would like to advocate the application of the conceptual framework as a methodology for the development of new immune-inspired systems. The conceptual framework facilitates a truly interdisciplinary approach where expertise from engineering can inform immunology and immunology can inform engineering. We have used this framework as a structure for our paper to highlight the interdisciplinary nature of AIS and through interactions across a variety of disciplines we should begin to harvest the complexity of the immune system into our engineering and, at the same time, develop new insights into the operation and functionality of the immune system. Indeed we concur with Cohen [19] in that a great deal can be learnt on all sides and maybe through the use of the conceptual framework the “literal” and “metaphorical” school may gain a greater understanding and appreciation of the underlying immunology so as to build better immune-inspired systems and the “modelling” school may develop richer and more informative models so as to further our understanding of this amazing complex system. This is not easy and will take the effort of many people over many years, but it is one that we will learn many lessons along the way in our quest to create truly artificial immune systems.

Acknowledgements The authors would like to thank many people in the AIS community who have engaged in many interesting discussions over the years, and in particular at the recent 2007 ICARIS conference who include: Emma Hart, Jorge Carnerio, Hugues Bersini, Rob de Boer, Uwe Aickelin, Julie Greensmith and Leandro de Castro. We would like to thank the reviewers for their many comments and suggestions.

References

1. Aickelin, U., Bentley, P., Cayzer, S., Kim, J., McLeod, J.: Danger theory: The link between AIS and IDS? In: Timmis et al. [104], pp. 147–155
2. Anderson, C.C., Matzinger, P.: Danger: the view from the bottom of the cliff. *Seminars in Immunology* **12**(3), 231–238 (2000)
3. Andrews, P.S., Timmis, J.: Inspiration for the next generation of artificial immune systems. In: Jacob et al. [63], pp. 126–138
4. Andrews, P.S., Timmis, J.: On diversity and artificial immune systems: Incorporating a diversity operator into aiNet. In: Proceedings of the International Conference on Natural and Artificial Immune Systems (NAIS05), *LNCS*, vol. 391, pp. 293–306. Springer (2005)
5. Andrews, P.S., Timmis, J.: A computational model of degeneracy in a lymph node. *LNCS*, pp. 164–177. Springer (2006)
6. Balthrop, J., Esponda, F., Forrest, S., Glickman, M.: Coverage and generalisation in an artificial immune system. In: Genetic and Evolutionary Computation, pp. 3–10 (2002)
7. Bersini, H.: Self-assertion versus self-recognition: A tribute to Francisco Varela. In: Timmis and Bentley [105], pp. 107–112
8. Bersini, H.: Immune system modeling: The OO way. In: Bersini and Carneiro [9], pp. 150–163

9. Bersini, H., Carneiro, J. (eds.): Proceedings of the 5th International Conference on Artificial Immune Systems, *LNCS*, vol. 4163. Springer (2006)
10. Bezerra, G., Barra, T., de Castro, L.N., Von Zuben, F.: Adaptive Radius Immune Algorithm for Data Clustering. In: Jacob et al. [63], pp. 290–303
11. Bretscher, P.: Contemporary models for peripheral tolerance and the classical ‘historical postulate’. *Seminars in Immunology* **12**(3), 221–229 (2000)
12. Burnet, F.M.: *The Clonal Selection Theory of Acquired Immunity*. Cambridge University Press (1959)
13. de Castro, L.N., Von Zuben, F.J.: An evolutionary immune network for data clustering. In: Proceeding of the IEEE Brazilian Symposium on Artificial Neural Networks, pp. 84–89 (2000)
