Towards an Artificial Immune System for Network Intrusion Detection: An Investigation of Dynamic Clonal Selection

Jungwon Kim
Department of Computer Science
King's College London
Strand, London, WC2R 2LS, UK
jungwon@dcs.kcl.ac.uk

Peter J. Bentley
Department of Computer Science
University College London
Gower St, London, WC1E 6BT, UK
P.Bentley@cs.ucl.ac.uk

Abstract: One significant feature of artificial immune systems is their ability to adapt to continuously changing environments, dynamically learning the fluid patterns of 'self' and predicting new patterns of 'non-self'. This paper introduces and investigates the behaviour of *dynamiCS*, a dynamic clonal selection algorithm, designed to have such properties of self-adaptation. The effects of three important system parameters: tolerisation period, activation threshold, and life span are explored. The abilities of dynamiCS to perform incremental learning on converged data, and to adapt to novel data are also demonstrated.

I INTRODUCTION

In this paper, we continue our effort to understand the role of important components of artificial immune systems especially for providing an appropriate artificial immune response against network intrusions. Following our previous work investigating two different evolutionary stages of AIS: negative selection (Kim and Bentley, 2001a) and static clonal selection (Kim and Bentley, 2001b), this paper focuses on the investigation of what we term here *dynamic* clonal selection.

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 Secondly its detectors should be replaced whenever previously observed normal behaviours no longer represent current normal behaviours.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 paper describes a *dynamic clonal selection* algorithm that has the above two properties.

II DYNAMIC NON-SELF ANTIGEN DETECTION BY HUMAN IMMUNE SYSTEMS

As our cells are constantly being made and dying, new immune cells are also endlessly being made and dying. The birth and death of immune cells is 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.

III DYNAMIC ANOMALY DETECTION BY AIS

Hofmeyr (1999) extended a negative selection algorithm in order to allow it to adapt to a continuously changing environment. Hofmeyr's 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, his extended AIS creates new detectors every day after the system experiences new network traffic which has not been presented before. The replacement of detectors is achieved by introducing new parameters to determine the usefulness of detectors in order to detect current non-self antigens.

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 the tests introduced in the system. These 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.

IV DYNAMIC CLONAL SELECTION (DynamiCS) ALGORITHM

The new AIS introduced in this paper follows the basic concept of the AIS proposed by Hofmeyr (1999). 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, we introduce an artificial immune algorithm, called the *dynamic clonal selection algorithm* (DynamiCS). Although Hofmeyr proposed various new features in order to effect great adaptability and distributed detection, DynamiCS attempts to

distill only the crucial components that yield adaptability to the system (and reduce the number of system parameters to ensure the algorithm is usable). The following pseudo code provides an overview of DynamiCS.

```
Initialise Dynamic Clonal Selection Algorithm
   Create an initial immature detector population with random detectors;

Generation_Number = 1;
Do
{
   If (Generation_Number = N)
      Select a new antigen cluster.

   Select 80% of self and non-self antigens from a chosen antigen cluster;

   Reset Parameters
        Generation_Number++;
        Memory Detector Age++;
        Mature Detector Age++;
        Immature Detector Age++;

   Monitor Antigens
   {
      Monitor Antigens by Memory Detectors
         Check whether any memory detector detects any non-self antigen;
         Check whether any memory detector detects any self antigen;

      Monitor Antigens by Mature Detectors
         Check whether any mature detector detects any non-self antigen;
         Check whether any mature detector detects any self antigen;
         Create new memory detectors;
         Old mature detectors are killed;

       Monitor Antigens by Immature Detectors
         Check whether any immature detector detects any self antigen;
         Delete any immature detector matching any self antigen;
         Create new mature detectors;
   }

   If (immature detector population size + mature detector population size
      < non-memory detector pop size)
   {
     Do
     {
       Generate a random detector;
       Add a random detector to an immature detector population;
     } Until  (immature detector population size +
                mature detector population size =
                non-memory detector pop size);
   }
} While (generation Number < max Generation)
```

DynamiCS Overview

DynamiCS starts by seeding initial immature detectors with random genotypes. DynamiCS then employs negative selection by comparing immature detectors to the given antigen set. As the result, immature detectors that bind to any antigens are deleted from the immature detector population and new immature detectors are generated until the number of immature detectors becomes the maximum size of the non-memory detector population. These same processes continue for the tolerisation period (T) number of generations. When the total number of generations reaches T, those immature detectors whose age reaches T (born at generation 1), become mature detectors.

At generation $T + 1$, a new antigen set is presented to the mature detectors to be monitored. Whenever a mature detector matches an antigen, the match count of a mature detector increases by one. After all the given antigens have been compared to all the existing mature detectors, the system checks: 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). 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, if the ages of mature detectors meet L, those mature detectors are deleted from the mature detector population.

