

Modelling Danger and Anergy in Artificial Immune Systems

Steve Cayzer
Hewlett-Packard Laboratories
Filton Road, Stoke Gifford
Bristol, UK
steve.cayzer@hp.com

Julie Sullivan
Department of Computer Science
University of Bristol
Bristol, UK
julie.sullivan@gmail.com

ABSTRACT

Artificial Immune Systems are engineering systems which have been inspired from the functioning of the biological immune system. We present an immune system model which incorporates two biologically motivated mechanisms to protect against autoimmune reactions, or false positives. The first, anergy, has been subject to the intense focus of immunologists as a possible key to autoimmune disease. The second is danger theory, which has attracted much interest as a possible alternative to traditional self-nonsel self selection models.

We adopt a published immunological model, validate and extend it. Using the same calculations and assumptions as the original model, we integrate danger theory into the software.

Without anergy, both models – the original and the danger model – produce similar results. When anergy is added, both models' performance improves. However, there seems to be some synergy between the mechanisms; anergy has a greater effect on the danger model than the original model.

These findings should be of interest both to AIS practitioners and to the immunological community.

Categories and Subject Descriptors

J.3 [COMPUTER APPLICATIONS] Life and Medical Sciences
I.6.4 [SIMULATION AND MODELING] Model Validation and Analysis

General Terms

Algorithms, Design, Experimentation

Keywords

Immune systems, artificial immune systems, ais, danger theory, anergy

1. INTRODUCTION

Artificial Immune Systems (AIS) are engineering systems

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

GECCO'07, July 7–11, 2007, London, England, United Kingdom.
Copyright 2007 ACM 978-1-59593-697-4/07/0007...\$5.00.

inspired from the functioning of the biological immune system [3]. AIS (there are many varieties) comprise a set of lymphocytes (white blood cells) which recognize damaging pathogen (bacteria, viruses and the like). There are several features that have a distinctive evolutionary flavour. Clonal selection uses recognition ability to choose the lymphocytes that will reproduce. Affinity maturation is where successful lymphocytes reproduce and mutate. Other features are more distinctive. Negative selection is the initial screening of potentially self reactive lymphocytes. Network algorithms take into account the possibility of lymphocytes recognizing each other. Gene libraries are used to provide building blocks for lymphocyte creation, and this provides a sort of species level learning.

A key function of the biological immune system is protection from infection. The recognition of infectious agents is through characteristic protein fragments called *antigens*. However, it is also possible for the immune system to recognize the antigens in the host's own body and thus mount a self destructive attack [12]. This is called *autoimmunity*, and more than eighty illnesses have been categorised as such. Transplant patients are also at risk due to unfamiliar antigens from a transplanted organ.

How can the immune system protect itself against autoimmune attacks? One mechanism, *anergy* [8,10] is invoked when a potentially self destructive, but weakly activated, response leads to a suppression of that response in the future. Danger Theory [9] on the other hand relies on signals from damaged tissue to differentiate between malignant and benign events, and thus to orchestrate the immune response. In this paper, we implement and extend an existing immunological model as a way of investigating the interplay of anergy and danger mechanisms in the immune response.

As a result multiple implications can be drawn for AIS, both from the view of the immunologist and the computer scientist. First, this investigation strengthens the synergy between the two fields, allowing those whose expertise may lie in one realm to communicate with the other in a way unheard of in the pre-in silico world. Secondly, the logical, mathematical relationship between anergy and T-Cell behaviour is proved. Furthermore, the flexibility of this AIS system was successfully tested by the introduction of Danger Theory. In fact anergy was shown to improve the performance of the danger model. Lastly, and on a broader scale, this investigation is a progression of the continuing efforts to apply computer science to biology; a field whose complexity resists quantification.

The paper is structured as follows. We introduce the immunology terminology necessary to understanding the model used in this paper. The original and danger-enhanced models are both described before we show how danger and anergy interact.

2. IMMUNOLOGY

The following is a brief overview of the immune system. It is not intended to be comprehensive, instead it is meant to introduce the concepts relevant to our model. For more detail see standard texts such as [5] and [7].