14. de Castro, L.N., Von Zuben, F.J.: Learning and optimization using the clonal selection principle. *IEEE Transactions on Evolutionary Computation* **6**(3), 239–251 (2002)
15. Clark, E., Hone, A., Timmis, J.: A markov chain model of the b-cell algorithm. In: Jacob et al. [63], pp. 318–330
16. Coelho, G., Von Zuben, F.J.: omni-ainet: An immune-inspired approach for omni optimization. In: Bersini and Carneiro [9], pp. 294–308
17. Cohen, I.R.: Discrimination and dialogue in the immune system. *Seminars in Immunology* **12**(3), 215–219 (2000)
18. Cohen, I.R.: *Tending Adam’s Garden: Evolving the Cognitive Immune Self*. Elsevier Academic Press (2000)
19. Cohen, I.R.: Real and artificial immune systems: Computing the state of the body. *Imm. Rev.* **7**, 569–574 (2007)
20. Cohen, I.R., Harel, D.: Explaining a complex living system: dynamics, multi-scaling and emergence [doi:10.1098/rsif.2006.0173](https://doi.org/10.1098/rsif.2006.0173) (2006)
21. Coutinho, A.: The network theory: 21 years later. *Scandinavian Journal of Immunology* **42**, 3–8 (1995)
22. Cruz Cortes, N., Coello Coello, C.: Multiobjective optimisation using ideas from the clonal selection principle. In: E. Cantu-Paz (ed.) *Genetic and Evolutionary Computation (GECCO)*, vol. 1, pp. 158–170 (2003)
23. Cutello, V., Nicosia, G.: The clonal selection principle for in silico and in vitro computing. In: L.N. de Castro, F.J. von Zuben (eds.) *Recent Developments in Biologically Inspired Computing*. Idea Group Publishing, Hershey, PA (2004)
24. Cutello, V., Nicosia, G., Oliveto, P., Romeo, M.: On the convergence of immune algorithms. In: *Proc. of Foundations of Computational Intelligence*, pp. 409–416 (2007)
25. Dasgupta, D. (ed.): *Artificial Immune Systems and their Applications*. Springer (1999)
26. Dasgupta, D., Krishna Kumar, K., Wong, D., Berry, M.: Negative selection algorithm for aircraft fault detection. In: Nicosia et al. [85], pp. 1–14
27. Davoudani, D., Hart, E., Paechter, B.: An immune-inspired approach to speckled computing. In: de Castro et al. [31], pp. 288–299
28. de Castro, L.N., Timmis, J.: *Artificial Immune Systems: A New Computational Intelligence Approach*. Springer (2002)
29. de Castro, L.N., Von Zuben, F.J.: Artificial immune systems: Part I—basic theory and applications. Tech. Rep. DCA-RT 01/99, School of Computing and Electrical Engineering, State University of Campinas, Brazil (1999)
30. de Castro, L.N., Von Zuben, F.J.: Artificial immune systems: Part II—a survey of applications. Tech. Rep. DCA-RT 02/00, School of Computing and Electrical Engineering, State University of Campinas, Brazil (2000)
31. de Castro, L.N., Von Zuben, F.J., Knidel, H. (eds.): Proceedings of the 6th International Conference on Artificial Immune Systems, *Lecture Notes in Computer Science*, vol. 4628. Springer (2007)
32. de Lemos, R., Timmis, J., Forrest, S., Ayara, M.: Immune-inspired adaptable error detection for automated teller machines. *IEEE Trans. on Systems, Man and Cybernetics Part C: Applications and Reviews* **37**(5), 873–886 (2007)
33. Edelman, L., Rosen, R.: Enzyme-substrate recognition. *Journal of Theoretical Biology* **73**(1), 181–204 (1978)

34. Esponda, F.: Negative representations of information. Ph.D. thesis, University of New Mexico (2005)
35. Esponda, F., Ackley, E.S., Forrest, S., Helman, P.: On-line negative databases (with experimental results). *International Journal of Unconventional Computing* **1**(3), 201–220 (2005)
36. Farmer, J.D.: A rosetta stone for connectionism. *Physica D* **42**, 153–187 (1990)
37. Farmer, J.D., Packard, N.H., Perelson, A.S.: The immune system, adaptation, and machine learning. *Physica D* **22**, 187–204 (1986)
38. Folwer, M.: *UML Distilled*. Addison Wesley (2004)
