Application Areas of AIS: The Past, The Present and The Future

Emma Hart[1] and Jonathan Timmis[2]

[1] School of Computing, Napier University
e.hart@napier.ac.uk
[2] Departments of Electronics and Computer Science, University of York
jt@york.ac.uk

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.

1 Introduction

The 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 naive metaphors are 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 no exemplars that really stand out as instances of successfully applying an AIS to a hard, real-world problems, or of AIS being used 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 [5], there were 38 papers describing the applications of EAs

C. Jacob et al. (Eds.): ICARIS 2005, LNCS 3627, pp. 483–497, 2005.

to real-world problems, and the EVONET repository [3] 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 focussing 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 [47], 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 thus make it distinctive. 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 [38,22,10], 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.

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 [14], used in a tutorial at ICARIS 2004 [15] on Engineering Application of AIS. The information contained in this tutorial has been expanded to include references from ICARIS 2004 [4] and is available from [1]. A useful summary of application areas can also be found in [16] 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.

Figure 1 therefore shows a summary of 97 papers which have been classified into 12 headings. Note that 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 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. Also, in several cases there are multiple papers published over a period of time by the same authors on the same application; in this case, only one paper per author is included in the list, as the intention is to reflect the diversity of applications and give some indication of the effort being directed towards a particular application area.

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 [41] to aircraft fault detection [13]. 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 sub-topics 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 used 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 [26,33].

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 algorithms — again the data being clustered is static in nature.

Combinatoric Optimisation covers a number of real-world application areas such as travelling salesman problems, scheduling (including inventory and job-shop scheduling), and routing problems. Typically, the publications report results on benchmark problem instances rather than real-world problem instances.

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 [28] attempts to lay a foundation for using an AIS to provide the basis of an architecture for a robot to acquire

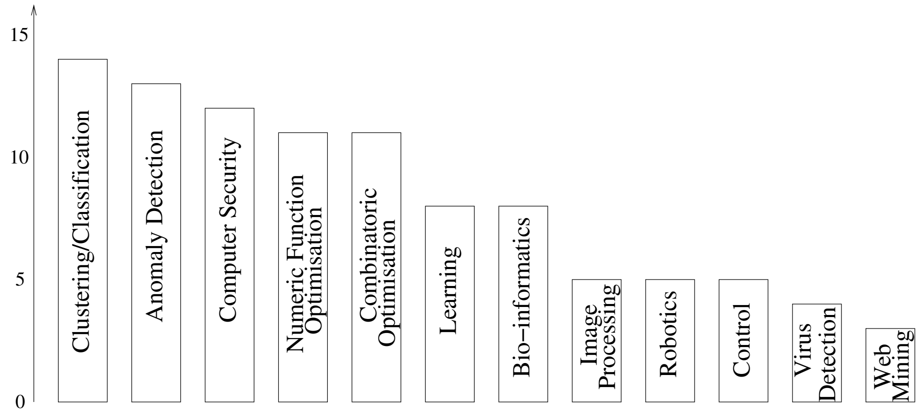

Fig. 1. Summary of Application Areas of AIS

new, more complex skills throughout its lifetime. *Adaptive control systems* form a related category of papers, for example pertaining to controlling a robotic arm [32]. The small topics of *Image Processing* and *Web-Mining* are self-evident.

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. [9], there has been a tendency to *reason by metaphor* [38], 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 [12] 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 in 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 [21] led to a great deal of research and proposal of immune inspired anomaly detection systems [20]. 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 [31], 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 [25,29], 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 [19], 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 [39,40] 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) [34]. When comparing the work of [29], (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 [29] 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.

So, from a "value added" perspective, at present it is not clear from the literature that the immune approach offers anything. It is necessary to use two classes of data to train and tune the system, a high false positive rate seems to blight systems and the computational complexity of generating detectors seems prohibitive in large dimensional data sets. In order to overcome some of these shortfalls, work proposed in [6] and later expanded on in [7] 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, however these ideas have yet to be proven in practice. Therefore, despite that the fact that at first glance, anomaly detection does appear to map to many of the features in the list given at the start of this section; i.e the problems are often distributed in nature, require feature extraction, recognition, memory and continuous learning, immunology has not yet provided all the answers.

3.2 Optimisation

A number of publications relate to 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 [17], opt-AINET [18] and the B-Cell algorithm [42]. Thus, for example, [11] applies Clonalg with a variety of modified hyper-mutation operators to solving static 'trap functions" — complex but toy problems often used in evolutionary algorithm trap investigations, and [42] compare versions of opt-AINET and the B-Cell algorithm to a variety of optimisation functions of various dimensions found in the literature.

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. We further conjecture that the only two features of Dasgupta's list that recommended immune-algorithms as a mechanism for performing function optimisation are that the algorithms require a diversity mechanism and a memory mechanism — however, these features are common components of many other algorithms. Therefore, we conjecture that there is no added value in applying an immune algorithm to static function optimisation problems. Admittedly, the B-Cell algorithm described has been found to use significantly fewer evaluations than a hybrid GA on some problems [42], however, we hypothesise that static function optimisation will not prove to be the Holy Grail of immune algorithms. Similar arguments apply to the use of AIS in Combinatorial Optimisation Problems. Thus, although AIS algorithms have provided superior results on benchmark job-shop scheduling problems when compared to other state-of-the art optimisation algorithms such as GRASP, these are again static problems, in which there is no obvious benefit to be gained from applying an 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 [23] used a network-based AIS to perform dynamic function optimisation. Walker *et al* [45] 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* [30] have adapted the B-Cell algorithm 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 also hypothesise that continuing research effort will reveal little of value; the immune system is not a natural model for extracting metaphors to perform optimisation.

