An Overview of Artificial Immune Systems

J. Timmis[1*], T. Knight[1], L.N. de Castro[2] and E. Hart[3]

[1]Computing Laboratory,
University of Kent.
Canterbury.
UK.
{jt6,tpk1}@ukc.ac.uk

[2]School of Electrical and
Computer Engineering,
State University of Campinas,
Brazil.
lnunes@dca.fee.unicamp.br

[3]School of Computing,
Napier University.
Edinburgh. Scotland.
UK.
e.hart@napier.ac.uk

ABSTRACT

The immune system is highly distributed, highly adaptive, self-organising in nature, maintains a memory of past encounters and has the ability to continually learn about new encounters. From a computational point of view, the immune system has much to offer by way of inspiration to computer scientists and engineers alike. As computational problems become more complex, increasingly, people are seeking out novel approaches to these problems, often turning to nature for inspiration. A great deal of attention is now being paid to the vertebrae immune system as a potential source of inspiration, where it is thought that different insights and alternative solutions can be gleaned, over and above other biologically inspired methods.

Given this rise in attention to the immune system, it seems appropriate to explore this area in some detail. This survey explores the salient features of the immune system that are inspiring computer scientists and engineers to build Artificial Immune Systems (AIS). An extensive survey of applications is presented, ranging from network security to optimisation and machine learning. However, this is not complete, as no survey ever is, but it is hoped this will go some way to illustrate the potential of this exciting and novel area of research.

1 Introduction

This contribution examines the growing field of Artificial Immune Systems (AIS). Artificial immune systems can be defined as computational systems inspired by theoretical immunology and observed immune functions, principles and models, which are applied to problem solving (de Castro & Timmis, 2002). The field of AIS is relativity new and draws upon work done by many theoretical immunologists (e.g., Jerne, 1974; Perelson, 1989; and Bersini & Varela, 1990) to name a few. What is of interest to researchers developing AIS is not the modelling of the immune system, but extracting or gleaning of useful mechanisms that can be used as metaphors or inspiration to help in the development of (computational) tools for solving particular problems. Within biologically inspired computing, it is quite common to see gross simplifications of the biological systems, on which the artificial systems are based: AIS is no exception. However, it should be remembered, that although a good understanding of the biological system is essential in this domain, it is inspiration from nature that is sought, rather than the creation of accurate models.

Through reading the literature, it can be observed that AIS have been applied to a wide range of application domains. Some of the first work in applying immune system metaphors was undertaken in the area of fault diagnosis (Ishida, 1990). Later work applied immune system metaphors to the field of computer security and virus detection (Forrest *et al*, 1994), which seemed to act as a catalyst for further investigation of the immune system as a metaphor in many areas. However, as yet, there seems to be no *niche* area for AIS. Some people have commented that this may be a weakness, or a gap in the field (Bersini, 2002) and that there needs to be a serious undertaking to find such a *niche* area and this will in turn go to strengthen the area. It has also be argued that AIS are incredibly flexible, as are many biologically inspired techniques, suitable for a number of applications and can be thought of

as a novel soft computing paradigm, suitable for integration with many more *traditional* techniques (de Castro and Timmis, 2003).

The growing interest in AIS is reflected in the growing number of special sessions and invited tracks at a number of well-established international conferences, such as the IEEE SMC and GECCO conferences. The first international conference on artificial immune systems (ICARIS) took place at the University of Kent at Canterbury (UKC) in September 2002[1]. Its great success in terms of organization and quality of papers presented motivated the second ICARIS to be held in Edinburgh in September 2003.

This chapter is organised in the following manner. First, reasons for why the immune system has generated such interest within the computing and engineering community. This is followed by a simple review of relevant immunology that has served as a foundation for much of the work reviewed in this contribution. Immunology is a vast topic and no effort has been made to cover the whole area, suitable citations are provided in the text to further direct the reader. The area of AIS is then presented, in terms of a general framework proposed in (de Castro and Timmis, 2002). A review of AIS applications is then presented, however, providing a general overview of a number of different application areas. Finally, comments on the perceived future of this emerging technology are then presented.

2 The Immune System: Metaphorically Speaking

When considered from a computational point of view, the immune system can be considered to be a rich source of inspiration as it displays learning, adaptability, is self-organising, highly distributed and displays a memory There are many reasons why the immune system is of interest to computing (Dasgupta, 1998b; de Castro & Timmis, 2002); these can be summarised as follows:

- *Recognition*: The immune system has the ability to recognise, identify and respond to a vast number of different patterns. Additionally, the immune system can differentiate between malfunctioning *self-cells* and harmful *nonself* cells, therefore maintaining some sense of self.
- *Feature Extraction*: Through the use of Antigen Presenting Cells (APC) the immune system has the ability to extract features of the antigen by filtering molecular noise from disease causing agents called an *antigen,* before being presented to other immune cells, including the lymphocytes.
- *Diversity*: There are two major processes involved in the generation and maintenance of diversity in the immune system. First, is the generation of receptor molecules through the recombination of gene segments from gene libraries. By recombining genes from a finite set, the immune system is capable of generating an almost infinite number of varying types of receptors, thus endowing the immune system with a large coverage of the universe of antigens. The second process, which assists with diversity in the immune system, is known as somatic hypermutation. Immune cells reproduce themselves in response to invading antigens. During reproduction, they are subjected to a somatic mutation process with high rates that allow the creation of novel patterns of receptors molecules, thus increasing the diversity of the immune receptors (Kepler & Perelson, 1993).
- *Learning*: The mechanism of somatic hypermutation followed by a strong selective pressure also allows the immune system to fine-tune its response to an invading pathogen; a process termed *affinity maturation* (Berek & Ziegner, 1993). Affinity maturation guarantees that the immune system becomes increasingly better at the task of recognising

* Corresponding author
[1] http://www.aber.ac.uk/icaris-2002

patterns. The immune network theory is another powerful example of learning in the immune system. It suggests that the immune system has a dynamic set of mutually recognizing cells and molecules, and the presence of an invading antigen causes a perturbation in this network. As a result, the dynamic immune network, which presents an intrinsic steady state in the absence of antigens, has to self-organise its pattern of behaviour again, so as to accommodate the disturbance (Bersini, 2002). Therefore, invading antigens require the immune network to adapt itself to this new element.

- *Memory*: After an immune response to a given antigen, some sets of cells and molecules are endowed with increased life spans in order to provide faster and more powerful immune responses to future infections by the same or similar antigens. This process, known as the maturation of the immune response, allows the maintenance of those cells and molecules successful at recognizing antigens. This is the major principle behind vaccination procedures in medicine and immunotherapy. A weakened or dead sample of an antigen (e.g., a virus) is inoculated into an individual so as to promote an immune response (with no disease symptoms) in order to generate memory cells and molecules to that antigen.
- *Distributed detection*: There is inherent distribution within the immune system. There is no one point of overall control; each immune cell is specifically stimulated and responds to new antigens that can invade the organism in any location.
- *Self-regulation*: Immune systems dynamics are such that the immune system population is controlled by local interactions and not by a central point of control. After a disease has been successfully combated by the immune system, it returns to its normal steady state, until it is needed in response to another antigen. The immune network theory explicitly accounts for this type of self-regulatory mechanism.
- *Metadynamics*: The immune system is constantly creating new cells and molecules, and eliminating those that are too old or are not being of great use. Metadynamics is the name given to this continuous production, recruitment and death of immune cells and molecules (Varela *et al*, 1988).
- *Immune Network*: In 1974 N. Jerne proposed the immune network theory as an alternative to explain how the immune system works. He suggested that the immune system is a dynamic system whose cells and molecules are capable of recognizing each other, thus forming an internal network of communication within the organism. This network provides the basis for immunological memory to be achieved, via a self-supporting and self-organising network.

The remainder of this chapter outlines some of the salient features of the immune system that have been employed in the development of artificial immune systems. Attention is then drawn to significant applications of the immune system as a metaphor for computational systems.

3 The Vertebrate Immune System

The vertebrate immune system is composed of diverse sets of cells and molecules that work in collaboration with other bodily systems in order to maintain a steady state within the host. A role of the immune system is to protect our bodies from infectious agents such as viruses, bacteria, fungi and other parasites. On the surface of these agents are antigens that allow the identification of the invading agents (pathogens) by the immune cells and molecules, thus provoking an immune response. There are two basic types of immunity, innate and adaptive. Innate immunity (Janeway, 1993) is not directed towards specific invaders into the body, but against any pathogens that enter the body. The innate immune system plays a vital role in the initiation and regulation of immune responses, including adaptive immune responses. Specialized cells of the innate immune system evolved so as to recognize and bind to common molecular patterns found only in microorganisms, but the innate immune system is by no means a complete solution to protecting the body.

Adaptive or acquired immunity (Roitt, 1997), however, allows the immune system to launch an attack against any invader that the innate system cannot remove. The adaptive system is directed against specific invaders, and is modified by exposure to such invaders. The adaptive immune system mainly consists of lymphocytes, which are white blood cells, more specifically B and T-cells. These cells aid in the process of recognizing and destroying specific substances. Any substance that is capable of generating such a response from the lymphocytes is called an antigen or immunogen. Antigens are not the invading microorganisms themselves; they are substances such as toxins or enzymes in the microorganisms that the immune system considers foreign. Adaptive immune responses are normally directed against the antigen that provoked them and are said to be antigen-specific. The immune system generalizes by virtue of the presence of the same antigens in more than one infectious agent. Many immunizations exploit this by presenting the immune system with an innocuous organism, which carries antigens present in more dangerous organisms. Thus the immune system learns to react to a particular pattern of antigen.

The immune system is said to be adaptive, in that when an adaptive immune response is elicited B-cells undergo cloning in an attempt to produce sufficient antibodies to remove the infectious agent (Burnet, 1959; Jerne, 1974). When cloning, B-cells undergo a stochastic process of somatic hypermutation (Kepler & Perelson, 1993) where an attempt is made by the immune system to generate a wider antibody repertoire so as to be able to remove the infectious agent from the body and prepare the body for infection from a similar but different infection at some point in the future.

After the primary immune response, when the immune system first encounters a foreign substance and the substance has been removed from the system, a certain quantity of B-cells remain in the immune system and acts as an immunological memory (Smith *et al*, 1998; Jerne, 1974). This is to allow for the immune system to launch a faster and stronger attack against the infecting agent, called the secondary immune response.

3.1 Primary and Secondary Immune Responses

A primary response (Tizard, 1988a) is provoked when the immune system encounters an antigen for the first time. A number of antibodies will be produced by the immune system in response to the infection, which will help to eliminate the antigen from the body. However, after a period of days the levels of antibody begin to degrade, until the time when the antigen is encountered again. This secondary immune response is said to be specific to the antigen that first initiated the immune response and causes a very rapid growth in the quantity of B-cells and antibodies. This second, faster response is attributed to *memory* cells remaining in the immune system, so that when the antigen, or similar antigen, is encountered, a new immunity does not need to be built up, it is already there. This means that the body is ready to combat any re-infection. Figure 1 illustrates this process.

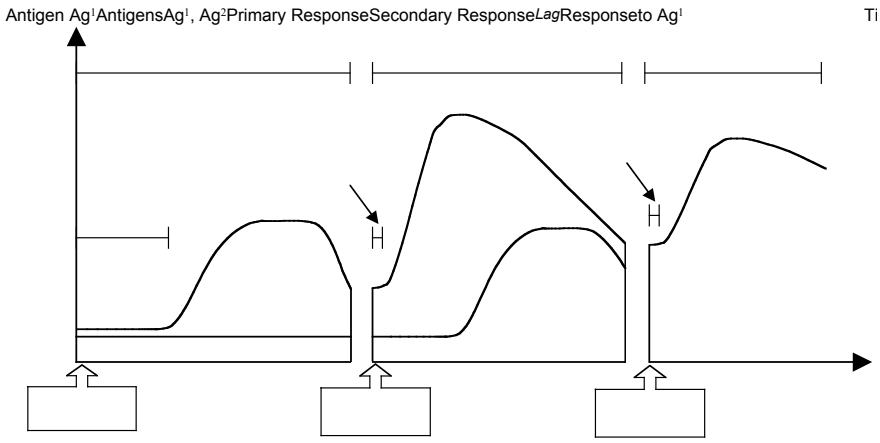


Figure 1 - Primary and Secondary Immune Response. Ag1 infects the system and a lag occurs before a primary immune response is initiated. The host is then re-infected with Ag_1 and a different antigen Ag_2. A fast secondary response is elicited against Ag_1, whilst a primary response is initiated against Ag_2. At some point in the future, the host is then infected with $Ag_{1'}$, which is a slight variation on Ag_1. Due to the generalist capability of the immune system, a secondary response is elicited against the antigen. © De Castro and Timmis, 2002

The amount of antibody is increased by the immune system generating a massive number of B-cells through a process called clonal selection (Burnet, 1959), this is now discussed in relation to the B-cell in the immune system.

3.2 B-cells and Antibodies

The B-cell is an integral part of the immune system. Through a process of recognition and stimulation, B-cells will clone and mutate to produce a diverse set of antibodies in an attempt to remove the infection from the body (Timmis, 2000). The antibodies are specific proteins that recognize and bind to another protein. The production and binding of antibodies is usually a way of signalling other cells to kill, ingest or remove the bound substance (de Castro & Von Zuben, 1999). Each antibody has two paratopes and two epitopes that are the specialised parts of the antibody that identify other molecules (Hunt & Cooke, 1996). Binding between antigens and antibodies is governed by how well the paratopes on the antibody matches the epitope of the antigen, the closer this match, the stronger the bind. Although it is the antibodies that surround the B-cell, which are responsible for recognising and attaching to antigen invaders, it is the B-cell itself that has one of the most important roles in the immune system.

