

# Artificial Immune Systems - Today and Tomorrow

**Jon Timmis**

Department of Computer Science and Department of Electronics, University of  
York, Heslington, York. YO10 5DD. UK

Received: Jan 10th 2006 / Revised version: April 7th 2006

**Abstract** In this position paper, we argue that the field of Artificial Immune Systems (AIS) has reached an impasse. For many years, immune inspired algorithms, whilst having some degree of success, have been limited by the lack of theoretical advances, the adoption of a naive immune inspired approach and the limited application of AIS to challenging problems. We review the current state of the AIS approach, and suggest a number of *challenges* to the AIS community that can be undertaken to help move the area forward.

## 1 Introduction

The UK research community has proposed a number of Grand Challenges for Computer Science research<sup>1</sup> and ambitious plans for the development of a variety of research areas. Grand Challenge 7 (GC-7)<sup>2</sup> [55] addresses the area of *Non-Classical Computation*, which includes exploring areas of both biologically inspired paradigms, and the exploitation of the natural world (for example, DNA computing and quantum computing) in order to develop new areas of computation. GC-7 consists of a number of *journeys* of which one is concerned with Artificial Immune Systems. In the spirit of GC-7, this position paper proposes a number of challenges to the AIS community, and initial thoughts on how we might go about addressing those challenges. We begin by providing a brief description of AIS in terms of a simple framework (as defined by [13]). We then explore the current *immunological mind set* of the AIS practitioner with a simple overview of the immunology that has served as inspiration to date, combining that with brief discussions on how those theories have helped to shape the development of AIS over the past 10 years. We then take a moment to reflect on the area of AIS as a whole and argue that there has been a lack of consideration of a number of issues, specifically: a lack of thought regarding the application areas of AIS, a lack of theoretical work, a limited view of the immune system, and we comment on the methodology employed in the development of AIS. We then conclude

---

<sup>1</sup> [http://www.nesc.ac.uk/esi/events/Grand\\_Challenges/](http://www.nesc.ac.uk/esi/events/Grand_Challenges/)

<sup>2</sup> <http://www.cs.york.ac.uk/nature/gc7/index.htm>

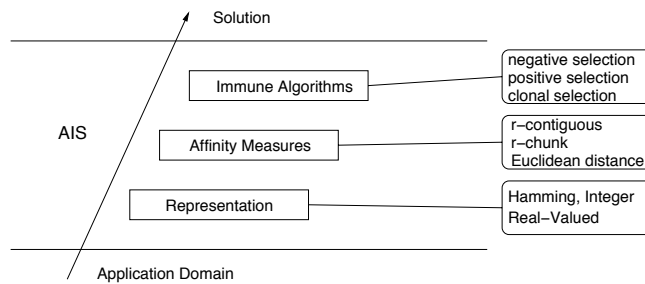
with a number of *challenges* to the AIS community. It should be said, that many of these ideas are not new, nor all my own, but have come from many discussions with people in the AIS community. This paper is intended to be both a review and position paper that will hopefully draw together many ideas and stimulate further discussion and research.

## 2 Artificial Immune Systems

In an attempt to create a common basis for AIS, work in [13] proposed the idea of a framework for AIS. The authors argued that in the case of other biologically inspired approaches, such as artificial neural networks (ANN) and evolutionary algorithms (EAs) such a basic idea of a framework exists, and helps considerably with the understanding and construction of such systems. For example, [13] consider a set of artificial neurons which can be arranged together so as to form an artificial neural network. In order to acquire knowledge, these neural networks undergo an adaptive process, known as learning or training, which alters (some of) the parameters within the network. Therefore, the authors argued that in a simplified form, a framework to design an ANN is composed of: a set of artificial neurons, a pattern of interconnection for these neurons, and a learning algorithm. Similarly, the authors argued that in evolutionary algorithms there is a set of "artificial chromosomes", representing a population of individuals, that iteratively suffer a process of reproduction, genetic variation, and selection. As a result of this process, a population of evolved artificial individuals arises. A

framework, in this case, would correspond to the genetic representation of the individuals of the population, plus the procedures for reproduction, genetic variation, and selection. Therefore, the authors adopted the viewpoint that a framework to design a biologically inspired algorithm requires, at least, the following basic elements:

- A representation for the components of the system
- A set of mechanisms to evaluate the interaction of individuals with the environment and each other. The environment is usually simulated by a set of input stimuli, one or more fitness function(s), or other mean(s) and
- Procedures of adaptation that govern the dynamics of the system, i.e. how its behavior varies over time



**Fig. 1 AIS Layered Framework adapted from [13]**

The framework for AIS can be thought of as a layered approach (Figure 1). In order to build a system, one typically requires an application domain or target function. From this basis, the way in which the components of the system will be represented will be considered. For example, the representa-

tion of network traffic may well be different than the representation of a real time embedded system. In AIS, the way in which something is represented is known as *shape space* [49]. There are many kinds of shape space, such as Hamming, Real valued and so on, each of which carries its own bias and should be selected with care [22]. Once the representation has been chosen, one or more affinity measures are used to quantify the interactions of the elements of the system. There are many possible affinity measures (which are partially dependent upon the representation adopted), such as Hamming and Euclidean distance metrics. Again, each has its own bias, and the affinity function must be selected with great care as it can affect the overall performance (and ultimately the result) of the system [22]. The final layer involves the use of algorithms, which govern the behavior (dynamics) of the system. Here, in the original framework proposal, algorithms based on the following immune processes were presented: negative and positive selection, clonal selection, bone marrow, and immune network algorithms. Work by Garrett [24] outlines criteria for assessing AIS in terms of effectiveness, usefulness and distinctiveness, and attempts to draw some conclusions about how AIS can be said to differ from other biologically inspired approaches. In this paper, we will not revisit that discussion as this is done elsewhere [30].

### 3 The Immune System: From an Artificial Immune Systems Perspective

The vast majority of developments within AIS focussed on three main immunological theories: clonal selection [8], immune networks [39] and negative selection. Researchers in AIS have concentrated, for the most part, on the *learning* and *memory* mechanisms of the immune system (typically taking clonal selection and immune network theories as a basis) and the selection of detectors for identifying anomalous entities (typically undertaken with negative selection theory).

