Nat Comput (2007) 6:1–18
DOI 10.1007/s11047-006-9029-1

ORIGINAL PAPER

Artificial immune systems—today and tomorrow

Jon Timmis

Received: 6 December 2005 / Accepted: 8 May 2006 /
Published online: 8 December 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.

Keywords Artificial immune system · Future directions · Position paper

Introduction

The UK research community has proposed a number of Grand Challenges for Computer Science research[1] and ambitious plans for the development of a variety of research areas. Grand Challenge 7 (GC-7)[2] (Stepney et al. 2003) 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 (AIS). 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 de Castro and Timmis (2002)). We then

[1] http://www.nesc.ac.uk/esi/events/Grand_Challenges/

[2] http://www.cs.york.ac.uk/nature/gc7/index.htm

J. Timmis (✉)
Departments of Computer Science and Electronics,
University of York, Heslington, York YO10 5DD, UK
e-mail: jtimmis@cs.york.ac.uk

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 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.

Artificial immune systems

In an attempt to create a common basis for AIS, work in de Castro and Timmis (2002) 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, (de Castro and Timmis 2002) 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 EAs 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 (de Castro and Timmis 2002).''

The framework for AIS can be thought of as a layered approach (Fig.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 representation 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* (Perelson 1989). There are

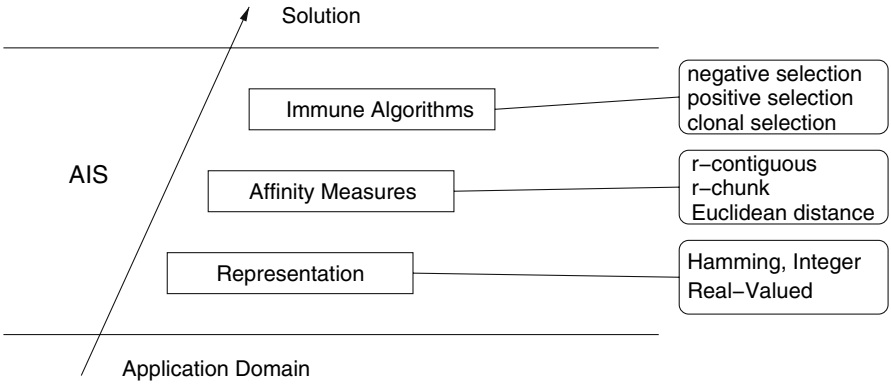

Fig. 1 AIS layered framework adapted from de Castro and Timmis (2002)

many kinds of shape space, such as Hamming, Real valued and so on, each of which carries it own bias and should be selected with care (Freitas and Timmis 2003). 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 (Freitas and Timmis 2003). 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 (2005) 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 (Hart and Timmis 2005).

The immune system: from an AIS perspective

The vast majority of developments within AIS focussed on three main immunological theories: clonal selection (Burnet 1959), immune networks (Jerne 1974) 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 Stepney et al. (2005) 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 de Castro and Timmis (2002).

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 (Cohen 2000). 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 (Janeway and Medzhitov 2002). 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 (Janeway and Travers 1997). 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.

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 (Berek and Ziegner 1993) and produce antibodies that are specific to the invading antigen. The clonal selection principle (Burnet 1959) 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 optimization approaches, such as the work on CLONALG by de Castro and Von Zuben (2002), the work by Cutello et al. (2004) and the work by Garrett (2004) on parameter free clonal selection. Clonal selection has also formed the basis of learning algorithms (mainly supervised) such as AIRS by Watkins (2001) and a parallel and a distributed version of AIRS (Watkins and Timmis 2004), a distributed version of CLONALG by Watkins et al. (2003), the work with DynamicCS applied to computer security by Kim and Bentley (2002), work by Secker et al. (2003) 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 optimization, 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 optimized). 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 Kim and Bentley (2002), Secker et al. (2003)) these solutions can be maintained over long periods of time. However, as argued by Newborough and Stepney (2005), at a certain level are such population-based algorithms like clonal selection ones that can be considered the same as genetic algorithms and other EAs. Figure 2 outlines the basic process in a clonal selection algorithm (CLONALG).

