

On Artificial Immune Systems and Swarm Intelligence

Jon Timmis · Paul Andrews · Emma Hart

Received date / Accepted date

Abstract This position paper explores the nature and role of two bio-inspired paradigms, namely Artificial Immune Systems (AIS) and Swarm Intelligence (SI). We argue that there are many aspects of AIS that have direct parallels with SI and examine the role of AIS and SI in science and also in engineering, with the primary focus being on the immune system. We explore how in some ways, algorithms from each area are similar, but we also advocate, and explain, that rather than being competitors, AIS and SI are complementary tools and can be used effectively together to solve complex engineering problems.

Keywords artificial immune systems · swarm intelligence · swarm robotics · immunology

1 Introduction

Swarms in nature are generally associated with groupings of insects or animals that collectively exhibit complex behaviours, for example, bees, ants or birds. Perhaps less obviously, it has been suggested (see for example Hoffmeyer 1997; Jacob et al 2006) that the immune system might also be considered a type of swarm system that shares similar properties with other swarm-like systems. This article examines the immune system from the perspective of a swarm system, and in particular the relationship of the fields of swarm intelligence (SI) and artificial immune systems (AIS), disciplines

Jon Timmis
Department of Computer Science and Department of Electronics
University of York, Heslington, York. UK.
E-mail: jtimmis@cs.york.ac.uk

Paul Andrews
Department of Computer Science
University of York, Heslington, York. UK.
E-mail: psa@cs.york.ac.uk

Emma Hart
School of Computing
Edinburgh Napier University, Merchiston Campus, Edinburgh, UK
E-mail: e.hart@napier.ac.uk

which both model and exploit biological systems for engineering. The article takes a position that *artificial immune systems (AIS) and swarm intelligence (SI) on the one hand share many similarities, but proposes they be utilised in a complementary, rather than competitive manner, to help solve complex engineering problems*. This position is reached by examining disciplines from both a scientific and engineering viewpoint. We draw various parallels between the disciplines which support the underlying hypothesis that the immune system should be viewed as an example of a swarm system. However, rather than attempting to place the two paradigms on a competitive footing, we conclude the paper by exploring areas in which the two disciplines can complement and learn from each other, particularly in an engineering context.

The term *swarm intelligence* is used widely to describe the collective behaviour of decentralized, self-organized systems, both natural and artificial, therefore it is useful to ground the discussion in a single definition. We base our discussion on the definition of Dorigo and Birattari (2007) who define SI as:

“the discipline that deals with natural and artificial systems composed of many individuals that coordinate using decentralised control and self-organisation. In particular, the discipline focuses on the collective behaviours that result from the local interactions of the individuals with each other and with their environment”

This definition encapsulates the key properties of a swarm system. Dorigo and Birattari (2007) go on to note that swarm systems occur in both natural and artificial systems. They consider that in the natural world, swarm intelligence studies a wide variety of systems ranging from ant colonies to flocks of birds. From an engineering perspective, they consider swarm intelligence covering systems from multi-robot systems to optimisation. The important concepts in this definition are clarified by Millonas (1994) who further proposes that a swarm system must encapsulate five important principles:

1. The principle of *proximity*
2. The principle of *quality*
3. The principle of a *diverse response*
4. The principle of *stability*
5. The principle of *adaptability*

We use these principles together with the definition of Dorigo and Birattari (2007) to describe the natural immune system and examine the extent to which it can be considered a swarm system.

Like swarm intelligence, *artificial immune systems* is a widely used term within bio-inspired computing research. Although a canonical definition of an AIS does not exist, the domain is well described by the definition of de Castro and Timmis (2002):

“adaptive systems, inspired by theoretical immunology and observed immune functions, principles and models, which are applied to problem solving”

This definition is more generic than the corresponding swarm definition in that it does not specify particular properties of an immune system; however, the natural immune system and immune-inspired artificial systems are generally understood to exhibit learning, memory, be decentralised and distributed, exploit imperfect recognition mechanisms, and operate in a collaborative manner.

According to Dorigo and Birattari (2007), SI is a broad area that can be discussed under two orthogonal classifications: natural vs. artificial (the study of biological systems or human engineered artefact); and science vs. engineering (see table 1). Whilst

	Swarm Intelligence	Artificial Immune Systems
Science	Understand how local individual behaviours result in coordinated population behaviours	Use models to explain phenomena and guide experimental work
Engineering	Exploit the understanding of natural swarms in designing problem solving systems	Apply systems inspired by immune functions, principles and models to problem solving

Table 1 Classification of the role of Swarm Intelligence and Artificial Immune Systems in Science and Engineering

the field of AIS focuses on the artificial engineered artefacts, a similar science vs. engineering view of AIS is proposed by Timmis et al (2008b) who explore diversifications in the field. A comparison of the science vs. engineering classification for SI and AIS is shown in table 1, which exposes a natural relationship between the goals of the two fields. We build on this relationship in this article exploring similarities and differences between SI and AIS, highlighting areas where, despite their similarities, the two fields can complement each other.

The paper is structured as follows. In section 2 we outline the basic concepts behind the immune system and discuss how we can view the immune system in the context of a swarm system taking into account work by Millonas (1994). In section 3 we then go on to discuss an engineering perspective on AIS. We discuss the core properties that underpin all AIS algorithms, and then compare at a *conceptual* level swarm and AIS algorithms under two headings: directly interacting swarms, and indirectly acting (stigmergic) swarms. We use a generic framework for studying population based algorithms proposed by Newborough and Stepney (2005), which enables us to readily discuss similarities and differences between algorithms from both fields. In section 4 we compare and contrast the use of traditional SI algorithms and immune swarms in science, before moving on to examine methodological issues in the development of AIS and SI algorithms in section 5 from an engineering perspective. In section 6, in order to illustrate how SI and AIS might complement each other in producing systems which address complex engineering problems, we focus on a simple case study taken from an EU funded research project, SYMBRION¹, where AIS techniques are being used to improve the fault tolerant capabilities of swarm and collective robotic systems. The paper then concludes in section 7.

It should perhaps be noted that although aimed primarily at an audience more familiar with SI than AIS, the article is not intended to be a primer on immunology or a detailed introduction to AIS algorithms. A number of review papers and text books exist with such detail to which the reader is referred (see Janeway et al 2001; de Castro and Timmis 2002; Timmis et al 2008a; Hart and Timmis 2008). Rather, the goal of the article is to examine synergies and differences between the two fields from a conceptual perspective, in order to identify how the two fields might learn from one another and work together.

2 Swarms and Real Immune Systems

The natural immune system consists of many evolved subsystems, that interact between various levels and over multiple timescales. Agents within the immune system work

¹ <http://www.symbion.eu>

together, to afford a level of protection to the host that maintains a steady operation state (Cohen 2000). Various mechanisms within the immune system, which we discuss briefly below, exhibit self-organising properties associated with SI type systems. We provide a very brief overview at a conceptual level of the components of the immune system, then using the work of Millonas (1994) we explore how the immune system can be considered to exhibit swarm-like properties.

2.1 The Immune System: Simply Stated

The vertebrate immune system is composed of a range of cells and molecules that work together with other systems, such as neural and endocrine, in order to maintain a *steady state* within the host and protect our bodies from infectious agents such as viruses, bacteria, fungi and other parasites. Figure 1 shows a simplistic view of the “levels” within the immune system and basic components therein. Various immune agents (cells and molecules) are capable of identifying other invading agents (pathogens) and, always in collaboration with a number of other immune agents, instigate an *inflammatory* response. At a simple level, we can consider there to be two types of immunity: innate and adaptive. Innate immunity is not directed towards specific invaders, but against general properties of pathogens that enter the body (Janeway and Medzhitov 2002). The innate immune system plays a vital role in the initiation and regulation of immune responses, including adaptive immune responses. Specialised cells of the innate immune system evolved so as to recognise and bind to common molecular patterns found only in microorganisms.

The adaptive immune system is directed against specific invaders, and immune agents within the system are modified by exposure to such invaders (Janeway et al 2001). The adaptive immune system mainly consists of immune agents called *lymphocytes*, more specifically *B-cells* and *T-cells*, which are engaged in the process of identifying and removing pathogenic infections. A mechanism known as clonal selection takes place when specific antigen (any molecule that is recognised by the immune system) is recognised jointly by both B and T-cells; this is a collaborative recognition event, facilitated by both direct recognition between receptors on cells (often referred to as a lock-and-key process) and additionally, indirect or stigmergic communication between cells. During clonal selection, cells undergo a *cloning* process with some differentiating into different functions, for example long-lived *memory cells*, or, in the case of B-cells, *antibody* producing *plasma cells*. B-cells also undergo *mutation* to their antigen recognising receptors, adding variability to their recognition abilities. Memory cells are maintained by the immune system, affording a level of memory of previous antigen-recognition encounters, which provides increased protection to the host over time. However, without appropriate stimulation, this memory will decay over time, leaving the host vulnerable to possible infection by the same, or similar, pathogen.

When considered together, the innate and adaptive immune systems operate over different timescales. The innate immune system operates on a short time scale often initiating a reaction either instantly or within a matter of minutes, whilst the adaptive immune system operates over a longer time period, taking of the order of days to initiate a reaction. It is the combination and interaction of both the innate and adaptive immune mechanisms that provides us with an effective immune system.




Fig. 1 Different layers of the immune system, innate and adaptive, after de Castro and Timmis (2002)

2.2 Immune Swarms

Within the immunology literature, the notion of swarms and swarm intelligence has been used to describe both behaviour at the conceptual (Hoffmeyer 1997) and cellular levels (Egen et al 2008; Chtanova et al 2008). At the conceptual level, the work of Hoffmeyer (1997) draws an analogy between swarm intelligence and the (human) body. Unlike swarms of insects, the body swarm consists of overlapping heterogeneous parts, each of which can itself be viewed as a swarm. The picture is that of a hierarchy of swarms, each engaging in distributed problem solving. The comparison of the body swarm to the swarm behaviour observed in social insects is simply an analogy with the detailed dynamics of each being quite different. One of the swarms in the body is identified as the immune system, which is considered to be a communicative process whereby receptors translate molecular messages to activity patterns inside immune cells. These cells continually interpret the environment based on the collective state of its receptors, with any response based on the history of the cell. The definition of SI presented in section 1 can be naturally applied to many mechanisms observed in the immune system. The main class of actors (agents) of the immune systems are the many populations of immune cells and messenger molecules that interact to achieve collective behaviours through self-organising mechanisms. Interactions of immune cell agents occurs between each other and their environment consisting of other bodily cells and any invading pathogens.