2.1 Innate and Adaptive Immunity

The first line of a body's defence is the *innate immune system* which individuals have from birth. This system is non-specific; it does not target any particular invader. Instead it is a broad layer of protection which includes skin, airway cilia, tears, saliva and cells such as macrophages which migrate to an infection area and neutralise threats.

Some foreign invaders are able to circumvent these basic defences. In this case, they encounter the body's second line of defence, the *adaptive immune system*. Adaptive immunity, sometimes called acquired immunity, is more specialised than innate immunity. The adaptive immune response is triggered by identification of the foreign invader which leads to the activation of certain cells that engulf, kill or remove the foreign agent from the body.

2.2 Cells of the Adaptive System

Our model focuses on T cells, a type of lymphocyte (white blood cell) involved in the adaptive response. T cells have a multistaged lifecycle, including an early maturation phase in the thymus. They are responsible for the 'cell-mediated response'. Naïve T cells theoretically can last for a lifetime (in our model also). A recent study measured the half-life at about two years [11]. They are 'presented' with a foreign agent for which they have specialised receptors.

T cells are created in the bone marrow and migrate to the thymus for their 'education'. Only 5 to 10% of immature T cells, lymphocytes, survive this process and leave the thymus. Thymic selection can be categorised into two phases: positive and negative selection. The first process, positive selection, deletes cells that fail to recognise foreign agents. This increases the sensitivity of the T cells – positive selection ensures T cells will correctly identify foreign antigen. Next a negative selection process deletes cells that have the potential to incorrectly activate when encountering 'self' molecules, molecules that are a normal part of a healthy body. This in turn ensures specificity, thus assuring that T cells will not signal an attack when encountering self antigen. Cells surviving selection are exported to the circulatory system.

Antigen-presenting cells (APCs) are cells such as dendritic cells, typically assigned to the innate system but capable of directing an adaptive immune response. These cells consume the antigen, display the molecules on their surface, and present them to the receptors on the T cell. APCs also send a costimulatory signal required for T cell activation. The label APC usually refers to dendritic cells, macrophages, or B cells.

2.3 Immune Response

The antigen T cell receptors (TCRs) are attached to the T cell membranes. A TCR is a protein that binds to another protein it recognises. T cells can only recognise proteins that are presented to it via APCs. If a TCR binds a protein and receives a costimulatory signal, signal 2, from the APC, the T cell activates. A TCR only recognises antigen that fit into its receptor

The immune response begins with an APC presenting a T cell with an antigen. Activation of the T cell requires more than just presentation, however. Activation requires two signals, one signal from the APC to the T cell and one signal from the T cell. If both signals are strong, the T cell activates. If signal 2 is absent or weak, the T cell is likely to become *anergic* (see next section). Activated T cells enlarge and proliferate, dividing into effector and memory cells. Effector cells seek to destroy the antigen and after the threat is neutralised will die away. Memory cells will remain in the system ready to meet another potential attack by the same antigen.

2.4 Anergy

The adaptive immune system is required to respond to any foreign invader. To achieve this, the body creates millions of T cells. Among these vast numbers of T cells are some that are potentially self-reactive, that is they might activate upon encountering self antigen. Although negative selection will protect the body from many such cells, there is a risk that self reactive cells will escape to the periphery. If such a cell were to activate, it would cause destruction of healthy tissue. Anergy is a regulatory mechanism to prevent this from happening. When a T cell recognises its cognate antigen, it sends a strong signal 1. In turn, the APC sends a costimulatory signal to the T cell. If that signal is absent or weak, the T cell is put in a quiescent state during which it ignores all signals. The T cell is 'turned off' and thus an inappropriate immune system response is prevented. This is anergy.

Anergy is an important immunological phenomenon. Further study of anergy could lead to therapies for autoimmune disease because a flaw in this system may be responsible for inappropriate immune system responses. Anergy is potentially important for artificial immune systems, where autoimmune responses correlate to false positives.

The question remains though: How important is anergy in preventing autoimmune responses? Are other mechanisms at play? How do they interact? The Chan-Stark-George model [4] explores anergy by manipulating certain variables: the threshold for activation, the number of T cells, the number and type of T cell encounters, and whether anergy is absent or present.