Immune networks

In a landmark paper for the time Jerne (1974) 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 to the idiotopes on similar B-cells, as it

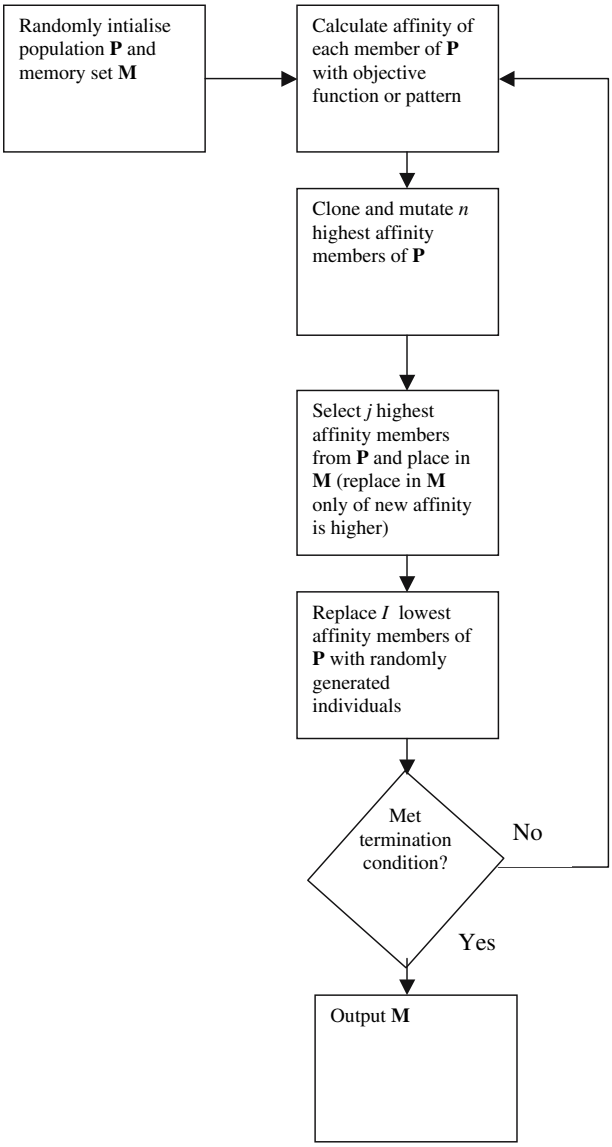

Fig. 2 Flowchart of a typical clonal selection algorithm

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 formalized in successive works by Farmer et al. (1986), Perelson (1989) and Bersini and Varela (1994). It is worth noting, however, that the immune network theory has not found a great deal of favor with other immunologists, being labeled as strongly as *absurd* by some (Langman and Cohn 1986). 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 Ishida (1997) on the application of immune network to diagnostic problems, work by Cooke and Hunt (1995), Hunt and Cooke (1996), Hunt et al. (1998) on the development of an immune network approach to DNA classification and mortgage fraud detection, and work on aiNET by de Castro and Von Zuben (2001), de Castro and Timmis (2002) all attempt to extract properties such as adaptation, self-organization and plasticity from the immune network theory. The work by Hunt et al. (1998) proved a useful basis for the refining and extending of their work into an unsupervised system by Timmis et al. (2000), Timmis and Neal (2001), which itself was then adopted for continuous learning by Neal (2002), Wierzchon and Kuzelewska (2002). There is little to unify these immune network algorithms, but an abstracted outline of one is presented in Fig. 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 Garrett (2003) where the author proposed a general artificial immune network, however, this was mainly conceptual and was never 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 (Hart and Ross 2004).