Many examples of swarm like behaviour in the immune system have been observed and documented by immunologists. For example, consider granuloma formation, which occurs in diseases such as tuberculosis. Granulomas are localised inflammatory responses that form around macrophages (a type of immune cell specialising in pathogen ingestion) when the macrophage is unable to completely remove the ingested pathogen and becomes infected (Janeway et al 2001). Other types of immune cells are attracted to the infected macrophage by messenger molecules emitted into the local neighbourhood by the macrophage. These other immune cells surround the infected macrophage constantly moving and interacting with each other and the infected macrophage. The various interactions result in self-organised positive and negative feedback mechanisms that regulate the level of attraction of additional immune cells to the site of infection. This leads to the formation of a stable population of aggregated cells: the granuloma. The purpose of the granuloma is to generate a physical barrier to prevent pathogen dissemination, and facilitate interactions between innate and adaptive immune system components (Egen et al 2008). Eventually the interaction between immune cells surrounding the granuloma results in a collective behaviour that removes the pathogen from the infected macrophage.

Taking the swarm analogy further, Chtanova et al (2008) describe the collective action of neutrophils in terms of swarm dynamics. Neutrophils are a type of immune cell present in large numbers in the blood that respond quickly to the presence of pathogen. Upon infection with an intracellular parasite (*Toxoplasma gondii*), many parasites accumulate in macrophages found in lymph nodes. Shortly after this event, neutrophils aggregate in the lymph nodes, resulting in the removal of the infected macrophages from the lymph node. Chtanova et al (2008) describe the dynamics of the neutrophils in terms of swarming. The neutrophil swarms first develop from the initial arrest of a small number of neutrophils followed minutes later by massive influx. This directed migration (where individual neutrophils move persistently toward the swarm over time) is likely to be caused by direct communication between cells via signalling molecules. These signalling molecules are produced by, and are attracted to, neutrophils. Another distinct behaviour of neutrophils that is identified by the authors is “streaming”: many neutrophils follow parallel paths in and out of the swarm. Based on visual inspection, Chtanova et al (2008) identify two distinct types of neutrophil swarm:

- *Transient*: these contain fewer than 150 cells and last for on average 20 minutes. They grow via rapid co-ordinated migration of neutrophils into the swarm followed by quick dissolution as neutrophils migrate out to join other growing swarms nearby.
- *Persistent*: these contain over 300 neutrophils and grew throughout the period in which they were observed (around 40 minutes). Growth can occur from both individual neutrophil migration and merging with nearby swarms.

Based on observations of each type of swarm, Chtanova et al (2008) hypothesise that neutrophils generate signals to induce swarming and once a swarm reaches a certain size, a large enough signalling centre exists to overwhelm the competing signals of nearby smaller swarms.

Many more examples of specific swarm like behaviours have been observed in the immune systems; instead of describing each individually, we attempt to highlight generic concepts of swarm-like behaviour in the immune system by taking the five principles of swarm behaviour stated by Millonas (1994) and noted in section 1, using them as a framework to explore swarm like behaviours in the immune system.

2.2.1 Principle of Proximity

Millonas (1994) states that swarms must be able to compute, in both time and space. As computation inevitably relates to energy expenditure by the swarm, he notes that that:

“computation must be performed as a direct behavioural response to environmental stimuli which in some sense maximises the utility of the group as a whole of some type of activity”

In typical swarm systems, activities might include nest-building, food-searching or group defence. In the immune system, we observe activities such as tissue-healing, body-maintenance, and response to pathogens. The overriding utility function in each case is survival, and in the case of the immune system this would be the survival of the host organism. This utility function is maximised in the immune system by a collective approach arising from interactions of a diverse range of cells. Dembic (2004) compares the approach to computation of the immune system to that of the neural system, drawing an analogy with neurons in the brain, which can only perform cognition as group rather than individuals. Thus, in the immune system we observe a collective approach to initiating an immune response, based on collaborative interactions of three types of cells: antigen-presenting cells (APCs), T-cells and B-cells. An example of this is the clonal selection process. When a B-cell binds with a suitable antigen, activation of that cell occurs, if a co-stimulatory signal is received from a T-cell that is processing the same type of antigen (which is undertaken with the help of an APC). Once activated, clones of the B-cell are produced expressing identical receptors to the original cell that encountered the antigen. Thus a clonal expansion of the original cell occurs, known as a primary immune response, see figure 2. This ensures that only cells specific to an activating antigen are produced in large numbers. A certain level of cells is maintained over time to act as a form of *memory*, thus affording a fast response to the infection, known as a secondary response. During the expansion process, receptors on cells are mutated, allowing the coverage of pathogenic material to increase, and thus potentially mount a response to an infection that the cell population has not encountered before - this is known as a cross-reactive response (figure 2).

2.2.2 Principle of Quality

The second principle described by Millonas (1994) is that of *quality*. It is noted that a group should be able to respond to quality factors, such as the quality of food source, or the safety of a location. With regard to the immune system, *quality* factors might refer to the seriousness or strength of an attack. The level of attack is indicated by signalling networks (e.g., cytokine and danger signals) which quickly up-regulate to amplify the “news” of an attack, and drive the immune system to respond more rapidly, for example through expansion of cell populations through proliferation. Conversely, attacks which are considered of low quality, that is not detrimental to the overall goals of the system, cause down-regulation of the immune response. The immune system as a collective system is therefore aware of, and driven by, principles of quality.

2.2.3 Principle of a Diverse Response

The principle of diverse response advocates that a group should:

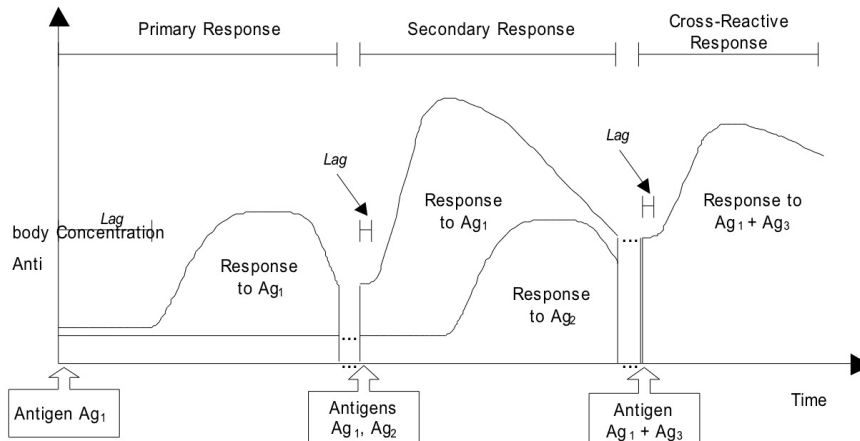


Fig. 2 Varying responses of the immune system over time, after de Castro and Timmis (2002)

“seek to distribute its resources along many modes as insurance against the sudden change in any one of them due to environmental fluctuations” (Millonas 1994)

The immune system can be considered both physically and functionally distributed. In the former case, it provides robustness by distributing its resources throughout the body of the organism it protects, and has evolved such that it has no central point of control. From the latter perspective, functionality is also distributed such that it comprises a diverse range of cells, each of which contribute to, but do not control, the immune response.

Further evidence of the immune system’s approach to “insurance” lies in its *degenerate* nature. This term was initially defined by Edelman and Gally (2001) as:

“the ability of elements that are structurally different to perform the same function or yield the same output”

Within the immune system, degeneracy is most notably typified by the non-specific nature of lymphocyte receptors, in that a single lymphocyte can recognise many different antigens and conversely, that a single antigen can activate different lymphocytes. Another form of degeneracy in the immune response is the two-signal mechanism required for activation of effector mechanisms as discussed above. This much celebrated “fail-safe” procedure is the result of two distinct signals, from structurally and functionally different cells (APC and helper T-cells), derived from the same raw antigen “data” present in the lymph nodes. Of particular relevance here is how the immune system has evolved coupled components with differing, limited views – rather than a single “smart” component – to deploy its lethal effector mechanisms. It is not difficult to speculate on why this would be evolutionary advantageous while in a constant and subtle arms-race with pathogens.

2.2.4 Principle of Stability

The fourth principle of Millonas (1994) notes that:

“no large shifts in behaviour in the system should occur as such shifts take energy”

The immune system has a series of regulatory networks, to attempt (in part) to ensure that large changes in behaviour are somewhat controlled. Immune responses are controlled and regulated by interacting cells, operating in various locations throughout the body and over multiple timescales. It is unusual, but not impossible, to see major shifts in immune behaviour. Of course, catastrophic infection can occur and overwhelm the host, killing it in a very short period of time, but this is not typical. A good example of a regulatory network in the immune system can be seen in response to Experimental Autoimmune Encephalomyelitis (EAE). This is an autoimmune disease in mice which serves as a suitable proxy for multiple sclerosis in humans. The disease causes damage to the central nervous system (CNS) and can lead to paralysis and death. In mice, EAE can be artificially induced by an intervention procedure (the details of which are not relevant here, see Read et al (2009) for more information), and a network of immune cell interactions operates to counter the onset of EAE. The regulatory network consists of various populations of regulatory T-cells, which interact to reduce promotion of a positive feedback loop that would lead to severe damage to the CNS.