We extend this model by testing parameter changes and by adding Danger Theory. While any outcomes of Chan-Stark-George would have in themselves been of great value, it was felt that testing the flexibility of the program would add yet another layer of interest. Too, quite often the final value measurement of a system is in its ability to work in tandem with other systems. From an AIS standpoint, the model offers insights about the use of anergy, Danger Theory and the combination of the two in practical systems.

2.5 Danger Theory

A central tenet of classical immunology is the distinction between self and non-self. Danger Theory [9] takes issue with this assumption, and instead asserts that the immune system is concerned with what is dangerous and not dangerous.

For instance, what about food? Food is not self, but neither is it inherently dangerous. What about foetuses, or foreign bacteria so useful in the digestive system? These are all foreign cells in the body that are not attacked by the body. Conversely, some harmful tumours are indeed correctly attacked even though they are clearly 'self'.

Danger theory posits that the coordinator of the immune system response is not the cells patrolling the body, instead the tissues themselves (ie APCs) are the entities responsible. The immune system response starts when a tissue is damaged and proceeds to send out a signal, the 'danger' signal. This sets up a perimeter around the tissue within which all APCs capture antigen and present them to the T cells.

As Figure 1 shows, if the T cell-APC encounter happens outside of the 'danger zone', the T cell is too far away and activation is unlikely. For encounters within the danger zone, activation is more likely. However, APCs are busying themselves with capturing all antigen in the area. This all-encompassing dragnet will inevitably ensnare those antigen that are present throughout the body, but who are not responsible for the damage to the tissue that originally sent out the danger signal. However, since these antigen are so ubiquitous, T cells will likely have already encountered these 'safe' antigen and tolerance will have been achieved. These mechanisms have obvious application to AIS [1], in particular in areas like computer security where the problem of false positives is rife [25].

From the point of view of our description above, danger mechanisms can be viewed as alternative to, or complementary with, anergy. Danger theory will be integrated into the Chan-Stark-George model, and simulations will be run to compare its performance to the original model.

3. Model Details

There are two major theoretical techniques for simulating the immune system: differential equations and computer models. Differential equations are adept at handling large populations, but are less adept at handling the special cases that are so common in the immune system. Conversely, agent based models excel at incorporating individual behaviour, non-linearities are handled with ease, and irregularities are easy to spot. In general, differential equations are best at answering 'what and when', whereas computer models answer 'why and how'.

We adopt an agent based approach where T cells are modeled as a population of entities. Each T cell has a series of rules governing its behaviour in response to the environment. These rules are described in the next section, and incorporate random events (ie they are stochastic).

The Chan-Stark-George model is an agent-based model of the immune system focusing on anergy and repeated T cell-APC encounters. This model explores the roles and interaction of different thresholds, costimulation and anergy.

We implement the model, investigate its sensitivity to parameters and use it as a way of investigating the interplay of anergy and danger mechanisms in the immune response.

3.1 The Chan-Stark-George model

In this section we summarise the main features of the Chan-Stark-George model. The key features of this model are: negative selection (thymic education); anergy (peripheral effect on T cells that survive negative selection) and costimulation (signal 2).

s	Signal 1
m	Slope of activation function
θ_1	Activation threshold
θ_2	Energy threshold (lower)
θ_3	Energy threshold (upper)
n	Number of T cell-APC encounters
p	Probability of activation per encounter
q	Probability of anergy per encounter
rf	Recognition fraction
cf	Co-stimulation fraction
if	Infection fraction