This is not the full story, as B-cells are also affected by Helper T-cells during the immune response (Tizard, 1988b). T-cell paratopes are different from those on B-cells in that they recognise fragments of antigens that have been combined with molecules found on the surfaces of the other cells. These molecules are called MHC molecules (Major Histocompatibility Complex). As T-cells circulate through the body they scan the surfaces of body cells for the presence of foreign antigens that have been picked up by the MHC molecules. This function is sometimes called immune surveillance. These helper T-cells when bound to an antigen secrete interleukines that act on B-cells helping to stimulate them.

3.3 Immune Memory

It is possible to identify two main philosophical avenues that try to explain how immune memory is acquired and maintained (Tew & Mandel, 1979), (Tew *et al*, 1980), (Ada &

Nossal, 1987) and (Matzinger, 1994): 1) clonal expansion and selection and 2) immune network.

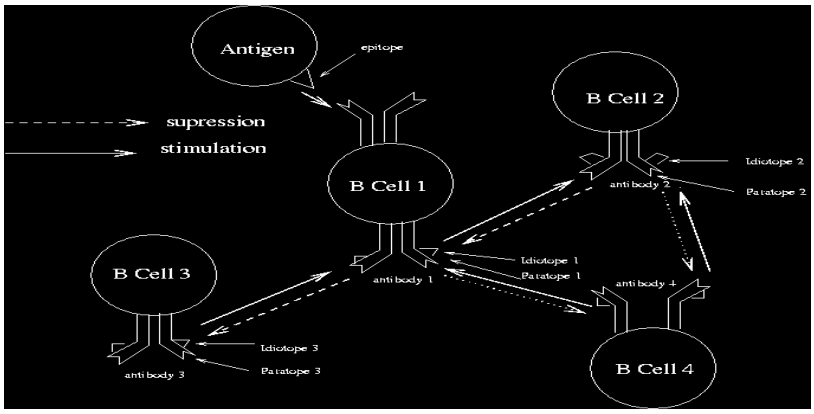
Throughout the lifetime of an individual, it is expected to encounter a given antigen repeatedly. The initial exposure to an antigen that stimulates an adaptive immune response is handled by a spectrum of small clones of B-cells, each producing antibodies of different affinity. The effectiveness of the immune response to secondary encounters is considerably enhanced by storing some high affinity antibody producing cells from the first infection, named memory cells, so as to form a large initial clone for subsequent encounters. Thus memory, in the context of secondary immune responses, is a clonal property (Coutinho, 1989).

Another theory that has been used in AIS for inspiration is the theory first proposed by Jerne (Jerne, 1974) and reviewed in (Perelson, 1989) called the Immune Network Theory. This theory states that B-cells co-stimulate each other via portions of their receptor molecules (idiotopes) in such a way as to mimic antigens. An idiotope is made up of amino acids within the variable region of an antibody or T-cell. A network of B-cells is thus formed and highly stimulated B-cells survive and less stimulated B-cells are removed from the system. It is further proposed that this network yields useful topological information about the relationship between antigens. For these reasons, this section focuses on this theory.

3.3.1 Immunological Memory via the Immune Network

Work in (Jerne, 1974) proposed that the immune system is capable of achieving immunological memory by the existence of a mutually reinforcing network of B-cells. These cells not only stimulate each other but also suppress connected B-cells, though to a lesser degree. This suppression function is a mechanism by which to regulate the over stimulation of B-cells in order to maintain a stable memory.

This network of B-cells occurs due to the ability of paratopes, located on B-cells, to match against idiotopes on other B-cells. The binding between idiotopes and paratopes has the effect of stimulating the B-cells. This is because the paratopes on B-cells react to the idiotopes on similar B-cells, as it would an antigen. However, to counter the reaction there is a certain amount of suppression between B-cells to act as a regulatory mechanism. Figure 2 shows the basic principles of the immune network theory. Here B-cell 1 stimulates three other cells, B-cells 2, 3 and 4, and also receives a certain amount of suppression from each one. This creates a network type structure that provides a regulatory effect on neighbouring B-cells. The immune network acts as a self-organising and self-regulatory system that captures



antigen information ready to launch an attack against any similar antigens.

Figure 2 - Jernes' idiotypic network hypothesis

3.3.1.1 A Proposed Immune Network Model

Attempts have been made at creating immune network models (Farmer *et al*, 1986), (Carneiro & Stewart, 1995) so as to better understand its complex interactions. Work in (Farmer *et al*, 1986) proposed a model to capture the essential characteristics of the immune network as described in (Jerne, 1974) and identify memory mechanisms in it, whereas the work in (Carneiro & Stewart, 1995) observed how the immune system identifies self and non-self. Both work by (Farmer *et al*, 1986) and (Perelson, 1989) investigated Jernes' work in more depth and provided insights into some of the mechanisms involved in the production and dynamics of the immune network. There are a number of immune network models, however it is impossible to review them all. This section will summarise the salient features of the Farmer *et al* model, as a form of case study to illustrate the potential power of such a model for computation.

Work in (Farmer *et al*, 1986) created a simplistic model to simulate the immune system. The model ignored the effect of T cells and of macrophages in an attempt to capture the essential characteristics of the immune network. Central to their work was the calculation of the dynamics of B-cell population related to a B-cell's stimulation level. The authors proposed a simple equation that they consider takes into account the three main contributing factors to B-cell stimulation level, these are: (i) the contribution of the antigen binding (ii) the contribution of neighbouring B-cells and (iii) the suppression of neighbouring B-cells. The rate of change of antibody concentration is given by

$$\frac{dx_i}{dt} = c\left[\sum_{j=1}^{N} m_{j,i}x_i x_j - k_1 \sum_{j=1}^{N} m_{i,j}x_i x_j + \sum_{j=1}^{M} m_{j,i}x_i y_j\right] - k_2 x_i,$$

Equation 1 – Change in antibody concentration (Farmer *et al*, 1986)

where the first term represents the stimulation of the paratope of an antibody type i by the epitope of an antibody j. The second term represents the suppression of antibody of type i when its epitope is recognized by the paratope of type j. The parameter c is a rate constant that depends on the number of collisions per unit time and the rate of antibody production stimulated by a collision. Constant k_1 represents a possible inequality between stimulation and suppression.

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

3.4 Repertoire and Shape Space

Coutinho (1980) first postulated the idea of repertoire completeness. He stated that if the immune systems antibody repertoire is complete, that is, present receptor molecules capable of recognizing any molecular shape, then antibodies with immunogenic idiotopes can be recognised by other antibodies, and therefore an idiotypic network would be created.

However, in order to understand completeness, it is first necessary to understand the concept of shape space. Shape space has been an important mechanism to create and represent abstract models of immune cells and molecules (de Castro and Timmis, 2002). The basic idea is that

all the features of a receptor molecule necessary to characterise its binding region with an antigen is called its generalised shape. The generalised shape of any receptor molecule can be represented by an attribute string of a given length L in a generic L-dimensional space, called shape space.

To illustrate this idea, consider a bi-dimensional space as illustrated in Figure 3. The set of all possible shapes lie within a finite volume V in this bi-dimensional shape space. The antibodies are represented by the letter A (black dots) and the antigens are depicted by the 'x'. Each antibody (A) can recognise a given number of antigens within an affinity threshold ε and therefore can recognise a volume (V_e) of antigens (x) in shape space. Therefore, a finite set of antibodies appropriately placed in the shape space and with appropriate affinity thresholds are sufficient to cover the whole shape space; thus being capable of recognizing any molecular shape that can be presented to the immune system.

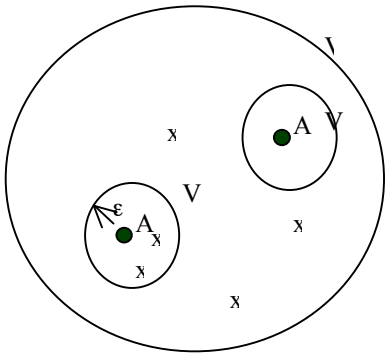


Figure 3 - A diagrammatic representation of shape space. Adapted from (Perelson, 1989).

3.5 Learning within the Immune Network

It has been proposed that the immune network can be thought of as being cognitive (Varela *et al*, 1988) and exhibits learning capabilities. The authors proposed four reasons as to why they consider immune systems to be cognitive: (i) they can recognise molecular shapes; (ii) they remember history of encounters; (iii) they define the boundaries of self, and (iv) they can make inferences about antigenic patterns they have yet to encounter. Taking these points, the paper explores cognitive mechanisms of the immune system and proposes that the immune network can be thought of as a cognitive network, in a similar way to a neural network.

The work suggests that the immune network is capable of producing dynamic patterns of activity over the entire network and that there is a self-regulatory mechanism working that helps to maintain this network structure. These emerging patterns within the immune network are characterised by varying numbers of B-cells that when in a response to an antigen undergo clonal selection. The authors use the term *metadynamics* of the immune system; see also (Bersini & Valera, 1994). This can essentially be taken to mean the continual production and death of immune cells and molecules. A large variety of new B-cells will be produced, but not all will be a useful addition to the immune system and many will never enter into the dynamics of the immune system (interact with other B-cells in the network) and will eventually die. The authors produced a simple model using these ideas and found that there are oscillations in many of the variables within their system, in particular the number of B-cells that are produced. There would often be rapid production of B-cells, followed by a sharp decline in number, which the authors argue, is what you expect to see in the natural immune system. Coupled with this oscillatory pattern, the authors observed that a certain core and stable network structure does emerge over time. This structure emerges due to a topological self-organisation within the network, with the resulting network acting to record the history of encounters with antigens. Therefore, the authors concluded that the immune

system is an excellent system for learning about new items and can support a memory of encounters by the use of complex pattern matching and a self-organising network structure, and can thus be thought of as being cognitive.

There is other research that goes to support the ideas presented above. (Bersini & Varela, 1990) implemented the model proposed in (Varela et al, 1988) and suggested that mechanisms such as immune memory, adaptability and the immune system's ability to perform distributed processing could be of potential use to engineering problem solving, in particular adaptive control (Bersini, 1991) and computational problem solving.
Following on from work in (Bersini & Valera, 1990), work in (Bersini & Valera, 1994) provides an effective summary of work done on exploring the dynamics and metadynamics of the immune system. They claim that the metadynamics of the immune system allows the identity of the immune system to be preserved over time, but still allows itself to adapt to new situations. Simulations of an immune network confirmed this. The reader is also directed to (Bersini, 2002) where further arguments for this position are proposed.

As a way to model the immune system metadynamics the authors proposed the use of the immune recruitment mechanism (IRM). The IRM is a mechanism by which the best new cells and molecules in the system are incorporated into the network. This can be translated as saying that one should only incorporate the best new items that are produced into the network. Therefore the selection of new items is based on the state of the surrounding network: any other items that are produced are lost. This gives rise to the metadynamical system that is believed to occur in the vertebrate immune system. In this paper, the authors proposed seven general principles that can be extracted from the immune system and applied to creating a controlling system for the area of adaptive control, but they hope, to other fields as well. These principles are:

- *Principle 1*: The control of any process is distributed around many operators in a network structure. This allows for the development of a self-organising system that can display emerging properties.
- *Principle 2*: The controller should maintain the viability of the process being controlled. This is keeping the system within certain limits and preventing the system from being driven in one particular way.
- *Principle 3*: While there may be perturbations that can affect the process, the controller learns to maintain the viability of the process through adaptation. This learning and adaptation requires two kinds of *plasticity*: a parametric plasticity, which keeps a constant population of operators in the process, but modifies parameters associated with them; and a structural plasticity which is based on the recruitment mechanism which can modify the current population of operators.
- *Principle 4*: The learning and adaptation are achieved by using a reinforcement mechanism between operators. Operators interact to support common operations or controls.
- *Principle 5*: The dynamics and metadynamics of the system can be affected by the sensitivity of the network.
- *Principle 6*: The immune recruitment mechanism can be considered to be a stand-alone optimisation algorithm.
- *Principle 7*: The controller retains a population-based memory, which can maintain a stable level in a changing environment.

The authors suggest that these principles, while being very general, could prove useful to many domains of learning, engineering control and so on. Indeed, in their paper they present a way of applying these general principles to the areas of adaptive control and to the creation of other immune inspired algorithms.

3.6 The Clonal Selection Principle

When antibodies on a B-cell bind with an antigen, the B-cell becomes activated and begins to proliferate. New B-cell clones are produced that are an exact copy of the parent B-cell, but then undergo somatic hypermutation (Berek & Ziegner, 1993) and produce antibodies that are specific to the invading antigen. The clonal selection principle (Burnet, 1959) is the term used to describe the basic properties of an adaptive immune response to an antigenic stimulus and is an alternative view to the position presented in the previous section. It establishes the idea that only those cells capable of recognizing an antigenic stimulus will proliferate, thus being selected against those that do not. Clonal selection operates on both T-cells and B-cells.

The B-cells, in addition to proliferating or differentiating into plasma cells, can differentiate into long-lived B *memory cells*. Memory cells circulate through the blood, lymph and tissues, probably not manufacturing antibodies (Perelson *et al*, 1978). However, when exposed to a second antigenic stimulus they commence differentiating into large lymphocytes capable of producing high affinity antibody.

