Application areas of AIS: The Past, The Present and the Future

Emma Hart [a] Jon Timmis [b]

[a] *School of Computing,Napier University*

[b] *Department of Electronics and Department of Computer Science, University of York*

Abstract

After a decade of research into the area of Artificial Immune Systems, it is worthwhile to take a step back and reflect on the contributions that the paradigm has brought to the application areas to which it has been applied. Undeniably, there have been a lot of successful stories — however, if the field is to advance in the future and really carve out its own distinctive niche, then it is necessary to be able to illustrate that there are clear benefits to be obtained by applying this paradigm rather than others. This paper attempts to take stock of the application areas that have been tackled in the past, and ask the difficult question "was it worth it ?". We then attempt to suggest a set of problem features that we believe will allow the true potential of the immunological system to be exploited in computational systems, and define a unique niche for AIS.

Key words: Artificial Immune Systems,, applications

1 Introduction

The Artificial Immune System (AIS) community has been vibrant and active for a number of years now, producing a prolific amount of research ranging from modelling the natural immune system, solving artificial or bench-mark problems, to tackling real-world applications, using an equally diverse set of immune-inspired algorithms. Whilst it is natural, and indeed healthy, for a somewhat scattergun approach to be taken in the early days of developing any new paradigm, in the sense that high-level, often naïve metaphors are

Email addresses: `e.hart@napier.ac.uk` (Emma Hart),
`jtimmis@cs.york.ac.uk` (Jon Timmis).

selected and applied to problem areas that have often been tackled with other paradigms, there comes a point at which research effort needs to have a more coherent focus in order to more clearly define the field, and allow it to go forward and be fully exploited. We argue that this point has now been reached in the AIS world — with a solid foundation of published work to build on, the time has come to try and define the role that AIS can play and the type of applications that will really allow its potential to be realised.

Without a doubt there have been a lot of successful applications of AIS, and these should not be ignored. However, at this point, there are still few exemplars that really stand out as instances of successfully applying an AIS to a hard, real-world problems, or of AIS being used in earnest in industry. This is in contrast for example to the field of Evolutionary Algorithms, where at the most recent flagship conference in the field, GECCO 2004 [4], there were 38 papers describing the applications of EAs to real-world problems, and the EVONET repository [2] is able to list 39 examples of *Evolution at Work*, i.e practical applications of EAs. On the one hand, this is somewhat of an unfair comparison, given the relative time-periods that the two fields have been active, however it illustrates the importance of focusing research effort in the next few years in order to provide hard evidence of a distinctive niche for AIS.

For any new paradigm to prove itself is always a difficult task — there is a lot of good competition from existing tried and tested algorithms. There has perhaps been a natural tendency for AIS to be compared to other biologically inspired paradigms such as Evolutionary Algorithms, Neural-networks, and to other more traditional classification or clustering algorithms. Scientifically, it is essential that such comparisons to be made; however, we argue that it is not sufficient for AIS simply to outperform other algorithms on any given set of problem instances to be declared useful. For a start, test instances (particularly benchmarks) are not necessarily difficult, and any number of other problem instances can be generated on which performance will be unknown. Secondly, in the light of the no-free lunch theorem [81], we cannot expect any one algorithm to outperform all others given all possible problem instances. We argue that for a paradigm to be truly successful, it should contain features that *are not present* in other paradigms *and/or* it should give rise to emergent system properties which are not obtained through other paradigms. Thus, the algorithm can be described as distinctive. This is in contrast to work in [33] which suggests that an appropriate manner in which to assess AIS algorithms is by their *distinctiveness* and *effectiveness*. [33] define distinctiveness in terms of an algorithm containing novel symbols, expressions or processes. We counter that this definition is not strong enough; some of the essential characteristics of immune algorithms are *emergent* properties that arise from the interaction of a number of processes which in themselves may not be described as distinct. In addition, whilst it is important that an algorithm is *effective*, we disagree with the definition proposed in [33] which measures effectiveness in terms of an algorithm performing *better* than other algorithms on shared benchmarks tests, *quicker* than other algorithms or providing a *unique* method of obtaining results. The problems associated with benchmarking data-sets have already been outlined and hence we do not believe that only being *better* or *quicker* will allow AIS to prove itself. Furthermore, in a climate in which there are any number of algorithms (biological and otherwise) inspired by an equally diverse range of processes, it seems critical that an algorithm *must* achieve something that cannot be achieved by any other means in order to earn its place in history.

In this position paper, we hope to extract some general features of problems that we believe will allow AIS to really bring some benefit, and thus distinguish it from other techniques. We suggest that the way forward for AIS is in part a focussed attempt to carefully select application areas based on mapping problem features to mechanisms exhibited by the IS, taking the problem-oriented perspective outlined by example in [69,31,13], and discussed further in section 4.2. However, we emphasise that application development needs to be under-pinned with a continuing line of research into the theoretical basis of AIS and with the overriding need for extraction of novel and accurate metaphors from immunology [1].

2 Survey of Existing Application Areas

In order to place the following discussions in context, we first present a general review of application areas to which AIS has currently been applied. The following brief summary is based in part on a bibliography produced by De Castro [22], used in a tutorial at ICARIS 2004 [23] on Engineering Applications of AIS. The information contained in this tutorial has been expanded to include references from ICARIS 2004 [3] and ICARIS 2005 [48]. A useful summary of application areas can also be found in [24] though as this was produced in 2000 it is slightly outdated. Whilst we stress that it does not represent all publications in the AIS domain, we believe it is reflective of the general picture. Note that this section does not describe in detail the application areas that AIS has been applied to. The reader is referred to the above publications for further information — the section is intended to provide an overview of the field as a whole and provide a basis for the following discussion. Based on this survey, we find that application areas can be roughly classified into 12 headings which are shown in table 1. The categories are chosen simply to reflect the natural grouping of papers and in some cases are rather broad, and in others very narrow. For example, computer security and virus detection could be classified as examples of anomaly detection, and the majority

[1] This is an extended version of [43]

Table 1
Application Areas of AIS

Major	Minor
Clustering/Classification	Bio-informatics
Anomaly Detection	Image Processing
Computer Security	Control
Numeric Function Optimisation	Robotics
Combinatoric Optimisation	Virus Detection
Learning	Web Mining

of the bio-informatics papers are essentially performing classification or clustering. However, where more than one paper has been written on a particular application area, these papers have been grouped together.

In brief, papers falling under the heading *Anomaly Detection* include a diverse range of topic areas, ranging for example from detection of temperature fluctuations in refrigeration units [74], aircraft fault detection [21], fault detection in automated teller machines [7] and in telephone networks [62]. As previously mentioned, computer security and virus detection applications could also be classified under this heading; these sub-headings speak for themselves as to the type of application covered. Some specific features of anomaly detection applications are discussed in more detail in section 3.1.