Table 1: Variables used in mathematical model

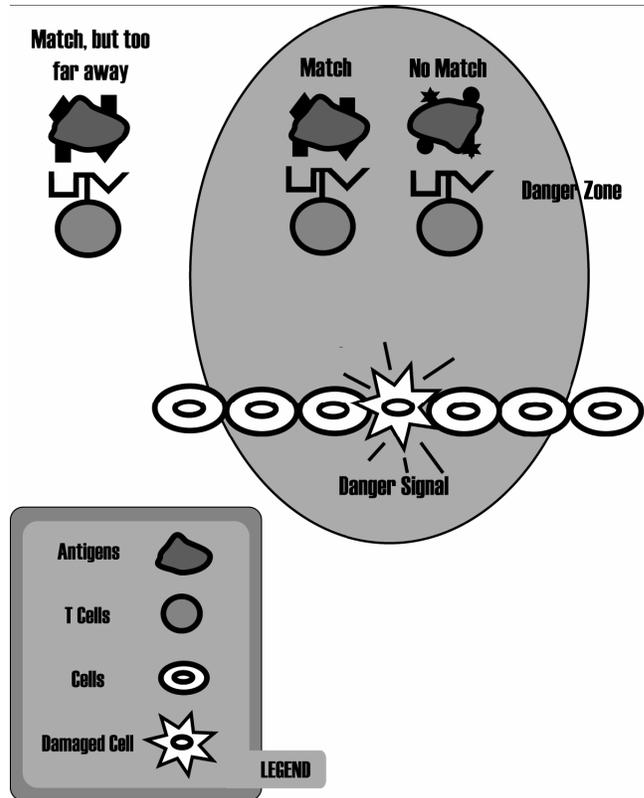


Figure 1: T Cell interactions in and around the 'danger zone'. Inside the danger zone, only T cells that match an antigen are activated. There may be many such antigens. Matches outside the danger zone induce tolerance. Adapted from [1]

3.1.1 Negative selection in the thymus

Negative selection occurs in the thymus, where immature T cells (thymocytes) are assumed to encounter only self antigen. It is modelled by generating a signal and based on that signal, probabilistically the T cell either activates or does nothing. If the T cell is activated, it is deleted. This is repeated until the T cell is deleted or survives a set number of APC encounters. The surviving cells then move out of the thymus to the periphery where they encounter APCs presenting foreign and self antigen

The activation of T cells is stochastic, so the probability of being activated per encounter (p) is calculated using a biologically motivated [14] sigmoid-shaped Hill function:

$$p = s^m / (s^m + \theta_1^m) \quad (1)$$

Other experimental findings [4] are used to provide values for the slope of the activation curve ($m=2$) and its threshold ($\theta_1 = 10e^{-1}$). This only leaves the value of s , signal 1, which naturally will vary from thymocyte to thymocyte due to binding affinity.

Keeping faithful to the Chan-Stark-George model [4], s is chosen from an exponential distribution, the mean of which is chosen so that about 50% of thymocytes survive negative selection. It can be noted in passing that this is far above the survival ration for most AIS implementations of negative selection, which implies that prior mechanisms (gene libraries, positive selection) are effective in producing an initial population of candidate thymocytes.

Equation 1 is then used to drive the negative selection process:

```
// for each T cell
While i < MAX_ENCOUNTERS
if RAND < probability of activation
    cell deleted
else
    cell survived
.
```

In the human immune system, negative selection takes about 2 weeks during which each T cell has about 1 APC encounter per minute. So we set MAX_ENCOUNTERS=20000 (60*24*14), after which the surviving cells are moved to the periphery.

It should be noted that the Chan-Stark-George model does not allow for interactions between the thymic and peripheral compartments. In other words, each individual T cell is not tracked through thymic maturation and moved to the periphery. Instead the compartments are studied separately.

3.1.2 Peripheral activation

In the periphery, the T cells encounter multiple APCs with the same activation function as in the thymus (Equation 1).

Signal 1 (s) is calculated the same as in the thymus (ie sampled from an exponential distribution), except for a certain number of T cells, set by the user (α : typically 1%), which recognise their cognate antigen. This is simulated by setting s to $2\theta_1$, which gives an 80% chance of activation.

3.1.3 Anergy

Anergy is a weak signal 1; the probability of which is the governed by two thresholds θ_2 and θ_3 (both less than θ_1):

$$q = s^m / (s^m + \theta_2^m) - s^m / (s^m + \theta_3^m) \quad (2)$$

$$\theta_2 = 1/20 \text{ of } \theta_1 ; \theta_3 = 1/10 \text{ of } \theta_1$$

While in the anergic state, a T cell simply ignores all signals.

