

Chapter 6. Dynamic Clonal Selection Algorithm

6.1 Introduction

The results described in chapter five show that static clonal selection can be used to build a misuse detector in an efficient way. The static clonal selection algorithm learned self and non-self antigen patterns hidden within a given static antigen data set. This work was performed by repeatedly experiencing an identical and static antigen data set. Although this type of static clonal selection algorithm is useful for building misuse detectors from previously collected data, a real environment (which a network-based IDS needs to monitor) produces new network traffic continuously in real-time. Thus antigens faced by the AIS will be different every day. More importantly, normal behaviours of network traffic on one day, which are considered as self antigens, can be different from normal behaviours of network traffic on another day. Therefore, the AIS needs to be extended, firstly to learn normal behaviours by undergoing only a small subset of self antigens at one time. Secondly its detectors should be replaced whenever previously observed normal behaviours no longer represent current normal behaviours.

This chapter describes a dynamic clonal selection algorithm that has the above two properties [Kim and Bentley, 2002a]. This algorithm is clearly distinguished from the static clonal selection algorithm studied in chapter five. The static clonal selection algorithm abstracts a niching mechanism of an antibody cloning process and was applied to static antigen data. In contrast, the dynamic clonal selection algorithm introduced in this chapter follows several procedures of the human immune system to provide the newly required properties of a dynamic clonal selection algorithm. This chapter starts by presenting those procedures of human immune systems that are directly used in the design of the dynamic clonal selection algorithm and previous work embedding these procedures in the AIS (Section 6.2). A detailed description of the dynamic clonal selection algorithm then follows (Section 6.3). This description introduces three important new parameters that control the degree of two properties required: tolerisation period, activation threshold and life span. The dynamic clonal selection algorithm is then investigated with respect to how these three parameters affect non-self and self antigen detection rates. In order for obtained non-self and self antigen detection rates to demonstrate the competence of the dynamic clonal selection algorithm in terms of the desired two properties, the algorithm is tested on two artificially created scenarios that reflect two common real circumstances of IDS's.

6.2 Related Work

6.2.1 Dynamic Non-Self Antigen Detection by Human Immune systems

Central Tolerisation

In previous chapters it has already been discussed how immune cells gain self tolerance via a negative selection process. Following this idea, a negative selection operator was adopted to reduce false positive errors in the static clonal selection algorithm. The negative selection operator of the static clonal selection algorithm conforms to the negative selection procedure which is used to generate helper T-cells in the thymus. Most epitopes of self cells are exposed to helper T-cells in the thymus and thus helper T-cells can gain relatively satisfactory self-tolerance. This is called Central Tolerisation (CT) [Sompayrac, 1999].

Distributed Tolerisation

Another type of immune cells, B-cells are also released from the bone marrow to the body for antigen detection after passing a negative selection process in the bone marrow. Both B-cells and helper T-cells continuously circulate around the body in the blood and encounter antigens for activation and evolution [Sompayrac, 1999]. The antibodies of B-cells, which recognise harmful antigens by binding to them, are activated directly or indirectly. When B-cell antibody receptors bind to antigen epitopes with strong affinity above a threshold, they are directly activated. On the other hand, B-cell antibody receptors can bind to antigen epitopes with weak affinity below a threshold. In this case, B-cells need the help of T-cells and Major-Histocompatibility Complex (MHC) molecules to be activated. MHC molecules have two important functions that help B-cell activation. Firstly, they bind to the fragments of antigens specially hidden inside B-cells (not visible on the cell surface), and secondly they transport these fragments to the B-cell surface. When B-cell antibody receptors bind to antigen epitopes with weak affinity, MHC molecules try to find some hidden antigen inside B-cells. When MHC molecules find them, they transport them on the surface of B-cells. The receptors of T-cells are genetically structured to recognise the MHC molecule on the B-cell surface. Thus, T-cells can bind to MHC molecules on B-cell surfaces. When the T-cell binds to MHC molecule with strong affinity, it sends a chemical signal to the B-cell which allows it to activate, grow and differentiate.

This raises the question of what makes the T-cells determine B-cell activation. One major difference between B-cells and T-cells is that only B-cells perform somatic hypermutation, which is a very high rate of mutation, to increase their diversity when they are matured in the germinal centres of lymph nodes. Hence, B-cells have more new and varied receptors than T-

cells. In addition, the thymus is centrally located while there are hundreds of distributed germinal centres of lymph nodes. Due to the central location of the thymus, helper T-cells can gain more reliable self-tolerance than B-cells. Therefore, the final decision of B-cell activation with weak affinity is made by the helper T-cells. The self-tolerance of B-cells gained by helper T-cell binding is called Distributed Tolerisation (DT) in contrast to CT [Sompayrac, 1999; Tizard, 1995].

Costimulation

However, there are still some self cells that have never been presented to T-cells in the thymus. Although T-cells have a much larger degree of self tolerance than B-cells, they do not show perfect self tolerance. Hence, T-cells and MHC molecule binding also needs an extra confirmation signal for activation. This signal is sent by innate immune cells when they sense actual tissue damage. This process is called costimulation [Sompayrac, 1999; Tizard, 1995]. Thus, if T-cells mistakenly bind self antigen fragments in MHC molecules and they are in an area having damaged tissues, they can trigger B-cell activation. However, if their MHC molecule binding affinities are not strong enough for activation before they move to different parts of the body that do not have damaged tissue, these T-cells do not allow B-cell activation.

Affinity Maturation

As described above, B-cells do not activate unless their affinities reach a threshold or receive a confirmation signal from helper T-cells. Affinity maturation is a process that ensures B-cells activate only if their affinities reach a threshold, and otherwise die off [Tizard, 1995; Life, 1993]. Similarly, T-cells also wait for activation signals from the innate immune system only when their binding MHC molecules on B-cell surfaces satisfy a threshold. Thus, affinity maturation is another scheme to supply a degree of self-tolerance by allowing T-cells and B-cells activation only when they match sufficient antigen epitopes.

Life Span

Since human body self cell births and deaths constantly occur, T-cells and B-cells also need to refresh their obtained self-tolerance. Mature T-cells and B-cells that monitor non-self antigens at the germinal centres of the lymph nodes have limited life spans as all other human body cells do [Sompayrac, 1999]. For B-cells, if they cannot activate within their life spans, they die off. Similarly, if T-cells do not contribute B-cell activation within their life spans because their binding MHC molecules on B-cell surfaces cannot receive signals from the innate immune system, they are initially unable to respond an antigen and eventually die [Sompayrac, 1999]. Thus, based on the existing self and non-self antigens, only useful B-cells and T-cells survive.

Memory Detectors

The life of B-cells and their mutated clones, the plasma cells, is relatively short. However, some of the B-cell clones survive as memory cells [Sompayrac, 1999; Tizard, 1995]. They have a special gene called *bcl-2* which is absent in short-lived B-cells and plasma cells. The *Bcl-2* gene enables the memory cells to survive for a longer time, such as several years. Some of the memory cells are exposed to antigens and differentiated into plasma cells without undergoing somatic mutation. Other memory cells undergo somatic mutation to be differentiated into plasma cells. Particularly plasma cells generated without somatic mutation allow the secondary response of immune systems. It generally takes some time for new pathogens to be detected by B-cells. This is called the primary response. Compared to the primary response, secondary responses by memory cells are very fast and efficient. Moreover, memory cells provide an associate memory property [Sompayrac, 1999; Paul, 1993]. New antigens having a structure that is not the same but similar to the structure of previously detected antigens can be detected by memory cells. This is because the binding of antibody and antigen is approximate. For example, when a body is infected by cowpox, the immune system takes time to detect and eliminate it. However, if somebody is infected by smallpox after being infected by cowpox, he/she is rapidly cured by the secondary response of immune system. This is because cowpox and smallpox are similar enough to induce a secondary response by memory cells.

Summary

In summary, as other human body cells are constantly born and die, new human immune cells are also endlessly born and die. The continuous birth and death of immune cells are vital to assure satisfactory self-tolerance and non-self antigen detection. Several mechanisms are used to achieve these properties. In particular, central tolerisation, distributed tolerisation, costimulation and affinity maturation contribute to ensure an adequate extent of self-tolerance. Memory cells provide quicker and more efficient non-self antigen detection. The limited life spans of immune cells improve both self-tolerance and non-self antigen detection by maintaining immune cells that reflect currently existing antigens.

6.2.2 Dynamic Anomaly Detection by AIS

Hofmeyr [1999] extended a negative selection algorithm in order to allow it to adapt to a continuously changing environment. In a real network environment, the distribution of self antigens, which is normal network traffic behaviour, often changes thus making it difficult for the AIS to observe a complete self and non-self antigen set at one time. This requires constant updating of detectors that can reflect current self and non-self antigen distribution. The extended AIS [Hofmeyr, 1999] generates detectors using negative selection with a currently obtained self

antigen set. In contrast to other AIS's producing detectors by monitoring a static antigen set, the extended AIS creates new detectors every day after the system experiences new network traffic which has not been presented before. The detectors that have been generated already are continuously compared to new antigens and their matching results of new antigens determine whether they have to be replaced by new detectors or not. More precisely, the replacement of detectors is achieved by introducing new parameters to determine the usefulness of detectors in order to detect current non-self antigens. These new parameters are defined mainly by mimicking human immune mechanisms introduced in the previous section that can contribute to maintain self-tolerance and non-self antigen detection in a dynamic environment. The main feature of these mechanisms is control of the life cycle of a detector according to its antigen matching results. Figure 6.1 shows the life cycle of a detector generated in Hofmeyr's AIS, which adopts life spans, activation threshold, memory detectors, central tolerisation and costimulation of the human immune system¹.