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, their 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 Stibor et al. (2005) 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 Forrest et al. (1994, 1997) to propose and develop a negative selection algorithm to detect data manipulation 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 **S** of elements of length l in shape-space. Then generate a set **D** of detectors, such that each fails to match any element in **S**. With these detectors, monitor a continual data stream for any changes, by continually matching

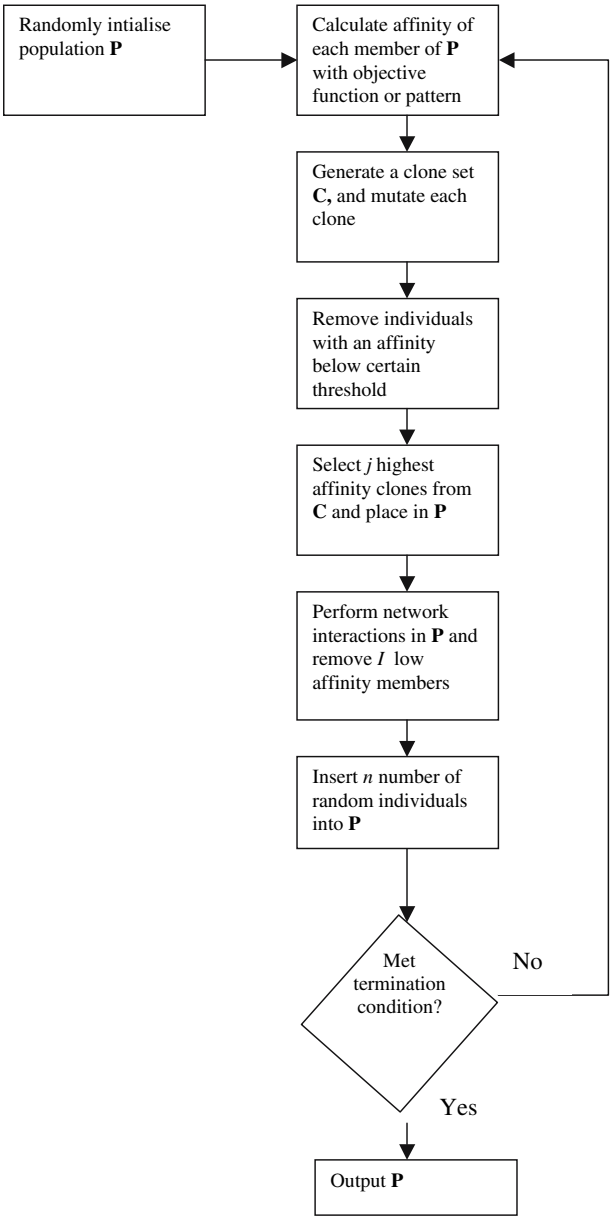

Fig. 3 Flowchart of a typical immune network algorithm

the detectors in **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 (Forrest et al. 1997), and more recently (Esponda et al. 2004). Figure 4 outlines the basic negative selection approach. However, building on the work of Esponda et al. (2004), work by Stibor et al. (2004), Stibor et al. (2005) begins to show (from a theoretical and practical perspective) potential problems of

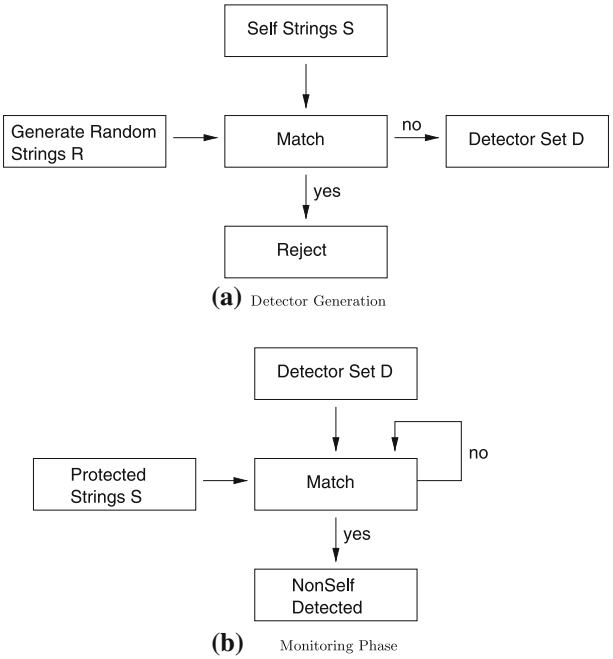

Fig. 4 Negative selection algorithm by Forrest et al. (1994)

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 problematic. 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.