The interplay between activation and anergy can best be described by figure 2:

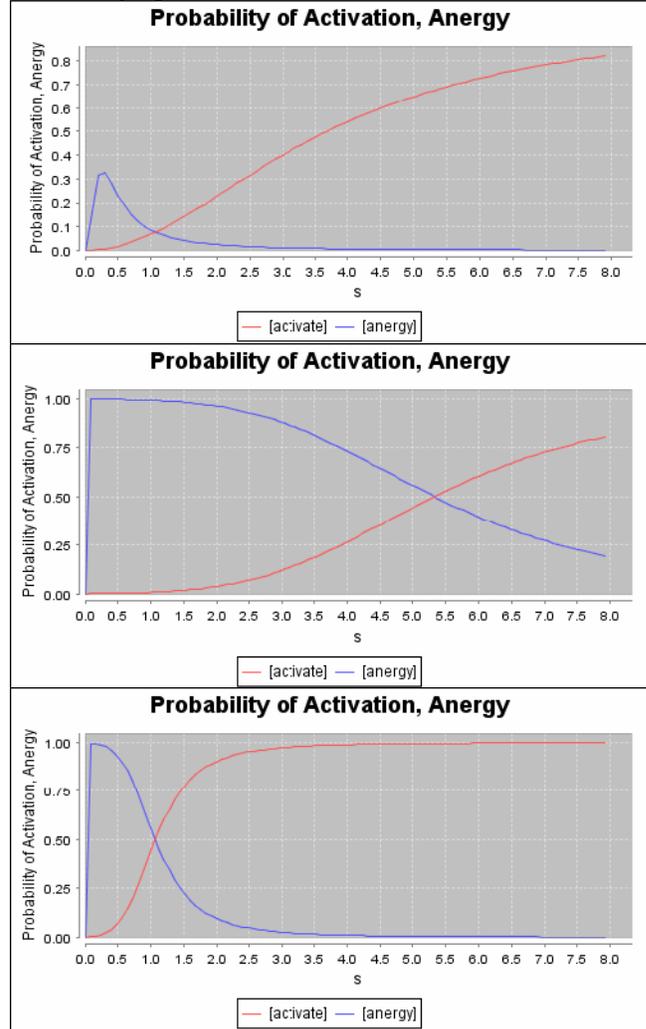


Figure 2: The effect of repeated encounters and costimulation on anergy. In the top graph, a single peripheral encounter will tend to lead to anergy at very low levels of signal one, and stimulation otherwise. In the middle graph we can see that the probability of anergy has greatly increased given multiple (here 100,000) encounters. In the bottom graph, the presence of costimulation shifts the balance away from anergy.

Putting this all together, we can see that the probability of activation after n encounters depends on the probability of activation p and the probability of anergy q at each encounter:

$$\text{Probability of activation} = p + p(1 - p - q) + p(1 - p - q)^2 + \dots + p(1 - p - q)^{n-1} \quad (3)$$

An activated cell never deactivates, but an anergised cell may revert back to its original state after a fixed number of encounters (*reversible anergy*).

The activation/anergy process can be programmed as follows:

```
r <- random()
While i < MAX_ENCOUNTERS
  // for each T cell
  if (r < p
    cell becomes activated
  elseif (p < r < (p + q) )
    cell becomes anergic
  else
    no change
  endif
endif
next i
```

3.1.4 Costimulation

In the original model [4], infection is simulated by creating two sets of APCs: those that present foreign antigen and those that present self antigen. Thus, the T cell can have three types of encounters:

1. Infected APC sends full costimulatory signal
2. Uninfected APC sends costimulatory signal
3. Uninfected APC sends no costimulatory signal

The type of encounter the T cell has is determined by the infection fraction (*if*) and the activation threshold is also determined by the presence/absence of costimulation (*cf*). *if* and *cf* are both set by the user.

```
r1 <- random()
r2 <- random()

if (r1 < if) //Infected APC
  // 'right' costimulatory signal
  θ1 = 10e-1

else if (r2 < cf) //Uninfected APC
  // 'wrong' costimulatory signal
  θ1 = 10e-1

else // Uninfected APC
  //no costimulatory signal
  θ1 = 5 * 10e-1

endif
```

3.1.5 Modelling Danger

In contrast to the previous section, in our model we used danger to control the level of costimulation.