It is becoming increasingly apparent that the biological inspiration behind AIS has been somewhat naive and that the mechanisms and processes within the immune system (not to mention the role of the immune system) exploited by the AIS community has taken a limited perspective. Indeed, as argued by [56] AIS, and the majority of bio-inspired paradigms, have been guilty of a *reasoning by metaphor* approach. We would agree with this, and despite the success to date of AIS (and this should not be ignored) the restricted view of the immune system adopted by the AIS practitioners will, in our opinion, limit the success of AIS.

In this section, we review the immunology that has served as the foundations for much of AIS. We outline the main immunological theories that have acted as a source of inspiration, notably: clonal selection, immune networks and negative selection theories, and provide a brief explanation of how these have been exploited within AIS. It is not our intention to provide

a detailed description of these algorithms here, a full review of these can be found in [13].

### 3.1 Immunity

The vertebrate immune system is composed of diverse sets of cells and molecules that work together with other systems, such as neural and endocrine, in order to maintain a *steady state* within the host. One traditionally held view on the role of the immune system is to protect our bodies from infectious agents such as viruses, bacteria, fungi and other parasites, however, not everyone would agree with that statement [10]. On the surface of these agents are antigens that allow the identification of the invading agents (pathogens) by the immune cells and molecules, thus provoking an immune response. There are two basic types of immunity, innate and adaptive. Innate immunity is not directed towards specific invaders, but against general pathogens that enter the body [38]. The innate immune system plays a vital role in the initiation and regulation of immune responses, including adaptive immune responses. Specialized cells of the innate immune system evolved so as to recognize and bind to common molecular patterns found only in microorganisms. However, the innate immune system is by no means a complete solution to protecting the body.

Adaptive or acquired immunity, allows the immune system to launch an attack against any invader that the innate system cannot remove [37]. The adaptive system is directed against specific invaders and is modified

by exposure to such invaders. The adaptive immune system mainly consists of lymphocytes, which are white blood cells, more specifically B and T-cells. These cells aid in the process of recognizing and destroying specific substances. Any substance that is capable of generating a response from the lymphocytes is called an antigen, or immunogen. Antigens are not the invading microorganisms themselves, they are substances such as toxins or enzymes in the microorganisms that the immune system considers foreign. Adaptive immune responses are normally directed against the antigen that provoked them and are said to be antigen-specific.

### *3.2 Clonal Selection*

A large part of AIS work has been based on the clonal selection theory. When antibodies on a B-cell bind with an antigen, the B-cell becomes activated and begins to proliferate. New B-cell clones are produced that are an exact copy of the parent B-cell, that then undergo somatic hypermutation [2] and produce antibodies that are specific to the invading antigen. The clonal selection principle [8] is the term used to describe the basic properties of an adaptive immune response to an antigenic stimulus. It established the idea that only those cells capable of recognizing an antigenic stimulus will proliferate, thus being selected against those that do not. Clonal selection operates on both T-cells and B-cells. The B-cells, in addition to proliferating or differentiating into *plasma cells*, can differentiate into long-lived B *memory cells*. Memory cells circulate through the blood, lymph and tissues.

In order for the immune system to be protective over periods of time, antigen recognition is insufficient. In the normal course of the evolution of the immune system, an organism would be expected to encounter a given antigen repeatedly during its lifetime. The initial exposure to an antigen that stimulates an adaptive immune response (an immunogen) is handled by a small number of B-cells, each producing antibodies of different affinity. Storing some high affinity antibody producing cells from the first infection, so as to form a large initial specific B-cell sub-population (clone) for subsequent encounters, considerably enhances the effectiveness and speed of the immune response to secondary encounters. Rather than starting from a *tabula rasa* every time, such a strategy ensures that both the speed and accuracy of the immune response becomes successively stronger after each infection.

Computationally, this has led to the development of population based algorithms inspired by this clonal selection process. A great number of the approaches in the literature have focussed on developing optimisation approaches, such as the work on CLONALG by [15], the work by [12] and the work by [23] on parameter free clonal selection. Clonal selection has also formed the basis of learning algorithms (mainly supervised) such as AIRS by [64] and a parallel and a distributed version of AIRS [65], a distributed version of CLONALG by [66], the work with DynamicCS applied to computer security by [41], work by [51] on email filtering, and many, many more. From a computational perspective, application of the clonal selection

theory leads to algorithms that evolve (through a cloning, mutation and selection phase), candidate solutions in terms of optimisation, or pattern detectors in terms of learning. Each of these algorithms have populations of *B-cells* (candidate solution) that match against *antigens* (function to be optimised). These B-cells then undergo cloning (usually in proportion to the strength of the match) and mutation (usually, inversely proportional to the strength of the match). High affinity B-cells are then selected to remain in the population, some low affinity cells are removed and new random cells are generated. In essence, this is a high level abstraction of the clonal selection process. Through this process, good solutions can be found, and in terms of dynamic environments (such as [41,51]) these solutions can be maintained over long periods of time. However, as argued by [48], at a certain level are such population based algorithms like clonal selection ones that can be considered the same as genetic algorithms and other evolutionary algorithms. Figure 2 outlines the basic process in a clonal selection algorithm.

### *3.3 Immune Networks*

In a landmark paper for the time Jerne [39] proposed that the immune system is capable of achieving immunological memory by the existence of a mutually reinforcing network of B-cells. This network of B-cells occurs due to the ability of paratopes, located on B-cells, to match against idiotopes on other B-cells. The binding between idiotopes and paratopes has the effect of stimulating the B-cells. This is because the paratopes on B-cells react