At generation $T + 2$, when memory detectors match any antigen, confirmation is sought immediately from a human security officer. In this case, if the detected antigen patterns are confirmed as intrusion signatures, 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). After the antigens have been monitored by the mature detectors, they are passed to immature detectors to perform negative selection. From generation $T + 3$ onwards, the same monitoring procedures that operated at generation $T + 2$ continue in order to monitor constantly changing antigen sets until the system terminates.

V DYNAMIC CLONAL SELECTION EXPERIMENTS

Objective

As introduced above, 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.

Data and Parameter Setting

The experiments performed for this paper used the Wisconsin breast cancer data set that was employed for the study of the static clonal selection algorithm in (Kim and Bentley, 2001b). 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.

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 systems' new behaviours, it was necessary to understand the degree of differences between various distributions of antigen sets in advance. Therefore, we adopted a following method for providing antigen data to DynamiCS.

In order to be sure of providing antigens of novel distributions, self and non-self antigen data was clustered into several groups and antigen data randomly selected from one cluster was presented for N generations. The Expectation Maximization (EM) clustering algorithm (Mitchell, 1997) was applied to cluster antigen data into three groups. 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.

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. 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.

All experiments were run for 2000 generations and repeated five times. A non-memory detector population size of 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 of a mature detector.

Experiment Design

Two series of experiments were performed by varying the distributions of the provided antigen data. The first series of experiments was carried out by providing DynamiCS with antigen data of a different distribution at every generation (i.e., the value of N was 1, ensuring that the system was able to experience the complete antigen data set). 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. 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.

VI EXPERIMENT RESULTS 1: EXAMINATION OF COMPLETE ANTIGEN DATA

Effect of the Tolerisation Period

Figure 1 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. 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 values decrease as T increases. The effects are illustrated in figure 1. Another key result revealed by figure 1 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 is 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. Although 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, FP can be dramatically reduced to an almost perfect 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. Since 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 table 1.

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

A=5	Total No. of Memory Detectors	Av. Mature Detector Pop. Size per Generation	Av. Immature Detector Pop. Size per Generation	Av. No. of New Mature Detector 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 1 Proportion of Three Different Types of Detectors when T varies and A = 5, L = 10, N =1. The values in parentheses are variances

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.

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 1. This can also be explained with the same reason. The 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. 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.

Effect of Activation Threshold

The second series of experiments were carried out with tolerisation period (T) equal to 5 and Life Span (L) equal to 10 with four different activation thresholds (A): {5, 10, 20, 50}. The figures 2 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 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 2. 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. As described in the previous section, when detectors are mainly generated by a negative selection without a niching mechanism, higher TP rates are expected when mature detectors can detect diverse niches existing in a non-self antigen set. Thus, the smaller amount of antigen detection detects only a subset of non-self antigen niches randomly scattered and induces lower TP and FP rates.

	Total No. of Memory Detectors	Av. Mature Detector Pop. Size per gen.	Av.Immature Detector Pop. Size per generation	Av. No. of New Mature Detector 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=50	14 (6)	149.13 (0.079)	90.87 (0.079)	14.96 (0.001)

Table 2 Proportion of Three Different Types of Detector when A varies and $T = 5$, $L = 10$, $N = 1$. The values in parentheses are variances.

Effect of Life Span

The third series of experiments were executed by varying the life span (L) of mature detectors from 5 to 10, 20 and 50, with the tolerisation period (T) fixed at 5 and the activation threshold (A) set at 150. Figure 3 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 after 2000 generations.

	Total No.of Memory Detectors	Av. Mature Detector Pop. Size per gen.	Av. Immature Det. Pop. Size / gen.	Av. No. of New Mature Detectors / gen.	Av. No. of Mature Detector deleted / gen.
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 3 Proportion of Three Different Types of Detector when L varies and $T = 5$, $A = 150$, $N = 1$. The values in parentheses are variances.

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 3. 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 discussed in the previous sections. Large L allows a mature detector to remain longer and thus lets it to experience more diverse non-self antigen clusters, more evenly. 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 cluster better and also learn different distributions more evenly.

Analysis

The above experiments showed clearly that the three parameters investigated in this paper 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.

The results suggest that lowering A and increasing L should be considered together in order to get an optimal result from DynamiCS. The appropriate decisions about lowering A or increasing L, or both, will be different according to the given environment. For instance, if we know that the distribution of an antigen subset presented at each generation will appear again in a near future, lowering A can be a good idea that can boost TP rates by detecting small niches. However, if any situation shows that the distribution of antigen subsets presented over time changes substantially, lowering A and keeping memory detectors cannot be such a good idea. Likewise, increasing T and L can also be equally bad for similar reasons. Since larger T and L imply keeping a larger number of immature and mature detectors that are not qualified to activate yet, increasing T and L can be an impractical idea, although they can reduce FP rate and generate more general and efficient detectors. The artificial scenario created for the experiments in this section follows the former case. We shall see that choosing arbitrarily large values for T can be a bad idea for the latter case in the next section.