A very large number of papers fall under the general heading of *Learning*. Learning can generally be understood to be the process of acquiring knowledge from experience and being able to re-apply that knowledge to previously unseen problem instances — this generic title applies to a variety of subtopics such as pattern recognition, concept-learning, and supervised and unsupervised versions of clustering data and classifying data. Papers relating to *clustering and classification* have been separated out from the general learning topic as a sub-topic where they relate specifically to clustering or classifying a particular data-set and have been compared to conventional classification techniques, and have been benchmarked using the standard accepted quality tests in data-mining such as classification accuracy. Almost all clustering applications which have gone beyond the conceptual stage focus on benchmark sets of data such as those available from the UCI repository which are static in nature, although there are few attempts to apply immune-based algorithms to dynamic data, e.g [39,57].

As previously mentioned, papers relating to *bio-informatics* have also been separated a distinct topic, as these form a natural group; however, it is important to realise that this topic essentially is just another set of applications of clustering or classification algorithms — again the data being clustered or classified is static in nature.

Optimisation covers a number of real-world application areas involving combinatoric optimisation such as travelling salesman problems, scheduling (including inventory and job-shop scheduling), and routing problems, and also optimisation of numeric functions. Typically, the publications report results on benchmark problem instances rather than real-world problem instances, although this picture is slowly changing. For example, Campelo *et al* discuss the use of an immune-based optimization algorithm when optimising the design of electromagnetic devices, and Kalinli *et al* apply an AIS algorithm to the design of infinite impulse response (IIR) filters.

Robotic applications tend to be based on controlling simulated robots around small, artificial environments, generally addressing the problem of behaviour arbitration and autonomous navigation, although work by [44] attempts to lay a foundation for using an AIS to provide the basis of an architecture for a robot to acquire new, more complex skills throughout its lifetime. *Adaptive control systems* form a related category of papers, for example pertaining to controlling a robotic arm [55]. The small topics of *Image Processing* and *Web-Mining* are self-evident, however, an increasing amount of literature reporting successful applications of immune algorithms in these domains is becoming apparent. For example, Zhong *et al* [83] report the use of an unsupervised AIS based classification algorithm to perform remote-sensing classification of satellite images. Su *et al* investigate approaches for compensating exposure in images where there is back-lighting, and use a supervised neuro-immune system to estimate the amount of compensation required. Web-mining applications are also becoming more prevalent — Nasraoui *et al* [56] use an AIS algorithm for tracking clusters in a dynamic data-stream of click-sequence data generated by web-users in order to construct user-profiles.

2.1 Summary of Application Areas

Having presented the above categorisations of application areas, it seems that application areas that have been addressed by AIS techniques can be broadly summarised as (1) Learning (2) Anomaly Detection and (3) Optimisation. Thus, *learning* includes clustering, classification and pattern recognition, robotic and control applications; *Anomaly Detection* includes fault detection and computer and network security applications, and *Optimisation* includes real-world problems which essentially boil down to combinatoric and also numeric function optimisation. To some extent, the fact that applications of AIS have fallen into the above categories is somewhat an accident of history. Early immune-based algorithms, proposed in the main by computer scientists with little if any immunological background, seized on what appeared to be be the obvious functions of the immune system as a defensive system, able to perform pattern recognition and learn over time. Hence, although very early work in the area was performed from an interdisciplinary slant, e.g. [11], there has been a tendency to *reason by metaphor* [69], and apply simplistic models such as clonal selection, immune-networks and negative-selection in isolation to problems which appear at first glance to be amenable to such techniques. Furthermore, again perhaps by accident, many of the AIS practitioners arrive in the field by way of working in other biologically inspired fields such as Evolutionary Computing, and thus there is a tendency to apply AIS algorithms to the same problems as have been tackled in other domains (e.g. optimisation), which often results in un-natural problem representations, and rather contrived mechanisms for mapping a problem to an AIS algorithm.

3 "Was it worth it " - a look at the added value of the AIS

It is now pertinent to re-evaluate the application of immune algorithms to the above application areas, and question whether there is really any added value in applying AIS to the three areas listed above. Again we re-iterate that there is no doubt that AIS has been successful in these areas; however, we question as to whether they AIS brings any benefits that could not have been gained from applying a different sort of algorithm. Recall the seminal list of features of an AIS, originally due to Dasgupta in [20] and so often quoted in AIS publications. This defines the features of an immune system that are relevant from a computational perspective as: recognition, feature extraction, diversity, learning, memory, distributed detection, self-regulation, thresholds, co-stimulation, dynamic protection and probabilistic detection. Although later we question as to whether the features on this list really distinguish an AIS from many other paradigms, it is useful to bear in mind during the following analysis of the three application areas.

3.1 Anomaly Detection

Anomaly detection has been an area of application that has found favor with the AIS practitioner. Such techniques are required to decide whether an unknown test sample is produced by the underlying probability distribution that corresponds to the training set of normal examples. Typically, only a single class is available on which to train the system. The goal of these immune inspired system was to take examples from one class (usually what was considered to be normal operational data) and generate a set of detectors that was capable of identifying when the *normal* or *known* system had changed, thus indicating a possible intrusion.

The early pioneering work of Forrest et al [30] led to a great deal of research and proposal of immune inspired anomaly detection systems [29]. Results reported in these works, did hint at the possibility that the immune approach was useful to some degree as both known and novel intrusions could be detected. This was extended by work of [53], who combined the clonal selection algorithm with a negative selection algorithm to help reduce the false positive rates. The interest of this immune approach was in part, due to the fact that it appeared possible to train a system with only a single class of examples and the intuitive link between the role of the natural immune system as the "great protector" and the development of intrusion detection systems. Notable work in [8] proposed the *r-chunk* matching rule which was to replace the computational expensive *r-contiguous bits* matching rule that had dogged the approaches to date. The *r-chunk* rule made it computationally more efficient to generate a set of detectors of the non-self space (in hamming shape space) and later computationally more efficient methods were developed in real-valued shape space [36,50], again based on only a single class of examples. This potentially made the use of the immune approach more attractive, as the main issue that had been raised to date was one of scalability with respect to the size of the *normal* data.