3.6.1 Learning and Memory via Clonal Selection

In order for the immune system to be protective over periods of time, antigen recognition is insufficient. The immune system must also have a sufficient number of cells and molecules so as to mount an effective response against antigens encountered at a later stage. The number of immune cells and molecules specific for the antigen with relation to the size of the antigen's population is crucial to determining the outcome of infection. Learning via clonal selection involves raising the population size and the affinity of those cells that have proven themselves to be valuable during the antigen recognition phase. Thus, the immune repertoire is biased from a random base to a repertoire that more clearly reflects the actual antigenic environment.

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

In summary, immune learning and memory are acquired through:

- Repeated exposure to an antigenic stimulus;
- Increase in size of specific immune cells and molecules;
- Affinity maturation of the antigenic receptors.
- Presence of long living cells that persist in a resting state until a second encounter with the antigen.

3.7 Self/Non-Self Discrimination

The immune system is said to be complete: it has the ability to recognize all antigens. Antibodies and T-cell receptors produced by the lymphocytes can recognize any foreign (or self) molecule. Antibody molecules have idiotopes and it follows from the idea of completeness that these will be recognized by other antibody molecules.

Therefore, all molecules (shapes) can be recognized including our own, which are also seen as antigens, or self-antigens. For the immune system to function properly, it needs to be able to distinguish between the molecules of our own cells (*self*) and foreign molecules (*non-self*), which are *a priori* indistinguishable (Perelson & Weisbuch, 1997). If the immune system is

not capable of performing this distinction, then an immune response will be triggered against the self-antigens, causing autoimmune diseases.

An encounter between an antibody and an antigen does not inevitably result in activation of the lymphocyte. It is possible that the encounter could actually cause the death of the lymphocyte. In order for this to happen, there must be some form of negative selection that prevents self-specific lymphocytes from becoming prevalent.

3.7.1 Negative Selection

The concept of a negative signal following certain lymphocyte-antigen interactions, allows for the control of those lymphocytes being anti-self. Negative selection of a lymphocyte describes the process whereby a lymphocyte-antigen interaction results in the death or *anergy* of that lymphocyte. The immune cell is simply purged from the repertoire. Location plays a role in negative selection: the primary lymphoid organs are designed to largely exclude foreign antigens and to preserve the self-antigens, whereas the secondary lymphoid organs are designed to filter out and concentrate foreign material, and to promote co-stimulatory intercellular immune reactions (Zinkernagel & Kelly, 1997).

The negative selection of T-cells has been broadly used by the AIS community as a model to perform anomaly detection. Basically, the negative selection of T-cells that occurs within the thymus is based on the following considerations. The thymus is comprised of a myriad of molecules that primarily present self-molecules to the naïve T-cells (immature T-cells just produced and with no function yet). The interactions of immature T-cells with the self-molecules results in the death of all those naïve T-cells that recognize the self-molecules. This means that only T-cells that do not recognize self-molecules are allowed to survive and become functional T-cells.

4 From Natural to Artificial Immune Systems

The immune system is a valuable metaphor as it is self-organising, highly distributed and has no central point of control. The theoretical aspects summarised above reveal interesting avenues for using the immune system as a metaphor for developing novel computational intelligence paradigms. These can potentially be applied to solve many problems in a wide range of domains, such as data mining, control and anomaly detection, to name a few. Some of these applications will be discussed in the following sections. Some of the interesting immunological aspects can be summarised as follows:

- *Using the idea of self-organisation*. Self-organisation is the ability of a system to adapt its internal structure to the environment without any external supervision. In the case of the immune system, clonal selection followed by affinity maturation and the immune network adapts to new antigens it comes across and ultimately can be said to represent the antigens. This fits in with the general principle 1 described above of having some inherent self-organising structure within a system that will exhibit emerging properties.
- *The primary and secondary immune responses*. It has been shown that more B-cells are produced in response to continual exposure to antigens. This suggests that to learn on data using the immune system metaphor, the data may have to be presented a number of times in order for the patterns to be captured.
- *Using the idea of clonal selection*. As B-cells become stimulated they reproduce in order to create more antibodies to remove the antigen from the system. This causes clusters of B-cells that are similar to appear. Clusters indicate similarity and could be useful in understanding common patterns in data, just as a large amount of specific B-cells in the immune system indicates a certain antigen.
- *Adaptation and diversification*. Some B-cell clones undergo somatic hypermutation. This is an attempt of the immune system to develop a set of B-cells and antibodies

that cannot only remove the specific antigen, but also similar antigens. By using the idea of mutation a more diverse representation of the data being learnt is gained than a simple mapping of the data could achieve. This may be of benefit and reveal subtle patterns in data that may be missed.

- *Knowledge extraction and generalisation*. Somatic hypermutation may be not only beneficial to generalise knowledge, i.e., to reveal subtle patterns in data but, together with a selective event, it might guarantee that those B-cells with increased affinities are selected and maintained as high affinity cells. The continuous processes of mutation and selection (affinity maturation) allow the immune system to extract information from the incoming antigens. Affinity maturation performs a better exploitation (greedy search) of the surrounding regions of the antibodies.
- *The use of a network structure*. The immune network represents an effective way of simulating a dynamic system and achieving memory. This idea could be exploited in helping to maintain a network of B-cells that are creating a model of some data being learnt. Indeed, visualising that network may reveal useful topological information about the network that leads to a greater understanding of the data being modelled.
- *Metadynamics*. The oscillations of immune system variables, such as antibody concentration and B-cell population, as discussed in (Varela *et al*, 1988) indicate that a stable network representative of the data being learnt could be possible. This would be very useful as once a pattern had been learnt, it would only be forgotten if it becomes useless in a far future. Additionally, the networks produced act as a life long learning mechanism, with B-cell population always in a state of flux, but representative of antigens it has been exposed to. This could be a useful metaphor for developing a system that could, in principle, learn a set of patterns in one data set, then go onto learn new patterns from other data sets, while still remembering the older ones.
- *Knowledge of self and non-self.* The immune system has a complete repertoire in its ability to recognise invading antigens. Additionally, the immune system is said to be tolerant to self, in that it can recognise the difference between self and non-self cells. This is a powerful metaphor when considering anomaly detection systems.

4.1 Summary

Immunology is a vast topic therefore; this chapter has introduced only those areas of immunology that are pertinent to this contribution. Through a process of matching between antibodies and antigens and the production of B-cells through clonal selection (Burnet, 1959) and somatic hypermutation (Kepler & Perelson, 1993), an immune response can be elicited against invading antigen so that it is removed from the system. In order to remember which antigens the immune system has encountered, some form of immunological memory must be present; this can be explained in part through theories such as the clonal selection theory or the more controversial immune network theories. Clearly, the immune system is performing a very important role within the body. The sheer complexity of the system is staggering, and current immunology only knows part of the story. Through complex interactions, the immune system protects out bodies from infection, interacts with other bodily systems to maintain a steady state (homeostasis). The focus of this chapter has been more on the immune network theory. This is not to lend more weight to that particular view point of the immune system, it has merely been presented in more depth to provide the reader with a deeper insight into one of the many complex ideas within immunology, that have helped computer scientists and engineers over the years. This area will now be examined in more detail.

5 The Immune System Metaphor

This section introduces the reader to the field of Artificial Immune Systems (AIS). There have been a number of attempts over the years to try and define exactly what is an AIS. For

example, (Timmis, 2000) defined AIS to be *"an AIS is a computational system based upon metaphors of the natural immune system"* and (Dasgupta, 1998b) defined them to be *"AIS are intelligent methodologies inspired by the immune system toward real-world problem solving"*. Feeling that neither of these definitions were complete, the most recent definition is taken from (de Castro and Timmis, 2002), where they define AIS to be *"AIS are adaptive systems, inspired by theoretical immunology and observed immune functions, principles and models, which are applied to problem solving"*. In this latest definition, the a more complete view of what an AIS has been captured: the fact they are inspired by the immune system, but the inspiration is not restricted to purely theoretical immunology, but also 'wet lab' type immunology, the systems are *adaptive* which means they must demonstrate some element of adaptability and are not restricted to pieces of software but could equally be implemented on hardware and that there is some form of application ultimately in mind – this allows for the distinction between the creation of pure models of the immune system (which indeed are useful for AIS, as has been discussed).

This section presents an overview of many different applications of AIS that can be seen in the literature. No attempt has been made on an exhaustive survey, for this, the readers are directed to (de Castro and Timmis, 2002), chapter 4, where such an exhaustive review is presented. The aim of this section is to merely illustrate the wide applicability of AIS. Very recently, (de Castro and Timmis, 2002) have proposed the idea of a framework for AIS which consists of basic components and processes, from which it is possible to both describe and build AIS. This framework is now presented – due the fact that it was only recently proposed however; the framework has not been used in this article when describing AIS literature published before the existence of this framework.

5.1 A Framework for AIS

In an attempt to create a common basis for AIS, work in (de Castro and Timmis, 2002) proposed the idea of a framework for AIS. The authors argued the case for proposing such as framework from the standpoint that in the case of other biologically inspired approaches, such as artificial neural networks (ANN) and evolutionary algorithms (EAs) such a basic idea exists and helps considerably with the understanding and construction of such systems. For example, (de Castro and Timmis, 2002) consider a set of artificial neurons, which can be arranged together so as to form an artificial neural network. In order to acquire *knowledge*, these neural networks undergo an adaptive process, known as learning or training, which alters (some of) the parameters within the network. Therefore, the authors argued that in a simplified form, a framework to design an ANN is composed of a set of artificial neurons, a pattern of interconnection for these neurons, and a learning algorithm. Similarly, the authors argued that in evolutionary algorithms, there is a set of "artificial chromosomes" representing a population of individuals that iteratively suffer a process of reproduction, genetic variation, and selection. As a result of this process, a population of evolved artificial individuals arises. A framework, in this case, would correspond to the genetic representation of the individuals of the population, plus the procedures for reproduction, genetic variation, and selection. Therefore, the authors adopted the viewpoint that a framework to design a biologically inspired algorithm requires, at least, the following basic elements:

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

Adopting this approach, (de Castro and Timmis, 2002) proposed such a framework for AIS. The basis of the proposed framework for is therefore a representation to create abstract

models of immune organs, cells, and molecules, a set of functions, termed affinity functions, to quantify the interactions of these "artificial elements", and a set of general-purpose algorithms to govern the dynamics of the AIS.

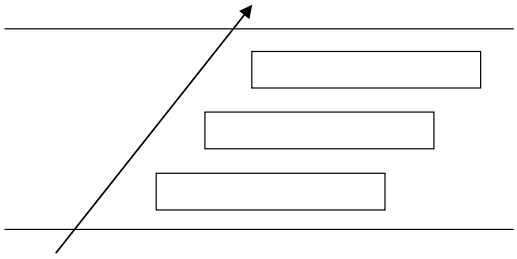


Figure 4 - A Framework for AIS © De Castro and Timmis, 2002

The framework can be thought of as a layered approach as shown in Figure 4. |In order to build a system, one typically requires an application domain or target function. From this basis, the way in which the components of the system will be represented will be considered. For example, the representation of network traffic may well be different that the representation of a real time embedded system. Once the representation has been chosen, one or more affinity measures are used to quantify the interactions of the elements of the system. There are many possible affinity measures (which are partially dependent upon the representation adopted), such as Hamming and Euclidean distances. The final layer involves the use of algorithms, which govern the behavior (dynamics) of the system. Here, in the original framework proposal, algorithms based on the following immune processes were presented: negative and positive selection, clonal selection, bone marrow, and immune network algorithms. It is not possible to explore these here in any detail, needless to say that each algorithm has its own particular use, or more than one use. For example, the immune network model proposed in the framework has been successfully applied to data mining (de Castro and von Zuben, 2000b) and with slight adaptations, multi-modal optimisation (de Castro and Timmis, 2002a)

5.2 Machine Learning

5.2.1 Recognising DNA

The past number of years has seen a steady increase in attempting to apply the immune metaphor to machine learning (de Castro & Von Zuben, 2000a). Amongst the first was that performed by (Cooke & Hunt, 1995) and (Hunt & Cooke, 1996). In these papers, the authors describe their attempts to create a supervised machine learning mechanism to classify DNA sequences as either promoter or non-promoter classes, by creating a set of antibody strings that could be used for this purpose. Work had already been done on this classification problem using different approaches such as C4.5 (Quinlan, 1993) standard neural networks and a nearest neighbour algorithm (Kolodner, 1993) The authors claimed that the AIS system they achieved an error rate of only 3% on classification, which, when compared to the other established techniques yielded superior performance. The system created used mechanisms such as B-cells and B-cell stimulation, immune network theory, gene libraries, mutation and antibodies to create a set of antibody strings that could be used for classification. Central to the work was the use of the Immune Network theory (Jerne, 1974).

Both work in (Hunt *et al*, 1995) and (Hunt & Fellows, 1996) attempted to apply this algorithm to the domain of case base reasoning. In this paper, the authors proposed creating a case memory organisation and case retrieval system based on the immune system. (Hunt *et al*, 1996) took the application to case base reasoning and attempted to apply it directly to data mining. In the previous work (Hunt *et al*, 1995), only cases were explicitly represented, with

no variations, but as indicated by the authors in (Hunt *et al*, 1996), a desirable property of any case base system is the ability to generalise; that is, to return a case that is a general solution if no specific solution is available. As the immune system creates generality in the fight against infection, the authors used this as inspiration to create the idea of a general case, which would attempt to identify trends in data, as opposed to simply the data themselves. By introducing the idea of a generalised case, the authors created a system that could help in the customer-profiling domain; specifically, identifying people who are likely to buy a Personal Equity Plan (PEPs) which were a tax-free investment available at the time.