We use infection fraction (*if*) as a proxy for the danger signal. If the random number is less than the infection fraction, the T cell is in the danger zone. Otherwise it's considered too far away. In the danger zone, all encounters are costimulatory. Therefore the thresholds are calculated as in the original model for a

costimulatory APC encounter, and the signal is calculated dependent on whether the T cell recognises the antigen (*rf*: typically set at 1%, see section 3.1.2). If the T cell-APC encounter happens outside the danger zone, then the threshold is calculated as it is in the original model for a non-costimulatory APC encounter:

```
r1 <- random()
r2 <- random()

//signal 1 by default is sampled from an
exponential distribution
s <- exponential()

if (r1 < if) // in danger zone

  if (r2 < rf) // match
    θ1 <- 10e-1
    s <- 2 * θ1 // 80% chance activation

  else // no match: ignore
    θ1 <- 10e-1

  endif

else // tolerance
  θ1 <- 5*10e-1

endif
```

Thus, danger, by controlling *s* and θ_1 , shifts the balance between activation and anergy in a way that makes the mechanisms synergistic.

4. Results

Experiments were run using the application to explore the model outlined in [4]. Further experiments were run comparing the new danger model with the original model.

First, the results in the original article [4] were replicated. Then various starting parameters were changed, experiments were run, and the results compared to the original results. Finally, 60 simulations were run: 30 simulations of the original model and 30 simulations of the new danger theory model. The findings of these experiments are detailed below.

4.1 Negative selection

Negative selection is effective at deleting those thymocytes that have a high probability of being activated by self antigen. Each cell has a certain probability of being activated at each encounter. As the cell encounters more and more antigens, the probability that a T cell that shows affinity for self antigen will activate and thus be deleted increases. Our model confirmed that as the number of encounters increases, the probability of export decreases even for thymocytes with a low activation threshold. Restated, multiple T cell-APC encounters in the thymus lowers the threshold for activation, so thymocytes with a high affinity for self antigen are likely to be deleted. Given 2000 encounters, if a cell is activated once out of 1000 encounters, $p=0.001$, it has a 14% chance of being exported. A more potentially self-reacting thymocytes with a probability of activation of 0.01 has only a 0.0000002% chance of surviving negative selection. Thus, given many T cell-APC encounters, negative selection is effective at deleting thymocytes that are potentially self-reactive. [4]

4.2 Peripheral energy

As the previous section demonstrated, multiple encounters lower the activation thresholds. Thus even a T cell with a low affinity for self antigen would likely activate in the absence of energy. However, in the presence of energy, energy is more likely than activation for cells with low avidity. Therefore an autoimmune response is only likely for T cells with high avidity. Since negative selection deleted cells with a high avidity, it's clear that negative selection and energy are complementary mechanisms which together lower the probability of an autoimmune response.

4.3 Sensitivity, specificity and reversible energy

Sensitivity is the ability of an immune system to correctly identify foreign antigen (true positives). Specificity is its ability to avoid autoimmune reactions (false positives). Both can be calculated by measuring true and false positives (TP, FP) and negatives (TN, FN):

$$\text{Sensitivity} = TP / (TP + FN) \quad (4)$$

$$\text{Specificity} = TN / (TN + FP) \quad (5)$$

In [4] it was reported that irreversible energy improves specificity while degrading sensitivity (this would make sense as it depletes the stock of T cells available to fight infection). Reversible energy restores sensitivity while maintaining the improved specificity levels.

We tested the model a number of ways – see [13] for full details. A higher infection rate, for example, gave the T cells more chances to identify foreign antigen, resulting in higher sensitivity, particularly for the irreversible energy scenario. We also tested doubling the number of thymic (to 20,000) and peripheral (to 200,000) encounters. Without energy, the increased number of encounters reduced sensitivity; ie increased the risk of autoimmune reactions (false positives). With energy, this risk abated and sensitivity (true positives) improved slightly. These changes were relatively small and the main thrust of [4] – that reversible energy improves specificity without sacrificing sensitivity – continued to hold.