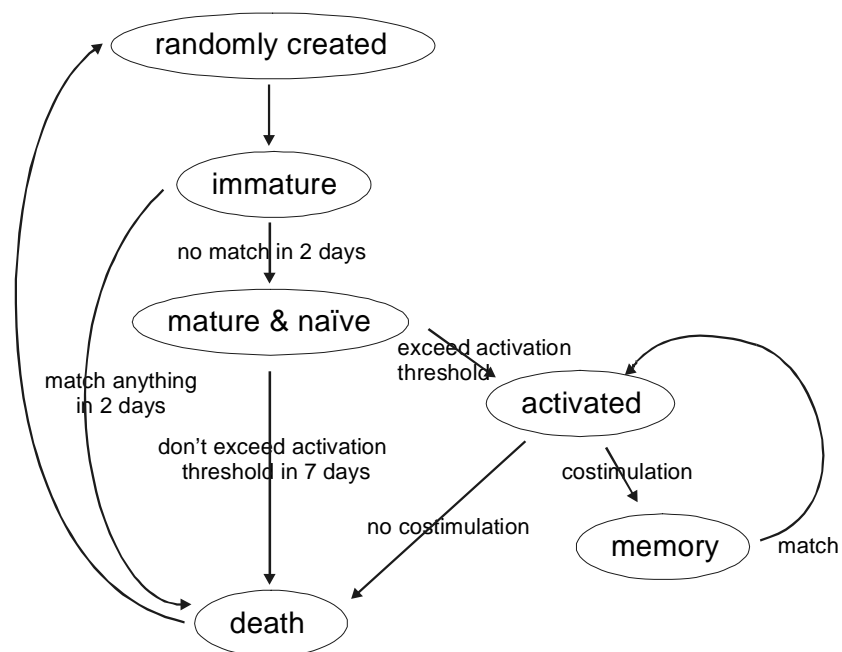


Figure 6.1 The Life of a Detector, [Hofmeyr, 1999]

A set of naïve detectors is randomly generated in each local host and these newly born detectors are exposed to currently transferred antigens (which are network packet headers). If these detectors, which are called immature detectors, are not deleted by a negative selection for a predefined time period (for instance, two days in figure 6.1), the immature detectors become mature detectors. Mature detectors are guaranteed to be self tolerant only against the self antigens having appeared during a past negative selection period. The mature detectors monitor

¹ Hofmeyr (1999) also implemented a distributed tolerisation function in his AIS. Since the dynamic clonal selection algorithm is investigated only in terms of its adaptability in this chapter, the distributed mechanisms used in

newly presented antigens and detect hidden anomalies. When each detector matches a sufficient number of anomalies, which is defined by activation threshold in figure 6.1, it sends a warning signal to a security officer for further checking. In this case, if a security officer confirms the detected anomaly pattern is indeed an intrusion pattern (costimulation in figure 6.1), the mature detector detecting this pattern becomes a memory detector. In contrast, if a security officer does not acknowledge the occurrence of an actual intrusion within a limited time, this mature detector is regarded as an auto-reactive detector and is killed. In addition, if a mature detector does not detect a sufficient number of anomalies before a pre-defined time (seven days in figure 6.1), it is eliminated. A memory detector transformed from a mature detector also monitors newly presented antigens, but it has two different features from a mature detector. It can live indefinitely even when it does not detect a sufficient number of anomalies and sends a warning signal to a security officer immediately when it detects a single anomaly.

In his system, Hofmeyr follows the life cycle of T-cells to implement dynamic network-based IDS. Each T-cell goes through various tests at different stages and it gains diverse features that are necessary to draw a complete immune response when it passes each test. Therefore, the success of this system greatly depends on controlling diverse tests introduced in the system. These tests can be characterised by the following questions: how long an immature detector has to be tested by negative selection, how long a mature detector can stay alive before it binds to a sufficient number of anomalies, what is a sufficient number of anomaly matches for a mature detector, etc. He also proposed interesting components that allow his system to provide answers to these questions. These components are tolerisation period, activation threshold, life span, decay rate, costimulation, sensitisation, distributed tolerisation, dynamic detectors and memory detector death rates. All of these new components were devised by following analogies of immune cell control mechanisms of human immune systems introduced in the previous section.

6.3 Dynamic Clonal Selection (DynamICS) Algorithm

The new AIS introduced in this chapter follows the basic concept of the AIS proposed by Hofmeyr (1999). Hofmeyr introduced various novel features of the AIS that allowed his system to be adaptive to continuously changing network traffic behaviour and perform distributed detection. From these two salient features, this section focuses on understanding how these novel features allow the AIS to be adaptive. The adaptability of Hofmeyr's AIS was achieved via co-ordinated dynamics of three different detector populations: immature, mature, and memory detector populations. In order to fully comprehend the co-ordinated dynamics of these three detector populations in terms of AIS adaptability, a new artificial immune algorithm, called the dynamic clonal selection algorithm (DynamICS) is introduced. Although Hofmeyr proposed various new features in order to effect great adaptability and distributed detection,

Hofmeyr's AIS are not reviewed in this section.

DynamiCS introduced in this section does not follow all the novel traits introduced in (Hofmeyr, 1999). Instead, DynamiCS employs only the most crucial components that can yield adequate adaptability to the system. The DynamiCS algorithm used for performing the experiments in this chapter is summarised in figure 6.2. The following sections introduce the details about DynamiCS with respect to three adopted detector populations and a non-self antigen detection mechanism using dynamics of these populations.

```

1. Initialise Dynamic Clonal Selection Algorithm
2. Create an initial immature detector population with random detectors
3.
4.  Generation_Number = 1;
5.  Do
6.  {
7.    If (Generation_Number /  $N^2$  == 1)
8.      Select a new antigen cluster.
9.
10.   Select 80% of self and non-self antigens randomly from a chosen antigen cluster;
11.
12.   Reset Parameters
13.     Memory Detector Age++;
14.     Mature Detector Age++;
15.     Immature Detector Age++;
16.
17.   Monitor Antigens
18.     Monitor Antigens by Memory Detectors if Memory Detectors are available
19.       Check whether any memory detector detects any non-self antigen
20.       Check whether any memory detector detects any self antigen
21.
22.     Monitor Antigens by Mature Detectors if Mature Detectors are available
23.       Check whether any mature detector detects any non-self antigen
24.       Check whether any mature detector detects any self antigen
25.       Create new memory detectors
26.       Old mature detectors are killed
27.
28.     Monitor Antigens by Immature Detectors if Immature Detectors are available
29.       Check whether any immature detector detects any self antigen
30.       Delete any immature detector matching any self antigen
31.       Create new mature detectors
32.
33.   If (immature detector population size + mature detector population size <
34.       maximum non-memory detector pop size)
35.   {
36.     Do
37.       Generate a random detector and add it to an immature detector population
38.     Until (immature detector population size + mature detector population size =
39.           non-memory detector pop size)
40.   }
41.   Generation_Number++;
42.
43. } While (Generation_Number < max Generation)
44.

```

Figure 6.2 Pseudo Code of the Dynamic Clonal Selection Algorithm

² N is the number of generations that each antigen cluster is used for when selecting antigen data. More details about N are presented in section 6.4.2 Data and Parameter Setting.

6.3.1 DynamiCS Overview

Before the detailed descriptions of each component in DynamiCS are presented, the overview of DynamiCS is stated in this section. DynamiCS starts with generating initial immature detectors by seeding with random genotypes. The generated immature detectors monitor a currently collected antigen set. Whereas two distinguished antigen sets, a self set and non-self set, are given to the StatiCS, DynamiCS considers the currently presented antigen set as a self antigen set and starts negative selection by comparing immature detectors to the given antigen set. As the result of negative selection, the immature detectors binding any antigen are deleted from the immature detector population and then new immature detectors are generated until the number of immature detectors becomes the maximum size of a non-memory detector population, see figure 6.3.