2.2.5 Principle of Adaptability

The final principle states that:

“The group should not shift its behaviour from one mode to another upon every fluctuation of the environment, since such changes take energy, and may not produce a worthwhile return for the investment”

An excellent example of this principle in practice observed in the immune system is its ability to reach a collective decision on whether or not to respond by analysing the *context* of any fluctuations observed in the environment. Thus, slight damage occurring to tissues alerts B-cells and T-cells that something is happening (via APCs as discussed in section 2.2.1). If there is an associated message that the function of the tissue has been compromised, then an immune response would ensue (Dembic 2004). On the other hand, should this functional-damage evidence be lacking, the immune system would assume that the matter should be either neglected or exploited; the immune system is able to ascertain whether or not responses to such situations are likely to produce a worthwhile return on investment.

2.2.6 Summary

Garnier et al (2007) argue that key to a decentralised collective decision making is *self-organisation*. This has four key *ingredients*: (1) positive feedback: this is where a change in the environment promotes certain activity, in the case of ants this might be the laying of pheromone; (2) negative feedback: used in order to counterbalance positive feedback, in the case of ants this might be the evaporation of pheromone; (3) amplifications: where a small change has a large effect and (4) multiple direct or stigmergic interactions: individuals must interact in some way.

We have now, briefly, examined both swarm and immune systems. There are clear parallels between the two types of systems. As with collective systems, immune systems are made up of many “units” (such as cells, molecules, tissues), Segal and Cohen (2001)

argue that they can also be viewed as a decentralised system with no overall control. Indeed, one could never say that a single cell has an overall view either the purpose, or goal, of the collection of cells during an inflammatory response. We can classify the immune system as a self-organising system that exhibits the same properties, making use of the same ingredients, as we have outlined above. Table 2 summarises some basic aspects of both swarm and immune systems in this context, outlining certain parallels between the two systems.

	Swarm System	Immune System
Positive feedback	laying pheromone	receptor recognition T-cell signalling
Negative feedback	pheromone evaporation	cell suppression
Amplification	locate new food source	clonal selection
Multiple direct or indirect signals	pheromone in ants visual communication in birds	lock & key recognition cytokine networks

Table 2 Self-organisation in the context of swarm and immune systems

3 Immune Swarms for Engineering

In order to derive engineering algorithms which take inspiration from the natural immune system, de Castro and Timmis (2002) propose a layered framework for engineering AIS. The layered framework takes the application domain of the AIS as its starting point, followed by three layers to be considered before the required AIS is engineered. These layers are:

- *Component representations*: how the components of the system are to be represented
- *Affinity measures*: how the interactions between the components of the system are to be quantified
- *Immune algorithms*: how the components of the systems are going to interact to determine the system dynamics

The first two components (representation and affinity measures) are most easily understood in the context of *shape-space*, which is described in the following section. We then briefly survey the range of algorithms which have been derived using this approach, focussing our attention on those that can be considered swarm-like.

3.1 Shape Space

Perelson and Oster (1979) first introduced the concept of shape-space in relation to research undertaken in theoretical immunology. Their goal was to discretise the physical aspects of antibody combining sites and antigenic determinants into a vector of n shape parameters which encapsulated physical interactions between cells such as geometric structure, electrostatic charge and hydrogen bonding. The resulting n -dimensional vector in the shape-space S represents the generalised shape of a molecule. A contiguous

recognition region surrounding each point accounts for imperfect matching between cells and ligands and receptors that have intersecting regions are said to have *affinity*.

Although biologically simplistic, for theoretical immunologists, the shape-space has a certain heuristic value in quantifying gross properties of the immune repertoire, away from the complex bio-chemical process of protein binding. The same shape-space formalism has also become the de facto method in AIS for representation of problem data or solutions; most AIS algorithms employ a population based approach, and represent data in some *instance based* mechanism. Components in a typical AIS will be represented as attribute strings of n parameters, for example a candidate solution would be represented as the vector $\mathbf{Ab} = \langle Ab_1, Ab_2, \dots, Ab_n \rangle$. The choice of these vector attributes will determine the shape space type. de Castro and Timmis (2002) highlight three main types that can be used in AIS (but this list is not exhaustive):

- *Real-valued shape-space*: attributes of all the components are real numbers
- *Hamming shape-space*: attributes of all the components are from a finite alphabet
- *Symbolic shape-space*: attributes can be of any type, including symbols such as age and name

The choice of attribute types used for candidate solutions will naturally determine the metric in S that can be used to measure the similarity (equivalent to the biological notion of affinity) between these components. It is normal in AIS that each of the n attributes in the component contribute equally to the affinity measure. For real valued shape spaces, this affinity measure is typically Euclidean distance, although similar measures such as Manhattan distance are equally applicable. The affinity measures used for the Hamming shape spaces are determined by the alphabet used. Two of the most popular affinity measures for the binary alphabet are the Hamming distance and the r -contiguous bit rule. The Hamming distance is simply calculated by applying the XOR operator to the two components that are being measured. The r -contiguous bit rule calculates the affinity as the length of the largest contiguous region between the two components. Defining the correct metric can be a non-trivial task and is clearly dependent on the nature of the symbols in S .

Although the key concept of mapping immunological shape-space to a vectorial representation of a data-set has driven much of successful AIS research, recent theoretical investigations have cast doubt on the tractability of the concept (see McEwan and Hart 2009; Stibor et al 2006). Instance based methods rely on two key assumptions; that there are dense regions in the space that can be generalised, compressed or sparsely represented, and that the distance between points is a meaningful proxy for comparison, discrimination and localisation. Unfortunately, these assumptions are based on low-dimensional intuitions, and the validity of both assumptions is a rapidly decreasing function of dimensionality; as the dimensionality of a space increases, its volume increases exponentially faster. In particular, this impacts the crucial affinity, or distance metric, which becomes increasingly meaningless as data-points become equidistant at high volume (Aggarwal et al 2001). This simple fact compounds several issues that undermine classical instance-based methods, and thus, any AIS built upon a shape-space abstraction. Recently, AIS researchers have proposed alternative representational abstractions, e.g., working in the space of feature-feature relations (Nanas et al 2004) or dictionary-based basis decomposition methods (McEwan and Hart 2009).

3.2 Immune Algorithms

Mechanisms occurring in the natural immune system have inspired a plethora of algorithms and models that have been applied to solving engineering problems with varying degrees of success (Hart and Timmis 2008). These algorithms take inspiration from a diverse range of immune functions, that occur across different timescales in the immune system and with varying levels of details. Using the definition of Dorigo and Bittari given in section 1, these immune inspired algorithms divide naturally into two categories; non-swarm and swarm-like algorithms. The most common algorithms falling into the former category are algorithms which are inspired by the immunological processes of *negative selection* and *clonal selection*. Negative selection algorithms model selective processes that take place in the thymus early in development and have been extensively used in anomaly detection; clonal selection algorithms mimic the adaptive immune response and have been applied to a wide range of optimisation and clustering problems (Garrett 2005; Hart and Timmis 2008); another class of algorithms based on *danger theory* has also inspired work in anomaly detection (Greensmith et al 2010). Detailed discussion of these algorithms is beyond the scope of this paper; however, the interested reader is referred to Timmis et al (2008a,c) for a detailed discussion of them.

Instead, we focus on the category of immune-inspired algorithms which can be described as swarm-like. The immune specific terminology commonly used in describing such algorithms can make a direct comparison of immune algorithms to those inspired from swarm systems difficult. We therefore ground the following discussion in a generic framework for population based algorithms, first proposed by Newborough and Stepney (2005), which enables us to abstract the underlying concepts of the algorithms in a unifying structure. Using this framework, we can readily highlight both differences and similarities between immune and swarm inspired algorithms and illustrate when and where the algorithms embody the defining principles of swarm intelligence as articulated by Dorigo and Birattari (2007) and discussed in section 1.

The framework of Newborough and Stepney (2005) was originally proposed as an abstraction of population based bio-inspired algorithms. To construct the framework, an empirical analysis of a range of population based algorithms inspired from different biological paradigms was carried out, investigating each algorithmic component and interaction. The derived framework consists of six generalised stages, which constitute a single algorithm *generation*. These stages, based on the description of Newborough and Stepney (2005), are shown in Figure 3.

It should be emphasised that to allow flexibility within the framework, a new population of candidate solution individuals is created at each generation. However, the framework allows for algorithms where individuals usually survive between generations using the select and spawn steps to create copies of these individuals from one generation to the next. In the next two sections, we use this framework to firstly compare and contrast SI and AIS algorithms which are based on *directly* interacting components, and secondly, those which utilise stigmergic or *indirect* interactions.

3.3 Directly Interacting Swarm Algorithms

One of the most commonly used swarm algorithms that exploits direct interactions between swarm members is the particle swarm optimisation (PSO) algorithm. The algorithm is based on a simplified social model of locally interacting agents, called

1. **Create** : a population of novel individuals is created that represent candidate solutions to the problem being optimised by the algorithm.
2. **Evaluate** : each individual is evaluated based on pre-defined criteria that determine how well it solves the optimisation problem.
3. **Test** : a condition is tested to establish whether the algorithm terminates, returning an individual solution or set of solutions upon termination.
4. **Select** : a set of candidate solution individuals is selected to be used as the basis for the creation of the next generation of individuals.
5. **Spawn** : the new population of candidate solution individuals is generated for use in the next generation.
6. **Mutate** : variability is introduced to the algorithm either via altering a number of individuals of the new population or some other aspect of the algorithm.

Fig. 3 A single algorithm *generation* of the framework by Newborough and Stepney (2005)

particles, that can be viewed simply as a set of vectors whose trajectories oscillate around a region defined by the particles' historical best position and the best position of other individuals Kennedy and Eberhart (2001). The topology of the swarm governs which particles from the swarm can influence a particle; usually the swarm is influenced by the global best particle (*gbest*) as well as a set of particles in the particles' local neighbourhood (*lbest*).

Using the framework, the algorithm can be viewed as shown in figure 4. The key step which differentiates swarm algorithms from population based algorithms such as evolutionary algorithms is step 5: *spawn* in which direct interactions between particles are used to calculate the position of new individuals.