Recent work in [27], proposed a formal framework for the negative selection approach, and when one examines this work, it is possible to see hints that the *r-chunk* may well suffer certain scaling problems. Indeed, this has now been confirmed by [70,71] who present an in-depth theoretical analysis of the negative selection algorithm over real and hamming shape spaces. The investigations reveal that defined over the hamming shape-space, the approach is not well suited for real-world anomaly detection problems. Problems arise with the generated detector set which under-fits exponentially for small values of r (where r is the size of the chunk. They suggest that in order avoid this under-fitting behavior, the matching threshold value r must lie near l (the length of the string). However, they point out that this has a consequence. This is that the detector generation process is once again infeasible, since all proposed detector generating algorithms have a runtime complexity which is exponential in r. In addition to their theoretical arguments, they undertook a simple study of comparison between the negative selection approaches on a one-class support vector machine (SVM) [63]. When comparing the work of [50], (the real-valued negative selection algorithm with variable-sized detectors) results revealed, that the classification performance of the method not only crucially depended on the size of the variable region, but results from the one-class SVM provides as good, if not better results. In addition, they noted that in order to tune the parameters of the system by [50] it was necessary to have the second class, as the probability distribution of this class impacted a great deal on the overall performance of the system. Further work in [72] presents a comparison between real-value negative selection, a one-class support vector machine and a Parzen-Window Estimator technique. In this paper, the authors used the well known KDD data set [46] for testing intrusion detection algorithms. Their analysis showed that on high-dimensional data such as the KDD set, the negative selection over real value shape space had very poor performance when compared to the other techniques (the best being the SVM). At the time, no reason for this poor performance was offered, but it would appear now that the reason for this is the properties of hyper-spheres that are used as a representation mechanism within these techniques. When the problem size scales with respect to the number of dimensions, the hyper-sphere representation (and combined affinity metric) the actual volume of the sphere decreases to almost zero. [2]

So, from a "value added" perspective, the contribution of the immune inspired approach is as yet unclear. The negative selection approach seems to have many drawbacks such as scaling issues, high false positive rates and complexity issues. One attractive feature of the approach was the ability to only use a single class of data on which to train the system, certain investigations would seem to indicate that the second class is needed to tune the system. However, in order to overcome some of these shortfalls, work proposed in [5] and later expanded on in [6] proposes the adoption of the *danger theory* approach. The authors claim that it should be possible to move away from the need to define what is *normal* for a system, and dynamically identify *normal* through the adoption of danger signals and context dependent responses.

Recent work by this *Danger team* [3] has revealed some interesting insights. This team are examining the interaction of the innate and adaptive immune system, rather than focusing purely in the adaptive immune system. For example, work in [37] presented initial investigations into the use of *dendritic cells* as inspiration for an anomaly detection system. In this paper, the authors present a biologically grounded algorithm that is able to detect changes in the input space. This is early work (and has been tested on the "benchmark" problems), but the potential for the reduced level of false positives is apparent. It is clear from this work, that the *dendritic cell algorithm* is part of a larger system, which was discussed in a companion paper [54]. In this paper, the authors outline an immune inspired system (which in part is made up of the work in [37]. In [54] an architecture is outlined of a worm detection system inspired by three central T-cell processes. The focus of this work is on the coordination of multiple responses, and it is the aim of the work to capture the coordination strategy of the T-cell response in their system. Therefore, the area of security is still being pursued by the AIS researcher. To date, results have been limited in success, but it is possible that through the adoption of the ideas from the innate immune system, significant breakthroughs may be possible.

[2] personal communication with T. Stibor, paper in preparation
[3] http://www.dangertheory.com

One might argue that the immune system does not do optimisation at all, at least in the manner that we use the term when solving traditional combinatoric or numerical "optimisation" problems. Most [4] computational optimisation problems have a single *goal* to obtain; the natural immune system on the other hand can be regarded as having multiple and possibly contradictory goals, and although it does improve its own response towards particular goal(s) as the result of feedback [66] it has no reason to evolve an *optimal* response. In fact, its network structure does not lead to the development of the best response, but results in the best possible response under exisiting conditions [60]. However, the AIS world has seized upon optmisation as a promising application area for AIS-inspired algorithms, leading to a number of publications relating to function optimisation problems, often declaring some success when compared against other state-of-the-art algorithms. The majority of these publications are based on the application of the clonal selection principle, resulting in a number of algorithms such as Clonalg algorithm [25], opt-AINET [26] and the B-Cell algorithm [75] and opt-IA [19]. For example, Freschi and Repetto provide an wide-ranging analysis of opt-IA and Clonalg on a robust test-bed that includes "toy" optimisation problems such as max-1s, trap functions and 23 numerical optimisation problems from [82] and find that immune algorithms are comparable to some of the most effective methods in the evolutionary algorithm literature such as Fast Evolutionary Programming (FEP). [75] compare versions of opt-AINET and the B-Cell algorithm to a variety of optimisation functions of various dimensions found in the literature, and [18] apply Clonalg to a range of *constrained* optmisation problems.

All of these algorithms essentially evolve solutions to problems via repeated application of a cloning, mutation and selection cycle to a population of candidate solutions (B Cells). A single antigen represents some function to be optimised, and good solutions are allowed to remain in the population, mimicking the memory cell mechanisms believed to exists in the natural immune system. The authors of optAINET state that it is characterised by the following features; it performs exploitation and exploration of the search space, it can determine the locations of multiple optima, it maintains many optimal solutions, and has defined stopping criteria. The main differences between this and Clonalg or the B-Cell algorithm lie in whether or not they maintain a static or adaptive population size, whether or not they include elitist mechanisms and in type of mutation operators they use. Anyone familiar with the EA literature will recognise all of these features as equally applicable to an EA, and even the differences between the immune algorithms are recognisable as differences between the various flavours of EA. Indeed, Newborough

[4] although not all, see later discussion

et al [58] even argue that the properties of most population-based algorithms — immune algorithms, genetic algorithms, swarms and ant-colonies— can be captured in a generic framework which factors out the commonalities of these algorithms and applies various properties uniformally across all of the classes, even though features such as niching and elitism may at first glance appear particular to just one class.

So have immune algorithms added value to the field of optimisation ? We are presented with a range of somewhat conflicting evidence in the literature. Cruz-Cortez [18], despite an in-depth evaluation of Clonalg, are unable to out-perform stochastic ranking techniques on constrained optimisation problems; [32] however finds that opt-IA performs comparably to FEP on 23 numeric optmisation problems, and the B-Cell algorithm [75] has been shown to use significantly fewer evaluations than a hybrid GA on some problems. In addition, [32] find that their multi-objective immune algorithm, VAIS, is comparable to NSGA2, the state-of-the-art for solving multi-objective optimisation problems. A note of caution however that their comparative study only involved three problems. This particular type of optimsation is more in-keeping with the goals of the natural IS as discussed earlier in this section, and hence proper exploitation of the IS metaphor may prove a promising avenue in relation to these problems in the future.

In general, we conclude that although AIS will prove a useful addition to the plethora of population-based algorithms that already exist, being able to solve some problems more easily than others, it will not *distinguish* itself as being significantly better over a wide-range of numerical optmisation problems than existing non-AIS techniques. Nonetheless, it is fruitful to continue with this line of work, to gain an understanding of *why* some AIS algorithms are better capable of searching certain landscapes than others.