4.4 Danger

Danger theory was added to the original model, then experiments were run on both models using the same parameters with one exception – infection rate. The tenet of danger theory is that costimulation originates primarily from the tissue sending out the danger signal. We use the infection fraction as a proxy for the danger signal. In these experiments, we compare the original model with a costimulation fraction of 1.0, 0.0, and 0.50 against the danger model with an infection fraction of 1.0, 0.0, and 0.50, respectively.

The graphs below show the results of 60 simulations in the absence and presence of energy, with energy being irreversible. Reversible energy is not shown since as in the original model, we obtained near perfect results for both sensitivity and specificity, and adding danger to the model did not change this. We also do not show results for infection fraction of 0.0 as with no infection there is no danger and hence no true positives. This means it's

impossible to measure sensitivity (in Equation 4) while specificity is 100%.

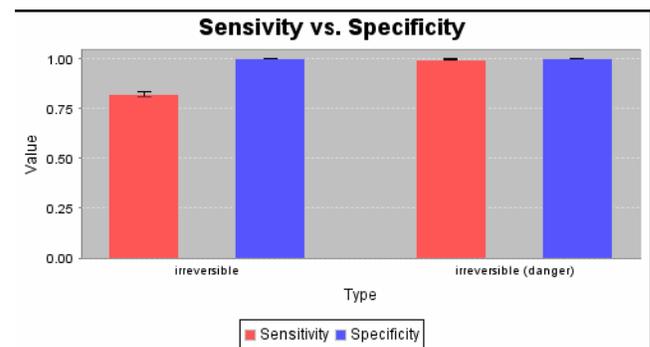
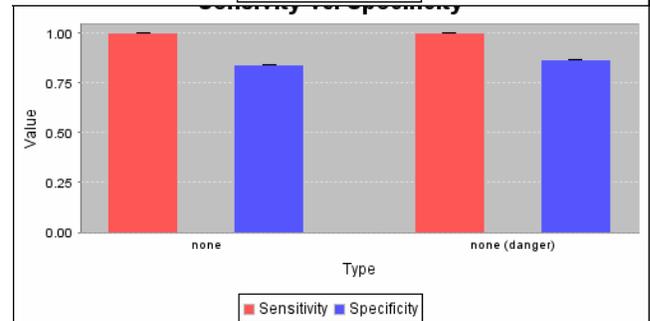
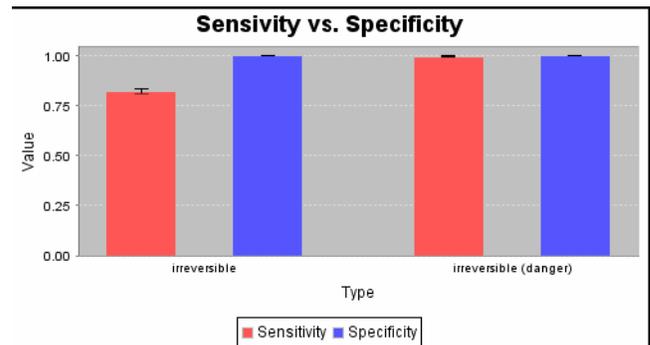
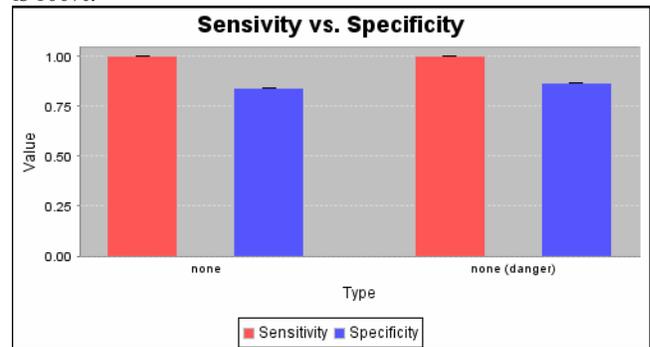


Figure 3 Effect of danger in different environments. Top graph shows response without (left hand side) and with (right hand side) danger in an environment with no energy but 100% costimulation (left) or 100% infection rate (our proxy for danger: right). 2nd graph shows same environment but irreversible energy. 3rd graph is like the 1st but with 50% costimulation (or infection rate). 4th graph is like the 3rd with irreversible energy.