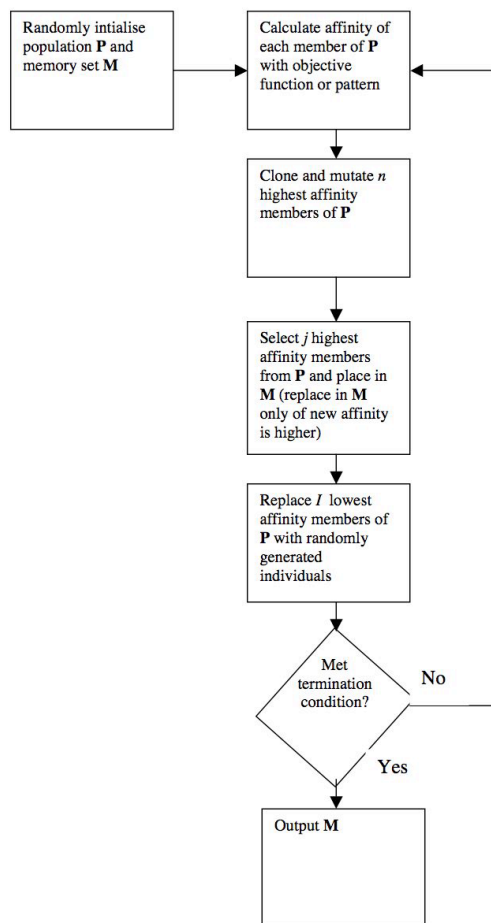


Fig. 2 Flowchart of a Typical Clonal Selection Algorithm

to the idiotopes on similar B-cells, as it would an antigen. However, to counter the reaction, there is a certain amount of suppression between B-cells, thus giving rise to a regulatory mechanism. This interaction of B-cells contributes to form a *stable* memory structure, and can account for the retainment of memory cells, even in the absence of antigen. This theory was

refined and formalised in successive works by [18,49] and [5]. It is worth noting, however, that the immune network theory has not found a great deal of favour with other immunologists, being labeled as strongly as *absurd* by some [44]. Indeed, talk to many experimental immunologists today and they are of the opinion that, whilst this theory is interesting, there is little experimental evidence to suggest that it is true.

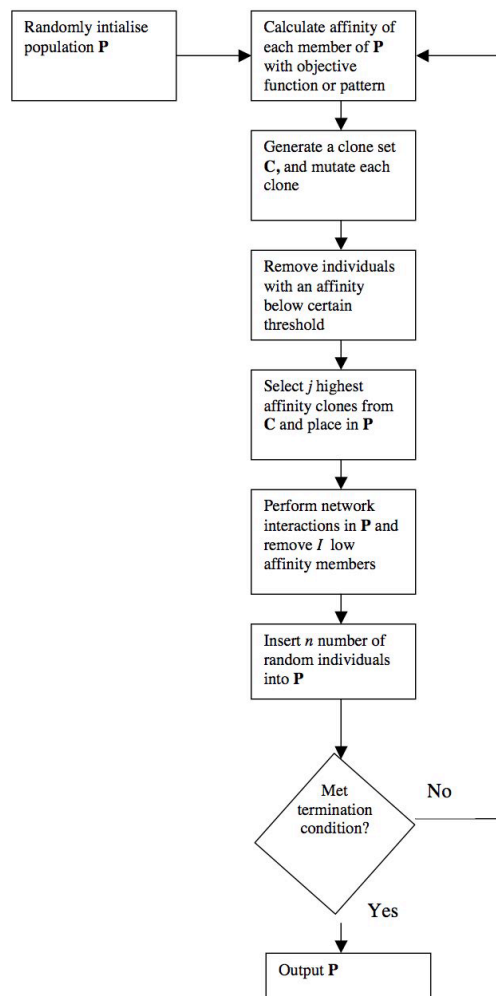
However, this lack of adoption in immunology circles of the immune network theory, has not stopped people from the area of AIS adapting it for their own use. Indeed, from a computational perspective, this has led to the development of a number of immune network algorithms, all of which vary somewhat, making general statements difficult. Work by [36] on the application of immune network to diagnostic problems, work by [11,34,35] on the development of an immune network approach to DNA classification and mortgage fraud detection, and work on aiNET by [14,16] all attempt to extract properties such as adaptation, self-organisation and plasticity from the immune network theory. The work by [35] proved a useful basis for the refining and extending of their work into an unsupervised system by [62,61], which itself was then adopted for continuous learning by [46,67]. There is little to unify these immune network algorithms, but an abstracted outline of one is presented in Figure 3. The main unifying theme is that they are all, in one way or another, based on the ideas of clonal selection (as outlined above). This was observed by [25] where the author proposed a general artificial immune network, however, this was mainly conceptual and has never

been extended further. The main addition to immune network algorithms (in general) is that the candidate solutions (or detectors) interact based on stimulation and suppression mechanisms employed from the immune network. Immune network algorithms have tended to suffer a great deal from a large overhead of computational complexity and a lack of understanding of their dynamics, which has made their applicability to date, rather limited [29].

### *3.4 Negative Selection*

Negative selection is a process of *selection* that takes place in the thymus gland. T-cells are produced in the bone marrow and before they are released into the lymphatic system, undergo a maturation process in the thymus gland. The maturation of the T-cells is conceptually very simple. T-cells are exposed to self-proteins in a binding process. If this binding activates the T-cell, then the T-cell is killed, otherwise it is allowed into the lymphatic system. In addition, there is an initial *positive selection* process that removes T-cells with the correct set of receptors that can recognize the MHC molecules responsible for self-recognition. This process has largely been ignored by AIS researchers, although some initial experiments have been undertaken in [59] where a tendency to overfit the data was reported, but reasons as to why were not given.

The negative selection principle inspired work such as [20, 19] to propose and develop a negative selection algorithm to detect data manipulation



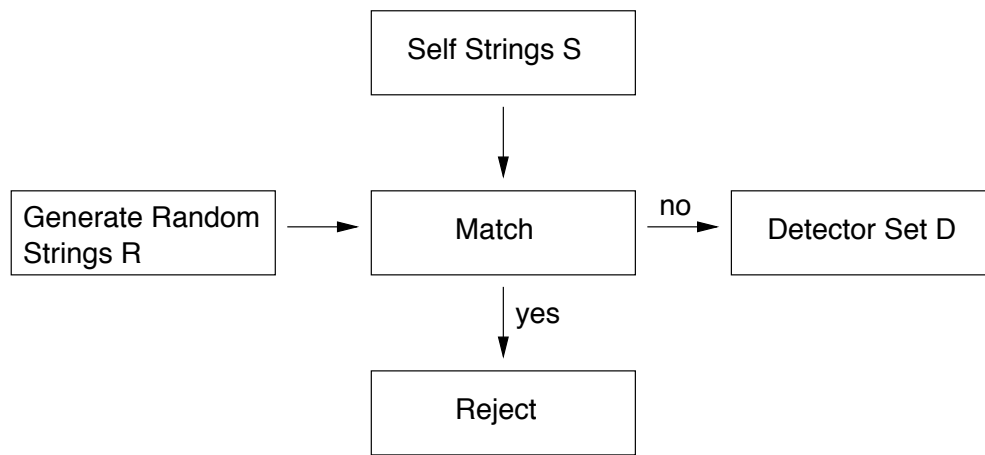
**Fig. 3 Flowchart of a Typical Immune Network Algorithm**

caused by computer viruses. The basic idea is to generate a number of detectors in the complementary space and then to apply these detectors to classify new (unseen) data as self (no data manipulation) or non-self (data manipulation). The negative selection algorithm proposed by Forrest et al.