A similar argument applies to the use of AIS in Combinatorial Optimisation Problems. [16] showed that an AIS algorithm utilising gene-libraries and a clonal selection mechanism was highly competitive with respect to GRASP [14] for solving a series of 31 benchmark job-shop scheduling problems taken from the OR-library [9]. Their algorithm showed improvements on GRASP on some problems and on all problems was less computationally expensive. Another clonal-selection based algorithm known as ClonaFlex was reported by [59] who apply it to a benchmark set of 12 flexible job-shop problems, comparing results against GENACE, a cultural EA [73]. They consider only one objective function — makespan — and find comparable results, although they report considerable practical difficulties in optimising algorithm parameters for individual problem instances. Despite these results however, we conjecture that will be no obvious benefit to be gained from applying an AIS to these essentially static problems when considering a wide-range of problems. However, note that [42,45,17] described some preliminary results applying an AIS to job-shop scheduling problems that attempted to capitalise on the properties of the IS to produce *robust* rather than optimal schedules. Optimal solutions to problems are often incredibly fragile — if the original problem changes slightly, then an optimal solution cannot be massaged a little to match, a new one must be produced. In the real-world of course, such changes happen all the time and hence there is a great deal of interest in the scheduling communities in generating robust, good-enough schedules rather than optimal ones, e.g. [49]. This may prove an interesting avenue of research of AIS.

Perhaps a more obvious optimisation area is that of *dynamic* function optimisation. In these problems, the goal is to find and track a continuously moving target — this at least fits better with the view of the immune system as a dynamic, and continuously adapting system. Gaspar and Collard [34] used a network-based AIS to perform dynamic function optimisation. Walker *et al* [78] have applied a version of Clonalg to a number of dynamic optimisation problems which they compare to an evolutionary strategy and find that generally an evolutionary strategy can optimise more quickly than the clonal selection algorithm. Recently, Kelsey *et al* [52] have adapted the B-Cell algorithm [51] to perform dynamic optimisation, and found that the fast adaptable nature of the algorithm enabled the tracking of multiple moving optima. Although there is little other work in this area, we hypothesise that continuing research effort will not lead to a dynamic-optimisation algorithm with distinctive performance, although it is reasonable to assume that comparable performance to other optimisation algorithms can be obtained on some problems.

3.3 Clustering and Classification

Immune-based algorithms which perform clustering make up a large number of the application areas shown in table 1. These range from supervised algorithms such as AIRS [80] and Carter, to aiNET [26] and algorithms based on idiotypic network models such as those of Neal and Timmis [57]. However, as already stated, the application areas to which these models are to clustering or classifying *static* data sets, where comparable or improved performance is achieved on many data-sets, when compared to traditional algorithms. Classification/clustering require *feature extraction*, *recognition* and *learning* — key features of the AIS — however, we conjecture that these are also key features of any machine-learning algorithms, and that there are no unique features of the *problem domain* that indicate an AIS based algorithm can offer anything over and above the more traditional machine learning algorithms. One potential distinguishing feature of the IS which *has* been exploited in classification is its *distributed* nature, which is used to advantage by Watkins [80] in a parallel version of AIRS. In his thesis, Watkins [79] takes these ideas even further and develops a *distributed* version of AIRS, which is significantly more biologically plausible than a parallel version. With the parallel version, speeds ups were observed with respect to the amount of data, with the distributed version the benefits were more subtle. Whilst speed ups were observed with respect to the amount of data, interestingly, regions of the network became more specialised at identifying certain classes than others. Therefore, there were regions of the network that were hopeless at classifying items, other regions excelled. This area is certainly interesting, and potentially could add something *distinct* to AIS classifiers.

A more promising application area for AIS may lie in the area of *dynamic* clustering or classification. Advances in technology now make it incredibly straightforward for huge amounts of data to be collected and stored cheaply and easily, and hence many companies and researchers now routinely collect data on a daily or even hourly basis. By tracking patterns and trends in the data, companies may be able to gain a competitive advantage. There are some existing learning algorithms which can cluster dynamic data — however, in an era of ever increasing computational processing power coupled with continually decreasing costs, it is pertinent to question why dynamic algorithms need even to be considered for time-varying problems. It is trivial for example to re-apply established "static" algorithms at each time-instant in a dynamic problem to the data in-hand; however, this type of approach totally disregards any information captured in either the current information or in previous time-series, therefore may miss vital clues. Therefore, we propose that AIS algorithms by definition, incorporate some form of memory, and can therefore outperform other state-of-the-art learning systems which are purely reactive. Most learning systems have very limited memory and hence no mechanism to balance the need to keep a record of currently under-used knowledge acquired in the past against the need to store newly-acquired knowledge that is valuable in the current climate.

Note that there is some existing, although limited, work in this area. Neals algorithm [57] is meta-stable in that it can in theory be continuously applied to a data-set. The work of Hart [39] models a self-organising system which is able to dynamically cluster moving data, whilst maintaining some memory of the past, but has only been tested with artificial data-sets. Work by Secker et al [64] developed a dynamic supervised learning algorithm for the filter of emails, and work by Kim and Bentley [53] a dynamic classification algorithm for use in intrusion detection.

Based upon the work presented by Secker *et al* [64], Ayara *et al* [7] propose a technique that aims to prevent system down-time by detecting states that are precursors of system failure in Automated Teller Machines (ATM). This is achieved via an immune inspired algorithm that produces a set of error detectors that are capable of identifying potential system failure, hours before the actual event (up to 12 hours in some cases). Once a potential failure is identified, an alarm can be raised, and an engineer despatched. Unlike the typical anomaly detection techniques discussed in section 3.1, this technique relies on the existence of sequences of states that represent the operational status of an ATM when errors are occurring (so not when the ATM is operating within normal bounds). The *adaptable error detection process* is able to identify those sequences that might contain fatal states and identify potential sequences that might lead to system failure. Of critical importance in this work, was the low false positive rate achieved by the system, which has seemed to blight other systems. However, this is a single application, so generalities can not be drawn from this.

4 A New Approach to AIS

The above discussion has shed a rather gloomy light on future of AIS in solving real-world applications. Perhaps this is a suitable point to take a step backwards and first re-evaluate our approach to designing AIS algorithms, as well as attempting to define what kind of applications they may be suitable for. With this in mind, we take brief look at both sides of the coin and take first an algorithm-oriented and then a problem-oriented view of the situation.