Without anergy, both models produce identical results (1st and 3rd graph in figure 3). As explained above, when anergy is introduced, specificity is improved, but sensitivity is sacrificed in the original model. However, when danger is added to the model, sensitivity is restored (2nd and 4th graphs). The results are statistically significant at the 1% level (Wilcoxon). Specificity is not significantly affected.

5. Discussion and Future Work

Anergy is an interesting biological phenomenon not often incorporated into AIS. Our work confirms that anergy has some interesting properties which would be worth exploring.

Moreover, we show that danger theory provides a mechanism which plays well with anergy by restoring sensitivity without compromising specificity. Danger theory can be, and has been incorporated into AIS [2] and so our findings should be of interest to the AIS community. However, since we build on an experimentally motivated (though simple) model, our results should also be of interest to immunologists.

In this paper we use the infection fraction as proxy for danger. In one sense, this is unsatisfying because one would like to measure the actual danger/damage caused. On the other hand, knowing the infection fraction does not tell you the cause of the infection. Nevertheless, we would like to take into account the idea of a spatial ‘danger zone’. The concept of space is key in the immune system. The immune response is by definition limited to those immune system cells that are in the physical neighbourhood of the infected site, and it would be helpful to model this explicitly.

We would like to follow a T cell throughout its development – in other words, to link the thymic and peripheral compartments of this model. Finally, APCs were not explicitly modelled in this simulation for reasons of performance. It would be interesting to explore models of APC development, such as mature and semi mature dendritic cells [6].

6. References

- [1] U Aickelin and S Cayzer. The Danger Theory and Its Application to Artificial Immune Systems. International Conference on Artificial Immune Systems (2002).
- [2] Aickelin, U., Bentley, P., Cayzer, S., Kim, J. And Mcleod, J. 2003. Danger Theory: The Link between AIS and IDS? In Proceedings of the Second International Conference on Artificial Immune Systems (ICARIS 2003), Sept 1-3, Edinburgh, Scotland
- [3] Castro, Leandro Nunes de, Timmis J. Artificial Immune Systems: A New Computational Intelligence Approach. Springer-Verlag. 2002: 380pp
- [4] C Chan, J Stark, and A George. The Impact of Multiple T cell—APC Encounters and the Role of Anergy. Journal of Computational and Applied Mathematics. 184 (2005) 101-120.
- [5] R Goldsby T Kindt, B Osborne, and J Kuby. Immunology, 5th edition. 2003
- [6] Greensmith J. Introducing the Dendritic Cell Algorithm International Conference on Artificial Immune Systems (ICARIS 2006), Lisbon Portugal September 2006
- [7] Janeway, Travers, Walport, and Shlomchik. Immunobiology. 6th Edition. 2005
- [8] LC Korb, S Mirshahidi, K Ramyar, AA Sidighi Ahka, and S Sadegh-Nasseri. Induction of T cell anergy by low numbers of agonist ligands. Journal of Immunology. 162 (1999) 6401-6409.
- [9] P Matzinger. The Danger Model: A Renewed Sense of Self. Science. 296 (12 April 2002) 301- 305.
- [10] S Mirshahidi, CT Huang, and S Sadegh-Nasseri. Anergy in peripheral memory CD4(+) T cells induced by low avidity engagement of a T cell receptor. Journal of Experimental Medicine. 194 (2001) 719-731.
- [11] R Neese, L Misell, S Turner, A Chu, J Kim, D Cesar, R Hoh, F Antelo, A Strawford, JM McCune, M Christiansen, and MK Hellerstein. Measurement in vivo of proliferation rates of slow turnover cells by 2H2O labeling of the deoxyribose moiety of DNA. Proc. Natl. Acad. Sci. USA 99 (2002) 15345–15350.
- [12] O Parnes. Autoimmune Disease. The Lancet. (4 February 02006-10 February 2006) 389.
- [13] Sullivan, J. The war within: A model of the immune system MSc dissertation, University of Bristol, UK. September 2006.
- [14] A Viola and A Lanzavecchia. T cell activation determined by T cell receptor number and tunable thresholds. Science 273 (1996) 104–106.