VII EXPERIMENT RESULTS 2: EXAMINATION OF INCOMPLETE ANTIGEN DATA

In contrast to the first experiment results, the detectors generated in the following experiments were presented with only 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.

Varying the Generation Numbers to Provide Antigens from a Same Cluster

In order to let generated detectors to 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 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 4 shows the results of four different experiments when four different values, {5, 10, 20, 30} were given to N. The other three parameters: T, A and L were set to 30, 100 and 10 respectively for these experiments.

However, subtle differences between the results of these two series of experiments can be found, as the overall TP and FP increases as N grows. Particularly, the previous results with a relatively large T value (see figure 1) always show nearly perfect FP rates of zero. However, the FP rates seen in figure 4 start increasing when $N = 20$ and $N = 30$. These are not surprising results. This is because although $T = 30$ 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.

VIII CONCLUSIONS

The dynamic clonal selection algorithm (DynamiCS) was introduced in this paper as a step towards an artificial immune system that is better able to deal with a real environment where self behaviours change after a certain period and only a small subset of self antigens is visible at one time. The significant features that allow the human immune system to provide these desired properties were identified. 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. 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.

A large tolerisation period directly lowered FP by allowing more immature detectors to remain and pushing mature detectors out. It was also found that both lowering the activation threshold and increasing life span could guide the system to attain a higher TP rate. From analysis, lowering A and increasing L should be considered together in order to obtain an effective application of DynamiCS. The appropriate

decision about lowering A or increasing L, or both, will be different according to a given environment.

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 memory detectors had never been exposed to a certain antigen cluster and thus they could not have perfect self-tolerance. This reason drives a further extension of DynamiCS, so that it can handle memory detectors based on their detection results. The modified dynamic clonal selection algorithm that employs this idea is currently under investigation.

References

Hofmeyr, S., (1999) *An Immunological Model of Distributed Detection and Its Application to Computer Security*, PhD Thesis, Dept of Computer Science, University of New Mexico.

Kim, J. and Bentley, P., (1999), "The Artificial Immune Model for Network Intrusion Detection, *7th European Conference on Intelligent Techniques and Soft Computing* (*EUFIT'99), Aachen, Germany*.

Kim, J. and Bentley, P. J. (2001a), "Evaluating Negative Selection in an Artificial Immune System for Network Intrusion Detection", *Genetic and Evolutionary Computation Conference 2001 (GECCO-2001), San Francisco,* pp.1330 - 1337, July 7-11.

Kim, J. and Bentley, P. (2001b), "The Artificial Immune System for Network Intrusion Detection: An Investigation of Clonal Selection with a Negative Selection Operator", *the Congress on Evolutionary Computation (CEC-2001),* Seoul, Korea, pp.1244-1252, May 27-30.

Mitchell, T. (1997), *Machine Learning*, McGraw-Hill.

Paul, W. E., (1993), "The Immune System: An Introduction", in *Fundamental Immunology* 3rd Ed., W. E. Paul (Ed), Raven Press Ltd.

Sompayrac, L., (1999), *How the Immune System Works*, Blackwell Sicence.

Tizard, I. R., (1995), *Immunology: Introduction*, 4th Ed, Saunders College Publishing.