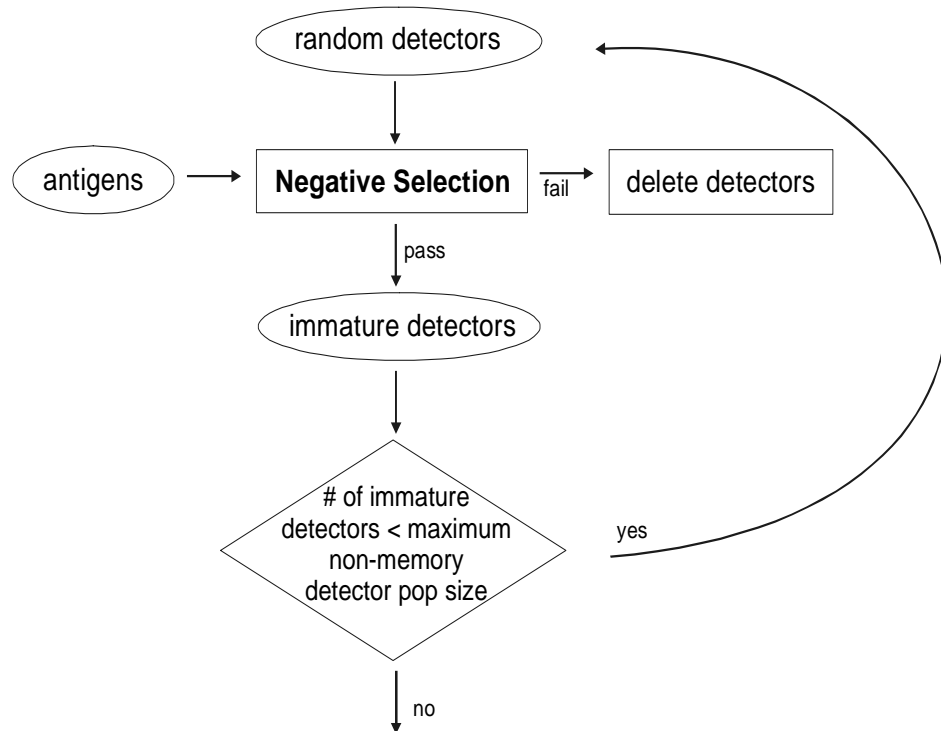


Figure 6.4

Figure 6.3 Immature Detector Generation

These same processes, which generate new immature detectors and perform negative selection with a new antigen set gathered only for a current generation, continue for the tolerisation period (T) number of generations. When the total number of generations reaches T , some immature detectors whose age reaches T , which were born at generation 1, become mature detectors (Figure 6.4). In other words, at generation T , one subset of immature detectors is deleted by negative selection and another subset of immature detectors becomes a mature detector set. The rest of the immature detectors remain in an immature detector population. The

generation T finishes by generating new immature detectors until the number of immature detectors plus with mature detectors meets the maximum size of non-memory detector population. The negative selection process used for generating immature detectors implements the central tolerisation of the human immune system.

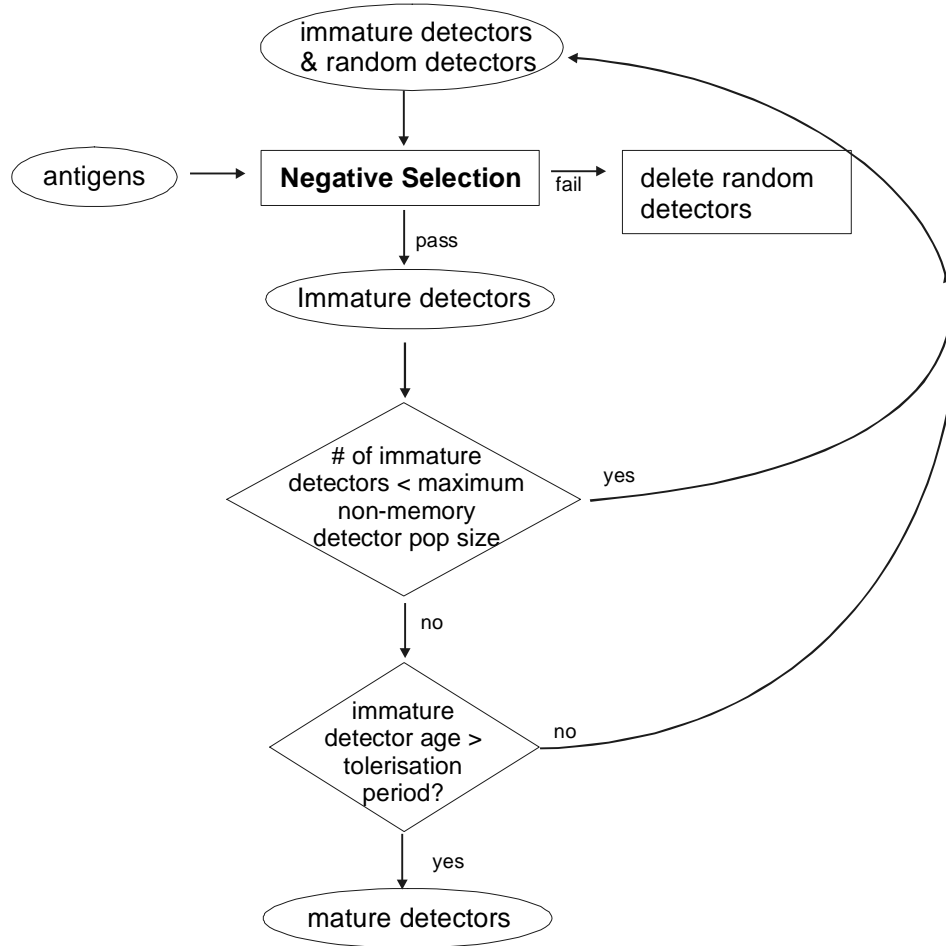


Figure 6.4 Mature Detector Generation at generation T

At generation $T + 1$, a new antigen set is presented to the mature detectors to be monitored (Figure 6.5). Whenever a mature detector matches an antigen, the match count of a mature detector increases by one. After all the given antigens are compared to all the existing mature detectors, the system checks two points: i) whether the match counts of mature detectors are larger than a pre-defined activation threshold (A) and ii) whether the ages of mature detectors meet a pre-defined life span (L). Firstly, if there is a mature detector with a match count that is larger than A , this mature detector becomes a memory detector only if it indeed detects an intrusion. When a human security officer acknowledges that this detector detects any intrusion signature (costimulation), the detector activates and eventually becomes a memory detector. In addition, the antigen patterns detected by this mature detector, which is confirmed to be a

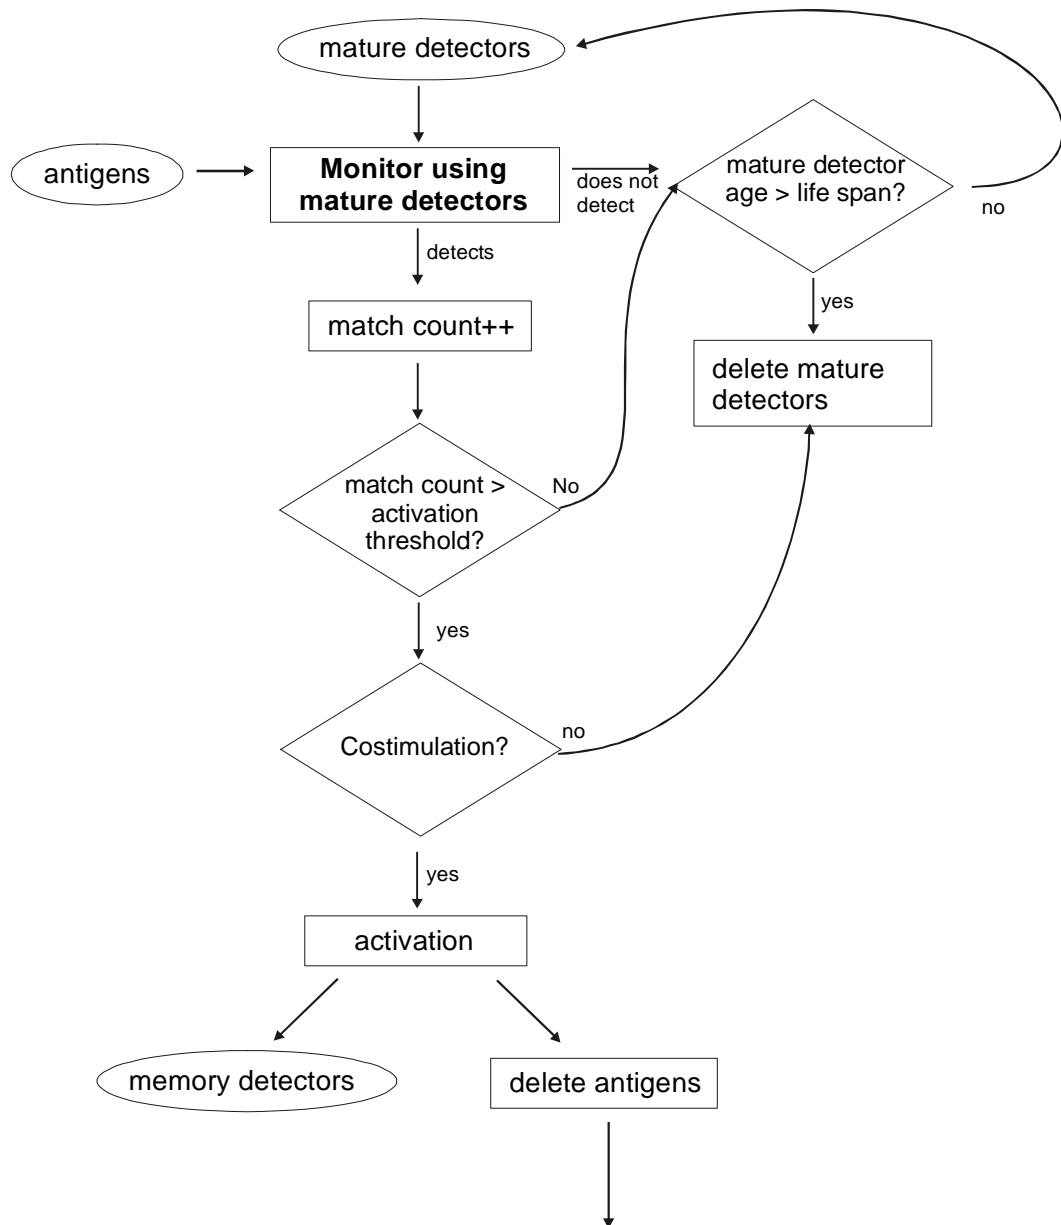


Figure 6.4

Figure 6.5 Mature Detector Activation

memory detector, are deleted from the antigen set. On the other hand, the other mature detectors having match counts less than A , still remain as mature detectors and wait, checking whether their ages reach given life span. In addition, the antigens detected by them are not discarded from the antigen set. Secondly, the system checks the ages of remaining mature detectors. If the ages of mature detectors meet L , those mature detectors are deleted from the mature detector population. After some antigens detected by activated mature detectors are deleted from the antigen set, the remaining antigens are presented to the immature detectors for negative selection. As stated before, some immature detectors are eliminated as a consequence of negative selection. In order to fully implement distributed tolerisation of the human immune system, both helper T-cells and costimulation are needed. However, DynamiCS does not