39. Forrest, S., Beauchemin, C.: Computer Immunology. *Immunol. Rev.* **216**(1), 176–197 (2007)
40. Forrest, S., Perelson, A., Allen, L., Cherukuri, R.: Self–nonself discrimination in a computer. In: *Proceedings of the IEEE Symposium on Research Security and Privacy*, pp. 202–212 (1994)
41. Freitas, A., Timmis, J.: Revisiting the foundations of artificial immune systems for data mining. *IEEE Trans. Evol. Comp.* **11**(4), 521–540 (2007)
42. Gamma, E., Helm, R., Johnson, R., Vlissides, J.: *Design Patterns*. Addison Wesley (1995)
43. Garrett, S.: How do we evaluate artificial immune systems? *Evolutionary Computation* **13**(2), 145–177 (2005)
44. Germain, R.N.: An innately interesting decade of research in immunology. *Nature Medicine* **10**, 1307–1320 (2004)
45. Gillespie, D.: Approximate accelerated stochastic simulation of chemically reacting systems. *J. Phys. Chem* **81**(25), 2340–2361 (1977)
46. Goldsby, R.A., Kindt, T.J., Osborne, B.A., Kuby, J.: *Immunology* 5th edition. W. H. Freeman and Company (2003)
47. González, F., Dasgupta, D., Gómez, J.: The effect of binary matching rules in negative selection. In: *Genetic and Evolutionary Computation – GECCO-2003, Lecture Notes in Computer Science*, vol. 2723, pp. 195–206. Springer-Verlag, Chicago (2003)
48. Gonzalez, F.A., Dasgupta, D.: Anomaly detection using real-valued negative selection. *Genetic Programming and Evolvable Machines* **4**(4), 383–403 (2003)
49. Greensmith, J., Aickelin, U., Cayzer, S.: Introducing dendritic cells as a novel immune-inspired algorithm for anomaly detection. In: Jacob et al. [63]
50. Greensmith, J., Aickelin, U., Twycross, J.: Articulation and clarification of the dendritic cell algorithm. In: Bersini and Carneiro [9], pp. 404–417
51. Grossman, Z., Paul, W.E.: Self-tolerance: context dependent tuning of T cell antigen recognition. *Seminars in Immunology* **12**(3), 197–203 (2000)
52. Guzella, T., Mota-Santos, T., Caminhas, W.: Towards a novel immune inspired approach to temporal anomaly detection. In: de Castro et al. [31], pp. 119–130
53. Harel, D.: Statecharts: a visual formalism for complex systems. *Sci. Computer Program* **8**(231-274) (1987)
54. Hart, E.: Not all balls are round: An investigation of alternative recognition-region shapes. In: Jacob et al. [63], pp. 29–42
55. Hart, E., Bersini, H., Santos, F.: Tolerance vs intolerance: How affinity defines topology in an idiotypic network. In: Bersini and Carneiro [9], pp. 109–121
56. Hart, E., Ross, P.: Studies on the implications of shape-space models for idiotypic networks. In: Nicosia et al. [85], pp. 413–426
57. Hart, E., Santos, F., Bersini, H.: Topological constraints in the evolution of idiotypic networks. In: de Castro et al. [31], pp. 252–263
58. Hart, E., Timmis, J.: Application areas of AIS: The past, the present and the future. In: Jacob et al. [63], pp. 483–497
59. Hart, E., Timmis, J.: Application areas of AIS: The past, the present and the future. To appear in *Applied Soft Computing* (2007). In Press, Corrected Proof, Available online 12 February 2007
60. Hofmeyr, S., Forrest, S.: Architecture for an artificial immune system. *Evolutionary Computation* **7**(1), 1289–1296 (2000)

61. Hone, A., van den Berg, H.A.: Modelling a cytokine network. In: Proc. of IEEE Workshop on Computational Intelligence, pp. 389–293 (2007)
62. Honorio, L., Leite da Silva, A., Barbosa, D.: A gradient-based artificial immune system applied to optimal power flow problems. In: de Castro et al. [31], pp. 1–12
63. Jacob, C., Pilat, M., Bentley, P., Timmis, J. (eds.): Proc. of the 4th International Conference on Artificial Immune Systems (ICARIS), *Lecture Notes in Computer Science*, vol. 3627. Springer (2005)