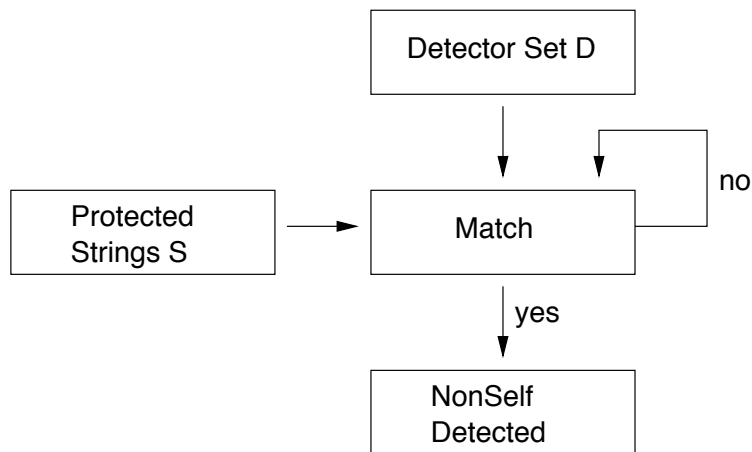
is summarized in the following steps. We can define self as a set  $\mathbf{S}$  of elements of length  $l$  in shape-space. Then generate a set  $\mathbf{D}$  of detectors, such that each fails to match any element in  $\mathbf{S}$ . With these detectors, monitor a continual data stream for any changes, by continually matching the detectors in  $\mathbf{D}$  against the stream. This work spawned a great deal of investigations into the use of negative selection for intrusion detection, with early work meeting with some success [19], and more recently [17]. Figure 1.4 outlines the basic negative selection approach. However, building on the work of [17], work by [57,58] begins to show (from a theoretical and practical perspective) potential problems of scalability and coverage of the problem space when employing the negative selection approach. However, we feel it is unclear from this research whether the negative selection approach, overall, is flawed. From the results, it would seem that the problem is directly associated with the representation and affinity metric issues. Affinity metrics such as  $r$ -contiguous bits and  $r$ -chunk have been tried with binary shape space, and Euclidean distance with real-valued shape space, it may be possible that alternative representations and metrics could be used with negative selection with more success.

### *3.5 Building Artificial Immune Systems*

When constructing an AIS, there are many computational and practical issues to consider. The first is computational complexity of the approach. This relates to the time and space required to generate a suitable number



(a) Detector Generation



(b) Monitoring Phase

**Fig. 4** Negative Selection Algorithm by Forrest [21]

of detectors (members of a population) that are required for the job [60]. For example, there are a number of works that outline the unacceptable computational complexity of the negative selection approach in AIS [42, 57, 58] as there is an exponential relationship between the size of the data set to be used and the number of detectors that it is possible to generate.

The second aspect to consider is the data to be used. When one abstracts away from the underlying data representation (e.g. real values from a sensor) then one has to be careful that there is an accurate mapping between the higher level representation and the actual system, and ensure that the representation adequately scopes the space to be immunised. Consideration here also has to be given to the way in which data is represented. The shape space paradigm proposes varying ways of data representation and interaction. However, when dealing with discrete values, such as those found in embedded systems, the method of defining affinity (i.e. seeing how similar one item is to another) is not as clear-cut as it may seem. This is coupled with the fact that mutation, even what might be thought of as a small amount, could have a huge impact on the meaning of the data. Should a binary shape space be employed, the mere flipping of one bit could indicate a huge shift in meaning of the state, rather than the small shift that may be desired. In both of these situations, domain knowledge can play a pivotal role in the success or failure of such as system [60]. For example, recent studies by [29,28] point out that the main effect on immune network algorithms may well be the way in which interaction is defined. 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 measure between two B-cells, which in turn affects how B-cells interact with each other. Whilst no concrete conclusions are drawn here, the message is clear: think before you design.

## 4 Challenges for AIS

In this section, we now present a number of *challenges* for the AIS community, we then explore how we might begin to address them.

### 4.1 Methodology and Beneficiaries

*Challenge: To Develop Novel and Accurate Metaphors and be a Benefit to Immunology.* Typically naive approaches to extracting metaphors from the immune system have been taken. This has occurred as an accident of history, and AIS has slowly drifted away from its immunological roots. Time is now ripe for greater interaction with immunologists and mathematicians to undertake specific experimentation and create useful models, all of which can be used as a basis for abstraction into powerful algorithms.

In the beginning, AIS were developed with an interdisciplinary slant. For example, Bersini [3–5] pays clear attention to the development of immune network models, and then applies these models to a control problem characterised by a discrete state vector in a state space. There are other examples of interdisciplinary work, such as the development of immune gene libraries and ultimately a bone marrow algorithm employed in AIS [32], and the development of the negative selection algorithm and the first application to computer security [21]. However, in more recent years, work on AIS has drifted away from the more biologically-appealing models and attention to biological detail, with the focus on more engineering-oriented approach. This has led to systems that are examples of *reasoning by metaphor* [56].

These include simple models of clonal selection, immune networks and negative selection algorithms as outlined above. For example, the clonal selection algorithm (CLONALG), whilst intuitively appealing, lacks any notion of interaction of B-cells with T-cells, MHC or cytokines. In addition, the large number of parameters associated with the algorithm, whilst well understood, make the algorithm less appealing from a computational perspective. aiNET, again, whilst somewhat affective, does not employ the immune network theory to a great extent. Only suppression between B-cells is employed, whereas in the immune network theory, there is suppression and stimulation between cells. With regards to negative selection, the simple random search strategy employed, combined with using a binary representation, makes the algorithm computationally so expensive that it is almost unusable in a real world setting [59].