generate helper T-cells. It only implements costimulation in order to obtain the distributed tolerisation effect.




Figure 6.6 Monitor using Memory Detectors

At generation $T + 2$, memory detectors monitor transferred antigens first if there are any memory detectors generated from generation $T + 1$ (Figure 6.6). When memory detectors match any antigen, the detected antigen patterns are instantly deleted from the antigen set. After monitoring of new antigens by memory detectors, the remaining antigens are shown to mature detectors if there are any mature detectors. Then, mature detectors follow the same monitoring steps as they did at generation $T + 1$. After the antigens are monitored by the mature detectors, they are passed to immature detectors to perform negative selection. Similarly, some detectors are killed from negative selection and new random detectors are created to fill up the numbers of the non-memory population. From $T + 3$ generation, the same monitoring procedures that have operated at generation $T + 2$ continue in order to monitor constantly changing antigen sets until the system terminates.

As seen in this section, three different detector populations and various parameters are important components in supplying effective non-self antigen detection and self tolerance. The detailed descriptions of these components are stated in the following sections.

6.3.2 Immature Detectors

DynamiCS starts by generating an immature detector population. An immature detector population is created by seeding it with random genotypes. This follows the same method of generating the initial detector populations in the negative selection and static clonal selection algorithms. In addition, detector genotypes and phenotypes used for the DynamiCS are the same as those used for the StatiCS. Newly created detectors in an immature detector population

are then tested by a negative selection operator immediately after their birth³. As seen in the previous chapter, a negative selection operator compares immature detectors to given self antigens. From this comparison, if an immature detector binds to any self antigen, this immature detector is discarded from an immature detector population.

The most important feature of DynamiCS is the monitoring of continuously changing self and non-self antigens. This is different from a static clonal selection algorithm, which analyses only static self and non-self antigens in a batch style. In DynamiCS, self antigens presented at one generation are different from self antigens given at another generation. An immature detector population at one generation experiences only a subset of the entire self antigen set. This different condition requires an extra mechanism to ensure sufficient self-tolerance, which is known as central tolerisation. If an immature detector passes a negative selection test only against self antigens presented at one generation, this detector will have tolerance only against a small subset of self antigens appearing in this generation. In order for an immature detector to gain satisfactory self-tolerance, it should be tested against a sufficient proportion of self antigens. However, it is very difficult to predict how much of the self antigen population are sufficient to provide adequate self-tolerance. Thus, a new parameter, called the *tolerisation period*, is defined. The tolerisation period can control the “sufficient proportion” of self antigens compared to an immature detector for negative selection.

A tolerisation period simply indicates the number of generations that an immature detector has to pass a negative selection test. If an immature detector binds no self antigen pattern for the tolerisation period, it becomes a mature detector and the mature detector starts monitoring newly-given antigens. On the other hand, if the number of generations which an immature detector has survived so far against a negative selection test is less than a pre-defined tolerisation period, an immature detector stays in the immature detector population and waits for a new self antigen set. Any immature detector can remain in an immature detector population for a tolerisation period as long as it survives against a negative selection test. Hence, different values for the tolerisation period will show varying degrees of self-tolerance.

6.3.3 Mature Detectors

Next, mature detectors, which have gained tolerance against self antigens, start monitoring new antigens. Each mature detector attempts to bind to new antigens but when a mature detector matches any new antigen, it does not immediately regard the detected antigen as non-self. Instead, it continues attempting to match more new antigens until it binds to a sufficient number

³ DynamiCS currently uses negative selection to generate immature detectors. This is because the purpose of the initial experiments performed in this chapter is to understand the system dynamics. Thus, the system is tested only on a relatively small data set that a negative selection algorithm can easily cope with and do not investigate the scalability of this system.

of new antigens. This mechanism is analogous to the affinity maturation of B-cells and T-cells in the human immune system – they have to bind to a sufficient number of antigen epitopes in order to activate. When B-cells activate, they send “non-self cell appearance” signals to other immune cells and kill detected non-self antigens in various ways. Since B-cell activation results directly in killing detected antigens, the human immune system is very cautious about triggering the activation of immune cells, in order to avoid killing mistakenly detected self antigens.

DynamiCS uses a similar idea in order to assure self-tolerance. Although mature detectors have obtained self-tolerance to an extent through negative selection tests performed when they were immature detectors, their degrees of self-tolerance vary greatly depending on the pre-defined tolerisation period. A relatively short tolerisation period will not allow mature detectors to develop a satisfactory degree of self-tolerance. Thus, DynamiCS embeds other parameters that inhibit false-positive errors: an *activation threshold* and a mature detector *life span* [Hofmeyr, 1999]. Each mature detector has a match count and an age. The match count of a mature detector is simply the number of matching antigens, and the age is the number of generations that the mature detector has remained in the mature detector population so far. The algorithm regards the antigen bindings of a mature detector as definite non-self antigen occurrences only when the match count is larger than a pre-defined activation threshold. To be more precise, DynamiCS provides no non-self antigen appearance signal until a mature detector binds to a sufficient number of anomalous patterns, defined by the value of the activation threshold.

However, if a mature detector stays in a mature detector population for a long period, it is highly likely that the match count of a mature detector will eventually meet the activation threshold. Therefore, in order to produce an appropriate effect from the activation threshold, it is necessary to restrict the maximum number of generations that a mature detector can stay in a mature detector population. The life span of a mature detector defines the maximum number of generations that a mature detector can remain in the mature detector population. To summarise, a mature detector stays in a mature detector population only for the number of generations corresponding to its life span. A mature detector accumulates its match count as it binds to more new antigens. If the match count does not reach an activation threshold before its age becomes the life span, this mature detector is eliminated from the mature detector population for good. In contrast, if the match count satisfies the activation threshold before the age of a mature detector reaches the life span, the mature detector provides a non-self detection signal and the detected antigens to a security officer for further checking. This mechanism is designed to detect intrusions that attempt to compromise target systems intensively and for a short period [Hofmeyr, 1999]. Those intrusions create a relatively large number of non-self patterns over a

short period. Therefore, the proposed criteria that allow mature detectors to activate work only when the system aims to detect this type of intrusions

The non-self detection signal and detected antigens sent by a mature detector are analysed by a human security officer [Hofmeyr, 1999]. This is called *costimulation* in the DynamiCS. As B-cells bound to antigens wait for the confirmation signal from helper T-cells and helper T-cells wait for the signal from the innate immune systems, DynamiCS also waits for a confirmation signal from a human security officer for a pre-defined period. If he or she acknowledges a detected pattern as a definite non-self antigen pattern (an intrusion pattern), the mature detector activates.

6.3.4 Memory Detectors

Activated mature detectors immediately become memory detectors. Memory cells in the human immune system have been proven to be useful in the detection of non-self antigens [Smith *et al.*, 1996]. This feature of memory detectors requires them to have two different traits from mature detectors: a lower activation threshold and a longer life span. In contrast to mature detectors, memory detectors activate immediately when they match any single antigen, that is to say that the activation threshold of memory detectors is one generation. This is because memory detectors have matched true non-self antigens in the past, and thus any antigen detected by these detectors is considered as a definite non-self antigen without extra checking.

This is analogous to the secondary response of memory cells in the human immune system. The secondary responses by memory cells are much quicker and more efficient than the primary responses by B-cells [Sompayrac, 1999; Paul, 1993]. Moreover, the memory cell provides an associate memory property. Similarly, memory detectors can quickly detect new antigens that are not the same as previously detected antigens, but are within the matching boundaries of the detector. Since these features of memory detectors are advantageous for the AIS, an infinite life span is set for them. Mature detectors are essentially undergoing a test of whether they can detect non-self antigens. Their limited life span is the maximum period they can participate in the test. However, memory detectors do not need such tests and thus do not require a limit on their life spans. Moreover, they are expected to detect non-self antigens that are associative to previously detected non-self antigens.