64. Janeway, C., Travers, P., Walport, M., Shlomchik, M.: Immunobiology: The Immune System is Health and Disease, 6th edition edn. Garland Science (2004)
65. Janeway, C.A.: The immune system evolved to discriminate infectious nonself from noninfectious self. *Immunology Today* **13**, 11–16 (1992)
66. Janeway, C.A., Travers, P., Walport, M., Shlomchik, M.: Immunobiology 5th edition. Garland Publishing (2001)
67. Jerne, N.K.: Towards a network theory of the immune system. *Ann. Immunol. (Inst. Pasteur)* **125C**, 373–389 (1974)
68. Ji, Z., Dasgupta, D.: Artificial immune system (AIS) research in the last five years. In: Congress on Evolutionary Computation, vol. 1, pp. 123–130. IEEE, Canberra, Australia (2003)
69. Ji, Z., Dasgupta, D., Yang, Z., Teng, H.: Analysis of dental images using artificial immune systems. In: Proceedings of Congress On Evolutionary Computation (CEC), pp. 528–535. IEEE Press (2006)
70. Kelsey, J., Timmis, J.: Immune inspired somatic contiguous hypermutation for function optimisation. In: Genetic and Evolutionary Computation Conference - GECCO 2003, pp. 207–218. Springer (2003)
71. 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)
72. Langman, R.E., Cohn, M.: Editorial summary. *Seminars in Immunology* **12**(3), 343–344 (2000)
73. Langman, R.E., Cohn, M.: A minimal model for the self–nonself discrimination: a return to basics. *Seminars in Immunology* **12**(3), 189–195 (2000)
74. Lutz, M.B., Schuler, G.: Immature, semi-mature and fully mature dendritic cells: which signals induce tolerance or immunity? *Trends in Immunology* **23**(9), 445–449 (2002)
75. Matzinger, P.: Tolerance, danger and the extended family. *Annual Review of Immunology* **12**, 991–1045 (1994)
76. Matzinger, P.: The danger model: A renewed sense of self. *Science* **296**, 301–305 (2002)
77. May, P., Timmis, J., Mander, K.: Immune and evolutionary approaches to software mutation testing. In: de Castro et al. [31], pp. 336–347
78. McEwan, C., Hart, E., Paechter, B.: Revisiting the central and peripheral immune system. In: de Castro et al. [31], pp. 240–251
79. Medzhitov, R., Janeway, C.A.: How does the immune system distinguish self from nonself? *Seminars in Immunology* **12**(3), 185–188 (2000)
80. Mendao, M., Timmis, J., Andrews, P.S., Davies, M.: The immune system in pieces: Computational lessons from degeneracy in the immune system. In: D.B. Fogel (ed.) Proc. of Foundations of Computational Intelligence, pp. 394–400 (2007)
81. Milner, R.: *Communicating and Mobile Systems: the π -Calculus*. Cambridge University Press (1999)
82. Mosmann, T.R., Livingstone, A.M.: Dendritic cells: the immune information management experts. *Science* **5**(6), 564–566 (2004)
83. Neal, M., Trapnel, B.: In Silico Immunology, chap. Go Dutch: Exploit Interactions and Environments with Artificial Immune Systems, pp. 313–330. Springer (2007)
84. Newborough, R., Stepney, S.: A generic framework for population based algorithms. In: Jacob et al. [63]

-
85. 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)
 86. Owens, N., Timmis, J., Greensted, A., Tyrrell, A.: On immune inspired homeostasis for electronic systems. In: de Castro et al. [31], pp. 216–227
 87. Perelson, A.S., Oster, G.F.: Theoretical studies of clonal selection: Minimal antibody repertoire size and reliability of self–non-self discrimination. *Journal of Theoretical Biology* **81**(4), 645–670 (1979)
 88. 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), vol. 4695, pp. 184–199 (2007)
 89. Priami, C.: Stochastic π -calculus. *The Computer Journal* **38**(7), 578–589 (1995)
 90. Priami, C., Regev, A., Shapiro, E.: Application of a stochastic name-passing calculus to representation for biological processes in the stochastic π -calculus. *Information Processing Letters* **80**, 25–31 (2001)