However, in the past year or so, work by the Danger Team<sup>3</sup> has started to address this imbalance. For example, recent work by [26] has begun initial explorations into the use of *dendritic cells* (which are a type of cell found in the innate immune system), as a mechanism for identifying *dangerous* (or anomalous) events in a data stream. Whilst that work is still preliminary and works only on static data at the moment, there appears to be some promise, and may go some ways towards making a real breakthrough in the intrusion detection area of AIS research. Work linked to that is by [1] proposes an *artificial tissue* which is a type of representation of the data space that can

---

<sup>3</sup> <http://www.dangertheory.com>

evolve and adapt over time. Again, this is very preliminary work, but could prove useful bridge between the data and the immune algorithm itself.

Work in [56] proposes a conceptual framework that allows for the development of more biologically grounded AIS, through the adoption of an interdisciplinary approach. Metaphors employed have typically been simple, but somewhat effective. However, as proposed in [56], through greater interaction between computer scientists, engineers, biologists and mathematicians, better insights into the workings of the immune system, and the applicability (or otherwise) of the AIS paradigm will be gained. These interactions should be rooted in a sound methodology in order to fully exploit the synergy. Modeling techniques can be employed such as Cellular automata [53] and process calculi such as  $\pi$ -calculus [43], and stochastic  $\pi$ -calculus [50]. Here, in order to develop any AIS, it is necessary to first develop mathematical and/or computational models of the immune system (focussing on aspects that you are interested in as an AIS developer). This process necessitates the interaction of a number of disciplines, which may naturally lead to an increased synergy of AIS and immunology: something that is lacking. There is great potential for AIS to make a contribution to the world of immunology through this process. Of course, modelling immune systems is nothing new, in fact, AIS arose to some degree from this area itself. What is new here is that by following this process, first, questions may be asked on the immunology that would not necessarily be asked - computer scientists like to ask naive questions about immunology that can

cause immunologists to think in a slightly different way, and second, algorithms that have been developed through this process should, in theory at least, exploit more fully the underlying principles on which they have been developed. At present, one such project we are involved in is trying to adopt this approach <sup>4</sup>: time will tell to see if we can get it to work.

#### 4.2 Problem Oriented Perspective

*Challenge: Consider the Application of AIS.* Work to date in the realm of AIS has mainly concentrated on *what other paradigms do*, such as simple optimisation, learning and the like. This has happened as an accident of history and whilst productive, the time is here to look for the *killer application* of AIS, or, if not that radical, then applications where the benefit of adopting the immune approach is clear.

Work in [22] outlines the need to consider carefully the application domain when developing AIS. They review the role AIS have played in the development of a number of machine learning tasks, including that of classification. However, Freitas and Timmis point out that there is a lack of appreciation for possible inductive bias within algorithms and positional bias within the choice of representation and affinity measures. Indeed, this observation is reinforced by the work of [29, 28] with the development of their simple immune network simulator with various affinity metrics. They make the argument that seemingly generic AIS algorithms, are not so generic af-

---

<sup>4</sup> <http://www.cs.kent.ac.uk/projects/biasprofs/>

ter all, and each has to be tailored to specific application areas. This may be facilitated by the development of more theoretical aspects of AIS, which will help us to understand how, when and where to apply various AIS techniques.

It should be noted that there have been some previous attempts at providing *design principles* for immune systems, such as work by Segal et al. [52] and Bersini and Varela [5]. However, work by Segal, whilst extremely interesting, focussed primarily on network signaling and did not provide a comprehensive set of general design principles, or provide any test application areas for those principles. Work by Bersini, focussed on the immune network and *self assertion* ideas of the immune system to create his design principles and whilst being more concrete, are still quite high level:

- Principle 1: The control of any process is distributed around many operators in a network structure. This allows for the development of a self-organising system that can display emerging properties.
- Principle 2: The controller should maintain the viability of the process being controlled. This is keeping the system within certain limits and preventing the system from being driven in one particular way.
- Principle 3: While there may be perturbations that can affect the process, the controller learns to maintain the viability of the process through adaptation. This learning and adaptation requires two kinds of plasticity: a parametric plasticity, which keeps a constant population of operators in the process, but modifies parameters associated with them;

and a structural plasticity which is based on the recruitment mechanism which can modify the current population of operators.

- Principle 4: The learning and adaptation are achieved by using a reinforcement mechanism between operators. Operators interact to support common operations or controls.
- Principle 5: The dynamics and metadynamics of the system can be affected by the sensitivity of the population.
- Principle 6: The system retains a population-based memory, which can maintain a stable level in a changing environment.

These are potentially useful principles, that should be refined in light of immunological advances and possibly taken on board (to some degree) by the community. These need to be tested in various application areas, and refined to allow for the creation of not only a generic set of AIS design principles that are useful to the community, but also specific ones for specific application areas. With this may come a better understanding of how to apply AIS, and not fall into the traps highlighted by [22]. A recent paper by [31] highlight the fact that to date, the development of AIS has been *scattergun* i.e. many applications have been tried without a great deal of thought. Indeed, this paper provides a detailed overview of the many application areas that AIS have tried, and this will not be repeated here: the interested reader should consult that paper. The authors go on to propose a number of properties that they feel any AIS should have, and that these properties may help guide the type of application they could be applied to:

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

#### 4.3 Theoretical Aspects of AIS

*Challenge: To Develop a Theoretical basis for AIS.* Much work on AIS has concentrated on simple extraction of metaphors and direct application. Despite the creation of a framework for developing AIS, it still lacks significant formal and theoretical underpinning. AIS have been applied to a wide variety of problem domains, but a significant effort is still required to understand the nature of AIS and where they are best applied. For this, a more theoretical understanding is required.

There is very limited work on the more theoretical aspects of AIS. To our knowledge, only two works exist on any kind of formal proofs of AIS algorithms. In [63], the authors present a complete proof for a specific multi-objective clonal selection algorithm using markov chains [27]. Work by [9] has developed a Markov Chain model of the B-cell algorithm [40] showing a convergence proof (within infinite time), and also a mathematical model of the mutation operator, which allows for an assessment of the bias of the operator, which can prove to be exceptionally insightful. In a paper

by [33] a useful and valued avenue to explore would be into the dynamics of immune algorithms based on nonlinear dynamical systems inspired by biological models [18], and stochastic differential equations [7]. Given the use of clonal selection based algorithms within AIS, a great deal could be gained by the community with further theoretical investigations such as, the role of mutation operators, which could be used to provide information for optimal mutation rates for specific functions. There is some advancement in the theoretical aspects of negative selection, with a firmer understanding of the role of affinity measures such as *r-chunk* matching and the scalability issues relevant to certain shape-spaces employed with negative selection [17, 57, 59]. At a more conceptual level, work by [48] propose that all *population* based algorithms are essentially the same: immune or evolutionary. In that paper, the authors abstract to a high level to introduce various operators (such as selection and mutation), and demonstrate how these can be interchanged between various algorithms. All it lacks now is the theoretical underpinning.