4.1 A Conceptual Framework for Algorithm Development

Work by Stepney el. al [69] 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 [69], 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. The basic outline of the approach proposed by Stepney et al. is to first *probe* the biological system in question. When one probes such a system, one has to bear in mind what it is you want to extract or observe. For example, you may be interested in initiation of danger signals, so one would undertake experimentation to observe that. This process is then followed by the development of suitable mathematical models. Properties of the system can then be modelled at a mathematical level, which allows for possible insights into the biological model that are not possible with "wet lab" experiments. From this, it is then possible to construct a computation model, based on the mathematical model. The creation of the computational model allows for the execution of the model, to observe and gain insight into the workings of the model. This model can then more easily be abstracted into an algorithm, or set of algorithms for deployment in an application area. Clearly, this is an iterative process, that allows for a great deal of interaction between all stages. Arising from this may be various *computational frameworks* that are suitable for instantiation into applications.

Stepney *et al* then go onto propose that once such frameworks are developed, it is possible to ask suitably posed *meta-questions* about the framework, that may give attention to interesting properties. The questions are concerned with openness (e.g. how much continual growth or development is required within the system), diversity (e.g how many agents are required), interaction (e.g. level of communication between agents), structure (e.g are the different levels required between agents) and scale (e.g how many agents are required). These are known as the ODISS questions. The potential benefit of adopting this approach is clear not only do all disciplines benefit from such work, but the immune algorithms developed at the end of the process will, all being well, be more grounded in the immunology than the simple *observe*, *implement* approach so dominant in the AIS literature today.

4.2 A Problem Oriented Perspective

Freitas and Timmis [31] outline 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 [40] with the development of their simple immune network simulator with various affinity metrics. They make the argument that seemingly generic AIS algorithms, are maybe 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 et al. [65], Bersini and Varela [13] and Somayaji *et al* [67] (which was specifically focussed on design of computer immune systems). However, work by Segal, whilst extremely interesting, focussed primarily on network signalling, 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. We assert that these potentially useful principles 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.

4.3 Towards a Theoretical Understanding

As we have alluded to in this paper, there is a significant lack of theoretical work in the area of AIS. This, again, has more than likely been by an accident of events, rather than anything else. In the scramble to develop novel and useful algorithms, less attention was paid to the theoretical side. In order for any subject to mature, a solid theoretical underpinning is required. In the past few years, there has been some attempt to address the imbalance. Drawing on their experience of developing multi-objective clonal selection based algorithms ([16]), work by Arias *et al* [77] provides a complete proof for their multi-objective immune inspired algorithm. This is one of the first attempts at providing a complete proof (via a Markov Chain method) for an AIS algorithm. Interestingly, at the same conference, work by Hone and Kelsey [47] discussed the application of dynamic systems techniques to the analysis of AIS algorithms. In their paper, they discuss a particular example (the B-cell algorithm [51]) and how it may be possible to first, use this algorithm to solve complex root functions in mathematics (by recasting the problem as an optimisation problem), and secondly, analyse the *dynamic* behavior of the algorithm (and other algorithms) using dynamical systems theory. They argue, that as most of the AIS literature was born from ideas such as Perelson *at al* [61] and Farmer *et al* [28], which themselves have foundations in dynamical systems, adopting such theoretical ideas within AIS would seem a logical step forward, and may prove to be very insightful.

Other work by Clark *et al* [15] has produced a theoretical analysis of the B-Cell algorithm [51] above which provides a complete and exact model of the B-cell algorithm with a proof of convergence. In addition, from their model, it would appear that it is possible to locate the optimum mutation rate for a given function. Thus, as there have been no convincing theoretical analyses that enable performance prediction in the EA world (at least using realistically sized strings), there is perhaps value in applying a properly understood algorithm to a problem, regardless of the nature of the problem. Even if a world-beating optimisation algorithm does not emerge from the AIS world, such theoretical studies will still go a long way towards increasing our understanding of these algorithms and hence the research is valuable.

In addition to the work reported above addressing pure theory, there has also been some recent work which attempts to gain a deeper understanding of dynamic immune algorithms through empirical simulation. Thus [12,40,41,38] have simulated immune network algorithms in Hamming and real-valued shape-space in order to understand how the size and shape of the recognition region of a cell can impact the emergence of immune networks, their stability and dynamics, and their memory capacity. They argue performing this kind of simulation allows a proper understanding of *why* algorithms exhibit certain performance to be gained as well as an understanding of how performance of the algorithms can be modified in a controlled and principled manner by altering parameters. This will ultimately lead to better engineered algorithms which can be tailored to individual applications in a scientific manner, rather than by ad-hoc parameter tweaking.

5 Suggestions as to the Way Forward

We have outlined what we believe to be the problems with the current applications to which AIS has been applied, from the perspective that although reasonably successful on a narrow range of problems, they do not add sufficient value over and above that which is offered by other paradigms to make them anything other than another tool in the engineers application tool-box. Although from some points of view, any tool is a worthwhile addition, we believe there is still a wealth of unexploited potential in the AIS domain. Adopting the methodology and problem oriented perspectives outlined above rather than the scatter gun approaches taken to date will surely help us tap into this potential. However, there are some crucial missing ingredients in our current perspectives in AIS that limit our current progress. Here we suggest three of the areas that we feel will play some part in defining the future of AIS — note that there will of course be several others.

5.1 The Innate Immune System

The natural immune system is known to comprise of two sub-systems, working in tandem with each other; the *innate* immune system, and the *adaptive* immune system. Almost without exception, the AIS community has chosen to model the adaptive immune system. This may partially reflect the historical interest in the adaptive immune system in the immunological community, which over a period of years, dismissed the innate system as the minor partner in the functioning of the immune system. Recently however there has been a resurgence of interest in the innate immune system in immunological circles — witness for example the work described in [35], and the influence it may have on the adaptive system. Directing some attention therefore towards understanding and modelling the innate system maybe prove fruitful in producing better immune-models. For example, we may choose to focus on a certain aspect such as signalling mechanisms within the innate immune system and apply the conceptual framework model to abstract useful mechanisms based on this. A step forward in this direction is conceptual work presented at [48] introducing the concept of artificial tissue, [10,37]. The authors propose that a layer of artifical tissue can essentially represent the problem data and store the current state of the application. Communication between an AIS and a problem is then mediated via the tissue-layer. The tissue thus provides some of the functionality of the innate immune system. Whilst these ideas are currently in their infancy they present a alternative mechanism that may potentially alter the way in which AIS applications are viewed and implemented in the future.

Strikingly, one of the key problems identified in section 3 with optimisation and clustering applications is that immune algorithms are applied to *static* systems without any justification. Yet, the inspiration behind the algorithms applied to such systems is the *adaptive* immune system, where we model clonal selection and learning on relatively fast time-scales. Perhaps such applications areas should be re-evaluated in the light of what we can learn from modelling the *innate* immune system. Many creatures, e.g. the nematode worm have *only* an innate immune system and yet function perfectly well — perhaps in many cases we have been too ambitious by trying to model the complete immune system and could achieve equally impressive results by abstracting mechanisms from a more simplistic yet still incredible system.