There is perhaps a caveat to the above statements. We are aware of work by Clark *et al* who have produced a theoretical analysis of the B-Cell algorithm discussed above. We believe this paper is in review for ICARIS 2005. This work 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. In addition, work by [44] provides a complete proof for their multi-objective immune inspired algorithm. Thus, as there have been no convincing theoretical analyses that enable performance prediction in the EA world, there is perhaps value in applying a properly understood algorithm to a problem, regardless of the nature of the problem.

3.3 Clustering and Classification

Immune-based algorithms which perform clustering make up a large number of the application areas shown in figure 1. These range from supervised algorithms such as AIRS [46] and Carter, to aiNET [18] and algorithms based on idiotypic network models such as those of Neal and Timmis [33]. 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 datasets, 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 [46] in a parallel version of AIRS.

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 [33] is meta-stable in that it can in theory be continuously applied to a data-set. The work of Hart [26] 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 [35] developed a dynamic supervised learning algorithm for the filter of emails, and work by Kim and Bentley [31] a dynamic classification algorithm for use in intrusion detection.

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 [38] proposes a conceptual framework that allows for the development of more biologically grounded AIS, through the adoption of an inter-disciplinary approach. Metaphors employed have typically been simple, but somewhat effective. However, as proposed in [38], 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 [22] 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 [27] 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. [36], Bersini and Varela [10] and Somayaji *et al* [37] (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.

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.

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 [24], 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.

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.

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 — [43] — 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.

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

An exciting example which represents a step forward in this direction is work currently in progress at the University of Kent, which proposes 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 through the development of an immune inspired continuous learning approach for updating the set of error detectors in a system. 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. The system 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.

Acknowledgements. 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 network [2].

References

1. `http://www.dcs.napier.ac.uk/~emmah/research_interests/AIS.html`.
2. Artist - uk based network, funded by esprc.
3. `http://evonet.lri.fr`.

4. *Artificial Immune Systems: Proceedings of ICARIS 2004*. Springer, 2004.
5. *Proceedings of Genetic and Evolutionary Computation Conference*. Springer, 2004.
6. Uwe Aickelin and Steve 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.
7. 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.
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.*, page 310, 2002.
9. H. Bersini. Immune network and adaptive control. In *Proceedings of 1st Conference on Artificial Life*, pages 217–226. MITpress, 1991.
10. H Bersini and F Varela. *The Immune Learning Mechansims: Recruitment, Reinforcement and their Applications*. Chapman Hall, 1994.
11. 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.
12. D. Dasgupta, editor. *Artifificial Immune Systems and their Applications*. Springer, 1999.
13. 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.
14. L. De Castro. Artificial immune systems bibliography. `http://www.dca.fee.unicamp.br/ lnunes/AIS.html`.
15. L. De Castro. Engineering applications of artificial immune systems. Tutorial at ICARIS 2004, available from `http://artificial-immune-systems.org/ICARIS2004/icaris2004.htm`, 2004.
16. 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.
17. L. De Castro and F.J. Von Zuben. The clonal selection algorithm with engineering applications. In *GECCO 2002 - Workshop Proceedings*, pages 37–37, 2000.
18. 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.
19. 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.
20. S Forrest, S Hofmeyr, and A Somayaji. Computer Immunology. *Communications of the ACM*, 40(10):88–96, 1997.
21. 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.
22. A Freitas and J Timmis. Revisiting the Foundations of Artificial Immune Systems: A Problem Oriented Perspective. In *LNCS 2787*, pages 229–241. Springer, 2003.
23. 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.

24. R. Germain. An innately interesting decade of research in immunology. *Nature Medicine*, 10(12):1307–1320, 2004.
25. 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*, page 261272, 2003.
26. 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.
27. E Hart and P Ross. Studies on the implications of shape-space models for idiotypic networks. In *LNCS 3239*, pages 413–427. Springer, 2004.
28. 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.
29. 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*, page 287298, 2004.
30. J Kelsey, J Timmis, and A Hone. Chasing chaos. In *Congress on Evolutionary Computation*, pages 89–98, Canberra, Australia., December 2003. IEEE.
31. 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.
32. 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.
33. M. Neal. Meta-stable memory in an artificial immune network. In *Artificial Immune Systems: Proceedings of ICARIS 2003*, pages 168–181. Springer, 2003.
34. 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.
35. 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.
36. L Segal and I Cohen, editors. *Design Principles for the Immune System and Other Distributed Systems.* Oxford University Press, 2001.
37. 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.
38. 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.
39. T. Stibor, K. Bayarou, and C. Eckert. An investigation of R-chunk detector generation on higher alphabets. In *LNCS 3102*, pages 26–30, 2004.
40. T. Stibor, P. Mohr, J. Timmis, and C. Eckert. Is negative selection algorithm suitable for anomaly detection? In *Proceedings of GECCO 2005. LNCS*, 2005.
41. 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.
42. 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.
43. J. Timmis and M. Neal. Once more unto the breach: Towards artificial homeostasis. *Recent Developments in Biologically Inspired Computing*, pages 340–365, 2005.

44. 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.
45. 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.
46. 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.
47. D. Wolpert and W. Macready. No free lunch theorems for optimisation. *IEEE Transactions on Evolutionary Computing*, 1997.