5.2.2 Fraud Detection

This algorithm was then applied to fraud detection, (Hunt *et al*, 1996), (Hunt *et al* 1998) and (Neal *et al* 1998). Work in (Hunt *et al*, 1996) simply proposed the idea that an AIS could be used to create a visual representation of loan and mortgage application data that could in some way aid the process of locating fraudulent behaviour. An attempt at creating such a system was proposed in (Hunt *et al*, 1998). This system, called JISYS, did not differ substantially from that described in (Hunt *et al*, 1996) apart from the application and the inclusion of more sophisticated string matching techniques, such as trigram matching and the inclusion of weighting in order of importance various fields in the B-cell object, taken from the weighted nearest neighbour idea (Kolodner, 1993).

5.2.3 Back to Basics

Work in (Timmis *et al*, 2000) developed an AIS inspired by the immune network theory, based on work undertaken by (Hunt *et al*, 1996). The proposed AIS consisted of a set of B-cells, links between those B-cells, and cloning and mutation operations that are performed on the B-cell objects. The AIS is tested on the well-known Fisher Iris data set. This data set contains three classes, of which two are not linearly separable. Each B-cell in the AIS represents an individual data item that could be matched (by Euclidean distance) to an antigen or another B-cell in the network (according to Jerne's immune network theory). The links between the B-cells were calculated by a measure of affinity between the two matching cells. If this affinity is above the *network affinity threshold* (NAT) it could be said that there is enough similarity between the two cells for a link to exist. The strength of this link is proportional to the affinity between them. A B-cell also has a certain level of stimulation that is related to the number and to the strength of links a cell has. The AIS also had a cloning mechanism that produced randomly mutated B-cells from B-cells that became stimulated above a certain threshold. The cloning mechanism is inspired by *somatic hypermutation* that produces mutated cells in the human body. The network is trained by repeatedly presenting the training set to the network. The AIS produced some encouraging results when tested on the Fisher Iris data set (Fisher, 1936). The proposed system successfully produced three distinct clusters, which when presented with a known data item could be classified. However, although the clusters were distinct there was still a certain amount connection between Iris Virginica and Iris Versicolor. The AIS also experienced an uncontrolled population explosion after only a few iterations, suggesting that the suppression mechanism (culling 5 % of the B-cell) could be improved. This work was compared to other traditional cluster analysis techniques and Kohonen Networks (Kohonen, 1997a) and found to compare favourably (Timmis *et al*, 1999).

This work was then taken further in (Timmis & Neal, 2001). In this paper the authors raise and address a number of problems concerning the work in (Timmis *et al*, 2000). A number of initial observations were clear: The network underwent exponential population explosion; the NAT eventually became so low that only very similar, if not identical clones can ever be connected; the number of B-cells removed from the system lags behind the number created to such an extent that the population control mechanism was not effective in keeping the network population at a sensible level; the network grew so large that they become difficult to compute each iteration with respect to time; the resultant networks were so large, they were

difficult to interpret, and were really too big to be a sensible representation of the data. With these concerns in mind, the authors proposed a new system called RLAIS (Resource limited artificial immune system). This was later renamed AINE (Artificial Immune Network). To summarize work in (Timmis & Neal 2001) AINE is initialised as a network of ARB objects (Artificial Recognition Balls); T-cells, again, are currently ignored. Links between ARBs are created if they are below the Network Affinity Threshold (NAT), which is the average Euclidean distance between each item in the data set. The initial network is a cross section of the data set to be learnt, the remainder makes up the antigen training set. Each member of this set is matched against each ARB in the network, again, with the similarity being calculated on Euclidean distance. ARBs are stimulated by this matching process and by neighbouring ARBs in the network. Again, a certain amount of suppression is included in the ARB stimulation level calculation. The equation used as a basis for B-cell stimulation calculation was based on Equation 1. The stimulation level of an ARB determines the survival of the B-cell. The stimulation level also indicates if the ARB should be cloned and the number of clones that are produced for that ARB. Clones undergo a stochastic process of mutation in order to create a diverse network that can represent the antigen that caused the cloning as well as slight variations. There exist a number of parameters to the algorithm, those being: network affinity scalar; mutation rate and number of times the training data is presented to the network. Each one of these can be used to alter algorithm performance. The population control mechanism that replaced the 5% culling mechanism, forces ARBs to compete for survival based on a finite number of resources that AINE contains; the more stimulated an ARB, the more resources it can claim. Once an ARB no longer claims any B-cells, it is removed from the AINE. Previously, always 5% was removed, with AINE this is not the case, a predetermined number is not set for removal and the amount removed depends on the performance of the algorithm. This gives rise to a meta-dynamical system that which will extract patterns or clusters from data being learnt. The authors propose that AINE is a very effective learning algorithm, and on test data so far, very encouraging results have been obtained. The authors test the system on a simulated data set and the Iris data set. With the Iris data set, three distinct clusters can be obtained, unlike the original AIS proposed. Additionally, the networks produced by AINE are much smaller than the original system. In effect, AINE is acting as a compression facility, reducing the complexity of the networks, as to highlight the important information, or knowledge, that can be extracted from the data. This is achieved by a special visualisation tool outlined in (Timmis, 2001). More details of these algorithms can be found in (Timmis, 2000), (Timmis & Neal, 2001). However, more recent work has shown that the networks produced by AINE suffer strong evolutionary pressure and converge to the strongest class represented in the data (Knight & Timmis, 2001). Whilst this is an interesting development that could potentially be applied to optimization, however, with regard to data mining it would not be preferential. From a continuous learning point-of-view it is more desirable if all patterns persist over time rather than the strongest. Work in (Neal, 2002) has developed a form of the original algorithm that is capable of finding stable clusters. Here, a different population control mechanism based on exponential decay of stimulation level calculations and the system allows for the continual learning of clusters of information, even in the absence of antigenic input.

5.2.4 Multi-Layered Immune Inspired Learning

In parallel to this work (Knight and Timmis, 2002) have developed a multi layered immune inspired algorithm for data mining. The motivation for this work was to take a step back from existing work and attempt to take a more holistic approach to the development of an immune inspired algorithm. It was noted that a more holistic approach might provide a better solution in the search for an immune inspired data-mining algorithm capable of continuous learning. Rather that focusing on the immune network theory the authors adopted aspects of the primary and secondary responses seen in the adaptive immune system. This new approach incorporates interactions between free-antibodies, B-cells, and memory cells, using the clonal selection processes as the core element of the algorithm. This three-layered approach consists of a free-antibody layer, B-cell layer and a memory layer. The free-antibody layer provides a

general search and pattern recognition function. The B-cell layer provides a more refined pattern recognition function, with the memory layer providing a stable memory structure that is no longer influenced by strong evolutionary pressure. Central to the algorithm is feedback that occurs between B-cells and is part of the secondary immune response in the algorithm. Novel data is incorporated into the B-cell layer and is given a chance to thrive, thus providing a primary immune response. Initial testing of this algorithm has shown good performance at static clustering.

5.2.5 Data Clustering

Similar work to that of (Timmis and Neal, 2001) has been undertaken in (de Castro & Von Zuben, 2000b). In this work the authors propose a system called aiNet, the driving force of which is data clustering and filtering redundant data. Again, for inspiration the authors utilise the immune network theory and the idea of shape space. The proposed aiNet is likened to a weighted disconnected graph, where each cell represents a set of variables (attributes or characteristics) which is said to characterise a molecular configuration, hence a point in p-dimensional space (shape space). Cells are allowed connections between them based on some similarity measure. Suppression within aiNet is achieved by eliminating self-similar cells under a given threshold (defined by the user). Cells within aiNet compete with each other for recognition of antigens (training data) and if successful proliferate and are incorporated into the emerging network. The algorithm is as follows: the training data are presented to an initial randomly generated network. Affinity between antigens and network cells is calculated and the highest matched cells are cloned and mutated. A heuristic is placed in the algorithm that increases the weighting of well-matched cells by decreasing their distance between the antigen items; this is akin to a greedy search. The affinity between these cells in this new matrix is then calculated with the lowest matched cells being removed (this is based on a predetermined threshold set by the user). A certain number of cells are then removed from the network; again, based on a threshold value predetermined by the user, the new clones are then integrated into the network. The cells in the network then have their affinities with each other recalculated, with again a certain number being removed, that fall under the user defined threshold. After the learning phase, the network can be said to be a representation of the data set that is being learnt. Clusters and patterns will emerge within the network and can be used for knowledge extraction. Once the networks have been created, the authors then use a variety of statistical techniques for interpreting the networks. The authors' main goal for aiNet is two-fold: identify the number of clusters within the data and determine which network cell belongs to which cluster. To achieve this, the authors apply the minimal spanning tree algorithm to the network. The authors test their system on two data sets, a simple five linearly separable data set and the famous Donut problem. Good results are obtained for each of the experiments, aiNet identifies the clusters within the data and manages to represent those clusters with a reduced number of points; thus reducing the complexity of the data. Work in (de Castro & von Zuben, 2001) explores the possibility of using immunological metaphors for Boolean competitive networks.

5.2.6 Inductive Learning

Research by (Slavov & Nikoleav, 1998) attempted to create an inductive computation algorithm based upon metaphors taken from immunology. In their paper, they describe an evolutionary search algorithm based on a model of immune network dynamics. By imitating the behaviour of constantly creating and removing good solutions, coupled with attempts to create a diverse range of solutions, the algorithm achieved high diversity and efficient search navigation. These dynamic features were incorporated in the fitness function of the immune algorithm in order to achieve high diversity and efficient search navigation. The authors claim an efficient and effective solution when compared to more traditional GAs.

5.2.7 Sparse Distributed Memory

Hart and Ross (Hart & Ross, 2001, Hart and Ross, 2002a, Hart & Ross 2002b) have used an immune system metaphor to address the problem of finding and tracking clusters in non-static databases. They note that in order to be ultimately useful in the real world, a successful machine-learning (ML) algorithm should address the following characteristics observed in very large, real-world databases:

- Databases are non-static; data is continually added and deleted
- Trends in the data change over time
- The data may be distributed across several servers
- The data may contain a lot of 'noise'
- A significant proportion of the data may contain missing fields or records

The biological immune system performs remarkably well in a dynamic environment; the system is continuously exposed to a variety of ever changing pathogens, and it must adapt quickly and efficiently in order to counteract them. Moreover, the biological immune system is robust to noisy and incomplete information. Therefore the metaphor embodies exactly those characteristics that it is proposed a good ML algorithm must contain. Hart and Ross's work combines an immune system metaphor with that of another class of associative memories: - the sparse distributed memory (SDM). This type of memory was first suggested by (Kanerva, 1988), and since then Smith (Smith *et al*, 1998) has shown that the IS and SDM can be considered analogous. The SDM is a robust memory that derives its properties from the manner in which it performs sparse sampling of huge input spaces by a small number of recognition units (equivalent to B cells and T cells in the immune system), and from the fact that that the memory is distributed amongst many independent units. This is analogous to the memory population of the IS which again consists of B-cells and T-cells.

In brief, an SDM is composed of a set of physical or *hard locations*, each of which recognises data within a specified distance of itself - this distance is known as the recognition *radius* of the location, and all data recognised is said to lie within the access *circle* of the location. In the case of storing binary data, then distance is simply interpreted as Hamming Distance. Each location also has an associated set of counters, one for each bit in its length, which it uses to 'vote' on whether a bit recalled from the memory should be set to 1 or 0. An item of data is stored in the memory by distributing it to every location which recognises it - if recognition occurs, then the counters at the recognising locations are updated by either incrementing the counter by 1 if the bit being stored is 1, or decrementing the counter by 1 if the bit being stored is 0. To recall data from the memory, all locations, which recognise an address from which recall is being attempted vote by summing their counters at each bit position; a positive sum results in the recalled bit being set to 1, a negative sum in the bit being set to 0. This results in a memory, which is particularly robust to noisy data due to its distributed nature and inexact method of storing data.

These properties make it an ideal candidate as a basis for building an immune system based model for addressing clustering problems in large, dynamic, databases. For example, we can consider each physical location along with its recognition radius to define a cluster of data; the location itself can be considered to be a concise representation or description of that cluster, and the recognition radius specifies the size of the cluster. Clusters can overlap --- indeed, it is this precisely this property that allows all data to be recognised with high precision whilst maintaining a relatively low number of clusters. This has a direct parallel in the biological immune system in which antibodies exhibit cross-reactivity. If no overlap was allowed in an SDM, then a large number of locations would be required to cluster the data, the system would become overly specific, and hence general trends in the data would be lost. The analogy between then IS and the SDM class of associative memories is detailed in Table 1, taken from (Smith *et al*, 1998).