However, we also suspect that the true value of the immune metaphor will be only revealed in systems which exploit the full richness of the natural immune system which is gained through the synergistic interaction between the innate and adaptive immune systems. Systems which can exploit this interaction have huge potential to benefit from application of the immune metaphor and categorically distinguish themselves from other biologically inspired paradigms. For example, we know of no other biological metaphor which genuinely is based on heterogeneous components and operates over multiple time-scales; applications which through their underlying physiology (whether that is physical or virtual) can encapsulate innate and adaptive immune components may hold the key to finding the elusive killer application.

5.2 The immune system does not operate in isolation

Living organisms show a remarkable ability to maintain homeostasis, that is, achieve a steady-state of internal body function in a varying environment. This is precisely what we wish to achieve in many practical anomaly detection systems, for example in maintaining a secure computing environment. In nature, this is made possible via the —em interaction of both a number of systems, for including the immune system, neural system and endocrine system, and via multiple components within each of these systems. Any one of these systems cannot and does not operate in isolation — this suggests that perhaps the true potential of modelling immune systems might only be achieved via combining them with other sub-systems. This is clearly an exciting new area of research to which attention should be paid. There has been some exploratory work in this area — [76] — yet much remains unknown. Furthermore, the fact that the immune system does not act in isolation gives us yet another important pointer; the immune system must be *embodied*. This fact has been acknowledged in robotic research for a long time, where it is well known that "there can be no intelligence without embodiment", however it is largely ignored in AIS research.

Indeed, work in progress by Stepney [68] describes in great detail the notion that *all* systems are embodied, either physically (such as a robot) or virtually (such as a software agent). Importantly, the embodied system exploits the environment in which it is in. However, Stepney argues that this can only be achieved through careful design of both system and environment, to allow for sufficient amounts of interactions between both, thus allowing for the necessary dynamics to arise.

5.3 Life-long learning

Although many application papers allude to this aspect of the immune system in their introductory text, few systems have really attempted to capture this feature of the IS, and those that have exhibit only a weak version of this. For instance, some optimisation and clustering algorithms have been applied in dynamic environments. However, there has been no published work on problems which *naturally* require a system to *improve its own performance* over the course of a life-time, as a result of its own experience. As this feature of the IS clearly distinguishes it from most other biologically inspired paradigms such as EAs or neural-nets which produce a fixed solution (or solutions) to a problem and then terminate, choice of application areas should focus on those problems which naturally require continuous learning.

6 Conclusions: Features of AIS Applications

We summarise by proposing a list of features that draw together some of the preceding discussion and that we believe point to the way forward for

AIS. Some of these features are currently absent in any of the AIS literature. Others, such as life-long learning, have been modelled in a limited sense. We emphasise that it is by the *combination* of these principles that a distinctive niche is carved for AIS.

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

For example, the system described by Ayara *et al* [7] (section 3.3) is embodied, distributed, has multiple components and its purpose is to maintain homeostasis in a distributed ATM network, therefore must exhibit life-long learning, and therefore exactly encapsulates the principles just outlined.

One ambitious idea might be to develop an *immune controlled operating system* [5] . The idea is that a set of heterogenous immune components would interact (and control) a combined neural system which may make decisions such as to conserve battery power by not running virus checking software when the computer is not connected to a network, to time certain diagnostics on the computer, or give priority to tools which are useful *at a specific moment in time* to the user. One may also consider buiding in *natural defense* mechanisms to the operating system (or one would hope for these to be an emergent property of the AIS in the computer!). Clearly, this system will be embodied, and will have to make use of multiple components, over multiple timescales (the time to decide on switching on and off power, will be different to that of deciding when to run a diskscan routine).

However, all this futuristic discussion is interesting, but what is needed is well grounded immune inspired techniques, that are applied in a logical and coherent matter. Maybe it is time to move away from the familiar scattergun approach we alluded to at the start of this paper, and think carefully about how, why and where we apply these techniques. Through a principled combination of good science (development of algorithms in a scientific manner) and careful engineering (the application of these ideas), both the state of the art in AIS will progress, but also we will be able to engineer solutions that we can not do with other techniques.

[5] This idea is not our own, but has come via personal communications with Dr. Mark Neal, University of Wales, Aberystwyth, UK

Acknowledgments

Many of the ideas in this paper have evolved from useful and stimulating discussions at the ICARIS conferences with many people in the field, and in particular at meetings of the UK based ARTIST [1]. We would like to thank ARTIST for financial support provided during the development of ideas in this paper.

References

[1] `http://artificial-immune-systems.org/artist.htm`.

[2] `http://evonet.lri.fr`.

[3] *Artificial Immune Systems: Proceedings of ICARIS 2004*, LNCS 3239. Springer, 2004.

[4] *Proceedings of Genetic and Evolutionary Computation Conference*. Springer, 2004.

[5] U. Aickelin and S. Cayzer. The danger theory and its application to artificial immune systems. In Jonathan Timmis and Peter J. Bentley, editors, *Proceedings of the 1st International Conference on Artificial Immune Systems ICARIS*, pages 141–148, University of Kent at Canterbury, September 2002. University of Kent at Canterbury Printing Unit.

[6] U Aicklen, P Bentley, S Cayzer, J Kim, and J McLeod. Danger Theory: The Link Between AIS and IDS? In *LNCS 2787*, pages 147–155. Springer, 2003.

[7] M. Ayara, J. Timmis, R. de Lemos, and S. Forrest. Immunising automated teller machines. In Jacob et al. [48], pages 404–417.

[8] J. Balthrop, F. Esponda, S. Forrest, and M. Glickman. Coverage and generalization in an artificial immune system. In *GECCO 2002: Proc. of the Genetic and Evolutionary Computation Conf.*, pages 3–10, 2002.

[9] J.E. Beasley. Or-library: Distributing test problems by electronic mail. *Journal of the Operations Research Society*, 41(11):1069–1072, 1990.

[10] P. Bentley, J. Greensmith, and S. Ujjin. Two ways to grow tissue for artificial immune systems. In Jacob et al. [48], pages 139–153.

[11] H. Bersini. Immune network and adaptive control. In *Proceedings of 1st Conference on Artificial Life*, pages 217–226. MITpress, 1991.

[12] H. Bersini. Revisiting idiotypic immune networks. In W. Banzhaf and *et al*, editors, *Advances in Artificial Life, 7th European Conference, Proceedings of ECAL 2003*, Lecture Notes in Computer Science, pages 164–174. Springer, 2003.

[13] H Bersini and F Varela. *The Immune Learning Mechansims: Recruitment, Reinforcement and their Applications.* Chapman Hall, 1994.