1. **Create** : particles are either initially created with random positions and velocities in the search space. Neighbourhoods can be defined with various topologies such as ring, grid or star.
2. **Evaluate** : usefulness of potential solutions are based on current position coordinate of a particle in solution space.
3. **Test** : upon triggering the termination condition, the single best individual solution is returned as the output of the algorithm.
4. **Select** : all particles are chosen to form the population for the next generation.
5. **Spawn** : new individuals (position and velocity) created from parent and highest affinity neighbour so that particle moves towards best neighbour.
6. **Mutate** : no mutation of an individual typically occurs, but the velocity of an individual undergoes an amount of random alteration which may be considered a type of mutation.

Fig. 4 A representation of a generic Particle Swarm Optimisation algorithm using the framework of Newborough and Stepney (2005)

One immune mechanism in particular has influenced the derivation of a family of swarm-like algorithms since the inception of the AIS field – immune network theory. Immune network theory stems from the notion that immune cells not only recognise foreign cells but also recognise each other, creating a structurally and functionally

plastic network of cells that dynamically adapts to stimuli over time. It is thus the *interactions* between cells that give rise to the emergence of complex phenomena such as memory (Farmer et al 1986) and other functionalities such as tolerance and reactivity (Hart et al 2007, 2009). Whilst controversial from an immunological perspective (see for example Langman and Cohn 1986) these models began to give rise to an interest from the computing community and led to significant development of immune network based algorithms (Timmis 2007). We illustrate the concept of an immune network with the aiNET algorithm, first developed by de Castro and Timmis (2002) and specialised into a series of algorithms for optimisation and data-mining in a variety of domains over the following years (de Castro and Zuben 2001; Coelho and Von Zuben 2006).

The generic aiNET algorithm is shown in figure 5. Solutions are represented in an appropriate shape-space as discussed in section 3.1 and are typically a vector of real-values or a string of bits. Unlike in PSO (figure 4), spawning is directly proportional to solution fitness, with fitter solutions spawning more offspring. Mutation of solutions plays a key role, in contrast to PSO, and is achieved via a *hypermutation* operator, which mutates a solution in a manner inversely proportional to solution fitness. Similarities between PSO and aiNET in terms of swarm-like behaviour are seen in step 6: *mutate* in which the interactions between solutions play an important role. In aiNET, solutions are suppressed based on their *similarity* to other solutions in the swarm; note that similarity is defined in terms of the solution representation and *not* in terms of solution quality. Similarity is calculated through the use of an *affinity* function which measures the *distance* between two solutions. For example, in a vector shape-space, a Euclidean metric is often used. The pairwise affinity between all vectors in the swarm is considered and in cases where the affinity is less than a defined threshold, the solution with lower fitness is removed. In contrast to PSO, this results in a dynamically adjusting population size, alleviating the need to determine a suitable population size by empirical testing. The interaction step functions as a mechanism for maintaining *diversity*, the net effect of which is that the algorithm returns a *set* of solutions, in which local optima can be simultaneously conserved.

1. **Create** : solutions (antibodies) are created with random shape-space receptors, or from those spawned in the previous generation.
2. **Evaluate** : potential solutions are evaluated based on the problem-specific quality function.
3. **Test** : upon triggering the termination condition, the entire population is returned as the output of the algorithm.
4. **Select** : the N best solutions are selected from the total population of existing solutions + any cloned solutions.
5. **Spawn** : clones of the selected individuals are spawned, where the number of clones produced by each individual is proportional to the quality of the individual.
6. **Mutate** : clones are mutated with a probability inversely proportional to their solution quality. *Diversity in the population is increased by considering interactions between all clone pairs; pairwise distances between clone vectors are calculated; if the distance is less than a predefined threshold, the less fit clone is deleted.*

Fig. 5 aiNET algorithm represented using the generic population framework. The key steps which model swarm like behaviours are highlighted in italics.

In both PSO and aiNET, swarm behaviour results from interactions between members of the swarm. In aiNET, the affinity metric and threshold together define a *neighbourhood* of interactions, and thus *implicitly* the topology of the immune network. In contrast, in PSO, the topology of the local neighbourhood of a particle is explicitly defined as part of the algorithm. Kennedy and Mendes (2002) note that the key parameters *lbest* and *gbest* which influence interactions in a swarm can be viewed as social neighbourhoods (Grosan et al 2006). The topology of these neighbourhoods influences information exchange between members of the swarm. Various swarm topologies are investigated for function optimisation problems, including pyramid, small world networks and von Neumann topologies. Their results suggests that von Neumann type neighbourhoods (where neighbours above, below and on each side of a two-dimensional lattice are connected) may be advantageous compared to the topologies commonly found in current practice.

A similar relationship between topology and function has also been noted both in studies of theoretical immunology and in AIS. Hart et al (2007, 2009) investigate how interactions between immune cells in a swarm-like system can result in self-organisation of the shape-space, such that the space is partitioned in tolerant and reactive regions, that is, those regions in which data will be ignored and those in which data will be viewed as anomalous and hence reacted against. Using a simple model in which cells become stimulated and hence persist via interactions with neighbours they show that the likelihood of the space being divided into distinct tolerant and reactive regions depends critically on the choice of neighbourhood. They investigate neighbourhoods defined through various affinity measures based on complementarity, similarity and Hamming distance in bitstring spaces, and show, for example, that in Euclidean space a complementarity affinity measure results in distinct tolerant regions, whereas none can be obtained using a similarity based measure.

3.4 Algorithms based on Indirect Swarm Interactions

Ant Colony Optimisation (ACO) algorithms exploit, as inspiration, stigmergic interactions which occur in natural swarms. A representation of an ACO algorithm using the population framework is shown in figure 6. The individuals that form the population of potential solutions in an ACO algorithm are complete solution paths consisting of a number of path steps. Each path step has an associated length and pheromone level. The ant agents are considered to be solution construction procedures that construct the paths, exploiting pheromone concentrations laid on the paths by other ants in the swarm. Communication between ants is therefore indirect, taking place only via the pheromone trail. In this case, the key steps to the swarm interaction are the creation of solutions through interaction with the pheromone trail in step 1: create, and step 6: mutate, which controls the laying of pheromone according to solution quality and also the gradual decaying of pheromone levels.

An immune mechanism apparent in the innate immune system concerned with “dendritic cell trafficking” also exploits *indirect* interactions in a swarm via chemical signals. In this case, the chemicals are chemokines that form chemical gradients in the body, which can be followed by immune cells, much in the same way that ant colonies utilise pheromone gradients.

Dendritic cells are often referred to as the sentinels of the immune system in that they circulate through the body, scouting for antigenic debris and chemical signals

1. **Create** : *a population of potential solution individuals is created each generation. A potential solution is constructed by an ant agent iteratively following a series of path steps based on pheromone levels until a complete potential solution is generated.*
2. **Evaluate** : the best individual solution is the one with the shortest path.
3. **Test** : upon triggering the termination condition, the single best individual solution is returned as the output of the algorithm.
4. **Select** : no individuals from the current generation are selected for the next as each generation creates its own population from scratch.
5. **Spawn** : no individuals spawned for next generation as none are selected.
6. **Mutate** : *additional pheromone is laid at each path step of solution individual proportionally to how good the solution is, whilst pheromone is also reduced by a decay function.*

Fig. 6 ACO algorithm represented using the generic population framework. The key swarm steps are highlighted in italics.

present in the tissues, and then return that information to the lymph node which aggregates the returned information and determines whether or not an immune response is appropriate. Receptors present on the surfaces of macrophages resident in tissue cells are activated by bacterial pathogen and stimulate the release of chemokines which attract dendritic cells to the potentially infected area. Dendritic cells (DCs) collect antigenic material and also various other contextual signals (often termed danger or safe signals) present in the tissue. This causes them to differentiate into mature, or semi-mature states, at which point they migrate back to the lymph node to present their information, and importantly the context in which the information was found. The migration process releases further chemokines which recruit further DCs to the infected site. Once at the lymph node, the cells present their information to T-cells which induces either tolerance or a response. Activated T-cells return to the infected site again via chemokine gradients induced by the differentiated DCs. In this case, the swarm interacts both *indirectly*, that is due to stigmergy, through chemokines and *directly* due to T-cell/dendritic-cell interactions in the lymph node.

This mechanism has inspired work in the domain of self-organising wireless sensor networks² (WSNs). Although very much work in progress, a developing line of work by Davoudani et al (2007) draws an analogy between the requirement for the body to be able to sense its state and react appropriately to acquired information, and the necessity for self-organising wireless sensor networks to be able to monitor their state and react accordingly. Davoudani et al (2007) focus on Specknets (Arvind and Wong 2004), which are self-organising networks of minute sensor devices which are capable of individually performing limited computation and processing but in which there is no central controller. Hart and Davoudani (2009); Davoudani and Hart (2008); Davoudani et al (2008, 2007) utilise circulating radio messages that mimic the functionality of dendritic cells in collecting information from nodes in the network. This information is returned to specks designated as lymph nodes by exploiting artificial chemical gradients for path-finding. In the natural immune system, dendritic cells migrating from areas of

² Note that this mechanism should not be confused with the *dendritic cell algorithm (DCA)* of Greensmith et al (2010) which is also inspired by dendritic cell behaviour but does not consider any interactions between dendritic cells and therefore cannot be considered swarm-like.

infection back to the lymph node release a further chemical signal which attracts new immature dendritic cells to the area thus creating a positive feedback loop; although this mechanism has yet to be implemented in the WSN work, it clearly remains ripe to be exploited. Information collected at the lymph node is interpreted through direct recognition of the information by artificial T-cells, and acted upon appropriately.