6.3.5 Controlling Detector Birth and Death

As emphasised earlier, one distinguishing feature of DynamiCS is that selection processes operate through experience of different antigen sets at different generations. It has to constantly monitor dynamically changing self and non-self antigens. This requires the algorithm to keep replacing existing detectors according to newly observed antigens. The dynamic replacement of

different levels of detector is achieved through the parameters introduced in above sections. These parameters enable the system to decide automatically the deaths and births of detectors in each population. The births and deaths of detectors in each detector population and related parameters are summarised in the table 6.1.

Detector Population	Size	Birth	Death	Detection
Memory	Fixed	From mature detectors	Live infinitely	✓
Mature	Half-Fixed	From immature detectors.	1) Become memory detectors when $matchCount > actThreshold$ 2) Deleted when the age $> LifeSpan$	✓
Immature	Half-Fixed	Filled up by randomly generated detectors	1) Deleted when they fail negative selection 2) Become mature detectors when age $> tolerisation\ Period$	

Table 6.1 Detector Birth and Death in the Dynamic Clonal Selection Algorithm

Most detector births and deaths are determined by detector-antigen binding and values of various parameters. Deaths of some detectors trigger births of other detectors. For instance, some deaths of immature detectors cause births of mature detectors, and some deaths of mature detectors lead to births of memory detectors. It should be noted that births of mature detectors and births of memory detectors only result from deaths of immature detectors and deaths of mature detectors respectively. Accordingly, only births of immature detectors and deaths of memory detectors are not pre-determined by the algorithm – other births and deaths of detectors are controlled by detector-antigen binding and other parameter values.

In each generation, some of the immature detectors will be deleted by a negative selection test, or else become mature detectors. The system generates new immature detectors until the size of the non-memory detector population (immature detectors plus mature detectors) reaches the pre-defined maximum size. The motivation for restricting the number of immature detector births in this way is to prevent excessive resource usage⁴. Here, DynamiCS restricts the number of immature detector births by defining the maximum size of non-memory detector population instead of the immature detector population. In contrast to immature detectors, mature detectors are born and die as an inherent part of the algorithm and so the mature population size cannot be directly controlled. Thus, the number of new immature births at each generation is limited not only by the number of immature detectors but also by the number of mature detectors.

⁴ One important requirement for IDS's is that they must be lightweight, as stated in chapter 3. The primary motivation for adopting a dynamic clonal selection algorithm can be to provide a lightweight AIS. It can be assumed that the system simply allows an unbounded accumulation of immature detectors until they might cover the non-self universe. However, it clearly requires the use of an excessive amount of resource for storage. For this reason, the dynamic clonal selection algorithm uses dynamic replacement of three different types of detectors.

Therefore, the number of immature detectors per generation varies according to the number of existing mature detectors. The number of mature detectors is also directly affected by their lifespan. The approach taken allows the lifespan of mature detectors (L) to be varied without causing a direct effect on the total number of immature and mature detectors. This allows L to be varied experimentally without a proportionate effect on resource utilisation caused by a change in the combined population size.

The memory detector population size also needs to be restricted for the same reason. However, DynamiCS as analysed in this chapter does not limit the memory detector population size. This is because defining the maximum number of memory detectors implies deaths of memory detectors. However, the elimination of memory detectors is not straightforward due to their verified usefulness. An appropriate strategy to kill memory detectors will be examined in the next chapter.

6.3.6 Self and Non-Self Antigen Detection

Whereas the static clonal selection algorithm is presented with two distinguished antigen sets: a non-self antigen set and a self antigen set, the antigens that are given to DynamiCS do not have a label indicating whether they are self or non-self. Therefore, the system treats transferred antigens as self and non-self at the same time depending on which type of detectors handle them. More precisely, among the three different types of detector, memory detectors monitor antigens first when antigens are presented to the system. Since memory detectors remember non-self antigen patterns that are already proven to be intrusions, it is expected that memory detectors will detect real intrusion patterns hidden in a presented antigen set. In this case, the detected antigen patterns should be deleted from the collected antigen set in order to avoid the loss of potential mature detectors.

The antigen set is presented to mature detectors after being filtered by memory detectors. Although memory detectors detect some non-self antigen patterns, there can be further non-self antigen patterns that have never occurred before, and thus no memory detector binds them. Mature detectors are expected to detect those new non-self antigen patterns. However, the antigen patterns that match mature detectors can be self antigens that have never occurred for a tolerisation period, when mature detectors were immature detectors. Therefore, the antigen patterns matched by mature detectors are not deleted from the antigen set until mature detectors can activate by receiving costimulation confirmation. In other words, mature detectors binding antigen patterns only increase their match counts, and matched antigen patterns are not eliminated from the antigen set. The antigen patterns matched are deleted from the antigen set only if the accumulated mature detector match counts meet the activation threshold and they

receive costimulation from a security officer. Finally, the antigen set filtered by memory detectors and mature detectors is sent to immature detectors for negative selection.

6.4 Dynamic Clonal Selection Algorithm Experiments

6.4.1 Objective

As introduced in section 3, there are several parameters that will control the performance of DynamiCS. Among them, three parameters, tolerisation period (T), activation threshold (A) and life span (L) are newly introduced in order to provide the adaptability of the AIS with a constantly changing antigen set. Although these parameters were introduced from previous work (Hofmeyr, 1999), the behaviours of the AIS directed by the various values of these parameters were not thoroughly analysed. The following experiments focus on understanding system behaviours under different values of these three parameters. The experimental results are investigated primarily in terms of how each parameter affects the adaptability of the AIS.

6.4.2 Data and Parameter Setting

The work in this chapter aims to provide an understanding of the nature of DynamiCS. The experiments performed for this chapter used the Wisconsin breast cancer data set that was employed for the study of the static clonal selection algorithm in chapter 5. The detailed features of this data set are already presented in chapter 5 and these features still apply to this data set. The cancer data has two classes, ‘Malignant’ and ‘Benign’. ‘Malignant’ has 240 examples and ‘Benign’ has 460 examples. The system treated ‘Malignant’ as non-self and ‘Benign’ as self (See Table 6.2).

Class	Malignant			Benign		
Number of examples	240			460		
Antigen Class	Non-self			Self		
Cluster-ID	1	2	3	1	2	3
Number of examples	45	117	78	42	335	63

Table 6.2 The features of Cancer Data

Although the same cancer data set used in chapter 5 was applied in the coming experiments, the new experiments required a different strategy for the presentation of antigen data to DynamiCS. Since the main benchmarking measure for the new experiments was the adaptability of the new algorithm, one criterion for the provision of antigen data to the AIS was that antigen data sets given in each generation should have varied distributions. Furthermore, in order to comprehend the system’s new behaviours, it was necessary to understand the degree of differences between various distributions of antigen sets in advance. Therefore, a different method is adopted for providing antigen data to DynamiCS.

In order to be sure of providing antigens with a new distribution each N generations, self and non-self antigen data was clustered into several groups and antigen data randomly selected from one cluster was presented for N generations. Thus, the system was provided with antigen data whose distribution was switched at every N th generation. The Expectation Maximization (EM) clustering algorithm was applied to cluster antigen data. The EM algorithm is widely-used as the basis for various unsupervised learning algorithms [Mitchell, 1997]. The EM algorithm clustered antigen data into three clusters. As the result of clustering, 240 ‘Malignant’ examples were divided into three clusters of 45, 117 and 78 examples. Similarly, 460 ‘Benign’ examples were grouped into three clusters 42, 355 and 63 examples (see table 6.2).

80% of the self and non-self antigen data belonging to each cluster were randomly selected for N generations. N , the number of generations that each cluster was used for selecting antigen data, was pre-defined. Therefore, DynamiCS was provided with different antigen data at each generation and the distributions of these data changed at every N generations. In addition, the antigen clusters used for providing antigen data were selected in a regular cyclical order, with the first cluster re-used after $3 * N$ generations. Thus, non-self and self antigen data were always selected from the same cluster for N generations and the same cluster pairs were re-used again after a break of $2 * N$ generations (see lines 7-10 in figure 6.2).

In addition, the costimulation mechanism involving a security officer was implemented by simply increasing the match count only when a detector detects non-self antigens. This is a simplification from the description in section 6.3. If the costimulation procedure is strictly implemented, detector match counts should increase whenever detectors detect an antigen, and the AIS should confirm whether they indeed detect non-self antigens only when their match counts reach activation thresholds. However, in order to reduce the running time of the system, a simpler method is chosen for the same job. Since a detector that gained its match count by detecting self antigens is not ultimately useful, it is expected to achieve the same result by incrementing the match count only when a detector detects a non-self antigen. This modification neglects one of the significant issues to be investigated in respect of DynamiCS: the frequency of false alarms caused by increased match counts due to self antigen detection. However, the investigation of this issue is delayed for future work and instead focus on understanding system behaviours under the influence of various parameters.