Immunological Memory	SDM
Antigen	Address/Data
B/T Cell	Hard Location
Ball of Stimulation	Access Circle
Affinity	Hamming Distance
Primary Response	Write and Read
Secondary Response	Read
Cross-Reactive Response	Associative Recall

Table 1: Analogy between the immune system memory and SDM

In its original form however, the SDM is a static form of memory, and is built on several assumptions that make it unsuitable to use directly as a model for data clustering. In brief, these assumptions are that the addresses of the hard locations are randomly chosen and fixed from the start, and that the recognition radii of each address are equal and constant. Hart and Ross first addressed these problems in a system named COSDM (Hart and Ross, 2001) in which they adapted a co-evolutionary genetic algorithm architecture first proposed by (Potter & De Jong, 2000), cGA, to form an immune system based model capable of clustering static and dynamic data-sets. cGA is another data-clustering algorithm which uses an immune-system metaphor to categorise a benchmark set of data, (Congress Voting records), and performs very well compared to more classical categorisation techniques such as ID3.

In COSDM, an *antigen* represented an item of data and an antibody defined a hard location and its recognition radius. The antibodies co-operate to form an SDM type of memory in which antigen data can be stored. The system consisted of a number of populations of potential antibodies - each population contributed one antibody to the memory A co-evolutionary GA was used to find quickly the 'location' of the antibodies and the size of their corresponding balls of stimulation in order to best cluster the data currently visible to the system. If an antibody recognised an antigen, the antigen was 'stored' by that antibody. The accuracy of clusters produced was determined by attempting to recall each antigen and then comparing the results to the actual data in the database. Antibody populations were added and deleted dynamically - if the best member of a population did not make a significant contribution to the memory, then the population was deleted. Similarly, if the system was not able to improve the clustering accuracy over a predetermined number of generations, then a new population was added. This system was tested on a number of benchmark static and dynamic data-sets – although it showed some promise on clustering dynamic data-sets, it was outperformed by the immune system of (Potter and De Jong, 2000) on large, static data-sets. The difficulties arose in evolving a suitable size for the ball of recognition of each antibody, which led to some antigen never being recognised by any of the antibodies in the system. Also, the system required large numbers of evaluations to find a reasonable SDM, due to the nature of the co-evolutionary architecture.

Hart and Ross thus tackled these issues in (Hart and Ross, 2002a, Hart, 2001b) in which they describe a system based on an SDM as in COSDM, but in which the architecture is akin to that used in a self-organising map, and thus the system is called SOSDM (Self-Organising SDM). A diagram of SOSDM is shown in figure 5. In this system, the recognition radius is replaced by a mechanism in which all antibodies in the system compete for antigen data; antigens bind to all those antibodies for which they have an affinity greater than some preset affinity threshold, with a strength proportional to their affinity. Thus, the binary counters in

the SDM are replaced with real-valued counters, and updated according to the strength of the binding. Each antibody accumulates a measure of its own error, that is, how distant are the antigens recognised by itself from its own description (based on Hamming Distance between the antibody and the antigen). This quantity is then used to allow the antibodies to self-organise, that is, antibodies gravitate towards regions of the space in which they best recognise antigen. The counters also move with each antibody, but decay over time, thus they contain a historical record of data that has been recognised by the antibody. As in COSDM, new antibodies are added periodically, and antibodies can also be deleted dynamically. SOSDM is thus truly adaptive and self-organising, and as such encapsulates some of the most important features of the biological immune system.

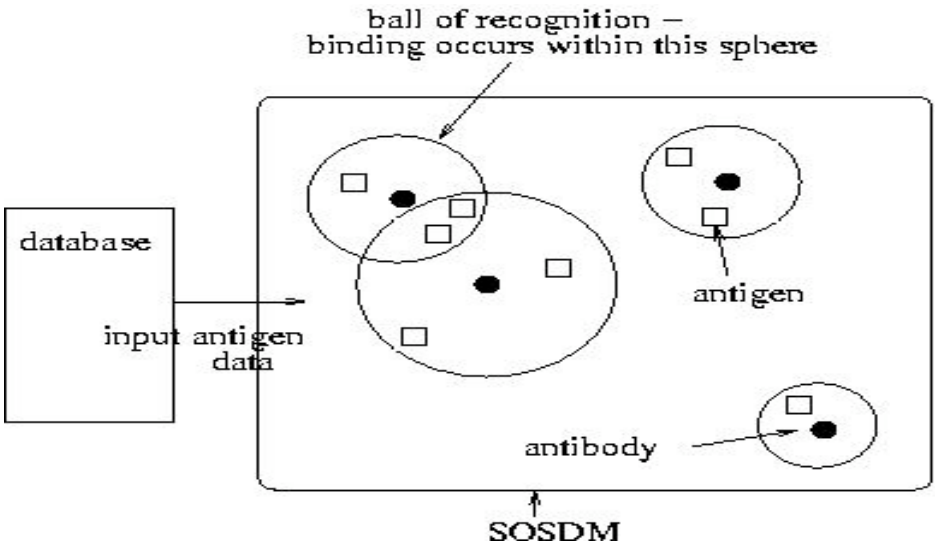


Figure 5: Diagrammatic representation of the SOSDM model

SOSDM has been shown to outperform other published immune algorithms on benchmark static data sets, and furthermore performance has been shown to scale both the size of the data set and with the length of the antigens within the data set. It was also tested on data sets, which contained known clusters of unequal sizes, and was shown to be satisfactory at detecting small clusters.

SOSDM was also tested on a number of time-varying data sets. The experiments tested scenarios, which are likely to represent the extremes of those scenarios, which might realistically occur in a real-world situation. Thus, one set examined scenarios in which data in one cluster was gradually replaced with new data, but still belonging to the same clusters, whereas the other set examined cases where whole clusters were suddenly deleted and replaced by entirely new clusters containing different data. SOSDM performed well at both tasks, though some loss in recall accuracy was observed as the number of clusters being replaced was increased. SOSDM was also shown to exhibit a basic form of memory; when re-exposed to familiar antigens, it reacted more rapidly than to previously unseen antigen. The system appeared relatively robust to the period of the memory. In summary, SOSDM provides a scalable, fast and accurate way of clustering data, but also builds on the analogy between the SDM and the immune system first presented by (Smith *et al,* 1998) to produce a system that is more faithful to the principles of the biological system than the original analogy suggested.

5.2.8 Supervised Learning with Immune Metaphors

(Carter, 2000) made use of the immune network theory to produce a pattern recognition and classification system. This system was known as Immunos-81. The author's aim was to produce a supervised learning system that was implemented based on high levels of abstraction on the workings of the immune system.

The model consisted of T-cells, B-cells, antibodies and an amino-acid library. Immunos-81 used the artificial T-cells to control the production of B-cells. The B-cells would then in turn compete for the recognition of the "unknowns". The amino-acid library acts as a library of epitopes (or variables) currently in the system. When a new antigen is introduced into the system, its variables are entered into this library. The T-cells then use the library to create their receptors that are used to identify the new antigen. During the recognition stage of the algorithm T-cell paratopes are matched against the epitopes of the antigen, and then a B-cell is created that has paratopes that match the epitopes of the antigen.

Immunos-81 was tested using two standard data sets, both of these from the medical field. The first set was the Cleveland data set, which consists of the results of a medical survey on 303 patients suspected of having coronary heart disease. This data set was then used as a training set for the second data set; a series of 200 unknown cases. Immunos-81 achieved an average classification rate of 83.2% on the Cleveland data set and approximately 73.5% on a second data set. When compared to other machine learning techniques, Immunos-81 performed very well. The bets rival was a *k*-nearest neighbour classifier (Wettschereck *et al*, 1997), which averaged 82.4% on the Cleveland data set, other clustering algorithms (Gennari *et al*, 1989) managed 78.9% and using C4.5 only 77.9% accuracy was obtained. The authors therefore argue that Immunos-81 is an effective classifier system, the algorithm is simple and the results are transparent to the user. Immunos-81 also has the potential for the ability to learn in real-time and be embeddable. It has proved to be a good example of using the immune system as a metaphor for supervised machine learning systems.

(Watkins, 2001) proposed a resource limited artificial immune system classifier model using as a basis work from (Timmis, 2000) and (De Castro & von Zuben, 2000c). Here the author extracted metaphors such as resource competition, clonal selection and memory cell retention to create a classification model named AIRS. Results presented in this work are very encouraging. Benchmark data sets such as Fisher Iris data set, Ionosphere data set and sonar data sets were used to test the effectiveness of the algorithm. AIRS was found to perform at the same level of accuracy as some other well established techniques, such has C4.5, CART etc. Recent work has highlighted several revisions that could be made to the original algorithm (Watkins and Timmis, 2002). The work highlighted that the internal data representation of the data items were over-complicated and by simplifying the evolutionary process it was possible to decrease the complexity whilst still maintaining accuracy. The authors also adopt an affinity aware somatic hypermutation mechanism to which they also attribute improved quality of memory cells and therefore greater data reduction and faster classification.

5.3 Robotics

Attempts have been made to apply the immune network idea to control large populations of robots to have some form of self-organising group behaviour. Work by (Mitsumoto *et al*, 1996) attempts to create a group of robots, which behave in a self-organising manner, to search for food without any global control mechanism. Central to their idea is the interaction between robots at the local level. The authors use three main immunological metaphors. The first is B-cells, where a robot represents a B-cell and each robot has a particular *strategy* on how to find food. The second is the immune network, allowing for interaction between robots. The third is the calculation of B-cell stimulation, where the more the robot is stimulated, and

then the better its strategy is considered to be. In order to calculate B-cell (robot) stimulation a modified version of Equation 1 is used, where the robot is stimulated and suppressed by neighbouring robots and stimulated by the outside environment. Each robot carries a record of its degree of success in collecting food, while neighbouring robots compare their success and strategies and stimulate and suppress each other accordingly. If a robot's stimulation level is considered low, then the strategy is considered too weak and, losing that strategy, randomly selects another. If the robot is well stimulated, the strategy is considered to be good and is preserved. Over time the robots interact and successfully achieve the food collection. The authors claim good results on their test data, but indicate the need for further research and testing.

This work is advanced by (Mitsumoto *et al*, 1997), where similar techniques were applied to create a group of robots to interact and achieve the transportation of multiple objects to multiple locations. The algorithm is very similar to the first: the B-cell is represented by a robot, the work to be done by the robots being analogous to antigens, and communication between robots is achieved via the network. The idea of B-cell cloning is also introduced into the algorithm, which is used to represent messages to other robots. Here, a robot is stimulated by interaction between other neighbouring robots and the work environment. If a robot is achieving the work, then it receives more stimulation. If that robot becomes well stimulated, it produces clone B-cells that contain information about the work it is doing, since it is considered to be good work. Other robots in the network then match these and, if they share similar work, they become stimulated and produce other similar work B-cells. If they do not match well, the robot will attempt to adapt its work to the most common work strategy it encounters. Both this interaction and passing of messages enables a group behaviour to emerge that can solve the transportation problem. It was also shown by the authors that this is successful if the work remains static or if the work requirement changes over time.

In very similar work by (Lee *et al*, 1997), the immune network metaphor is applied to creating *swarm* strategies for mobile robots. However, this work is virtually identical to that presented above. The authors do extend the concept in (Lee *et al,* 1999) who introduce the metaphor of the T-cell into the algorithm. They propose a modified version of Equation 1 with the addition of the T-cell metaphor. However, the authors fail to include the results of using the modified equation in their simulation results, presenting instead results of only using the equation without the T-cell interaction.

Work by (Watanabe *et al*, 1998) and (Kondo *et al*, 1998) attempts to create a mechanism by which a single, self-sufficient autonomous robot, the *immunoid*, can perform the task of collecting various amounts of garbage from a constantly changing environment. The environment for the immunoid consists of garbage to be collected, and a home base consisting of a wastebasket and a battery charger. The authors use the metaphors of antibodies, which are potential behaviours of the immunoid, antigens, which are the environmental inputs such as existence of garbage, wall and home bases and the immune network, which is used to support good behaviours of the immunoid. In order for the immunoid to make the best strategy decision, the immunoid detects antigens and matches the content of the antigen with a selection of all the antibodies that it possesses. For example, the immunoid may have antibodies that are suitable for when a wall is met head-on and therefore needs to turn right. Each antibody of the immunoid records its concentration level, which is calculated using Equation 1. A number of antigens (environmental inputs) are detected and the levels of antibodies are calculated and the antibody with the highest concentration is selected as the appropriate behaviour to employ. In experimental results, the authors prepared 24 antibodies for the immunoid (potential behaviours) and observed good results. The authors then extended this work. This was an attempt to create more emergent behaviour within the network of robots (Watanabe *et al*, 1998) by the introduction of genetic operators.

5.4 Fault Diagnosis and Tolerance

The field of diagnosis is a vast field driven by the requirement to accurately predict or recover from faults occurring in plant. One approach to detect abnormal sensors within a system (Kayama *et al*, 1995) has been to use the combination of Learning Vector Quantization (LVQ) (Kohonen, 1997b) and the immune network metaphor. The idea behind the system is to use LVQ to determine a correlation between two sensors from their outputs when they work properly, and then use an immune network to test sensors using extracted correlations. Within the system, each sensor corresponds to a B-cell and sensors test one another's outputs to see whether or not they are normal. Each sensor calculates a value based on an adapted version of Equation 1 where the inputs to the equation are reliability of the sensor, rather than similarity to the neighbour. A sensor that has a low value is considered to be faulty and can therefore be flagged for needing repair. Using this method has the advantage of having no overall control mechanism for checking for faulty sensors; they can detect for themselves when they are faulty. Simulations of their system showed the potential for good diagnostic results, and the paper points the way forward for more research and actual application to real plants.