Building AIS

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 of detectors (members of a population) that are required for the job (Timmis et al. 2002). For example, there are a number of works that outline the unacceptable computational complexity of the negative selection approach in AIS (Kim and Bentley 2002; Stibor et al. 2004, 2005) 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 immunized. 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 (Timmis et al. 2002). For example, recent studies by Hart and Ross (2004), Hart (2005) 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.

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.

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 (1991, 1992); Bersini and Varela (1994) pays clear attention to the development of immune network models, and then applies these models to a control problem characterized 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 (Hightower et al. 1995), and the development of the negative selection algorithm and the first application to computer security (Forrest et al. 1994). However, in more recent years, work on AIS has drifted away from the more biologically appealing models and attention to biological detail, with the focus on more engineering-oriented approach. This has led to systems that are examples of *reasoning by metaphor* (Stepney et al. 2005). These include simple models of clonal selection, immune networks and negative selection algorithms as outlined above. For example, the 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 (Stibor et al. 2005).

However, in the past year or so, work by the Danger Team[3] has started to address this imbalance. For example, recent work by Greensmith et al. (2005) 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 Bentley et al. (2005) proposes an *artificial tissue* which is a type of representation of the data space that can 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 Stepney et al. (2005) 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 Stepney et al. (2005), 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 (Sieburg and Clay 1991) and process calculi such as π-calculus (Kuttler et al. 2004), and stochastic π-calculus (Phillips and Cardelli 2004). 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, modeling 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[4]: time will tell to see if we can get it to work.

[3] http://www.dangertheory.com

[4] http://www.cs.kent.ac.uk/projects/biasprofs/

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 optimization, 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 Freitas and Timmis (2003) 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 Hart and Ross (2004), Hart (2005) 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 after 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 and Cohen (2001) and Bersini and Varela (1994). 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-organizing 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 Freitas and Timmis (2003). A recent paper by Hart and Timmis (2005) 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 (Hart and Timmis 2005)."

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 Villalobos-Arias et al. (2004), the authors present a complete proof for a specific multi-objective CLONALG using markov chains (Grimmett and Stirzaker 1982). Work by Clark et al. (2005) has developed a Markov Chain model of the B-cell algorithm (Kelsey and Timmis 2003) 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. According to a paper by Hone and Kelsey (2004) a useful and valued avenue to explore would be into the dynamics of immune algorithms based on nonlinear dynamical systems inspired by biological models (Farmer et al. 1986), and stochastic differential equations (Brzezniak and Zastawniak 1999). 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 (Esponda et al. 2004; Stibor et al. 2004, 2005). At a more

conceptual level, work by Newborough and Stepney (2005) 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.

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 (Besendovsky and del Ray 1996). 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[5]. However, for a more detailed overview and for some initial ideas see Neal and Timmis (2004).

Whilst the immune system is clearly an interesting system to investigate, if viewed in isolation, many key emergent properties arising from interactions with other systems will be missed. Such systems do not operate in isolation in biology, therefore, consideration should to be given to the interactions of the immune, neural and endocrine systems, and how, together, they allow for emergent properties to arise (Smith 1983; Liu and Deng 1991; Sieburg and Clay 1991). 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 (Neal and Timmis 2004).

[5] 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!

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 an artificial 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 fuel consumption and battery levels. Thus the artificial immune system components can make similar decisions within the artificial context.