Another important assumption made in the following experiments is that there is no noise in antigen sets provided for negative selection. When an antigen set including both self and non-self antigens is transferred to the system, if non-self antigens are not detected by memory detectors and activated mature detectors, those non-self antigens are treated as self antigens. Thus, these non-self antigens become noise for the purposes of a negative selection. This

immediately affects the system's likelihood of providing false self information and thus results in the elimination of sound immature detectors. In other words, it causes the loss of some potential mature detectors that could have detected non-self antigen patterns. Therefore, the lower non-self antigen detection rates caused by antigen noise greatly depend on how quickly memory detectors and activated mature detectors detect hidden non-self antigen patterns. This issue is also left as a future research until the behaviours of DynamiCS is sufficiently understood. Hence, the antigen data set provided for negative selection is always comprised solely of self antigens.

All experiments were run for 2000 generations and repeated five times. A non-memory detector population size of 240 was used. This number was the number of worst detectors replaced at each generation from the static clonal selection algorithm. Since the non-memory detectors keep being replaced, they played a similar job to the worst detectors that were replaced in the static clonal selection algorithm. Thus, the same population size, 240, was used. Experiments were run by taking various values of the three parameters; tolerisation period of an immature detector (T), activation threshold of a mature detector (A) and the life span (L) of a mature detector.

6.4.3 Experiment Design

Three series of experiments were performed by varying the distributions of the provided antigen data. The first series of experiments was carried out by giving the value one to N , the number of generations that the same cluster is selected for presenting antigen data. DynamiCS was provided with antigen data of a different distribution at every generation. This antigen presentation method gave the system the opportunity to experience the complete antigen data set. This is because the complete antigen set was divided into three groups based on distribution, and all of these three distributions of antigens were input equally into the system. Moreover, all of the three distributions of antigens were provided a sufficient number of times to be learned. In this way, these experiments aim to evaluate whether DynamiCS is able to learn the converged behaviours of three different antigen clusters after a certain period, when only a subset of antigens is presented at one time. In contrast, the values of N employed for the second series of experiments ranged from 5 to 50. As N increases, the system will overfit the distribution of only one antigen cluster more for N generations. Thus, the significant question for investigation in the second set of experiments is how quickly the system is able to learn the distribution of a new antigen cluster, when an antigen cluster is replaced.

The antigen presentation setting for the first series of experiments was defined in order to test whether the system can incrementally learn globally converged distributions, when only its one subset distribution is given at each generation. This is quite a common environment for IDS's

to face. Anomaly detectors adopted for IDS's have to learn the globally converged normal behaviours of monitored targets. However, the audit data sets collected to build normal profiles are usually enormous. Therefore, they have to learn the globally converged behaviours by learning incrementally from a small amount of subset data, which is practical for IDS's to store. Bearing this kind of real environment in mind, the section 6.5 examined the first series of experimental results to determine how much varying values of the three parameters influenced the incremental learning capability of the AIS.

The second series of experiments followed another common real circumstance. This circumstance was such that converged behaviours learned in an incremental way are suddenly altered due to a normal change in the self. Examples of normal changes in the self include the arrival of a new authorised user, the departure of established, authorised users, and some local system replacement. In this case, the IDS's should quickly forget what they have learned in the past (especially that which does not fit the new environment) and learn behaviours that emerge from the new environment. Thus, in the section 6.6, the second set of experiments was analysed with respect to the sensitivities of the following features to the three parameters: forgetting out-of-date behaviours and rapid learning of new behaviours.

6.5 Experiment Results 1: Examination of complete antigen data

6.5.1 Effect of the Tolerisation Period

The first set of experiments is performed for examining the effect of the tolerisation period. The parameter values are used for these experiments are shown in table 6.3.

Parameters	Values
Tolerisation Period (T)	{5, 10, 20, 50}
Life Span of Mature Detectors (L)	10
Activation Threshold of Mature Detectors (A)	{5, 100}
Number of Generations that Antigens are Selected from a Same Cluster (N)	1

Table 6.3 Parameter values used for the experiments performed in 6.5.1 Effect of the Tolerisation Period

Figure 6.7 illustrates the results of the first set of experiments, where tolerisation period (T) was varied from 5 to 10, 20 and 50 with activation threshold (A) equal to 100 and life span (L) equal to 10. The X-axes of these graphs represent the number of generations and the Y-axes indicate detection rates. Each graph has two lines, one displaying a True Positive (TP) rate and another showing a False Positive (FP) rate. As defined in the previous chapter, TP was the “non-self” detection rate and FP was the rate at which “self” was mistakenly detected by a generated detector set. First of all, it can be seen that TP values oscillate between two converged minimum and maximum values and these converged minimum and maximum values decrease

Figure 6.7 TP and FP rates when T varies and $A = 100, L = 10, N = 1$

Figure 6.8 here

Figure 6.8 TP and FP rates when T varies and $A = 5, L = 10, N = 1$

as T increases. For instance when two extreme cases, $T = 5$ and $T = 50$, are compared, TP displays a range from around 0.6 to near 0.9 when $T = 5$, but it oscillates from 0.3 to near 0.8 when $T = 50$. However, no change in the FP rates for the four experiments is found in figure 6.7. In all four cases, FP remains stable at the optimum value zero⁵. In order to see whether FP is indeed insensitive to T , another set of experiments was performed. This time, A is set to 5 instead of 100. Since a larger value of A makes the DynamiCS stricter when evaluating whether to activate mature detectors, it can cause less frequent detection, whether of self or non-self. By giving a smaller value of A , we encouraged mature detectors to activate more easily and thus expected to gain higher TP and FP rates. The second set of experimental results is given in figure 6.8.

As expected, the second set of experiment results show higher TP and FP rates than those exhibited in the first series of experiment results. All four results demonstrate comparatively high TP rates ranging between 0.9 and 1.0. In addition, converged maximum and minimum TP ranges oscillate within much smaller scopes than those displayed in figure 6.7. These results are different from the TP rates shown in figure 6.7. These different results are analysed by studying the effects of varying A in the next section, 6.5.2 Effect of the Activation Threshold.

In both figures 6.7 and 6.8, TP values oscillate between two converged minimum and maximum values and these converged minimum and maximum values decrease as T increases. These effects are more clearly illustrated in figure 6.7. Another key result revealed by figure 6.8 is a dramatic drop in FP when T increases from 5 to 10. When $T = 5$, FP steadily increases and reaches 0.6 by the time the number of generations becomes 2000. By contrast, when $T = 10$, FP immediately becomes zero from generation one and stays at this optimal value for the entire 2000 generations. Thus, two significant changes were effected by varying T : firstly both TP and FP rates decrease and secondly the drop in FP is much sharper than the drop in TP.

These results clearly illustrate the role that the *tolerisation period* plays in DynamiCS. They demonstrate that the employment of a *tolerisation period* directly benefits the system by complementing the expected low degree of self-tolerance when only a subset of antigens is provided at each generation. As long as immature detectors have an opportunity to experience various antigen distributions for a sufficient period, which is defined by the tolerisation period, the FP can be dramatically reduced to an almost perfect near-zero rate.

Moreover, these results confirm that having large value of T results in a high degree of self tolerisation at the expense of TP. This outcome can be scrutinised by examining the proportion of the population made up of non-memory detectors. Since the maximum number of non-

memory detectors, consisting of immature detectors and mature detectors, is fixed, one type of detector has to diminish when another type of detector expands. It means that when a large value of T forces detectors to remain immature longer, the average immature detector population size per generation gets larger and the average mature detector population size per generation becomes smaller. This is shown in tables 6.4 and 6.5. The smaller number of mature detectors implies that a smaller number of candidate detectors are qualified to activate. Consequently, this results in a smaller number of total detector activations. For the same reason, a large value of T leads the system to produce a smaller number of memory detectors in total. Since DynamiCS does not employ any niching mechanism like the one introduced in the static clonal selection algorithm, the smaller number of generated mature and memory detectors directly causes low TP and FP rates. This is because it is unlikely that any detector will match a significantly larger number of antigens than any other detector when all detectors are randomly generated. Random generation might produce a powerful detector by chance, but this will not occur consistently. Thus, a more consistently expected outcome, which is shown from the experiments, is that more mature detectors produce more frequent antigen detection.

$A=100$	Total Number of Memory Detectors	Average Mature Detector Population Size per Generation	Average Immature Detector Population Size per Generation	Average Number of Mature Detector Generated per Generation
$T=5$	8.5 (1.667)	149.11 (0.0021)	90.89 (0.0021)	14.95 (0.000009)
$T=10$	7 (3.333)	113.12 (0.007)	126.88 (0.0071)	11.35 (0.00006)
$T=20$	8 (6.667)	76.56 (0.003)	163.44 (0.003)	7.68 (0.00005)
$T=50$	5.5 (4.333)	38.52 (0.0006)	201.48 (0.0006)	3.87 (0.000005)

Table 6.4 Proportion of Three Different Types of Detector when T varies and $A = 100$, $L = 10$, $N = 1$. The values in parentheses are variances.

$A=5$	Total Number of Memory Detectors	Average Mature Detector Population Size per Generation	Average Immature Detector Population Size per Generation	Average Number of Mature Detector Generated per Generation
$T=5$	65.5 (13.67)	151.58 (0.141)	88.42 (0.141)	15.21 (0.0014)
$T=10$	42 (23.33)	122.99 (0.0024)	127.01 (0.0024)	11.343 (0.00001)
$T=20$	39 (24.66)	76.51 (0.0006)	163.49 (0.0006)	7.68 (0.000002)
$T=50$	37.25 (40.92)	38.45 (0.0007)	201.55 (0.0007)	3.87 (0.000006)

Table 6.5 Proportion of Three Different Types of Detectors when T varies and $A = 5$, $L = 10$, $N = 1$. The values in parentheses are variances.