#### 4.4 *The Immune System is not an Island*

*Challenge: To Consider the Integration of Immune and Other Systems.* The immune system does not work in isolation. Therefore, attention should not only be paid to the potential of the immune system as inspiration, but also other systems with which the immune system interacts, in particular the neural and endocrine systems. This will pave the way for a greater understanding of the role and function of the immune system and develop a new breed of immune inspired algorithms.

Homeostasis is the ability of an organism to achieve a steady state of internal body function in a varying environment [6]. This is achieved via complex interactions between a number of processes and systems within the organism, in particular the immune, neural and endocrine systems. We propose that by examining these systems, and their interactions, it should be possible to gain insight into how organisms achieve homeostasis, and therefore exploit these interactions in the realm of computer science and engineering. There is a large body of work investigating the interactions of these systems, but little work **specifically** in the area of AIS has paid any attention to this <sup>5</sup>. However, for a more detailed overview and for some initial ideas see [47].

Whilst the immune system is clearly an interesting system to investigate, if viewed in isolation, many key emergent properties arising from interac-

---

<sup>5</sup> This is not to say there does not exist a huge literature in the modeling community on this matter, it is just the AIS people have not seen it!

tions with other systems will be missed. Such systems do not operate in isolation in biology, therefore, consideration should be given to the interactions of the immune, neural and endocrine systems, and how, together, they allow for emergent properties to arise [54,45,53]. Immune, neural and endocrine cells express receptors for each other. This allows interaction and communication between cells and molecules in each direction. It appears that products from immune and neural systems can exist in lymphoid, endocrine and neural tissue at the same time. This indicates that there is a bi-directional link between the nervous system and the immune system. Therefore, it would seem that both endocrine and neural systems can affect the immune system. There is evidence to suggest that by stimulating areas of the brain it is possible to affect certain immune responses, and also that stress (which is regulated by the endocrine system) can suppress immune responses: this is also reciprocal in that immune cells can affect endocrine and neural systems. The action of various endocrine products on the neural system is accepted to be an important stimulus of a wide variety of behaviors. These range from behaviors such as flight and sexual activity to sleeping and eating [47].

Computationally then, what does this have to say to AIS? It should be possible to explore the role of interaction between these three systems. One interesting avenue would be to design an AIS to help select the types of components which will be most useful when added to a control system at any moment (differentiation) and to remove components when they are proving

harmful to the control system (apoptosis). The biological immune system cells select which action to perform by detecting properties of the cells and chemical environment through molecular interactions at membrane receptors. In an artificial system, similar properties can be detected by looking at activation states of artificial neurons and endocrine cells as well as global state information such as current consumption and battery levels. Thus the artificial immune system components can make similar decisions within the artificial context.

## 5 Conclusion

This paper has been an (incomplete) attempt at summarising the current state of research in the area of Artificial Immune Systems (AIS). By taking a holistic view, such as we have tried here, we hoped to gain a general impression of how well (or not) the area of AIS has been progressing. To be sure, there has been some significant developments over the years, and there is no doubt that the area has received a great deal of interest recently, and people have shown an interest in not only developing systems, but have started to think more carefully about why and how they both develop and apply these immune inspired ideas. We think a certain amount of reflection is good for any area, all too easily we can rush into developing novel, maybe even *cool* ideas, but we can get carried away! We have now had time to pause, think and reflect: now is the time to do something about it. It is clear, that it will not be possible to address these challenges in one day, or

even one year, but it is our hope, that this paper has set the seeds for new ideas and some exciting research in the future.

**Acknowledgments** This paper is a result of many useful interactions with a wide variety of people in the AIS community, and in particular during meetings that have taken place under the aegis of the EPSRC funded ARTIST<sup>6</sup> network in the UK. Particular thanks to: Mark Neal, Emma Hart, Susan Stepney, Andy Tyrrell, Andy Greenstead, Andy Hone, Hugo Van-den-Berg, Adrian Robins, Jamie Tycross, Al Lawson, Colin Johnson, Qi Chen and Paul Andrews for stimulating discussions. I would also like to thank the anonymous reviewers for their very helpful feedback and suggestions.

## References

1. P. Bentley, J. Greensmith, and S. Ujin. Two ways to grow artificial tissue. In C. Jacob, M. Pilat, P. Bentley, and J. Timmis, editors, *Proceedings of the 4th International Conference on Artificial Immune Systems*, volume 3627 of *LNCS*, pages 139–152. Springer, 2005.
2. C. Berek and M. Ziegner. The maturation of the immune response. *Immunology Today*, 14:200 – 402, 1993.
3. H. Bersini. Immune network and adaptive control. In *Proceedings of the 1st European Conference on Artificial Life (ECAL)*, pages 217–226. MIT Press, 1991.