4 Immune Swarms for Science

Although the focus of the swarm intelligence community is generally on developing systems and algorithms inspired by swarm behaviours that can be applied to solve engineering problems, it is also apparent that developing methods for better understanding biological systems can potentially lead to breakthroughs in our understanding of science and medicine. The swarm community – composed in the majority of engineers and computer scientists – can contribute to this effort in two ways:

1. By building computational models specifically to understand biological systems
2. By using artificial systems as a proxy for developing understanding of biological systems

With regard to the former approach, agent-based models (ABM) lend themselves well to modelling swarms. In a swarm ABM each entity, or agent, represents a single organism in the swarm; in an immune ABM, entities represent single cell or pathogens. In both cases, a set of simple rules encode the behaviour of entities and their interactions with other agents. In SI, modelling tools have been very successful in explaining observed phenomena in natural swarming systems. We refer the reader again to the book of Bonabeau et al (1999) for a plethora of examples of models which explain a range of diverse phenomena. Further resources, including bibliographies and software tools can also be found at SwarmWiki (ongoing). From an immunological perspective, computational modelling has been used to construct synthetic computer models that capture the relevant phenomena. This approach has the advantage of enabling experiments to be performed using a model which is much easier than experimenting on a natural system. Forrest and Beauchemin (2007) suggest that using ABM to model the immune system has the following advantages:

- It is much easier to isolate mechanisms and test hypotheses about how they function and therefore determine their significance to the overall system
- Synthetic models can be used to integrate specialised models for different phenomena into one system to see how they interact
- It enables researchers to identify gaps and inconsistencies in their knowledge by making assumptions explicit, allowing them to make predictions, generate new hypotheses, and suggest new experiments
- They can be used to predict the effects of therapeutic interventions
- They simplify the work of testing alternative hypotheses, designing experiments, and discovering both necessary and sufficient mechanisms to explain observed behaviour by trying out variations within a framework and adding complexity incrementally.

By way of illustration, we briefly mention a number of complex simulation tools being developed by computer scientists, all of which are far too large to be reviewed in any detail here. IMMS:VIGO:3D (Jacob et al 2006) was developed by the Evolutionary

and Swarm Design Laboratory at the University of Calgary. It is a swarm-based simulation environment used to demonstrate results for clonal selection and primary and secondary collective responses after viral infection, as well as the key response patterns encountered during bacterial infection. IMMSIM (Kleinstei 2000) employs a bitstring representation of cells and a 2D grid, and has been used for example to investigate phenomena such as the dynamics of HIV infection (Bernaschi and Castiglione 2001), vaccine efficiency (Kohler et al 2000) and the rheumatoid factor paradox (Kohler et al 2000). CyCells (Warrender 2004) is a freely available simulator which represents cells discretely on a 3D grid, but molecular concentrations continuously, and has been used to investigate hypotheses about the maintenance of peripheral macrophage population sizes in the lung (Warrender et al 2004). Work by Salazar-Bañuelos (2009) serves to illustrate that even the simplest of agent-based models can aid understanding of complex systems; the author, a practising surgeon, has used a simple agent-based model to shed light on observed patterns of inflammation that take place in the body. Results offer an explanation for why self-tissue is treated as non-self when injury occurs. The reader is referred to Forrest and Beauchemin (2007) for a detailed review of application of ABM modelling to biomedicine.

Artificial systems are also being used as a proxy for developing understanding of biological systems. Swarm systems are composed of large numbers of dynamically moving, interacting components which makes them notoriously difficult to both study and understand if focusing only on the natural systems themselves. Forrest and Beauchemin (2007) discuss how artificial systems on the other hand can be studied and analysed much more easily. They propose that the study of artificial systems might lead to uncovering of phenomena that might be difficult to uncover experimentally. Swarm robotic platforms have already been used to study insect behaviour. Halloy et al (2007) and Sempo et al (2006) both propose to make use of robotic cockroaches in societies of actual cockroaches to see if it is possible to develop a mixed society. Through the use of a chemical *lure* they were able to demonstrate that marked robots (with the *lure*) did not effect in any significant way the resting patterns of actual cockroaches, when compared to unmarked robots which did effect the behaviour. The work showed how important chemical recognition of the robots was to the success of being able to integrate robots into a nest of cockroaches and provided insight into the behaviour of cockroaches collective response to other agents in their society. This work is interesting, as it helped provide insight into the development of a robotic system that was able to interact with an actual cockroach population. To our knowledge, there is no example of such use of engineering to provide insights into immunology in the literature.

5 Approaches to Developing Swarm Systems

Artificial swarm systems derive their inspiration from processes that occur in nature, whether that be in the immune system or in systems of social insects or animals. In order to best exploit these inspirations in algorithms, it seems evident that developing a detailed understanding of the natural system under consideration should lead to better algorithms; while on the one hand it is clear that algorithms do not have to faithfully replicate the exact processes that occur in the natural system, on the other hand it is also necessary to avoid the tendency to “reason by metaphor” expressed by Stepney et al (2006), which they suggest can lead to uninspired algorithms.




Fig. 7 The Conceptual Framework, after Stepney et al (2006)

As argued by Stepney et al (2006) bio-inspired algorithms, such as AIS, are best developed in a more principled manner that better captures the properties of the biological system that one intends to mimic in some way. Their proposed conceptual framework attempts to capture biological richness and complexity but, at the same time, appreciates the need for sound engineered systems that need to work. They argue that this should help avoid the “reasoning by metaphor” approach prevalent in bio-inspired computing whereby algorithms are typically a weak analogy of the biology on which they are based, being developed directly from (often naïve) biological models and observations, (Timmis 2007).

A major issue, of course, in designing bio-inspired algorithms is understanding not only which aspects of the biology are necessary to generate the required behaviour and which aspects are surplus to requirements, but also the appropriate level of abstraction at which to work. Often people observe a biological system and identify a property they wish to include in their engineered system, but this property may well be an artefact of *physical* properties, and not *logical* ones, thus proving very difficult to replicate in an artificial domain (Freitas and Timmis 2007). To try and counter this, the conceptual framework advocates an interdisciplinary approach, involving the design of bio-inspired algorithms through a series of observational and modelling stages in order to identify the key characteristics of the underlying biological processes on which the algorithm will be based. The first stage of the conceptual framework, as outlined in figure 7, aims to probe the biology, utilising biological observations and experiments to provide a partial view of the biological system from which inspiration is being taken. This view is used to build abstract models of the biology. These models can be both mathematical and computational, and are open to validation techniques not available to the actual biological system. From the execution of the models and their validation, insight can be gained into the underlying biological process. It is this insight that leads to the construction of the bio-inspired algorithms. This whole process is iterative, and can also lead to the construction of computational frameworks that provide a suitable structure for specific application-oriented algorithms to be designed from.

As noted by Stepney et al (2006) each step in the standard conceptual framework is biased, be it modelling some particular biology mechanism or designing an algorithm for which there is an intended end product or specific concept. The first instantiations of the conceptual framework will produce models specific to certain biological systems and algorithms for solutions to specific problems. One could attempt to

produce a computational framework based on some biology without a particular end algorithm/application in mind, that is examining biology and hoping to come across something applicable to a generic computational problem. This, however, would seem to be a very difficult task and one has to ground the development of AIS in some form of application at some point. Therefore, it is far easier to orient these steps toward some particular problem giving necessary focus to the modelling work (Freitas and Timmis 2007). In addition, it should be noted that if this framework is to work to full effect, then some form of feedback into the biological area under study should be possible.

In the area of AIS, this has been taken up in the guise of *immuno-engineering*, proposed by Timmis et al (2008b). The authors adopt Orosz's definitions of immunology and immuno-informatics (Orosz 2001), as follows:

Immuno-ecology : “the study of immunological principles that permit effective immunological function within the context of the immensely complex immunological network . . . the principles serve mainly to provide an infrastructure for the immune system.”

Immuno-informatics : “the study of the immune system as a cognitive, decision-making device . . . addresses mechanisms by which the immune system converts stimuli into information, how it processes and communicates that information, and how the information is used to promote an effective immuno-ecology . . . how the immune system generates, posts, processes, and stores information about itself and its environment”

The authors then go on to define:

Immuno-engineering : the abstraction of immuno-ecological and immuno-informatics principles, and their adaptation and application to engineered artefacts (comprising hardware and software), so as to provide these artefacts with properties analogous to those provided to organisms by their natural immune systems.

The authors argue that immuno-engineering should take into account the differences between artificial systems and biological systems. They cite an example from Orosz (2001) who hypothesises that the major design features of the biological immune system that provides speed, flexibility and multiple response options rely on a parallel-processing system which has “wasteful” use of resources, countless back-up systems, and requires the ability to immediately and continuously monitor physical sites. Clearly, this is of crucial importance to an engineer, who may well be constrained by processing speeds, communication overheads, and physical resources: therefore great care needs to be taken in identifying the appropriate processes and abstraction to take into account when developing the immune-inspired system.

5.1 Methodologies for Deriving Swarm Algorithms

The above discussion highlights the importance of understanding the biological system to be used as inspiration for algorithm development. Usually this is achieved through the use of some kind of modelling tool. Models can inform artificial systems in many ways. At one end of the spectrum, models can be used to understand *concepts* which can then be mapped onto artificial systems. At the other end, models might be directly translated into algorithms. We find many examples of the former approach in the swarm literature; the same is also true in AIS literature. However, it is apparent that

the immune literature also contains some examples of the latter approach where models have explicitly driven algorithm development.

5.1.1 *Concept Driven Algorithm Development*

As an example of the former approach, consider stigmergic effects observed in swarms e.g., the manner in which ants deposit pheromone when foraging, or deposit items such as grains of sand at initial random locations. Many models have been constructed to investigate these phenomena, see for example (Bonabeau et al 1999), all of which illuminate a simple, but critical point for artificial systems, i.e., that *indirect communication* can be used efficiently to achieve collective goals. Similarly, the concept of incremental construction used by termites, and also modelled in (Bonabeau et al 1999) is readily mapped to the optimisation domain. Gazi and Passino (2003) note that the study of automated systems such as autonomous multi-robot or multi-agent systems is likely to benefit from modelling of biological principles such as swarms and coordination strategy specification, and further from the analysis of such models, e.g. to show that certain group dynamics can achieve group goals. They present a mathematical continuous time swarm model of foraging swarms in the presence of an attractant/repellent profile and analyse the stability properties of the model. They suggest that the model directly addresses the problem of coordination of agents and interactions with the environment based on simple potentials, and that their analysis is therefore a contribution to literature on multi-agent systems. They also suggest that their model might be directly relevant to applications such as undersea exploration by robots in which the agents need to follow a gradient of some substance.