Also in the field of diagnosis, there has been an interest in creating other distributed diagnostic systems. Initial work in (Ishida, 1990) and (Ishida & Mizessyn, 1992) proposed a parallel distributed diagnostic algorithm. However, the authors likened their algorithm to that of an immune network, due to its distributed operation, and the systems emergent co-operative behaviour between sensors. This work was then continued in (Ishida, 1996), (Ishida & Tokimasa, 1996), and *active diagnostic* mechanism (Ishida, 1997). The work in (Ishida1997) builds on foundations laid in the others so will be briefly examined here.

Active diagnosis continually monitors for consistency between the current states of the system with respect to the normal state. The authors argue that the immune system metaphor is a suitable idea for creating an effective active diagnostic system. Central to their idea is the immune network theory, where each sensor can be equated with a B-cell (Kayama *et al,* 1995). Sensors are connected via a network (the immune network), with each sensor maintaining a record of sensory reliability, which is continually changed over time - creating a dynamic system. Sensors in the network can test each other for reliability, but where this work differs from the above is the way in which the reliability of each sensor is calculated. This will not be explored here. The key features of the immune system that is used by this work are distributed agents that interact with each other in parallel (each agent only reacting on its own knowledge and not via a central controller), and the creation of memory of the sensor state formed via a network.

Hardware fault tolerant systems seek to provide a high degree of reliability and flexibility even in the presence of errors within the system. The said system must be protected from a variety of potential faults, manifesting in such forms as permanent stuck at faults or intermittent faults.

Bradley & Tyrrell (2000a,b) proposed what they called Immunotronics (immunological electronics) in order to implement a finite state machine based counter using immune principles. Their proposed system relied upon the negative selection algorithm that is responsible for creating a set of tolerance conditions to monitor changes in hardware states. They employed a binary Hamming shape-space to represent the tolerance conditions.

Recent work in (Timmis et al., 2002) discusses important issues when considering the design of immune inspired fault tolerant embedded systems. The authors highlight that one advantage of using a technique based on AIS in comparison to traditional fault tolerant approaches, is the possibility to exploit the evolutionary property of the immune system. While conventional fault tolerant techniques generate static detectors that have to be updated

offline, AIS-based techniques will enable the development of adaptable fault tolerant systems, in which error detectors may evolve during runtime. This feature will increase the availability of embedded systems since acceptable variations of non-erroneous states can be integrated to the self-system. For example, external factors (e.g. temperature) induce changes that might have significant effects on the system functionalities, while internal changes (e.g. component replacement) could give rise to variability in self that must be noticed. The authors also argue that AIS techniques however pose some challenges. One of them is the need to ensure that the detectors generated fully cover the non-self space (i.e. the erroneous states). This is determined by the mode of detector generation, which in turn affects the resulting detector set as well as the speed of the operation. However, the distribution of the self-data can be exploited to enhance the process. Other metaphors of the immune system are also tagged as potential avenues for research in this area such as the adaptability feature, which is inherent in the immune network metadynamics (Bersini, 2002).

5.5 Optimisation

In order to address the issue of designing a Genetic Algorithm (GA) with improved convergence characteristics, particularly in the field of design constraints, work in (Hajela *et al*, 1997) proposed a GA simulation of the immune system. The motivation for their work stems from the fact that genetic algorithms, when applied to design constraints, have been found to be very sensitive to the choice of algorithm parameters, which can ultimately affect the convergence rate of the algorithm. The authors use the idea of antibody -antigen binding to define a complex matching utility to define similarity between design solutions. This is based on work found in (Farmer *et al*, 1986), (section 3.3.1.1) and is simply a bit by bit match for continuous regions. The model created also simulates the dynamics of the immune system by creating and removing possible new solutions. Some solutions will be more specific to the problem areas, whereas others will be more generalised. However, the authors point out that both specialist and general solutions are important in the context of structural design, so they introduce a control parameter into the algorithm that enables them to control the production of specialist and general case solutions. The authors suggest their algorithm leads to a higher convergence rate when compared to a traditional GA, but indicate the need for further research and application. It should be noted, however, that while the authors claim to use the immune network as a metaphor, in reality they use the immune system, as there is no apparent network interaction going on during the algorithm.

The above work focused on a specific search problem in a particular domain, work in (Toma *et al*, 1999) adopts a more generic approach to adaptive problem solving by the use of the immune network metaphor. Again, the authors claim the use of a network structure, but do not present the work as such, but simply immune system metaphors including B-cells, T-cells, macrophages and the Major Histocompatibility Complex (MHC). The immune algorithm given in the paper is used to produce adaptive behaviours of agents, which are used to solve problems. The algorithm is then applied to the n-TSP problem, and for small-scale problems achieves good results. The authors also experiment with removing the interaction of the T-cell in the searching algorithm and present convincing results that the effect of the T-cell on performance is significant, as the solutions found with using the T-cell result in lower cost solutions overall. Other, similar application of the immune network metaphor for multi-modal function optimisation can be found in (Mori *et al*, 1996), (Fukuda *et al*, 1998), (Mori *et al*, 1998). Here the authors use somatic hypermutation and immune network theory to create and sustain a diverse set of possible solutions in the search space and combine it with traditional genetic algorithms. The authors propose that their algorithm possess two main characteristics: (i) the ability to create a diverse set of candidate solutions, and (ii) is a parallel-efficient search. Combined with the somatic mutation, the authors also employ standard genetic algorithm mutation operators of crossover and mutation. The authors apply their algorithm to finding optimal solutions to various functions and compare the result

obtained with a standard GA approach. They argue that the strength of their algorithm lies in its ability to maintain a higher diversity of candidate solutions compared to a standard GA, which is important when attempting to find the global maximum on any search surface.

De Castro & Von Zuben (2000c) focused on the clonal selection principle and affinity maturation process of an adaptive immune response to develop an algorithm suitable to perform tasks like machine learning, pattern recognition, and optimisation. Their algorithm was evaluated in a simple binary character recognition problem, multi-modal optimisation tasks and a combinatorial optimisation problem; more specifically the travelling salesman problem (TSP). The main immune aspects taken into account to develop the algorithm were: maintenance of a specific memory set, selection and cloning of the most stimulated cells, death of non-stimulated cells, affinity maturation and re-selection of the clones proportionally to their antigenic affinity and generation and maintenance of diversity. The performance of their algorithm was compared with a GA for multi-modal optimisation, and the author's claim their algorithm was capable of detecting a high number of sub-optimal solutions, including the global optimum of the function being optimised. This work was further extended with the use of the immune network metaphor for multi-modal optimisation in (de Castro and Timmis, 2002a).

5.6 Scheduling

Creating optimal schedules in a constantly changing environment is not easy. Work by (Mori *et al*, 1994), (Chun *et al*, 1997) and (Mori *et al*, 1998) proposes and develops an immune algorithm that can create adaptive scheduling system based on the metaphors of somatic hypermutation and the immune network theory. Work in (Mori *et al*, 1998) builds on the above by addressing the issue of batch sizes and combinations of sequence orders, which optimise objective functions. In these works, antigens are considered as input data or disturbances in the optimisation problem, and antibodies are considered as possible schedules. Proliferation of the antibodies is controlled via an immune network metaphor where stimulation and suppression are modelled in the algorithm. This assists in the control of antibody (or new solution) production. The T-cell effect in this algorithm is ignored. The authors claim that their algorithm is an effective optimisation algorithm for scheduling and has been shown to be good at finding optimal schedules. The authors indicated that further work could be undertaken in applying this algorithm to a dynamically changing environment. This work was undertaken in (Hart *et al*, 1998) and more recently in (Hart & Ross, 1999a).

(Hart *et al*, 1998) proposes a system that can create a diverse set of schedules, but not necessarily an optimal solution for the scheduling problem that can be easily adapted should the situation change. The authors consider antibodies as a single schedule and antigens to be possible changes to the schedule. Their system produces a set of antibodies (schedules) that can cover the whole range of possible changes in the antigen set. Using these metaphors, and that of gene libraries to create new antibodies, the authors have shown that they can create a set of schedules, using a GA, from an initial random state of possible changes. Their system can then successfully retrieve schedules corresponding to antigens existing in that set, and also new antigens (or changes in situations) previously unseen. In a later work, (Hart & Ross 1999b) proposed a scheduling application of an artificial immune system, called PRAIS (Pattern Recognising Artificial Immune System). In their system, sudden changes in the scheduling environment required the rapid production of new schedules. Their model operated in two phases. A first phase comprising the immune system analogy, in conjunction with a genetic algorithm, in order to detect common patterns amongst scheduling sequences frequently used by a factory. In phase II, some of the combinatorial features of the natural immune system were modelled to use the detected patterns to produce new schedules, either from scratch or by starting from a partially completed schedule.

5.7 Computer Security

The problem of protecting computers from viruses, unauthorised users, etc. constitutes a rich field of research for artificial immune systems. The existence of a natural immune system to fight against biological micro organisms like viruses and bacteria is probably the most appealing source of inspiration for the development of an artificial immune system to combat computer viruses and network intruders.

5.7.1 Network Security

The role of the immune system may be considered analogous to that of computer security systems (Dasgupta, 1999). Whilst there are many differences between living organisms and computer systems, researchers believe that the similarities are compelling and could point the way to improved computer security. Long-term research projects have been established in order to build a computer immune system (Forrest *et al*, 1994), (D'haeseleer *et al*, 1996), (Forrest *et al*, 1996) and (Forrest *et al*, 1997) which could augment a simple computer security system with more advanced and novel features. A good overview of the current work in this field is presented by (Somayaji *et al,* 1997), where an attempt is made to draw together various pieces of research in the field in order to derive some basic principles of computer immune systems.

There are a number of approaches to implementing a computer security system. Host based intrusion detection methods (Hofmeyr *et al*, 1998), (Warrender *et al*, 1999), construct a database that catalogues the normal behaviour pattern of a piece of software that is specific to a particular machine, software version etc. Construction of such a database would enable the programs' behaviour to be monitored.

In order to build up a pattern of normal behaviour for a particular database of software, system calls made by the software are monitored and recorded over time. As this record builds up, the database may be monitored for any system calls not found in the normal behaviour database, which are then flagged. The authors argue that, while simplistic, this approach is not computationally expensive and can be easily used in real time. It also has the advantage of being platform and software independent.

An alternative method is the network based intrusion detection approach. This tackles the issue of protecting networks of computers rather than an individual computer. This is achieved in a similar way in monitoring network services, traffic and user behaviour and attempts to detect misuse or intrusion by observing departures from normal behaviour. Work in both Hofmeyr & Forrest (1999, 2000) and (Kim & Bentley, 1998) lay foundations for a possible architecture and general requirements for a network based intrusion detection system based on immune system metaphors. (Kim & Bentley, 1999) propose a network intrusion detection algorithm based on metaphors presented in the previous paper. The algorithm is based on the negative selection algorithm, first proposed by (Forrest *et al*, 1994). Negative selection in the immune system is the immune system's ability to eliminate harmful antibodies while not attacking the self of the immune system (Section 3.7). The algorithm in (Forrest *et al*, 1994) consists of three phases: defining self, generating detectors and monitoring the occurrence of anomalies. In this paper, it was applied to the detection of computer viruses.

Recently, Dasgupta (1999, 2000) proposed an agent-based system for intrusion/anomaly detection and response in networked computers. In his approach, the immunity-based agents roamed around the nodes and routers monitoring the situation of the network. The most appealing properties of this system were mobility, adaptability and collaboration. The immune agents were able to interact freely and dynamically with the environment and each other.

5.7.2 Virus Detection

Much interest has been shown in applying immune system metaphors to virus detection. The first work done on this (Forrest *et al*, 1994), (Forrest *et al*, 1997) developed a simple algorithm using the negative selection metaphor to detect potential viruses in computer systems. This work was concerned with distinguishing normal computer resources and behaviour from abnormal. A different approach is taken in (Kephart, 1994), (Kephart *et al*, 1997) and (Kephart *et al*, 1998). Their initial approach is to use the metaphor of the innate immune system. This resides on the user's PC and applies virus-checking heuristics to .COM and .EXE files. If an unknown virus is detected, a sample is captured that contains information about the virus and is sent to a central processing system for further examination. This is analogous to how the innate immune system works, as the first line of defence. In the central processing service, the virus is encouraged, or baited, to produce itself in a controlled environment. This allows examination of the virus and extraction of its signature. An antidote can then be constructed, which may be sent out to the infected PC and the virus removed.

The signature extraction mechanism is based on immune system metaphors, such as clonal selection, by producing large numbers of possible code signatures in order to detect the virus code signature. This is achieved by generating large numbers of random signatures and checking each one of these signatures against the potential virus. A positive match indicates that a virus has been detected. (Marmelstein *et al*, 1998) proposed an alternative multi-layer approach which attempts to tackle the infection at varying levels of protection where ultimately if the infection cannot be identified an evolutionary algorithm is applied to create alternative decoy programs to trap the virus. This was extended in (Lamont *et al*, 1999), then in (Harmer & Lamont, 2000) as a distributed architecture for a computer virus system based on immune system principles.