---

<sup>6</sup> <http://www.artificial-immune-systems.org/artist.htm>

4. H. Bersini. Reinforcement and recruitment learning for adaptive process control. In *Proc. Int. Fuzzy Association Conference (IFAC/IFIP/IMACS) on Artificial Intelligence in Real Time Control*, pages 331–337, 1992.
5. H Bersini and F Varela. *The Immune Learning Mechanisms: Recruitment, Reinforcement and their Applications*. Chapman Hall, 1994.
6. H O Besendorfsky and A del Ray. Immune-Neuro-Endocrine Interactions: Facts and Hypotheses. *Nature*, 249:356–358, 1996.
7. Z Brzezniak and T Zastawniak. *Basic Stochastic Processes*. Springer, 1999.
8. F Burnet. *The Clonal Selection Theory of Acquired Immunity*. Cambridge University Press, 1959.
9. E. Clark, A. Hone, and J. Timmis. A markov chain model of the b-cell algorithm. In C. Jacob, M. Pilat, P. Bentley, and J. Timmis, editors, *Proceedings of the 4th International Conference on Artificial Immune Systems*, volume 3627 of *LNCS*, pages 318–330. Springer, 2005.
10. Irun R. Cohen. *Tending Adam's Garden: Evolving the Cognitive Immune Self*. Elsevier Academic Press, 2000.
11. D Cooke and J Hunt. Recognising Promoter Sequences using an Artificial Immune Systems. In *Proceedings of Intelligent Systems in Molecular Biology*, pages 89–97. AAAI Press, 1995.
12. V Cutello, G Nicosia, and M Parvone. Exploring the capability of immune algorithms: A characterisation of hypermutation operators. In *LNCS 3239*, pages 263–276. Springer, 2004.
13. L. N. de Castro and J. Timmis. *Artificial Immune Systems: A New Computational Intelligence Approach*. Springer, 2002.
14. L. N. de Castro and F. J. Von Zuben. *aiNet: An Artificial Immune Network for Data Analysis*, pages 231–259. Idea Group Publishing, USA, 2001.

15. L. N. de Castro and F. J. Von Zuben. Learning and optimization using the clonal selection principle. *IEEE Transactions on Evolutionary Computation*, 6(3):239–251, 2002.
16. L.N. de Castro and J. Timmis. Hierarchy and convergence of immune networks: Basic ideas and preliminary results. In J. Timmis and P. Bentley, editors, *Proceedings of the 1st International Conference on Artificial Immune Systems ICARIS*, pages 231–240, University of Kent at Canterbury, September 2002. University of Kent at Canterbury Printing Unit.
17. F. Esponda, S. Forrest, and P. Helman. A formal framework for positive and negative detection schemes. *IEEE Transactions on Systems, Man and Cybernetics Part B*, 34:357–373, 2004.
18. J. D. Farmer, N. H. Packard, and Perelson A. S. The immune system, adaptation, and machine learning. *Physica D*, 22:187–204, 1986.
19. S Forrest, S Hofmeyr, and A Somayaji. Computer Immunology. *Communications of the ACM*, 40(10):88–96, 1997.
20. S. Forrest, A. Perelson, L. Allen, and R. Cherukuri. Self–nonself discrimination in a computer. In *Proceedings of the IEEE Symposium on Research Security and Privacy*, pages 202–212, 1994.
21. S Forrest, A Perelson, L Allen, and R Cherukuri. Self-Nonsel Self Discrimination in a Computer. In *Proc. of the IEEE Symposium on Research in Security and Privacy*, pages 202–212, 1994.
22. A Freitas and J Timmis. Revisiting the Foundations of Artificial Immune Systems: A Problem Oriented Perspective. In *LNCS 2787*, pages 229–241. Springer, 2003.
23. S. Garrett. Parameter-free, adaptive clonal selection. In *Congress on Evolutionary Computing, CEC*. IEEE, 2004.

24. S. Garrett. How do we evaluate artificial immune systems? *Evolutionary Computation*, 13(2):145–177, 2005.
25. S. M. Garrett. A paratope is not an epitope: Implications for clonal selection and immune networks. In *Proceedings of the 2nd International Conference on Artificial Immune Systems*, volume 2787, 2003.
26. J. Greensmith, U. Aickelin, and S. Cayzer. Introducing Dendritic Cells as a Novel Immune-Inspired Algorithm for Anomaly Detection. In C Jacob, M Pilat, P Bentley, and J Timmis, editors, *Proceedings of the 4th International Conference on Artificial Immune Systems*, volume 3627, 2005.
27. G R Grimmett and D R Stirzaker. *Probability and Random Processes*. Oxford University Press, 1982.
28. E. Hart. Not all balls are round. In C. Jacob, M. Pilat, P. Bentley, and J. Timmis, editors, *Proceedings of the 4th International Conference on Artificial Immune Systems*, volume 3627 of *LNCS*. Springer, 2005.
29. E. Hart and P. Ross. Studies on the implications of shape-space models for idiotypic networks. In *Proceedings of the 3rd international Conference on Artificial Immune Systems (ICARIS 2004)*, *LNCS 3239*, pages 413–426. Springer, 2004.
30. E. Hart and J. Timmis. Application areas of ais: The past, the present and the future. *Applied Soft Computing*, In Review., 2005.
31. E. Hart and J. Timmis. Application areas of ais: The past, the present and the future. In C. Jacob, M. Pilat, P. Bentley, and J. Timmis, editors, *Proceedings of the 4th International Conference on Artificial Immune Systems (ICARIS 2005)*, *LNCS 3627*, volume 3627 of *LNCS*, pages 126–138. Springer, 2005.
32. R. R. Hightower, S. A. Forrest, and A. S. Perelson. The evolution of emergent organization in immune system gene libraries. In L. J. Eshelman, editor, *Proc.*

- 6th Int. Conf. Genetic Algorithms*, pages 344–350. Morgan Kaufmann, 1995.
33. A Hone and J Kelsey. Optima, extrema and artificial immune systems. In *Lecture Notes in Computer Science*, pages 89–98. Springer, September 2004.
34. J Hunt and D Cooke. Learning using an artificial immune system. *Journal of Network and Computer Applications*, 19:189–212, 1996.
35. J Hunt, J. Timmis, D Cooke, M Neal, and C King. *Artificial Immune Systems and their Applications*, chapter JISYS: Development of an Artificial Immune System for Real World Applications, pages 157–186. Springer, 1998.
36. Y Ishida. Active diagnosis by self-organisation : An approach by the immune network metaphor. In *Proceedings of the International Joint Conference on Artificial Intelligence*, pages 1084–1089, Nagoya, Japan., 1997. IEEE.
37. C. A. Janeway Jr. and P. Travers. *Immunobiology: The Immune System in Health and Disease*. Garland Publishing, New York, third edition, 1997.
38. C.A. Janeway Jr and R. Medzhitov. Innate immune recognition. *Ann. Rev. Immunol.*, 20:197–216, 2002.
39. N. K. Jerne. Towards a network theory of the immune system. *Ann. Immunol. (Inst. Pasteur)*, 125C:373–389, 1974.
40. J. Kelsey and J. Timmis. Immune inspired somatic contiguous hypermutation for function optimisation. In *Genetic and Evolutionary Computation Conference - GECCO 2003*, volume LNCS 2723, pages 207–218. Springer, 2003.
41. J. Kim and Peter J. Bentley. Immune memory in the dynamic clonal selection algorithm. In J. Timmis and P. Bentley, editors, *Proceedings of the 1st International Conference on Artificial Immune Systems ICARIS*, pages 59–67, University of Kent at Canterbury, September 2002. University of Kent at Canterbury Printing Unit.