In the domain of AIS, the seminal work of Farmer et al (1986) using mathematical models inspired a variety of immune-network algorithms, for example (Timmis and Neal 2001). Models of self-assertion in idiotypic networks by Bersini (2001); Hart (2005) are further developed by Dilger and Strangfeld (2006) who suggest that the emergent phenomena observed in separation of the shape-space can be exploited in machine-learning algorithms. Agent-based models are utilised by Hart and Davoudani (2009) to illustrate the properties that emerge from dendritic cell trafficking. The authors discuss the necessity of altering the agent-based model, originally based on biological principles, to adapt it to the requirements of a highly constrained artificial system. Thus, although the model becomes biologically implausible, it is used to successfully derive a protocol for sampling in WSNs (as discussed in section 3.4).

5.1.2 *Explicit Derivation of Algorithms from Models*

There is limited literature in both AIS and SI where algorithms have explicitly been derived from plausible biological models. In the context of SI, early work by Deneubourg et al (1991) proposed a Monte Carlo model to attempt to illustrate sorting behaviour in ants. Whilst their motivation was not to explain directly ant behaviour, they did compare their model with actual ant behaviour observed in the lab. Indeed, the authors state “*Our aim in this article is not to prove that the model is actually how the ants behave, but to show that such an algorithm both works and could be used by a team of robots*” (Deneubourg et al 1991). The observation was that the model reproduced, in a qualitative manner, similar types of sorting as observed by the actual ant colony. The authors discuss how this model may be exploited in a robotics context. Work by Lumer and Faieta (1994) further developed such an ant-like sorting approach to investigate the

role of diversity in a population of “ants”, and showed that diverse populations of “ants” perform clustering better than an homogeneous set of “ants”. From these models, a general algorithm could be derived for clustering. This modelling work formed the basis of algorithm development in work by Handl et al (2006) where the authors modified the general algorithm specifically to clustering and topographical mapping, and unlike previous work, make a detailed analysis of the performance of that algorithm and interesting find that the algorithm is not that well suited to topographical mapping.

Moving into the AIS domain, work by Owens et al (2009) has developed an anomaly detection system that was derived from computational models of T-cell signalling events developed in Owens et al (2008). Owens et al (2009) developed an AIS based on properties of T-cell receptors, and has the ability to identify anomalies in the data that deviate from some notion of normality and those where the structure of the data changes over time. Of key importance to their work, was the modelling phase, as discussed in Owens et al (2008). This phase provided key insight into the signalling properties of the T-cells, that not only highlighted deficiencies in early T-cell models, but highlighted the key *computational* aspects of the signalling process that would be most useful in an algorithmic context. From the model, Owens et al (2009) proposed a generalised receptor that is capable of identifying small changes in sensor data. This is achieved by a combination of positive and negative feedback mechanisms that result in generating an *activation* signal that exceeds a specified threshold, dependent on the strength of input. One strength of this system is the fact that it can be mapped to a kernel estimation function, so the theoretical aspects of the algorithm are well understood.

Work by Andrews (2008) undertook an extensive study of the use of modelling in the development of immune-inspired algorithms. They examined various immune theories in order to extract logical principles for algorithm development. A number of conclusions were taken from that work. In particular, an important conclusion was that whilst it is possible to develop algorithms by first constructing abstract models of biological phenomena, having an idea as to the possible algorithm application area would assist greatly by positively biasing model design. Without such bias it is easy to get lost in the depth and breadth of possibilities immunology gives you to model. Therefore taking into account the application area allows for the identification of particular properties of interest, which then permits the undertaking of a principled mapping between the said application area and an area of immunology that might be relevant. Timmis et al (2008b) proposed Immuno-engineering as one possible way of avoiding the issue of modelling without any kind of focus, and to take seriously the nature of the application area.

In summary, it is clear that the design of algorithms or artificial systems can be informed through modelling. However we have shown that a caveat applies to this statement; the literature also describes how this process can be fraught with difficulty. Stepney et al (2006) note that great care must be taken to know what you want to model, taking into account the application area, and to extract logical properties from the model that might be useful in the computational domain. Finally, Freitas and Timmis (2007) warn against assuming physical properties will map well into the logical domain.

6 Swarm vs. Immune: Complementing not Competing

Algorithms from AIS and SI have been applied to similar types of problems, from optimisation to machine learning (Hart and Timmis 2008), and in the many papers in each domain it is common to see an immune network algorithm compared with the performance of a PSO algorithm. However, of course we know that different algorithms will perform differently on different problems, and that there is no single algorithm that will outperform every other algorithm (Wolpert and Macready 1997). For the purposes of this paper, we prefer to focus, not on which types of problems one approach might perform better than another, but on where they might be able to work together, and complement each other in order to solve a complex engineering problem. We examine the area of swarm robotics, where we feel there is a large amount of scope for such complementarity.

From an engineering standpoint, the design of complex distributed systems based upon the swarm intelligence paradigm is compelling but problematic. Due to the characteristics of the system itself, no hierarchical command and control structures are needed and hence no common failure point, or vulnerability, can be identified. Each individual agent makes decisions autonomously, based upon local sensing and communication (Bonabeau et al 1999). Systems with these characteristics could, potentially, exhibit very high levels of robustness, in the sense of tolerance to failure of individual agents and much higher levels of robustness (Winfield et al 2006). Simple units can, in fact, collaborate in achieving their common goal without the need of being aware of the rest of the group. Resilience achieved in this way makes the paradigm very appealing in many applications; however one or more faulty robots may jeopardise the success of the overall mission. As noted by Winfield et al (2006), there are two reasons for undesirable behaviours in swarm robotics: random errors, or systematic (design) errors. Random errors are those due to hardware or component faults. The likelihood that random errors cause undesirable behaviours can be reduced, in the first instance, by employing high reliability components (Winfield et al 2006). However, these systems also need to be fault tolerant, through redundancy for example. In this section, we review work where a complementary approach combining ideas from SI and AIS has been developed in the context of swarm robotics research.

6.1 100 Robots, 100 Days

A major challenge in swarm robotics, is to achieve long-term autonomy. To this end, a series of challenges have been put forward as part of the SYMBRION/REPLICATOR projects³, one of which focusses on achieving long-term operation of a swarm of robots, size of 100, for a period of 100 days (Kernbach et al 2010). Central to the project, is the ability of the robots to join together to form *organisms* of robots, where many robots join together to form a single structure. For more information on the actual robots the reader is referred to Levi and Kernbach (2010). The underlying idea of this challenge is for a collection of 100 robots to be left in a room, unattended, for a period of 100 days, with the objective for as many robots to keep operating for as long as possible, whilst undertaking some simple task, such as exploration. Robots have the ability to recharge via sockets located in the wall of the room, and of course have facilities and the ability

³ <http://www.symbion.eu>

to dock with such sockets. In order to do this, the robots will have to self-configure and re-configure as appropriate to extract as much energy from the environment as possible. Clearly, this challenge will be unlikely to be met in the lifetime of the project, but it provides an interesting testbed to drive the development of SI and related issues.

Key to the survival of the swarm, is a tolerance to mechanical and power failure. AIS techniques are being used to detect errors in the robotic units, and also to predict errors before they occur, thus affording a greater level of operational time for the unit. If certain errors can be identified and circumvented, in principle, this will increase the longevity of operation of the swarm. This is an example where techniques from various areas, not only SI and AIS, can be brought together in a complementary manner for the solving of a very complex, engineering issue.

6.2 AIS Framework

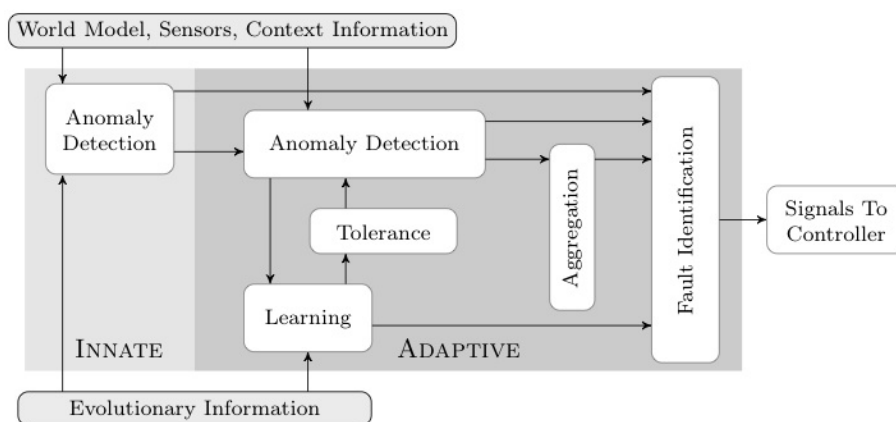


Fig. 8 Framework for immune-inspired fault tolerance in swarm robotic systems, after Timmis et al (2010)

Timmis et al (2010) outline a framework for an AIS to be developed on a SYMBRION robot, see figure 8. The framework has been divided into two separate levels: innate and adaptive, around the concept of an artificial lymph node. As we have outlined earlier, the innate level is akin to a pre-programmed ability to identify specific patterns which then results in specific responses. The adaptive level is akin to a non-specific response which adapts over the life-time of the host, which in some cases may improve over time. This analogy is carried over into this framework, with a combination of simple, pre-programmed type responses, and ones that adapt over time. It is also worth noting that the AIS is to be encoded in a genome that is held within the robotic unit, and this itself will adapt over time.

The innate component consists of a process of anomaly detection that takes data from internal and external sources. The innate system only adapts on an evolutionary timescale, which is directed by the genome. As such, the transfer of immunological genetic information between robotic units can be implemented. Timmis et al (2010)

make use of the innate immune system antigen presenting cell (APC), or Dendritic Cell (DC) analogy (Mokhtar et al 2009). This layer allows for an effective “filtering” of the data for detecting potential anomalies in the data stream. The innate layer has the ability to *correlate* a variety of input sources over time and identifies the presence of an error based on the correlations of these signals. The innate layer also has the ability for prioritising an alarm, by indicating the severity of the danger and *learning* by the incorporation of new knowledge over time on an evolutionary timescale via changes at the genome level.