6 Summary

Using the immune system as inspiration has proved very useful when trying to address many computational problems. The immune system is a remarkable learning system. Through the use of B-cells and T-cells the immune system can launch an attack against invading antigens and remove them from the system. This is achieved through a process of B-cell stimulation followed by cloning and mutation of new antibodies. This diversity that is generated allows the immune system to be adaptive to new, slightly different infections. The immune system is able to retain information about antigens; so that next time the body is infected a quicker, secondary immune response can be triggered to eliminate the infection. A number of theories exist about how the immune system retains this information, in a type of memory: the most popular being the clonal selection theory, and the idea of memory cells, and the alternative immune network theory, with the idiotypic interactions on antibodies.

From observing this natural system, researchers have identified many interesting processes and functions within the immune system that can provide a useful metaphor for computation. The review of the field of Artificial Immune Systems (AIS) has revealed many and varied applications of immune metaphors. The proposed framework for AIS was outlined, with the main ideas being that it is possible to think of AIS in terms of a *layered framework* that consists of representations, affinity measures and immune algorithms. The field of machine learning was the examined. Early work by (Cooke & Hunt, 1995) spawned a great deal of research, which led to a generic unsupervised learning algorithm proposed in (Timmis & Neal, 2001), which ultimately forms part of the proposed framework for AIS. Other approaches to learning have also been adopted by (de Castro & von Zuben, 2000b) to create a clustering algorithm. Work in (Hart & Ross, 2001) proposes a modified immune algorithm capable of clustering moving data and is adaptable to clustering large volumes of data. (Carter, 2000) proposed the use of immunological metaphors for supervised machine

learning, that had advantages over other supervised methods in that the results were transparent. Attempting to create collective behaviour in robots has also been a major field of study. By using ideas from the immune network theory, work in (Mitsumoto *et al*, 1996) and subsequent works built a small system of self-organising and autonomous robots that could be used for simple collection and navigational tasks. Some of the pioneering work in AIS was done in the field of fault diagnosis, primary following work from (Ishida, 1990). This led to the development of the active diagnosis field (Ishida, 1996) and more recently the idea of Immunotronics and hardware tolerance (Bradley & Tyrrell, 2000a). Attention has also been paid to using immune metaphors for solving optimisation problems (Hajela & Lee, 1997), by augmenting genetic algorithms. More recent work by (Hart *et al*, 1998) and (Hart & Ross, 1999a, 1999b) and (Hart, 2002) tackled the difficult problem of producing an adaptive scheduling system. Here work used on combining the use of genetic algorithms and immunological metaphors. A significant field of research is that of computer security and virus protection. From the early work in (Forrest *et al*, 1994) for computer security and (Kephart, 1994) a significant body of research has been generated and architectures for various security and virus detection systems proposed (Hofmeyr & Forrest, 2000) and (Lamont *et al*, 1999).

7 Comments on the Future for AIS

For each of the many contributions discussed in Section 5 it would be possible to talk at length regarding possible extensions and improvements. Instead, we will keep our discussion about future trends on AIS at a high level.

Although several papers have been discussed proposing the use AIS to solve problems in many areas of research, few of those have attempted to present a formal artificial immune system (e.g., Hunt & Cooke, 1996; de Castro & Von Zuben, 1999, 2000a; Hofmeyr & Forrest, 2000). In reviewing all these works on AIS, it becomes clear that the area is lacking the proposal and development of a general framework in which to design artificial immune systems. Observing other comparable computational intelligence (or soft computing) paradigms, such as artificial neural networks, evolutionary computation and fuzzy systems, it is clear that there is the presence of well described sets of components and/or mechanisms with which to design such algorithms. To this end, a framework has been proposed (de Castro and Timmis, 2002), however, much work remains to be done on this framework in terms of formalisation from a mathematical viewpoint and augmentation in terms of new shapes spaces and development of new algorithms which have been inspired by other areas of immunology, as yet unexplored by computer scientists and engineers.

This leads to an interesting avenue of research. To date, the concentration has mainly been on the more *basic* immunology, such as simple antibodies, clonal selection and so forth. Recent work in (Aickelin and Cayzer, 2002) postulated the use of the Danger Theory (Matzinger, 1994a) for AIS. This is an interesting idea, sadly beyond the scope of this chapter. However, it is quite possible the danger theory has much to offer AIS in terms of a paradigm shift in thinking – as yet unexplored. This shift is away from the idea that the immune system controls the response, typically adopted in AIS, to the idea that the tissue initialises and controls the response: in some way contextualises the response. This could provide a powerful metaphor, especially in terms of data mining in dynamic environments, where the context of what you might want to learn may change over time. Indeed, this debate was opened further by (Bersini, 2002), where it is argued that the danger theory idea is nothing more than a change of terminology from the idea of self/non-self to the idea that something is dangerous or not dangerous. It will be interesting to observe if this debate grows pace.

The debate also has to be had: is there a single immune algorithm? To date, this has not been addressed and is still an open question. One observation that is made regarding AIS is that

there are so many different algorithms around, it is so complicated and you never know which one to use. In answer to this, you may state that either more rigor and analysis has to be applied to the current algorithms to identify exactly their suitability for problems and therefore, what AIS will offer is a very rich suit of effective and well understood algorithms. Alternatively, you could pursue a single *unified* algorithm: but then would that either enhance or restrict the field of AIS? It may enhance it in the fact that there then exits a single commonly understood algorithm – so people then know what you mean when you say an immune algorithm, or it may restrict in the sense that the immune system is such a complex system, why try and limit that to one simple algorithm – why not exploit the many complexities therein?

It would certainly seem that there are many challenges ahead. Immunology has a great deal to teach people involved with computational problems: we have only just scratched the surface. A greater interaction between biology and computer science is needed if we are to fully exploit the richness of this marvellous biological system.

8 References

Ada, G. L. & Nossal, G. J. V. (1987). The Clonal Selection Theory, *Scientific American*, **257**(2), pp. 50-57.

Acklien, U and Cayzer, S. (2002). The Danger Theory and its Application to Artificial Immune Systems. *Proceedings of the 1st International Conference on Artificial Immune Systems (ICARIS)*. pp. 141-148. Canterbury, UK.

Berek, C. & Ziegner, M. (1993). The Maturation of the Immune Response, *Imm. Today*, **14**(8), pp. 400-402.

Bersini, H. (1991). Immune Network and Adaptive Control, *Proc. of the First European Conference on Artificial Life*, MIT Press, pp. 217-226.

Bersini, H and Varela, F. (1990). Hints for adaptive problem solving gleaned from immune networks. *Parallel Problem Solving from Nature*, 1st Workshop. PPSW1. Dortmund, Germany. Pub. Springer-Verlag, pp. 343-354.

Bersini, H and Varela, F. (1994). The immune learning mechanisms : Reinforcement and recruitment and their applications. *Computing with Biological Metaphors*. Pages 166-192. Pub. Chapman Hall

Bersini, H. (2002). A Tribute to ….. In proceedings of 1st International Conference on Artificial Immune Systems (ICARIS).

Bradly, D. W. & Tyrrell, A. M. (2000a), Immunotronics: Hardware Fault Tolerance Inspired by the Immune System, *Lecture Notes in Computer Science*, **1801**, pp. 11-20.

Bradly, D. W. & Tyrrell, A. M. (2000b). Hardware Fault Tolerance: An Immunological Approach, *Proc. of the IEEE SMC'00*, pp. 107-112.

Burnet, F. M. (1959). *The Clonal Selection Theory of Acquired Immunity*, Cambridge University Press.

Carneiro, J & Stewart, J. (1995). Self and nonself revisited: Lessons from modelling the immune network. *Third European Conference on Artificial Life*, Granada, Spain. Pp. 405-420.

Carter, J.H. (2000). The Immune System as a Model for Pattern Recognition and Classification. *Journal of the American Medical Informatics Assocation*.**7/1**. Pp. 28-41.

Chun, J, Kim, M & Jun, H. (1997). Shape Optimisation of Electromagnetic Devices Using Immune Algorithms. *IEEE Transactions on Magnetics*, **33**,(2).

Cooke, D and Hunt, J. (1995). Recognising Promoter Sequences Using an Artificial Immune System. *Proc. of Intelligent Systems in Molecular Biology*. AAAI Press, pp. 89-97.

Coutinho, A. 1980. The self non-self discrimination and the nature and acquisition of the antibody repertoire. *Annals of Immunology***.** (Inst. Past.) **131D**.

Coutinho, A. (1989). Beyond Clonal Selection and Network, *Imm. Rev.*, **110**, pp. 63-87.

Dasgupta, D (Ed.) (1998a). *Artificial Immune Systems and Their Applications*. Springer-Verlag.

Dasgupta, D (1998b). An overview of artificial immune systems. *Artificial Immune Systems and Their Applications*. pp. 3-19. Pub: Springer-Verlag.

Dasgupta, D (1999). Immunity based intrusion detection systems: A general framework. *Proceedings of the 22nd National Information Systems Security Conference (NISSC)*. Pp. 147-159

Dasgupta, D. (2000). An Immune Agent Architecture for Intrusion Detection, *Proc. of GECCO'00*, Workshop on Artificial Immune Systems and Their Applications, pp. .

De Castro, L. N. & Von Zuben, F. J. (1999). Artificial Immune Systems: Part I – Basic Theory and Applications, *Technical Report – RT DCA 01/99*, p. 95

De Castro, L. N. & Von Zuben, F. J. (2000a). Artificial Immune Systems: Part II – A Survey of Applications, *Technical Report – RT DCA 02/00*, p. 65.

De Castro, L. N., & Von Zuben, F. J., (2000b). An Evolutionary Immune Network for Data Clustering, *Proc. of the IEEE SBRN*, pp. 84-89.

De Castro, L. N. & Von Zuben, F. J. (2000c). The Clonal Selection Algorithm with Engineering Applications, *Proc. of GECCO'00 – Workshop Proceedings*, pp. 36-37.

De Castro, L. N., & Von Zuben, F. J., (2001). The Construction of a Boolean Competitive Neural Network Using Ideas From Immunology, *submitted*.

De Castro, L.N and Timmis, J (2002). Artificial Immune Systems: A New Computational Intelligence Approach. Springer-Veralg. London. ISBN 1-85233-594-7

De Castro, L.N and Timmis, J (2002a). An Artificial Immune Network for Multimodal Optimisation. In Proceedings of the *Congress on Evolutionary Computation. Part of the 2002 IEEE World Congress on Computational Intelligence.*, pages 699-704, Honolulu, Hawaii, USA. IEEE.

De Castro, L.N and Timmis, J. (2003). Artificial Immune Systems as a Novel Soft Computing Paradigm. *Soft Computing*.

D'haeseleer, P, Forrest, S and Helman, P (1996). An Immunological Approach To Change Detection: Algorithm and Analysis and Implications. *Proceedings of the 1996 IEEE Symposium on Computer Security and Privacy*. pp. 110-119

Farmer, J, Packard, N and Perelson, A. (1986). The Immune System, Adaptation and Machine Learning. *Physica D*. **22,** pp**.** 187-204.

Fisher, R (1936). The use of multiple measurements in taxonomic problems. *Annual Eugenics*. **7,** II. Pp. 179-188

Forrest, S, Perelson, A, Allen, L and Cherukuri, R (1994). Self-Nonself Discrimination in a Computer. *Proc. of IEEE Symposium on Research in Security and Privacy*. Oakland, USA, pp. 202-212.

Forrest, S, Hofmeyr & Somayaji, A & Longstaff, T. (1996). A sense of self for UNIX processes. *Proc. IEEE Symposium on Research in Security and Privacy*. Oakland, USA, pp. 120-128.

Forrest, S, Hofmeyr, S and Somayaji, A (1997). Computer Immunology. *Communications of the ACM*. **40/10**. pp. 88-96

Fukuda, T, Mori, K and Tsukiyama, M. (1998). Parallel Search for Multi-Modal Function Optimisation with Diversity and Learning of Immune Algorithm. *Artificial Immune Systems and Their Applications*. pp. 210-220. Pub. Springer-Verlag

Gennari, J.H. Langley, P and Fisher, D. (1989). Models of information concept formation. *Artificial Intelligence*; **40**:11-61.

Hajela, P., Yoo, J. & Lee, J. (1997). GA Based Simulation of Immune Networks - Applications in Structural Optimization, *Journal of Engineering Optimization*.

Harmer, P.K and Lamont, G.B. (2000). An Agent Based Architecture for a Computer Virus Immune System. *In Proceedings of Artificial Immune Systems Workshop*. pp. 45-46. GECCO 2000, Las Vegas, USA.

Hart, E. Ross, P. and Nelson, T (1998). Producing robust schedules via an artificial immune system. *Proc. of IEEE CEC'98*, pp. 464-469. IEEE.

Hart, E. & Ross, P. (1999a). The Evolution and Analysis of a Potential Antibody Library for Use in Job-Shop Scheduling, In *New Ideas in Optimisation*, D. Corne, M. Dorigo & F. Glover (Eds.), McGraw Hill, London, pp. 185-202.

Hart, E. & Ross, P. (1999b). An Immune System Approach to Scheduling in Changing Environments", *Proc. of GECCO'99*, pp. 1559-1566.

Hart, E & Ross, P. (2001). *Clustering Moving Data with a Modified Immune Algorithm*. EvoWorkshops 2001 - Real World Applications of Evolutionary Computing.

Hart, E & Ross, P. (2002a*)*. Exploiting the Analogy between Immunology and Sparse Distributed Memories*. Proc. of ICARIS-2002,* pp. 49-58.

Hart, E. (2002b) *Immunology as a Metaphor for Computational Information Processing: Fact of Fiction ?* PhD thesis. University of Edinburgh.