Conclusion

This paper has been an (incomplete) attempt at summarizing the current state of research in the area of 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[6] 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

Bentley P, Greensmith J, Ujin S (2005) Two ways to grow artificial tissue. In: Jacob C, Pilat M, Bentley P, Timmis J (eds) Proceedings of the 4th international conference on artificial immune systems, vol 3627 of LNCS. Springer, pp 139–152

Berek C, Ziegner M (1993) The maturation of the immune response. Immunol Today 14:200–402

Bersini H (1991) Immune network and adaptive control. In: Proceedings of the 1st European conference on artificial life (ECAL), MIT Press, pp 217–226

Bersini H (1992) Reinforcement and recruitment learning for adaptive process control. In: Proc int fuzzy Association Conference (IFAC/IFIP/IMACS) on artificial intelligence in real time control, pp 331–337

[6] http://www.artificial-immune-systems.org/artist.htm

Bersini H, Varela F (1994) The immune learning mechansims: recruitment, reinforcement and their applications. Chapman Hall
Besendovsky HO, del Ray A (1996) Immune-neuro-endocrine interactions: facts and hypotheses. Nature 249:356–358
Brzezniak Z, Zastawniak T (1999) Basic stochastic processes. Springer
Burnet F (1959) The clonal selection theory of acquired immunity. Cambridge University Press, Cambridge
Clark E, Hone A, Timmis J (2005) A markov chain model of the b-cell algorithm. In: Jacob C, Pilat M, Bentley P, Timmis J (eds) Proceedings of the 4th international conference on artificial immune systems, vol 3627 of LNCS. Springer, pp 318–330
Cohen IR (2000) Tending Adam's Garden: evolving the cognitive immune self. Elsevier Academic Press
Cooke D, Hunt J (1995) Recognising promoter sequences using an artificial immune systems. In: Proceedings of intelligent systems in molecular biology. AAAI Press, pp 89–97
Cutello V, Nicosia G, Parvone M (2004) Exploring the capability of immune algorithms: a characterization of hypermutation operators. In: LNCS vol 3239, Springer, pp 263–276
de Castro LN, Timmis J (2002) Artificial immune systems: a new computational intelligence approach. Springer
de Castro LN, Von Zuben FJ (2001) aiNet: an artificial immune network for data analysis. Idea Group Publishing, USA, pp 231–259
de Castro LN, Von Zuben FJ (2002) Learning and optimization using the clonal selection principle. IEEE Trans Evolut Comput 6(3):239–251
de Castro LN, Timmis J (2002) Hierarchy and convergence of immune networks: Basic ideas and preliminary results. In: Timmis J, Bentley P (eds) Proceedings of the 1st international conference on artificial immune systems ICARIS, University of Kent at Canterbury, University of Kent at Canterbury Printing Unit, pp 231–240
Esponda F, Forrest S, Helman P (2004) A formal framework for positive and negative detection schemes. IEEE Trans Systems Man Cybernet Part B 34:357–373
Farmer JD, Packard NH, Perelson AS (1986) The immune system, adaptation, and machine learning. Physica D 22:187–204
Forrest S, Hofmeyr S, Somayaji A (1997) Computer immunology. Commun ACM 40(10):88–96
Forrest S, Perelson A, Allen L, Cherukuri R (1994) Self-nonself discrimination in a computer. In: Proceedings of the IEEE symposium on research security and privacy, pp 202–212
Freitas A, Timmis J (2003) Revisiting the foundations of artificial immune systems: a problem oriented perspective. In: LNCS, vol 2787, Springer, pp 229–241
Garrett SM (2003) A paratope is not an epitope: implications for clonal selection and immune networks. In: Proceedings of the 2nd international conference on artificial immune systems, vol 2787
Garrett S (2004) Parameter-free, adaptive clonal selection. In Congress on evolutionary computing, CEC. IEEE
Garrett S (2005) How do we evaluate artificial immune systems? Evol Comput 13(2):145–177
Greensmith J, Aickelin U, Cayzer S (2005) Introducing dendritic cells as a novel immune-inspired algorithm for anomaly detection. In: Jacob C, Pilat M, Bentley P, Timmis J (eds) Proceedings of the 4th international conference on artificial immune systems, vol 3627
Grimmett GR, Stirzaker DR (1982) Probability and random processes. Oxford University Press, Oxford
Hart E (2005) Not all balls are round. In: Jacob C, Pilat M, Bentley P, Timmis J (eds) Proceedings of the 4th international conference on artificial immune systems, vol 3627 of LNCS. Springer
Hart E, Ross P (2004) 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, Springer, pp 413–426
Hart E, Timmis J (2005) Application areas of ais: the past, the present and the future. Applied Soft Computing, In Review
Hart E, Timmis J (2005) Application areas of ais: the past, the present and the future. In: Jacob C, Pilat M, Bentley P, Timmis J (eds) Proceedings of the 4th international conference on artificial immune systems (ICARIS 2005), LNCS 3627, vol 3627 of LNCS. Springer, pp 126–138
Hightower RR, Forrest SA, Perelson AS (1995) The evolution of emergent organization in immune system gene libraries. In: Eshelman LJ (ed) Proceedings of the 6th international conference on genetic algorithms. Morgan Kaufmann, pp 344–350