[14] S. Binato, W.J. Hery, D.M. Loewenstern, and M.G.C. Resende. A grasp for job shop scheduling. In C. Ribiero and P. Hansen, editors, *Essays and Surveys in Meta-heuristics*, pages 59–80. Kluwer, 2001.

[15] 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.

[16] C. Coello Coello, D. Cortez-Rivera, and N. Cruz Cortez. Use of an artificial immune system for job shop scheduling. In *Artificial Immune Systems: Proceedings of ICARIS 2003*, pages 1–10. Springer, 2003.

[17] D Corne, M Dorigo, and F Glover, editors. *New Methods in Optimisation*, chapter Artificial Immune Systems for Job-Shop Scheduling, E. Hart and P. Ross. McGraw-Hill, 1999.

[18] N. Cruz-Cortez, D. Trejo-Perez, and C. Coello Coello. Handling constraints in global optimization using an artificial immune system. In Jacob et al. [48], pages 234–247.

[19] V. Cutello, G. Nicosia, and M. Pavone. Exploring the capability of immune algorithms: A characterization of hypermutation operators. In *Artificial Immune Systems: Proceedings of ICARIS 2004*, pages 263–273. Springer, 2004.

[20] D. Dasgupta, editor. *Artifificial Immune Systems and their Applications.* Springer, 1999.

[21] D. Dasgupta, K. KrishnaKumar, D. Wong, and M. Berry. Negative selection algorithm for aircraft fault detection. In *Artificial Immune Systems: Proceedings of ICARIS 2004*, pages 1–14. Springer, 2004.

[22] L. De Castro. Artificial immune systems bibliography. `http://www.dca.fee.unicamp.br/~lnunes/AIS.html`.

[23] L. De Castro. Engineering applications of artificial immune systems. Tutorial at ICARIS 2004, available from `http://artificial-immune-systems.org/ICARIS2004/icaris2004.htm`, 2004.

[24] L. De Castro and F. Von Zuben. Artifical immune systems: Part ii - a survery of applications. Technical Report DCA-RT 02/00, Department of Computer Engineering and Industrial Automation, State University of Campinas, Brazil, 2000.

[25] L. De Castro and F.J. Von Zuben. The clonal selection algorithm with engineering applications. In *GECCO 2002 - Workshop Proceedings*, pages 37–37, 2000.

[26] L.N. De Castro and F.J. Von Zuben. Ainet: An artificial immune network for data analysis. In H.A. Abbass, R.A. Sarker, and C.S. Newton, editors, *Data Mining: A Heuristic Approach*. Idea Group Publishing, USA, 2001.

[27] 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.

[28] J D Farmer, N H Packard, and A S Perelson. The Immune System, Adaptation, and Machine Learning. *Physica*, D(22):187–204, 1986.

[29] S Forrest, S Hofmeyr, and A Somayaji. Computer Immunology. *Communications of the ACM*, 40(10):88–96, 1997.

[30] S Forrest, A Perelson, L Allen, and R Cherukuri. Self-Nonself Discrimination in a Computer. In *Proc. of the IEEE Symposium on Research in Security and Privacy*, pages 202–212, 1994.

[31] A Freitas and J Timmis. Revisiting the Foundations of Artificial Immune Systems: A Problem Oriented Perspective. In *LNCS 2787*, pages 229–241. Springer, 2003.

[32] F. Freschi and M. Repetto. Multiobjective optimisation by a modified artificial immune system. In Jacob et al. [48], pages 248–261.

[33] S. Garrett. How do we evaluate artificial immune systems ? *Evolutionary Computation*, 13(2):145–178, 2005.

[34] A. Gaspar and P. Collard. From GAs to artificial immune systems: Improving adaptation in time-dependent optimization. In *Proceedings of the Congress on Evolutionary Computation*, pages 1859–186. IEEE Press, 1999.

[35] R. Germain. An innately interesting decade of research in immunology. *Nature Medicine*, 10(12):1307–1320, 2004.

[36] F. Gonzalez, D. Dasgupta, and L.F. Nino. A randomized real-valued negative selection algorithm. In *Proceedings of the 2nd International Conference on Artificial Immune Systems (ICARIS). Lecture Notes in Computer Science*, pages 261–272, 2003.

[37] 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.

[38] E. Hart. Not all balls are round: An investigation of alternative recognition region shapes. In Jacob et al. [48], pages 29–42.

[39] E. Hart and P. Ross. Exploiting the analogy betweeen the immune system and sparse distributed memories. *Genetic Programming and Evolvable Machines*, 4:333–358, 2003.

[40] E Hart and P Ross. Studies on the implications of shape-space models for idiotypic networks. In *LNCS 3239*, pages 413–427. Springer, 2004.

[41] E. Hart and P. Ross. The impact of the shape of antibody recognition regions on the emergence of idiotypic networks. *Journal of Unconventional Computing*, 1(3):281–313, 2005.

[42] E. Hart, P. Ross, and J. Nelson. Producing robust schedules via an artificial immune system. In *IEEE World Congress on Computational Intelligence, International Conference on Evolutionary Computing*, 1998.

[43] E. Hart and J. Timmis. Application areas of AIS: The Past, Present and Future. 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 483–499. Springer, 2005.

[44] E. Hart, A. Webb, P. Ross, and A. Lawson. A role for immunology in next generation robot controllers. In *Artificial Immune Systems: Proceedings of ICARIS 2003*, pages 46–57. Springer, 2003.

[45] Emma Hart and Peter Ross. An immune system approach to scheduling in changing environments. In Wolfgang Banzhaf, Jason M. Daida, A. E. Eiben, Max H. Garzon, Vasant Honavar, Mark J. Jakiela, and Robert E. Smith, editors, *GECCO*, pages 1559–1566. Morgan Kaufmann, 1999.

[46] S. Hettich and S. D. Bay. `http://kdd.ics.uci.edu`.

[47] A. Hone and J. Kelsey. Optima, extrema and artificial immune systems. In G. Nicosia, V. Cutello, P. Bentley, and J. Timmis, editors, *Proceedings of the 3rd International Conference on Artificial Immune Systems*, volume 3239 of *LNCS*. Springer, 2004.

[48] C. Jacob, M. Pilat, P. Bentley, and J. Timmis, editors. *Artificial Immune Systems*, number 3627 in LNCS. Springer, 2005.

[49] M.T. Jensen. Generating robust and flexible jobshop schedules using genetic algorithms. *IEEE Transactions on Evolutionary Computing*, 7(3):275–288, 2003.

[50] Z. Ji and D Dasgupta. Real-valued negative selection algorithm with variable-sized detectors. In *Genetic and Evolutionary Computation. GECCO-2004, Part I. Volume 3102 of Lecture Notes in Computer Science*, pages 287–298, 2004.