The adaptive artificial layer is analogue to the collaborative effort of B-cells and T-cells, key actors in the natural immune system’s adaptive layer as we have outlined above in section 2.1. The system, therefore, employs two distinct and complementary methods of identifying anomalies, one taking inspiration from the behaviour of B-cells, the other from the operation of T-Cells. The B-cell approach maintains a B-cell population and is akin to an instance based learning approach where actual instances of data are stored in a feature vector and are used to compare against the current readings from the robot sensors. They employ a time window, similar to the approach taken in (de Lemos et al 2007) where we build up a population of detectors capable of identifying potential anomalies based on examples seen in the past and learnt over time. The T-cell approach takes inspiration from the T-cell receptor and performs estimates of the densities associated with the distributions of sensor values, based on work by Owens et al (2009).

Finally, not only will the system be able to detect errors and make recommendations to the control system, the robotic swarm is endowed with a collective behaviour of *aggregation* in order to isolate a faulty robot, and in some cases allow for the repair of the robot, both when operating in swarm mode or organism mode: this affords a *self-healing* property on the system. For this, inspiration is taken from the immunological process known as granuloma formation as described in section 2.2. Applying the granuloma formation concepts to collective robotics allows the authors to perform two related tasks. First, the isolation of faulty robots, which occurs after the initial identification of a potential fault. The authors propose a case when a permanent fault is located in a robot, and the robot ceases to move. They assume that certain visual signals can be sent by the robot which other functional robots nearby can recognise. These functional robots are then attracted towards the faulty robot, similar to how T-cells are attracted by cytokines emitted by an infected macrophage. A limited number of these robots then isolate the fault robot, akin to T-cells surrounding an infected macrophage, but still move around the fault robot so that other functional robots in the swarm are no longer drawn to the “anchor” point that could be the faulty robot. This approach would be ideally used when certain repairs could be initiated by the robots themselves, such as certain power failures where a re-charge may well be sufficient to help complete a task (Humza et al 2009).

7 Conclusions

We have argued that immune system can be viewed as a swarm like system: there are many interacting agents that operate on multiple timescales that collectively maintain the host, through a process of collaboration and competition. We have discussed how we can view both AIS and SI in the context of natural vs. artificial (the study of biological systems or human engineered artefact); and science vs. engineering. From our discus-

sions, we conclude that even though there are similarities both at the methodological level and at the algorithmic level, when viewed as two approaches that can complement each other, rather than compete with each other, new possibilities for designing and implementing artificial systems become possible. In addition, each area may well learn from the other, in terms of methodology, modelling approaches and engineering design. This paper was intended to highlight the potential for collaboration between the two areas of research, highlight similarities and encourage new interactions between disciplines, as there is more that joins these two fields together, than separates them.

Acknowledgements Paul Andrews is supported by EPSRC grant number EP/E053505/. Jon Timmis would like to acknowledge support from EU grant number FP7-ICT- 2007.8.2, EOARD grant number FA-8655-07-3061 and EPSRC grant number EP/E053505/.

References

- Aggarwal C, Hinneburg A, Keim D (2001) On the surprising behavior of distance metrics in high dimensional space. In: Van den Bussche J, Vianu V (eds) Database Theory – ICDT 2001, 8th International Conference, Proceedings, Lecture Notes in Computer Science, vol 1973, Springer, Berlin, Germany, pp 420–434
- Andrews PS (2008) An investigation of a methodology for the development of artificial immune systems: A case study in receptor degeneracy. PhD thesis, Department of Computer Science, University of York, UK
- Arvind D, Wong K (2004) Speckled computing: Disruptive technology for networked information appliances. In: IEEE International Symposium on Consumer Electronics, 2004, ISCE 2004, IEEE Press, Piscataway, NJ, pp 219–223
- Bernaschi M, Castiglione F (2001) Design and implementation of an immune system simulator. *Computers in Biology and Medicine* 31:303–331
- Bersini H (2001) Self-assertion versus self-recognition: A tribute to francisco varela. In: Timmis J, Bentley PJ (eds) Proceedings of the First International Conference on Artificial Immune Systems, ICARIS 2002, University of Kent Printing Unit, Kent, UK, pp 107–112.
- Bonabeau E, Dorigo M, Theraulaz G (1999) *Swarm Intelligence: From Natural to Artificial Systems*. Oxford University Press, New York, NY
- de Castro LN, Timmis J (2002) An artificial immune network for multimodal function optimization. In: IEEE Congress on Evolutionary Computation, 2002, CEC 2002, IEEE Press, Piscataway, NJ, pp 699–674
- Chtanova T, Schaeffer M, Han S, van Dooren G, Nollmann M, Herzmark P, Chan SW, Satija H, Camfield K, Aaron H, Striepen B, Robey E (2008) Dynamics of neutrophil migration in lymph nodes during infection. *Immunity* 29(3):487–496
- Coelho G, Von Zuben FJ (2006) omni-aiNet: An immune-inspired approach for omni optimization. In: Bersini H, Carneiro J (eds) Artificial Immune Systems, 5th International Conference, ICARIS 2006, Lecture Notes in Computer Science, vol 4163, Springer, Berlin, Germany, pp 294–308
- Cohen IR (2000) *Tending Adam’s Garden: Evolving the Cognitive Immune Self*. Elsevier Academic Press, London, UK
- Davoudani D, Hart E (2008) Computing the state of specknets: an immune inspired approach. In: International Symposium on Performance Evaluation of Computer and

-
- Telecommunication Systems, 2008, SPECTS 2008, IEEE Press, Piscataway, NJ, pp 52–59
- Davoudani D, Hart E, Paechter B (2007) An immune-inspired approach to speckled computing. In: de Castro LN, Von Zuben FJ, Knidel H (eds) *Artificial Immune Systems, 6th International Conference, ICARIS 2007*, Lecture Notes in Computer Science, vol 4628, Springer, Berlin, Germany, pp 288–299
- Davoudani D, Hart E, Paechter B (2008) Computing the state of specknets: Further analysis of an innate immune-inspired model. In: Bentley PJ, Lee D, Jung S (eds) *Artificial Immune Systems, 7th International Conference, ICARIS 2008*, Lecture Notes in Computer Science, vol 5132, Springer, Berlin, Germany, pp 95–106
- de Castro LN, Timmis J (2002) *Artificial Immune Systems: A New Computational Intelligence Approach*. Springer, London, UK
- de Castro LN, Zuben FJV (2001) aiNet: An artificial immune network for data analysis. In: Abbass HA, Sarker RA, Newton CS (eds) *Data Mining: A Heuristic Approach*, Idea Group Publishing, Hershey, PA, pp 231–259
- de Lemos R, Timmis J, Forrest S, Ayara M (2007) Immune-inspired adaptable error detection for automated teller machines. *IEEE Transactions on Systems, Man and Cybernetics Part C: Applications and Reviews* 37(5):873–886
- Dembic Z (2004) Response to Cohn: The immune system rejects the harmful, protects the useful and neglects the rest of microorganisms. *Scandinavian Journal of Immunology* 60:3–5
- Deneubourg JL, Goss S, Franks N, Sendova-Franks A, Detrain C, Chretien L (1991) The dynamics of collective sorting: Robot-like ants and ant-like robots. In: *Proceedings of the First International Conference on Simulation of Adaptive Behavior: From animals to animats*, MIT Press, Cambridge, MA., pp 356–365
- Dilger W, Strangfeld S (2006) Properties of the Bersini experiment on self-assertion. In: *Genetic and Evolutionary Computation Conference, GECCO 2006*, Proceedings, ACM, New York, NY, pp 95–102
- Dorigo M, Birattari M (2007) Swarm intelligence. *Scholarpedia* 2(9):1462
- Edelman GM, Gally JA (2001) Degeneracy and complexity in biological systems. *Proceedings of the National Academy of Science (PNAS)* 98(24):13,763–13,768
- Egen JG, Rothfuchs AG, Feng CG, Winter N, Sher A, Germain RN (2008) Macrophage and T cell dynamics during the development and disintegration of mycobacterial granulomas. *Immunity* 28:271–284
- Farmer JD, Packard NH, Perelson AS (1986) The immune system, adaptation, and machine learning. *Physica D* 22:187–204
- Forrest S, Beauchemin C (2007) Computer Immunology. *Immunol Rev* 216(1):176–197
- Freitas A, Timmis J (2007) Revisiting the foundations of artificial immune systems for data mining. *IEEE Transactions on Evolutionary Computing* 11(4):521–540
- Garnier S, Gautrais J, Theraulaz G (2007) The biological principles of swarm intelligence. *Swarm Intelligence* 1(1):3–31
- Garrett S (2005) How do we evaluate artificial immune systems? *Evolutionary Computation* 13(2):145–177
- Gazi V, Passino KM (2003) Stability analysis of swarms. *IEEE Transactions on Automatic Control* 48:692–697
- Greensmith J, Aickelin U, Tedesco J (2010) Information fusion for anomaly detection with the dendritic cell algorithm. *Information Fusion* 11(1):21–34
- Grosan C, Abraham A, Chis M (2006) Swarm intelligence in data-mining. *Studies in Computational Intelligence (SCI)* 34:1–20