Hofmeyr, S, Forrest, S & Somayaji, A. (1998). Intrusion detection using a sequence of system calls. *Journal of Computer Security*, **6**, pp. 151-180.

Hofmeyr, S and Forrest, S (1999). Immunity by Design: An artificial immune system. *Proc. of GECCO'99*, Pub. Morgan-Kaufman. pp. 1289-1296

Hofmeyr, S.A. and Forrest, S. (2000). Architecture for an Artificial Immune System. *Evolutionary Computation* **7(1)**:45-68.

Hunt, J. E. & Cooke, D. E. (1996). Learning Using an Artificial Immune System, *Journal of Network and Computer Applications*, **19**, pp. 189-212.

Hunt, J & Fellows, A (1996). Introducing an Immune Response into a CBR system for Data Mining. *BCS ESG'96 Conference and published as Research and Development in Expert Systems XIII*. Pp. 35-42. Springer-Verlag.

Hunt, J, Cooke, D and Holstein, H. (1995). Case Memory and Retrieval Based on the Immune System. *Case-Based Reasoning Research and Development, Lecture Notes in Artificial Intelligence.* **1010**. pp. 205-216

Hunt, J, King, C and Cooke, D. (1996). Immunising Against Fraud. *Proc. Knowledge Discovery and Data Mining.* IEE Colloquium. IEE., pp. 38-45.

Hunt, J, Timmis, J, Cooke, D, Neal, M and King, C. (1998). JISYS: Development of an Artificial Immune System for real world applications. In *Artificial Immune Systems and theory Applications*. Ed. D. Dasgupta. pp. 157-186.

Ishida, Y. (1990). Fully Distributed Diagnosis by PDP Learning Algorithm: Towards Immune Network PDP Model. *Proc. of the IEEE International Joint Conference on Neural Networks*. San Diego, USA, pp. 777-782.

Ishida, Y (1996). Distributed and autonomous sensing based on immune network. *Proc. of Artificial Life and Robotics. Beppu.* AAAI Press, pp. 214-217.

Ishida, Y. (1997). Active Diagnosis by Self-Organisation : An approach by the immune network metaphor. *Proceedings of the International Joint Conference on Artificial Intelligence*. pp. 1084-1089. Nagoya, Japan.

Ishida, Y & Mizessyn, F. (1992). Learning algorithms on immune network model: application to sensor diagnosis. *Proc. International Joint Conference on Neural Networks*, Beijing, China, pp. 33-38.

Ishida, Y & Tokimasa, T. (1996). Diagnosis by a dynamic network inspired by immune network. *Proc. World Congress of Neural Networks*, San Diego, USA, pp. 508-511.

Janeway, C. (1993). Life, Death and the Immune System. *Scientific American Special Issue. How the immune system recognises invaders*, pp. 27-36.

Jerne, N. (1974). Towards a network theory of the immune system. Annals of Immunology (Inst.Pasteur). **125C.** pp. 373-389.

Kanerva, P. (1998) *Sparse Distributed Memory*. MIT Press, Cambridge,MA.

Kayama, M, Sugita, Y, Morooka, Y & Fukuodka, S. (1995). Distributed diagnosis system combining the immune network and learning vector Quantization, pp. 1531-1536 *of Proc. IEEE 21st International Conference on Industrial Electronics and Control and Instrumentation*, Orlando, USA.

Kephart, J. Sorkin, B. Swimmer, M and White, S. (1998). Blueprint for a computer immune system. In *Artificial Immune Systems and their Applications*. Ed. D. Dasgupta. pp. 242-260. Springer-Verlag.

Kephart, J. O., Sorkin, G. B. & Swimmer, M. (1997), An Immune System for Cyberspace, *Proc. of the IEEE SMC'97*, pp. 879-884.

Kephart, J. O. (1994). A Biologically Inspired Immune System for Computers, R. A. Brooks & P. Maes (Eds.), *Artificial Life IV Proceedings of the Fourth International Workshop on the Synthesis and Simulation of Living Systems*, MIT Press, pp. 130-139.

Kepler, T and Perelson, A. (1993). Somatic Hypermutation in B-cells : An Optimal Control Treatment. *Journal of Theoretical Biology*. **164**. pp. 37-64.

Kim, J. & Bentley, P. (1999), Negative Selection and Niching by an Artificial Immune System for Network Intrusion Detection, *Proc. of GECCO'99*, pp. 149-158.

Kim, J and Bentley, P. (1998). The human immune system and network intrusion detection. *Proc. of 7th European Congress on Intelligent Techniques - Soft Computing*. Aachan and Germany

Kolodner, J. (1993). *Case Based Reasoning*. Pub. Morgan Kaufman.

Kohonen, T. (1997a). *Self-Organising Maps*. 2nd Edition.

Kohonen, T. (1997b). *Self-Organising Maps*. 2nd Edition. Chap. Learning Vector Quantization. pp. 203-217. Springer.

Kondo, T, Ishiguro, A, Watanabe, Y and Uchikawa, Y. (1998). Evolutionary construction of an immune network based behaviour arbitration mechanism for autonomous mobile robots. *Electrical Engineering in Japan*. **123/3**. pp. 1-10

Knight, T and Timmis, J. (2001). In N Cercone, T Lin, and Xindon Wu, editors, *IEEE International Conference on Data Mining*, pages 297-304, San Jose, CA. USA, December 2001. IEEE

Knight, T and Timmis, J. (2002). A Multi-Layered Immune Inspired Approach to Data Mining. *Recent Advances in Soft Computing*, Nottingham, UK. 2002

Lamont, G. B., Marmelstein, R. E. & Van Veldhuizen D. A. (1999), A Distributed Architecture for a Self-Adaptive Computer Virus Immune System, *New Ideas in Optimisation*, D. Corne, M. Dorigo & F. Glover (Eds.), McGraw Hill, London, pp. 167-183.

Lee, Dong-Wook and Sim, Kwee-Bo. (1997). Artificial immune network based co-operative control in collective autonomous mobile robots. *Proc. of IEEE International Workshop on Robot and Human Communication. Sendai, Japan.* IEEE. pp. 58-63.

Lee, Dong-Wook and Sim, Kwee-Bo. (1999). Artificial Immune Systems for realisation of co-operative autonomous mobile robots. *Proc. of IEEE International workshop on Robot and Human Communication*. Sendai, Japan. pp. 58-53.

Marmelstein, M, Veldhuizen & Lamont, G. (1998). A Distributed Architecture for an Adaptive Computer Virus System. *Proc. of the IEEE SMC*, San Diego, USA. pp. 3838-3843.

Matzinger, P. (1994). Immunological Memories Are Made of This? *Nature*, **369**, pp. 605-606.

Matzinger, (. (1994a). Tolerance, Danger and the Extended Family. *Annual Review of Immunology*. 12:991-1045.

Mitsumoto, N, Fukuda, T and Idogaki, T. (1996). Self-Organising Multiple Robotic System. *Proceedings of IEEE International Conference on Robotics and Automation*. Pp. 1614-1619. Minneapolis, USA. IEEE.

Mitsumoto, N, Fukuda, T, Arai, F & Ishihara, H (1997). Control of distributed autonomous robotic system based on the biologically inspired immunological architecture. *Proceedings of IEEE International Conference on Robotics and Automation*. Pp. 3551-3556. Albuquerque, USA . IEEE

Mori, K, Tsukiyama, M and Fukuda, T. (1994). Immune Algorithm and Its Application to Factory Load Dispatching Planning. pp. 1343-1346 of *Proc. Japan-USA Symposium on Flexible Automation.*

Mori, K, Tsukiyama, M and Fukuda, T. (1996). Multi-optimisation by immune algorithm with diversity and learning. *Proc. of the IEEE SMC*, pp. 118-123.

Mori, K, Tsukiyama, M and Fukuda, T (1998). Application of an immune algorithm to multi-optimisation problems. *Electrical Engineering in Japan*. **122/2**. pp. 30-37

Neal, M, Hunt, J and Timmis, J. (1998). Augmenting an artificial immune network. *Proc. of the IEEE SMC*, San Diego, USA. pp. 3821-3826.

Neal, M. (2002). An Artificial Immune System for Continuous Analysis of Time-Varying Data. In *1st International Conference on Artificial Immune Systems (ICARIS)*, pages 75-86, Canterbury, UK.

Perelson, A. S. & Weisbuch, G. (1997). Immunology for Physicists, *Rev. of Modern Physics*, 69(4), pp. 1219-1267.

Perelson, A. S. (1989). Immune Network Theory, *Imm. Rev.*, **110**, pp. 5-36.

Perelson, A. S., Mirmirani, M. & Oster, G. F. (1978). Optimal Strategies in Immunology II. B Memory Cell Production, *J. Math. Biol.*, **5**, pp. 213-256.

Potter, M.A & De Jong, K.A (1998). *The coevolution of antibodies for concept learning*. In Parallel Problem Solving From Nature - PPSN V, pages 530--540. Springer-Verlag.

Potter, M.A. & De Jong, K.A (2000) *Cooperative coevolution: An architecture for evolving co adapted subcomponents*. Evolutionary Computation, 8(1):1--29.

Quinlan, J. (1993). *C4.5: Programs for machine learning*. Pub. Morgan Kaufman.

Roitt, I. (1997). *Essential Immunology: 9th Edition*. Chap. Specific Acquired Immunity, pp. 22-39. Pub. Blackwell Science.

Slavov, V & Nikoleav, N (1998). Immune network dynamics for inductive problem solving. *Lecture Notes in Computer Science*, **1498**, pp. 712-721.

Somayaji, A., Hofmeyr, S. A. & Forrest, S. (1997), Principles of a Computer Immune System, *Proc. of the new Security Paradigms Workshop*, pp. 75-81.

Smith, D. J., S. Forrest & A. S. Perelson (1998). Immunological Memory is Associative. In *Artificial Immune Systems and their Applications*. Ed. D. Dasgupta. Springer-Verlag.

Tew, J & Mandel, T. (1979). Prolonged antigen half-life in the lymphoid follicles of antigen-specifically immunised mice. *Immunology*, **37**, pp. 69-76.

Tew, J, Phipps, P & Mandel, T. (1980). The maintenance and regulation of the humoral immune response. Persisting antigen and the role of follicular antigen-binding dendritic cells. *Immunological Review*, **53**, pp. 175-211.

Timmis, J, Neal, M and Hunt, J. (1999). Data Analysis with Artificial Immune Systems and Cluster Analysis and Kohonen Networks: Some Comparisons. *Proceedings of the IEEE SMC,* Tokyo, Japan. pp. 922-927.

Timmis, J, Neal, M and Hunt, J. (2000). An Artificial Immune System for Data Analysis. *Biosystems*. **55**(1/3), pp. 143-150

Timmis, J. (2000). *Artificial Immune Systems: A novel data analysis technique inspired by the immune network theory*. Ph.D. Thesis. University of Wales, Aberystwyth. 2000.

Timmis, J and Neal, M.(2001) A Resource Limited Artificial Immune System for Data Analysis. *Knowledge Based Systems*, 14(3-4):121-130, June 2001.

Timmis, J (2001). aiVIS: Artificial Immune Network Visualisation. *EuroGraphics UK 2001 Conference Proceedings*, pages 61-69, University College London., April 2001.

Timmis, J, de Lemos, R, Ayara, M and Duncan R. (2002) Towards Immune Inspired Fault Tolerance in Embedded Systems. To appear in the Proceedings of International Conference on Neural Information Processing. Singapore. November 2002.

Tizzard, I. (1988a). *Immunology: An Introduction. 2nd edition*. Chap. The response of B-cells to Antigen. Pp. 199-223. Pub. Saunders College.

Tizzard, I. (1988b). *Immunology: An Introduction. 2nd edition*. Chap. The response of T-cell to Antigen. Pp. 224-260. Pub. Saunders College.

Toma, N, Endo, S & Yamada, K (1999). Immune algorithm with immune network and MHC for adaptive problem solving. *Proc. IEEE SMC*. Tokyo, Japan, **IV**, pp. 271-276.

Varela, F, Coutinho, A, Dupire, B and Vaz, N. (1988). Cognitive Networks : Immune and Neural and Otherwise. *Theoretical Immunology: Part Two, SFI Studies in the Sciences of Complexity*, **2**, pp.359-371

Warrender, C, Forrest, S & Pearhmutter, B. (1999). Detecting intrusions using system calls: Alternative data models. *Proc. of Symposium on Security and privacy*. IEEE. pp. 133-145.

Watanabe, Y, Ishiguro and Uchikawa, Y. (1998). Decentralised behaviour arbitration mechanism for autonomous mobile robots using immune network. In *Artificial Immune Systems and their applications*. Ed. D. Dasgupta. pp. 187-209. Springer-Verlag.

Watkins, A. (2001). A resource limited artificial immune classifier. MS Thesis. Mississippi State University. USA.

Watkins, A and Timmis, J. (2002). Artificial Immune Recognition Systems (AIRS): Revisions and Refinements. In *Proceedings of the 1st International Conference on Artificial Immune Systems.* pages 173-181, University of Kent at Canterbury, September.

Wettschereck, D. Aha, D.W, and Mohri, T. 1997. A review and empirical evaluation of feature weighting methods for a class of lazy learning algorithms. *Artificial Intelligence Review*. 11:273-314.

Zinkernagel, R. M. & Kelly, J. (1997). How Antigen Influences Immunity, *The Immunologist*, **4/5**, pp. 114-120.