In addition, the range between maximum and minimum values of TP rates tends to get larger as T increases. This difference is more evidently presented when two cases, when $T = 5$ and $T = 10$, are compared in figure 6.7 and 6.8. This can also be explained with the same reason. The

⁵ Since FP values are constantly shown as zero for 2000 generations, FP values in figure 6.7 and 6.9 do not appear to be shown. Thick lines overlapped with X-axes are zero FP values in these figures.

smaller number of mature detectors tends to cover a smaller number of the niches that could exist in the non-self antigen set. Since three different distributions of antigen set were given in turn at each generation, oscillating TP rates indicate differing results of detection of non-self antigens between the three different clusters. The TP rates obtained between generation 800 and 899 when $T = 20$ in figure 6.7 are redrawn in figure 6.9 with a larger scale. It can be seen from figure 6.9 that three stabilised TP rates appear repeatedly every three generations. This verifies that the fluctuating TP rates result from three different antigen detection rates associated with the three different clusters. Therefore, the range of fluctuation in TP rates will be reduced if detectors which are qualified to perform antigen detection cover niches in non-self antigen clusters evenly. A smaller number of detectors will only cover a smaller number of randomly scattered niches in each non-self antigen cluster. This is again because detectors randomly generated via a negative selection have not had any chance to evolve to tightly matching antigen niches. In this case, activated mature detectors are detecting a random subset of non-self antigens. This causes the range between maximum and minimum values of TP rates to become larger as T increases.




Figure 6.9 TP rates between generation = 800 and generation = 899 when $T = 20$

Furthermore, when T increases, TP does not drop radically in contrast to the large drop in FP shown in figure 6.8. This again proves that negative selection with a tolerisation period drives the system to gain sufficient self-tolerance directly, whilst it affects non-self antigen detection only indirectly. Another noticeable result shown in figure 6.8 is steady growth of FP when $T = 5$. This is contrast to the other cases in which TP and FP rates somehow tend to stabilise. This outcome can be caused by using a relatively small value for T . In this case, mature detectors were generated by testing against a largely incomplete self-antigen set. Thus, the generated mature detectors are still more-or-less random detectors, which cannot distinguish between self and non-self. However, among these more-or-less random detectors, the only detectors to activate are those subsets of detectors which actually detect a number of non-self antigens.

Thus, initial random detectors have a chance to be selected based on non-self antigen detection performance, but do not have opportunities to be tested against sufficient self antigens. As a consequence, although the generated memory detectors have some information about previously detected non-self antigens, they did not have enough information about which patterns are not self. Thus, as this kind of memory detector is generated more and more, the system is expected to make more mistakes in detecting self antigens.

Finally, we can discern that all TP's suddenly jump up from zero just after the generation number reaches the end of the first tolerisation period. For example, when $T = 5$ in figure 6.7 and figure 6.8, TP starts jumping up to 0.5 just after the generation number is 5. Similarly, when $T = 50$, TP suddenly becomes around 0.5 just after the generation number reaches 50. This is simply because every detector, which can trigger activation, has been generated until the system run passes the first tolerisation period. Only after the first tolerisation period, the system starts an antigen detection process.

6.5.2 Effect of the Activation Threshold

The second series of experiments is performed for examining the effect of the activation threshold. The parameter values are used for these experiments are shown in table 6.6.

Parameters	Values
Tolerisation Period (T)	5
Life Span of Mature Detectors (L)	10
Activation Threshold of Mature Detectors (A)	{5, 10, 20, 40, 50, 100, 150, 200}
Number of Generations that Antigens are Selected from a Same Cluster (N)	1

Table 6.6 Parameter values used for the experiments performed in 6.5.2 Effect of the Activation Threshold

The figures 6.10 and 6.11 show the experimental results, displaying TP and FP rates obtained from these eight experiments. As seen in the previous set of experiments, the results gained from the new series of experiments also exhibit fluctuating TP and FP values between two converged minimum and maximum values. Both TP and FP rates tend to decrease as A increases. These results confirm that the activation threshold contributes to reduce FP further by making the system stricter in triggering activation. However, similar symptoms that were observed from the previous experiments in 6.5.1 Effect of the Tolerisation Period are also found: lowering FP causes decline of TP. The explanation of variations in TP and FP rates according to various values for A can be found in table 6.7.

Figure 10 here.

Figure 6.10 TP and FP rates when A varies with $T = 5, L = 100, N = 1$ (1)

Figure 6.11 TP and FP rates when A varies with $T = 5, L = 100, N = 1$ (2)

	Total Number of Memory Detectors	Average Mature Detector Population Size per generation	Average Immature Detector Population Size per generation	Average Number of Mature Detector generated per generation
A=5	63.75 (49.58)	151.59 (0.26)	88.41 (0.26)	15.21 (0.0026)
A=10	37.5 (33.67)	150.03 (0.119)	89.97 (0.119)	15.05 (0.0011)
A=20	22.5 (12.33)	149.30 (0.013)	90.70 (0.013)	14.98 (0.0002)
A=40	16.5 (19)	149.38 (0.117)	90.62 (0.181)	14.98 (0.0010)
A=50	14 (6)	149.13 (0.079)	90.87 (0.079)	14.96 (0.001)
A=100	8.5 (3.67)	149.19 (0.035)	90.81 (0.035)	14.96 (0.001)
A=150	6.75 (0.92)	149.01 (0.063)	90.99 (0.063)	14.95 (0.001)
A=200	5 (3.33)	148.91 (0.007)	91.09 (0.007)	14.93 (0.00007)

Table 6.7 Proportion of Three Different Types of Detector when A varies and $T = 5, L = 10, N = 1$
The values in parentheses are variances.

As can be seen, differing values for A do not affect the average mature and immature detector population sizes per generation. Unlike T , a large A does not reduce the number of candidate detectors activated. Instead, large A causes the activation of mature detectors to be much less frequent. Accordingly, a much smaller number of memory detectors were generated during the full 2000 generations. DynamiCS will obtain higher TP rates when mature detectors detect diverse niches existing in a non-self antigen set. As discussed in section 6.5.1 Effect of the Tolerisation Period, this is because detectors are mainly generated by a negative selection without a niching mechanism. Thus, the smaller amounts of memory detectors detect only a subset of non-self antigen niches randomly scattered and this induces lower TP and FP rates. This also explains the differences of TP rates shown in figure 6.7 and 6.8 in the previous section. The results with $A = 5$ in figure 6.8 exhibit higher TP rates than those displayed in figure 6.7 showing the results with $A = 100$. In addition, converged maximum and minimum ranges with $A = 5$ oscillate within a much smaller scope than those with $A = 100$. These results clearly follow the outcome shown in figure 6.10 and 6.11, which larger A lowers FP rates but makes an converged maximum and minimum TP rates lower and oscillate with much larger range.

6.5.3 Effect of the Life Span

The third series of experiments is performed for investigating the effect of the life span. The parameter values are used for these experiments are shown in table 6.8.

Parameters	Values
Tolerisation Period (T)	5
Life Span of Mature Detectors (L)	{5, 10, 20, 50}
Activation Threshold of Mature Detectors (A)	150
Number of Generations that Antigens are Selected from a Same Cluster (N)	1

Table 6.8 Parameter values used for the experiments performed in 6.5.3 Effect of the Life Span

Figure 6.12 here

Figure 6.12 TP and FP rates when L varies and $T = 5, A = 150, N = 1$

Figure 6.12 exhibits TP and FP rates gained from four different experiments. As was seen in the previous experiments, these results also show that TP rates oscillating between minimum and maximum values had stabilised during 2000 generations. As L gets larger, two similar tendencies of TP rate changes are perceived. Firstly, its minimum and maximum values get larger and secondly, the oscillating scopes between minimum and maximum values tend to be narrower.

These outcomes can also be interpreted by examining the proportion of the population that is made up of non-memory detectors, shown in table 6.9. The larger number of mature detectors again implies that a larger number of candidate detectors were to be activated. Consequently, the larger number of mature detectors triggers a higher frequency of detector activation and this results in higher TP rates. The second effect can also be interpreted by the same reason

	Total Number of Memory Detectors	Average Mature Detector Population Size per generation	Average Immature Detector Population Size per generation	Average Number of Mature Detector generated per generation	Average Number of Mature Detector deleted per generation
$L=5$	4.75 (2.92)	108.07 (0.007)	131.93 (0.007)	21.65 (0.0004)	21.59 (0.0003)
$L=10$	7 (3.33)	149.14 (0.024)	90.86 (0.024)	14.96 (0.0003)	14.88 (0.0002)
$L=20$	10.75 (4.25)	183.79 (0.005)	56.21 (0.005)	9.24 (0.00002)	9.14 (0.00001)
$L=50$	18.25 (8.25)	213.74 (0.006)	26.27 (0.006)	3.34 (0.006)	4.19 (0.0002)

Table 6.9 Proportion of Three Different Types of Detector when L varies and $T = 5$, $A = 150$, $N = 1$. The values in parentheses are variances.

discussed in the previous sections. Large L allows a mature detector to remain longer and thus permits it to experience more diverse non-self antigen clusters. When each TP rate represents a TP rate for each non-self antigen cluster, the TP rate differences among three non-self clusters are not large when L is sufficiently large. Conversely, the differences become wider as smaller values of L are given. In summary, a mature detector that meets more non-self antigens can learn the distributions of each non-self antigen cluster better.