Hone A, Kelsey J (2004) Optima, extrema and artificial immune systems. In: Lecture notes in computer science. Springer, pp 89–98
Hunt J, Cooke D (1996) Learning using an artificial immune system. J Network Comput Appl 19:189–212
Hunt J, Timmis J, Cooke D, Neal M, King C (1998) Artificial immune systems and their applications, chapter JISYS: development of an artificial immune system for real world applications. Springer, pp 157–186
Ishida Y (1997) Active diagnosis by self-organisation: an approach by the immune network metaphor. In: Proceedings of the international joint conference on artificial intelligence, IEEE, Nagoya, Japan, pp 1084–1089
Janeway CA Jr, Travers P (1997) Immunobiology: the immune system in health and disease, 3rd edn. Garland Publishing, New York
Janeway CA Jr, Medzhitov R (2002) Innate immune recognition. Ann Rev Immunol 20:197–216
Jerne NK (1974) Towards a network theory of the immune system. Ann Immunol (Inst Pasteur) 125C:373–389
Kelsey J, Timmis J (2003) Immune inspired somatic contiguous hypermutation for function optimisation. In: Genetic and evolutionary computation conference – GECCO 2003, vol LNCS 2723. Springer, pp 207–218
Kim J, Bentley PJ (2002) Immune memory in the dynamic clonal selection algorithm. In: Timmis J, Bentley P (eds) Proceedings of the 1st international conference on artificial immune systems ICARIS, University of Kent at Canterbury, University of Kent at Canterbury Printing Unit, pp 59–67
Kim J, Bentley PJ (2002) A model of gene library evolution in the dynamic clonal selection algorithm. In: J Timmis, Bentley P (eds) Proceedings of the 1st international conferece on artificial immune systems ICARIS, University of Kent at Canterbury, University of Kent at Canterbury Printing Unit, pp 182–189
Kuttler C, Niehren J, Blossey R (2004) Gene regulation in the pi calculus: Simulating cooperativity at the lambda switch. In: Proc concurrent models in molecular biology (Bioconcur 04)
Langman RE, Cohn M (1986) The complete idiotype network is an absurd immune system. Immunol Today 7(4):100–101
Liu BZ, Deng GM (1991) An improved mathematical model of hormone secretion in the hypothalamo-pituitary-gonadal axis in man. J Theor Biol 150:51–58
Neal M (2002) An artificial immune system for continuous analysis of time-varying data. In: Jonathan T, Bentley PJ (eds) Proceedings of the 1st international conference on artificial immune systems ICARIS, University of Kent at Canterbury, University of Kent at Canterbury Printing Unit, pp 76–85
Neal M, Timmis J (2004) Recent advances in biologically inspired computing, chapter: once more unto the breach... towards artificial homeostasis? IGP
Newborough R, Stepney S (2005) A generic framework for population based algorithms. In: Jacob C, Pilat M, Bentley P, Timmis J (eds) Proceedings of the 4th international conference on artificial immune systems, vol 3267 of LNCS. Springer
Perelson AS (1989) Immune network theory. Immunol Rev 110:5–36
Phillips A, Cardelli L (2004) A correct abstract machine for the stochastic pi-calculus. In: Proc concurrent models in molecular biology (Bioconcur'04). ENTCS
Secker A, Freitas A, Timmis J (2003) AISEC an artificial immune system for email classification. In: Proceedings of the congress on evolutionary computation, pp 131–139
Segal L, Cohen I (eds) (2001) Design principles for the immune system and other distributed systems. Oxford University Press, Oxford
Sieburg HB, Clay OK (1991) The cellular device machine development system for modeling biology on the computer. Complex Syst 5:575–601
Smith WR (1983) Qualitative mathematical models of endocrine systems. Am J Physiol 245:473–477
Stepney S, Clark JA, Tyrrell A, Johnson CG, Timmis J, Partridge D, Adamatsky A, Smith RE (2003) Journeys in non-classical computation: a grand challenge for computing research. Grand Challenge Report 7, National E-Science Centre, University of Edinburgh
Stepney S, Smith R, Timmis J, Tyrrell A, Neal M, Hone A (2005) Conceptual frameworks for artificial immune systems. Int J Unconvent Comput 1(3):315–338
Stibor T, Bayarou KM, Eckert C (2004) An investigation of r-chunk detector generation on higher alphabets. In: Proceedings of the genetic and evolutionary computation conference (GECCO 2004). Springer, pp 299–307