- Halloy J, Sempo G, Caprari G, Rivault C, Asadpour M, tache F, Said I, Durier V, Canonge S, Ame J, Detrain C, Correll N, Martinoli A, Mondada F, Siegwart R, Deneubourg J (2007) Social integration of robots into groups of cockroaches to control self-organized choices. *Science* 318(5853):1155–1158
- Handl J, Knowles J, Dorigo M (2006) Ant-based clustering and topographical mapping. *Artificial Life* 12:36–61
- Hart E (2005) Not all balls are round: An investigation of alternative recognition-region shapes. In: Jacob C, Pilat ML, Bentley PJ, Timmis J (eds) *Artificial Immune Systems, 4th International Conference, ICARIS 2005, Lecture Notes in Computer Science*, vol 3627, Springer, Berlin, Germany, pp 29–42
- Hart E, Davoudani D (2009) Dendritic cell trafficking: From immunology to engineering. In: Andrews PS, Timmis J, Owens N, Aickelin U, Hart E, Hone A, Tyrrell A (eds) *Artificial Immune Systems, 8th International Conference, ICARIS 2009, Lecture Notes in Computer Science*, vol 5666, Springer, Berlin, Germany, pp 11–13
- Hart E, Timmis J (2008) Application areas of AIS: The past, the present and the future. *Applied Soft Computing* 8(1):191–201
- Hart E, Bersini H, Santos F (2007) How affinity influences tolerance in an idiotypic network. *Journal of Theoretical Biology* 249(3):422–436
- Hart E, Bersini H, Santos F (2009) Structure vs function: a topological perspective on immune networks. *Natural Computing* doi:10.1007/s11047-009-9138-8:(on-line first)
- Hoffmeyer J (1997) The swarming body. In: *Semiotics Around the World, Proceedings of the Fifth Congress of the International Association for Semiotic Studies*, Mouton de Gruyter, Berkeley, CA, pp 937–940
- Humza R, Scholz O, Mokhtar M, Timmis J, Tyrrell A (2009) Towards energy homeostasis in an autonomous self-reconfigurable modular robotic organism. In: *Computation World 2009*, IEEE Press, Piscataway, NJ, pp 21–26
- Jacob C, Steil S, Bergmann KP (2006) The swarming body: Simulating the decentralized defenses of immunity. In: Bersini H, Carneiro J (eds) *Artificial Immune Systems, 5th International Conference, ICARIS 2006, Lecture Notes in Computer Science*, vol 4163, Springer, Berlin, Germany, pp 52–65
- Janeway C, Medzhitov R (2002) Innate immune recognition. *Ann Rev Immunol* 20:197–216
- Janeway CA, Travers P, Walport M, Shlomchik MJ (2001) *Immunobiology*, 5th edn. Garland Publishing, New York, NY
- Kennedy J, Eberhart R (2001) *Swarm Intelligence*. Morgan Kaufmann, San Francisco, CA
- Kennedy J, Mendes R (2002) Population structure and particle swarm performance. In: *IEEE Congress on Evolutionary Computation, 2002, CEC 2002*, IEEE Press, Piscataway, NJ, pp 1671–1676
- Kernbach S, Scholz O, Harada K, Popesku S, Leidke J, Raja H, Liu W, Caparrelli F, Jemai J, Havlik J, Meister E, Levi P (2010) Multi-robot organisms: state of the art. In: Koy K, Nagpal R, Shen W (eds) *IEEE International Conference on Robotics and Automation (Workshop on Modular Robotics)*, IEEE Press, Piscataway, NJ, pp 1–10
- Kleinstein P SH Seiden (2000) Simulating the immune system. *Computing in Science and Engineering* 2:67–77
- Kohler B, Puzone R, Seiden P, Celada F (2000) A systematic approach to vaccine complexity using an automaton model of the cellular and humoral immune system, I. Viral characteristics and polarized responses. *Vaccine* 19(7-8):862–76

-
- Langman RE, Cohn M (1986) The ‘complete’ idiootype network is an absurd immune system. *Immunology Today* 7(4):100–101
- Levi P, Kernback S (eds) (2010) *Symbiotic Multi-Robot Organisms: Reliability, Adaptability and Evolution*. Springer, Berlin, Germany
- Lumer E, Faieta B (1994) Diversity and adaptation in populations of clustering ants. In: *Proceedings of the Third International Conference on Simulation of Adaptive Behavior: From Animals to Animats 3*, MIT Press, pp 501–508
- McEwan C, Hart E (2009) Representation in the (artificial) immune system. *Journal of Mathematical Modelling and Algorithms* 8:125–149
- Millonas M (1994) Swarms, phase transitions, and collective intelligence. In: Langton CG (ed) *Artificial Life III*, Addison-Wesley, Redwood City, CA, pp 417–445
- Mokhtar M, Timmis J, Tyrrell A, Bi R (2009) A modified dendritic cell algorithm for on-line error detection in robotic system. In: *IEEE Congress on Evolutionary Computation, 2009, CEC 2009*, IEEE Press, Piscataway, NJ, pp 2055–2062
- Nanas N, Uren V, de Roeck A (2004) Nootropia: a user profiling model based on a self-organising nootropia: a user profiling model based on a self-organising term network. In: Nicosia G, Cutello V, Bentley PJ, Timmis J (eds) *Artificial Immune Systems, Third International Conference, ICARIS 2004, Lecture Notes in Computer Science*, vol 3239, Springer, Berlin, Germany, pp 146–160
- Newborough R, Stepney S (2005) A generic framework for population based algorithms. In: Bersini H, Carneiro J (eds) *Artificial Immune Systems, 5th International Conference, ICARIS 2006, Lecture Notes in Computer Science*, vol 4163, Springer, Berlin, Germany, pp 43–55
- Orosz M (2001) An introduction to immuno-ecology and immuno-informatics. In: Segal LA, Cohen IR (eds) *Design Principles from the Immune System*, Oxford University Press, New York, NY, pp 125–150
- Owens N, Timmis J, Tyrrell A, Greensted A (2008) Modelling the tunability of early T-cell activation events. In: Bentley PJ, Lee D, Jung S (eds) *Artificial Immune Systems, 7th International Conference, ICARIS 2008, Lecture Notes in Computer Science*, vol 5132, Springer, Berlin, Germany
- Owens N, Greensted A, Timmis J, Tyrrell A (2009) T cell receptor signalling inspired kernel density estimation and anomaly detection. In: Andrews PS, Timmis J, Owens N, Aickelin U, Hart E, Hone A, Tyrrell A (eds) *Artificial Immune Systems, 8th International Conference, ICARIS 2009, Lecture Notes in Computer Science*, vol 5666, Springer, Berlin, Germany, pp 122–155
- Perelson AS, Oster GF (1979) Theoretical studies of clonal selection: Minimal antibody repertoire size and reliability of self–non-self discrimination. *Journal of Theoretical Biology* 81(4):645–670
- Read M, Timmis J, Andrews PS, Kumar V (2009) A domain model of experimental autoimmune encephalomyelitis. In: *CoSMoS 2009, Proceedings of the 2009 International Workshop on Complex Systems Modelling and Simulation*, Luniver Press, Frome, UK, pp 3–39
- Salazar-Bañuelos A (2009) Non-deterministic explanation of immune responses: A computer model. In: Andrews PS, Timmis J, Owens N, Aickelin U, Hart E, Hone A, Tyrrell A (eds) *Artificial Immune Systems, 8th International Conference, ICARIS 2009, Lecture Notes in Computer Science*, vol 5666, Springer, Berlin, Germany, pp 7–10
- Segal L, Cohen I (eds) (2001) *Design Principles for the Immune System and Other Distributed Systems*. Oxford University Press, NY.

- Sempo G, Depickere S, Ame JM, Detrain C, Halloy J, Deneubourg J (2006) Integration of an autonomous artificial agent in an insect society: Experimental validation. In: From Animats to Animals 9, The Ninth International Conference on the Simulation of Adaptive Behavior, SAB 2006, Lecture Notes in Artificial Intelligence, vol 4095, Springer, Berlin, Germany, pp 703–712
- Stepney S, Smith R, Timmis J, Tyrrell A, Neal M, Hone A (2006) Conceptual frameworks for artificial immune systems. *International Journal of Unconventional Computing* 1(3):315–338
- Stibor T, Timmis J, Eckert C (2006) On the use of hyperspheres in artificial immune systems as antibody recognition regions. In: in Proceedings of 5th International Conference on Artificial Immune Systems (ICARIS), Springer, Lecture Notes in Computer Science, vol 4163, pp 215–228
- SwarmWiki (ongoing) Swarm wiki. <http://swarm.org>
- Timmis J (2007) Artificial immune systems: Today and tomorrow. *Natural Computing* 6(1):1–18
- Timmis J, Neal M (2001) A resource limited artificial immune system for data analysis. *Knowledge Based Systems* 14(3–4):121–130
- Timmis J, Andrews PS, Owens N, Clark E (2008a) An interdisciplinary perspective on artificial immune systems. *Evolutionary Intelligence* 1(1):5–26
- Timmis J, Hart E, Hone A, Neal M, Robins A, Stepney S, Tyrrell A (2008b) Immuno-Engineering. In: 2nd IFIP International Conference on Biologically Inspired Collaborative Computing, 20th IFIP World Computer Congress, IFIP International Federation for Information Processing, vol 268, Springer, Berlin, Germany, pp 3–17
- Timmis J, Hone A, Stibor T, Clark E (2008c) Theoretical advances in artificial immune systems. *Journal of Theoretical Computer Science* 403(1):11–32
- Timmis J, Tyrrell A, Mokhtar M, Ismail A, Owens N, Bi R (2010) An artificial immune system for robot organisms. In: Levi P, Kernbach S (eds) *Symbiotic Multi-Robot Organisms: Reliability, Adaptability and Evolution*, Springer, Berlin, Germany, pp 268–288
- Warrender C (2004) Modeling intercellular interactions in the peripheral immune system. PhD thesis, Computer Science Department, University of New Mexico, NM
- Warrender C, Forrest S, Segel L (2004) Homeostasis of peripheral immune effectors. *Bulletin of Mathematical Biology* 66:1493–1514
- Winfield AF, Harper CJ, Nembrini J (2006) Towards the application of swarm intelligence in safety critical systems. In: Proceeding of the 1st IET International Conference on System Safety, IEE/IET, Stevenage, Hertfordshire, pp 89–95
- Wolpert DH, Macready WG (1997) No free lunch theorems for optimization. *IEEE Transactions on Evolutionary Computing* 4:67–82