6.5.4 Analysis

The above experiments showed clearly that the three parameters investigated in this chapter influence the non-self antigen detection (TP) and self-tolerance (FP) rates significantly. A common trait found in the three different result sets is that TP and FP rates vary depending on the number of detectors which are qualified to activate. Since DynamiCS generated its initial immature detectors only through negative selection, the degree of antigen detection did not vary greatly between detectors. This was because no detector had evolved to match existing niches in the given antigen set. To summarise, the antigen detection capability of DynamiCS was governed by the total number of detector activations and this number was directly affected by three parameters.

This result indicates that effective application of DynamiCS requires a strategy to choose an ideal combination of the three parameter values, in order to get satisfactory TP and FP rates. Recalling the three sets of experimental results, it was observed that larger T led the system to produce a smaller number of detectors to activate and resulted in lower TP and FP rates. In contrast, different A values did not affect the number of candidate detectors to activate, but larger A made mature detectors activate much less frequently, and ultimately triggered a smaller number of detector activations, which also brought about lower TP and FP rates. In addition, larger L values forced mature detectors to live longer and let them to experience more diverse antigen clusters. This led the system to provoke more detector activation, which produced higher TP and FP rates.

From these outcomes, the effects caused by T and L are easily comparable: large T allows more immature detectors to remain and pushes mature detectors out, while large L compels more mature detectors to remain and pushes immature detectors out. Thus, the only difference between the effects of T and L is the direction of the change. T controls self-tolerance more directly by letting immature detectors be compared to more self antigens while L governs non-self antigen detection more directly by allowing mature detectors to have more chances of attempting to match non-self antigens. Therefore, if in a particular application the first priority is a low false positive error rate, and secondly the true positive detection rate, a T value that is large enough to show a fulfilling FP rate should be found first regardless of A and L values. After a satisfactory FP rate is obtained by tuning T , A and L can be changed in order to get a satisfactory TP rate while still maintaining the FP rate. More discussion about setting values of A and L after tuning T is presented in Appendix D Setting Activation Threshold and Life Span Values in the DynamiCS.

6.6 Experiment Results 2: Examining only antigen subsets

In contrast to Experiment 1 in section 6.5, the detectors generated in the following experiments were presented with a subset of antigens for a number of generations so that generated detectors overfit a certain antigen cluster. This experimental setting is defined in order to test whether DynamiCS is able to learn newly emerged behaviours of self antigens and forget old behaviours that are no longer parts of the self antigen behaviour.

6.6.1 Varying the Generation Numbers to Provide Antigens from a Same Cluster

In order to let generated detectors overfit a specific antigen cluster, a large value was given for N , the number of generations that antigens are selected from a same cluster. The parameters used for the experiments of this section are shown in table 6.10.

Figure 6.13.

Figure 6.13 TP and FP rates when N varies and $T = 30, A = 100, L = 10$

Parameters	Values
Tolerisation Period (T)	30
Life Span of Mature Detectors (L)	10
Activation Threshold of Mature Detectors (A)	100
Number of Generations that Antigens are Selected from a Same Cluster (N)	{5, 10, 20, 30}

Table 6.10 Parameter values used for experiments performed in 6.6.1 Varying the Generation Numbers to Provide Antigens from a Same Cluster.

The values of N employed for the second series of experiments ranged from 5 to 30. As N increases, the system will overfit the distribution of only one antigen cluster for N generations. Figure 6.13 shows the results of four different experiments when four different values were given to N and the other three parameters: T , A and L have fixed values. In figure 6.13, the grid lines on the X axis were placed at every 100 generations for $N = 5$ and $N = 10$ but they were given at every N generations for $N = 20$ and $N = 30$. As expected, the clear feature of these results is that large N leads the system to show three different converged TP and FP rates according to the given antigen cluster. For instance, in the case of the largest, $N = 30$, three converged TP rates repeatedly appear at every 30 generations and so do FP rates. This outcome does not look different from the graphs observed in the previous section 6.5 Experiment Results 1: Examining complete antigen data. From the experimental results obtained in section 6.5, the TP and FP rates also oscillated across three stable values meaning that current detectors produced three converged detection results for the three antigen clusters. However, subtle differences between the results of these two series of experiments can be found. Particularly, the results in section 6.5 with a relatively large T value (see figure 6.7 and figure 6.8) always show nearly perfect FP rates of zero. However, the FP rates seen in figure 6.13 start increasing when $N = 20$ and $N = 30$ even when a relatively large T value, 30, is given. These are not surprising results. This is because although T was large enough for detectors to activate only when they experience three antigen clusters evenly (because $N = 1$), this is no longer true when $N = 20$ and 30. For instance, when $N = 30$, mature detectors were generated by experiencing only one particular cluster or a maximum of two clusters. Then these detectors increased their match counts by matching antigens belonging to a different cluster which was not used for negative selection. Therefore, new memory detectors, which were generated as the results of activation, never had sufficient self-tolerance and easily made errors in detecting self-antigens, although they can increase TP rates by detecting small niches in each cluster. Thus, the increase of TP by new detector activation can cause an increase of FP at the same time.

Thus, the clarified question raised by this result is how the AIS can handle this kind of situation in reality, when N is not known in advance and thus it is difficult to select appropriate T . The

simplest method to resolve this problem is obviously to give a very large value to T . However, when N is also very large, a further question is immediately raised: how large a value of T is feasible. Since T is the period during which the system continues generating immature detectors, large T can cause waste of resources with no great detection effect. This is because immature detectors are not used for detection. Furthermore they limit the generation of mature detectors that are actually used. Therefore, blindly giving a very large value to T is not such a good solution for a real situation. Nevertheless, the original purpose of the activation threshold and life span is to compensate for the problem caused by insufficiently high T . Increasing A and decreasing L will help to reduce FP rates as has been already seen in sections 6.5.2 and 6.5.3. However, there still remains a question: is there any automated way to control an appropriate combination of the three parameters, depending on given antigens?

DynamiCS currently treats memory detectors generated from any activation equally. As long as detectors get activated, they are considered to be useful in detecting any non-self antigen occurring in the future. This approach worked well in the experiments of section 6.5 because there was no big difference between the distribution of antigens appearing during an initial activation procedure and the distribution of antigens being presented later. Thus, memory detectors generated as the consequence of activation are useful at any time in the future. In contrast, if the converged antigen distributions alter after a certain period and thus previously generated memory detectors do not reflect new self and non-self antigen distributions any more, these memory detectors have to be replaced as well. Therefore, another way of handling this kind of situation can be to monitor further detection results of generated memory detectors and use their detection results as feedback input to generate new immature detectors, or maintaining current memory detectors. This method might provide a better way to define an appropriate combination of the three parameters. DynamiCS could be extended by adding these features.

6.7 Summary

In order for the AIS to be able to deal with a real environment, whose self behaviours change after a certain period and shows only a small subset of self antigens at one time, the dynamic clonal selection algorithm (DynamiCS) was introduced. The significant features that allow the human immune system to provide these desired properties were introduced. They are central tolerisation, distributed tolerisation, costimulation, affinity maturation, life span and memory detectors. DynamiCS implemented these features by introducing three important parameters: tolerisation period, activation threshold and life span.

Two sets of experiments were performed in order to examine system behaviours under various values of the three parameters. The first series of experiments tested whether the system can incrementally learn the globally converged distributions when only its one subset distribution is

given at each generation. The experimental results showed that the AIS was able to incrementally learn the globally converged distributions when only one small subset of antigens was given at each generation. However, this was achieved only when appropriate combinations of the three parameter values were given. It was revealed that the system performance measured by TP and FP rates was primarily controlled by the number of detector activations in total, and that this number was directed by values of the three parameters. This is because DynamiCS generated initial immature detectors only through negative selection, and as a result the degree of antigen detection did not vary greatly between detectors. In addition, it was observed that a larger tolerisation period caused the system to produce a smaller number of detectors to activate and resulted in lower TP and FP rates. In contrast, different values of activation threshold did not affect the number of detectors activating, but a larger activation threshold made mature detectors activate much less frequently, and ultimately triggered a smaller number of detector activations, also bringing about lower TP and FP rates. In addition, larger values for the life span of mature detectors forced mature detectors to live longer and experience more antigens. This led the system to provoke more detector activation, which produced higher TP and FP rates.

In order to see different effects of parameter values depending in different scenarios, the second set of experiments simulated a situation in which converged behaviours learned in an incremental way are suddenly altered due to legal self change. The experimental results showed that large T values that were sufficient to show perfect FP rates in previous experiments no longer demonstrated perfect FP rates. This was because generated memory detectors had never been exposed to a certain antigen cluster and thus they could not have perfect self-tolerance. This reason drives further extension of DynamiCS, so that it can handle generated memory detectors based on their detection results. The modified dynamic clonal selection algorithm that employs this idea is studied in the next chapter.