[51] J. Kelsey and J. Timmis. Immune Inspired Somatic Contiguous Hypermutation for Function Optimisation. In *Proceedings of GECCO*, volume 2723 of *LNCS*. Springer, 2003.

[52] J Kelsey, J Timmis, and A Hone. Chasing chaos. In *Congress on Evolutionary Computation*, pages 89–98, Canberra, Australia., December 2003. IEEE.

[53] J Kim and P Bentley. Immune Memory in the Dynamic Clonal Selection Algorithm. In J Timmis and P Bentley, editors, *Proceedings of the First International Conference on Artificial Immune Systems ICARIS*, pages 59–67, 2002.

[54] J Kim, W. Wilson, U Aickelin, and J McLeod. Cooperative automated worm response and detection immune algorithm (cardinal) inspired by t-cell immunity and tolerance. In C Jacob, M. Pilat, P. Bentley, and Timmis J., editors, *Proceedings of 4th International Conference on Artificial Immune Systems*, volume 3627, pages 168–181, 2005.

[55] A. Ko, H.Y.K. Lau, and T.L. Lau. An immuno control framework for cedcentralized mechatronic control. In *Artificial Immune Systems: Proceedings of ICARIS 2004*, pages 46–57. Springer, 2004.

[56] O. Nasraoui, C. Rojas, and C. Cardona. A framework for mining evolving trends in web data streams using dynamic learning and retrospective validation. *Computer Networks*, 50(10):1425–1429, July 2006.

[57] M. Neal. Meta-stable memory in an artificial immune network. In *Artificial Immune Systems: Proceedings of ICARIS 2003*, pages 168–181. Springer, 2003.

[58] J. Newborough and S. Stepney. A generic framework for population-based algorithms, implemented on multiple fpgas. 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*, page 43:55. Springer, 2005.

[59] Z.X. Ong, J.C. Tay, and C.K. Kwoh. Applying the clonal selection principle to find flexible job-shop schedules. In Jacob et al. [48], pages 442–455.

[60] G. Orosz, C. An introduction to immuno-ecology and immuno-informatics. In L.A. Segel and I. Cohen, editors, *Design Principles for the Immune Systems and Other Distributed Systems*, pages 125–150. Oxford University Press, 2001.

[61] A. S. Perelson. Immune network theory. *Immunological Review*, 110:5–36, 1989.

[62] J.C.L. Pinto and F. Von Zuben. Fault detection algorithm for telephone systems based on the danger theory. In Jacob et al. [48], pages 418–431.

[63] B Scholkopf, J. Platt, A ShaweTaylor, A. Smola, and A Williamson. Estimating the support of a high-dimensional distribution. Technical Report MSR-TR-99-87, 1999.

[64] A. Secker, A. Freitas, and J. Timmis. AISEC: An Artificial Immune System for E-mail Classification. In *Proceedings of the Congress on Evolutionary Computation*, pages 131–139. IEEE, December 2003.

[65] L Segel and I Cohen, editors. *Design Principles for the Immune System and Other Distributed Systems*. Oxford University Press, 2001.

[66] L.A. Segel. Diffuse feedback from a diffue informational network: In the immune system and other distributed autonomous systems. In L.A. Segel and I. Cohen, editors, *Design Principles for the Immune Systems and Other Distributed Systems*, pages 203–226. Oxford University Press, 2001.

[67] A Somayaji, S Hofmeyr, and S. Forrest. *Design Principles for the Immune System and Other Distributed System*, chapter Principles of a Computer Immune System. 2001, Oxford University Press.

[68] S. Stepney. Embodiment. In D. Flower and J. Timmis, editors, *In Silco Immunology*, chapter Embodiment. Springer, 2006.

[69] S. Stepney, R. Smith, J. Timmis, and A. Tyrrell. Towards a Conceptual Framework for Artificial Immune Systems. In *LNCS 3239*, pages 53–64. Springer, 2004.

[70] T. Stibor, K. Bayarou, and C. Eckert. An investigation of R-chunk detector generation on higher alphabets. In *LNCS 3102*, pages 26–30, 2004.

[71] T. Stibor, P. Mohr, J. Timmis, and C. Eckert. Is negative selection algorithm suitable for anomaly detection? In *Proceedings of GECCO 2005. LNCS*, 2005.

[72] 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. Springer, 2005.

[73] J.C. Tay and D. Wibowo. Genace: An efficient cultural algorithm for solving the flexible job-shop problem. In *Proceedings of the IEEE Congress on Evolutionary Computation*, pages 1759–1766, 2004.

[74] D. Taylor and D. Corne. An investigation into negative selection algorithm for fault detction in refrigeration systems. In *Artificial Immune Systems: Proceedings of ICARIS 2003*, pages 34–45. Springer, 2003.

[75] J. Timmis, C. Edmonds, and J. Kelsey. Assessing the performance of two immune inspired algorithms and a hybrid genetic algorithm for optmisation. In *Proceedings of Genetic and Evolutionary Computation Conference, GECCO 2004*, pages 308–317. Springer, 2004.

[76] J. Timmis and M. Neal. Once more unto the breach: Towards artificial homeostasis. *Recent Developments in Biologically Inspired Computing*, pages 340–365, 2005.

[77] M. Villalobos-Arias, C.A. Coello Coello, and O. Hernandez-Lerma. Convergence analysis of a multiobjective artificial immune system algorithm. In *Lecture Noted in Computer Science 3239*, pages 226 – 235, 2004.

[78] J. Walker and S. Garrett. Dyanmic function optimisation: Comparing the performance of clonalg and evolution strategies. In *Artificial Immune Systems: Proceedings of ICARIS 2003*, pages 273–285. Springer, 2003.

[79] A. Watkins. *Exploiting Immunological Metaphors in the Development of Serial, Parallel and Distributed Learning Algorithms*. PhD thesis, University of Kent, Computing Laboratory, 2005.

[80] A. Watkins and J.. Timmis. Exploiting the parallelism inherent in airs, and artificial immune classifier. In *Artificial Immune Systems: Proceedings of ICARIS 2004*, pages 427–438. Springer, 2004.

[81] D. Wolpert and W. Macready. No free lunch theorems for optimisation. *IEEE Transactions on Evolutionary Computing*, 1997.

[82] X. Yao, Y. Liu, and G.M. Lin. Evolutionary programming made faster. *Evolutionary Computation*, pages 82–102, 1999.

[83] Y. Zhong, L. Zhang, B. Huang, and P. Li. An unsupervised artificial immune classifier for multi/hyperspectral remote sensing imagery. *IEEE Transactions on Geoscience and Remote Sensing*, 44(2):420–431, February 2006.