 91. Regev, A., Silverman, W., Shapiro, E.: Representation and simulation of biochemical processes using the pi-calculus process algebra. In: Pacific Symposium on Biocomputing, vol. 6, pp. 459–470 (2001)
 92. Rudolph, G.: Finite markov chain results in evolutionary computation: a tour d’horizon. *Fundam. Inf.* **35**(1-4), 67–89 (1998)
 93. Secker, A., Freitas, A.: Wairs: Improving classification accuracy by weighting attributes in the airs classifier. In: Proceedings of the Congress on Evolutionary Computation, pp. 3759–3765 (2007)
 94. Secker, A., Freitas, A., Timmis, J.: A danger theory inspired approach to web mining. In: Timmis et al. [104], pp. 156–167
 95. Silverstein, A.M., Rose, N.R.: There is only one immune system! The view from immunopathology. *Seminars in Immunology* **12**(3), 173–178 (2000)
 96. Stepney, S., Smith, R., Timmis, J., Tyrrell, A., Neal, M., Hone, A.: Conceptual frameworks for artificial immune systems. *Int. J. Unconventional Computing* **1**(3), 315–338 (2006)
 97. Stepney, S., Smith, R.E., Timmis, J., Tyrrell, A.M.: Towards a conceptual framework for artificial immune systems. In: Nicosia et al. [85], pp. 53–64
 98. Stibor, T.: On the appropriateness of negative selection for anomaly detection and network intrusion detection. Ph.D. thesis, Darmstadt University of Technology (2006)
 99. Stibor, T.: Phase transition and the computational complexity of generating r -contiguous detectors. In: de Castro et al. [31], pp. 142–155
 100. Stibor, T., Timmis, J.: An investigation into the compression quality of ainet. In: D. Fogel (ed.) *Proc. of Foundations of Computational Intelligence* (2007)
 101. Stibor, T., Timmis, J., Eckert, C.: The link between r -contiguous detectors and k -cnf satisfiability. In: Proceedings of the Congress on Evolutionary Computation, pp. 491–498 (2006)
 102. Tauber, A.I.: Moving beyond the immune self? *Seminars in Immunology* **12**(3), 241–248 (2000)
 103. Timmis, J.: Artificial immune systems: Today and tomorrow. *Natural Computing* **6**(1), 1–18 (2007)
 104. Timmis, J., Bentley, P., Hart, E. (eds.): Proceedings of the 2nd International Conference on Artificial Immune Systems (ICARIS 2003), LNCS 2787. Springer (2003)
 105. Timmis, J., Bentley, P.J. (eds.): Proceedings of the 1st International Conference on Artificial Immune Systems (ICARIS 2002). University of Kent Printing Unit (2002)
 106. Timmis, J., Hone, A., Stibor, T., Clark, E.: Theoretical advances in artificial immune systems. Submitted to *Journal of Theoretical Computer Science* (2007)
 107. Timmis, J., Knight, T.: Data Mining: A Heuristic Approach, chap. Artificial immune systems: Using the immune system as inspiration for data mining., pp. 209–230. Idea Group (2001)

108. Timmis, J., Neal, M.: A resource limited artificial immune system for data analysis. *Knowledge Based Systems* **14**(3-4), 121-130 (2001)
109. Twycross, J., Aickelin, U.: Towards a conceptual framework for innate immunity. In: Jacob et al. [63], pp. 112-125
110. Varela, F.J., Coutinho, A.: Second generation immune networks. *Immunology Today* **12**(5), 159-166 (1991)
111. Villalobos-Arias, M., Coello, C.A.C., Hernandez-Lerma, O.: Convergence analysis of a multiobjective artificial immune system algorithm. In: Nicosia et al. [85], pp. 226 - 235
112. Voigt, D., Wirth, H., Dilger, W.: A computational models for the cognitive immune system theory based on learning classifier systems. In: de Castro et al. [31], pp. 264-275
113. Watkins, A., Timmis, J., Boggess, L.: Artificial immune recognition system (AIRS): An immune-inspired supervised learning algorithm. *Genetic Programming and Evolvable Machines* **5**(1), 291-317 (2004)