42. J. Kim and Peter J. Bentley. A model of gene library evolution in the dynamic clonal selection algorithm. In J. Timmis and P. Bentley, editors, *Proceedings of the 1st International Conference on Artificial Immune Systems ICARIS*, pages 182–189, University of Kent at Canterbury, September 2002. University of Kent at Canterbury Printing Unit.
43. C. Kuttler, J. Niehren, and R. Blossey. Gene regulation in the pi calculus: Simulating cooperativity at the lambda switch. In *Proc Concurrent Models in Molecular Biology (Bioconcur 04)*, 2004.
44. R. E. Langman and M. Cohn. The ‘complete’ idiotypic network is an absurd immune system. *Immunology Today*, 7(4):100–101, 1986.
45. B.Z. Liu and G.M. Deng. An improved mathematical model of hormone secretion in the hypothalamo-pituitary-gonadal axis in man. *J Theor Biol*, 150:51–8, 1991.
46. M. Neal. An artificial immune system for continuous analysis of time-varying data. In Jonathan Timmis and Peter J. Bentley, editors, *Proceedings of the 1st International Conference on Artificial Immune Systems ICARIS*, pages 76–85, University of Kent at Canterbury, September 2002. University of Kent at Canterbury Printing Unit.
47. M Neal and J Timmis. *Recent Advances in Biologically Inspired Computing*, chapter : Once more unto the breach... towards artificial homeostasis? IGP, 2004.
48. R. Newborough and S. Stepney. A generic framework for population based algorithms. In C. Jacob, M. Pilat, P. Bentley, and J. Timmis, editors, *Proceedings of the 4th International Conference on Artificial Immune Systems*, volume 3267 of *LNCS*. Springer, 2005.

49. A. S. Perelson. Immune network theory. *Immunological Review*, 110:5–36, 1989.
50. A. Phillips and L. Cardelli. A correct abstract machine for the stochastic pi-calculus. In *Proc Concurrent Models in Molecular Biology (Bioconcur'04)*, ENTCS, 2004.
51. A. Secker, A. Freitas, and J. Timmis. AISEC an artificial immune system for email classification. In *Proceedings of the Congress on Evolutionary Computation*, pages 131–139, 2003.
52. L Segal and I Cohen, editors. *Design Principles for the Immune System and Other Distributed Systems*. Oxford University Press, 2001.
53. Hans B. Sieburg and Oliver K. Clay. The cellular device machine development system for modeling biology on the computer. *Complex Systems*, 5:575–601, 1991.
54. W.R. Smith. Qualitative mathematical models of endocrine systems. *Am. Journal of Physiology*, 245:473–477, 1983.
55. S. Stepney, J.A. Clark, A. Tyrrell, C.G. Johnson, J. Timmis, D. Partridge, A. Adamatsky, and R.E. Smith. Journeys in non-classical computation: A grand challenge for computing research. Grand Challenge Report 7, National E-Science Centre, University of Edinburgh, May 2003.
56. S. Stepney, R. Smith, J. Timmis, A. Tyrrell, M. Neal, and A. Hone. Conceptual frameworks for artificial immune systems. *International Journal of Unconventional Computing*, 1(3):315–338, 2005.
57. T. Stibor, K. M. Bayarou, and C. Eckert. An investigation of r-chunk detector generation on higher alphabets. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO 2004)*, pages 299–307. Springer, 2004.

58. T. Stibor, P. Mohr, J. Timmis, and C. Eckert. Is negative selection appropriate for anomaly detection? In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO 2005)*, pages 321–328. Springer, 2005.
59. T. Stibor, J. Timmis, and C. Eckert. A comparative study of real-valued negative selection to statistical anomaly detection techniques. In C. Jacob, M. Pilat, P. Bentley, and J. Timmis, editors, *Proceedings of the 4th International Conference on Artificial Immune Systems*, volume 3627 of *LNCS*, pages 262–275, 2005.
60. J. Timmis, R. de Lemos, M. Ayara, and R. Duncan. Towards immune inspired fault tolerance in embedded systems. In L Wang, J Rajapakse, K Fukushima, S Lee, and X Yao, editors, *Proceedings of 9th International Conference on Neural Information Processing*, pages 1459–1463. IEEE, November 2002.
61. J. Timmis and M. Neal. A Resource Limited Artificial Immune System. *Knowledge Based Systems*, 14(3/4):121–130, 2001.
62. J Timmis, M Neal, and J Hunt. An artificial immune system for data analysis. *Biosystems*, 55(1/3):143–150, unknown 2000.
63. M Villalobos-Arias, C A Coello Coello, and O Hernandez-Lerma. Convergence analysis of a multiobjective artificial immune system algorithm. In *Lecture Notes in Computer Science 3239*, pages 226 – 235, 2004.
64. A Watkins. *An Artificial Immune Recognition System*. Mississippi State University. MSc Thesis., 2001.
65. A. Watkins and J. Timmis. Exploiting Parallelism Inherent in AIRS, an Artificial Immune Classifier. In G. et al Nicosia, editor, *Third International Conference on Artificial Immune Systems*, number 3239 in *LNCS*, pages 427–438. Springer, September 2004.

66. A. Watkins, B. Xintong, and A. Phadke. Parallelizing an immune-inspired algorithm for efficient pattern recognition. In *Intelligent Engineering Systems through Artificial Neural Networks: Smart Engineering System Design: Neural Networks, Fuzzy Logic, Evolutionary Programming, Complex Systems and Artificial Life*, pages 224–230. ASME Press, 2003.
67. S. Wierzchon and U. Kuzelewska. Stable clusters formation in an artificial immune system. In Jonathan Timmis and Peter J. Bentley, editors, *Proceedings of the 1st International Conference on Artificial Immune Systems ICARIS*, pages 68–75, University of Kent at Canterbury, September 2002. University of Kent at Canterbury Printing Unit.