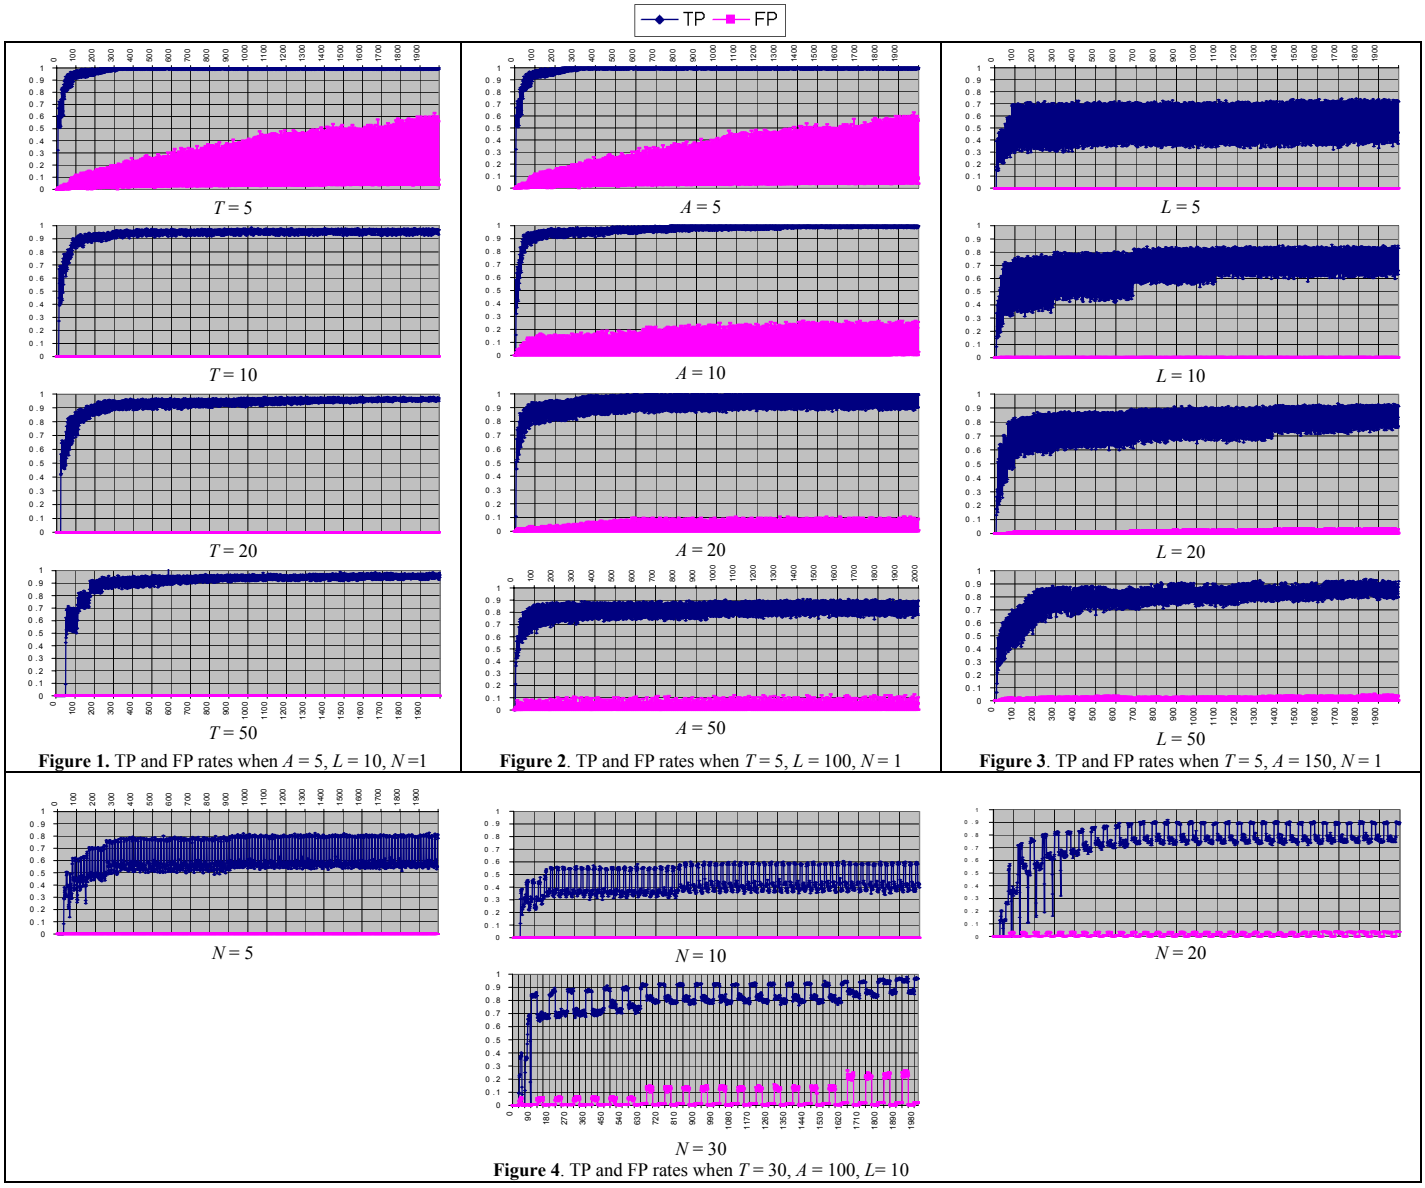

Figure 1. TP and FP rates when $A = 5$, $L = 10$, $N = 1$

Figure 2. TP and FP rates when $T = 5$, $L = 100$, $N = 1$

Figure 3. TP and FP rates when $T = 5$, $A = 150$, $N = 1$

Figure 4. TP and FP rates when $T = 30$, $A = 100$, $L = 10$