Stibor T, Mohr P, Timmis J, Eckert C (2005) Is negative selection appropriate for anomaly detection? In Proceedings of the genetic and evolutionary computation conference (GECCO 2005). Springer, pp 321–328

Stibor T, Timmis J, Eckert C (2005) A comparative study of real-valued negative selection to statistical anomaly detection techniques. In: Jacob C, Pilat M, Bentley P, Timmis J (eds) Proceedings of the 4th international conference on artificial immune systems, vol 3627 of LNCS, pp 262–275

Timmis J, de Lemos R, Ayara M, Duncan R (2002) Towards immune inspired fault tolerance in embedded systems. In: Wang L, Rajapakse J, Fukushima K, Lee S, Yao X (eds) Proceedings of 9th international conference on neural information processing. IEEE, pp 1459–1463

Timmis J, Neal M (2001) A resource limited artificial immune system. Knowl Based Syst 14(3/4):121–130

Timmis J, Neal M, Hunt J (2000) An artificial immune system for data analysis. Biosystems 55(1/3):143–150, unknown

Villalobos-Arias M, Coello Coello CA, Hernandez-Lerma O (2004) Convergence analysis of a multiobjective artificial immune system algorithm. In: Lecture notes in computer science, vol 3239, pp 226–235

Watkins A (2001) An artificial immune recognition system. Mississippi State University. MSc Thesis.

Watkins A, Timmis J (2004) Exploiting parallelism inherent in AIRS, an artificial immune classifier. In: Nicosia G et al (eds) Third international conference on artificial immune systems, vol 3239 in LNCS, Springer, pp 427–438

Watkins A, Xintong B, Phadke A (2003) Parallelizing an immune-inspired algorithm for efficient pattern recognition. In: Intelligent engineering systems through ANN: smart engineering system design: neural networks, fuzzy logic, evolutionary programming, complex systems and artificial life. ASME Press, pp 224–230

Wierzchon S, Kuzelewska U (2002) Stable clusters formation in an artificial immune system. In: Timms J, Bentley PJ (eds) Proceedings of the 1st international conference on artificial immune systems ICARIS, University of Kent at Canterbury, University of Kent at Canterbury Printing Unit, pp 68–75