International Journal of Medical, Medicine and Health Sciences
ISSN: 2517-9969
Vol:4, No:5, 2010

Negative Selection as a Means of Discovering Unknown Temporal Patterns

Wanli Ma, Dat Tran, and Dharmendra Sharma

***Abstract*—**The temporal nature of negative selection is an under exploited area. In a negative selection system, newly generated antibodies go through a maturing phase, and the survivors of the phase then wait to be activated by the incoming antigens after certain number of matches. These without having enough matches will age and die, while these with enough matches (i.e., being activated) will become active detectors. A currently active detector may also age and die if it cannot find any match in a pre-defined (lengthy) period of time. Therefore, what matters in a negative selection system is the dynamics of the involved parties in the current time window, not the whole time duration, which may be up to eternity. This property has the potential to define the uniqueness of negative selection in comparison with the other approaches. On the other hand, a negative selection system is only trained with "normal" data samples. It has to learn and discover unknown "abnormal" data patterns on the fly by itself. Consequently, it is more appreciate to utilize negation selection as a system for pattern discovery and recognition rather than just pattern recognition. In this paper, we study the potential of using negative selection in discovering unknown temporal patterns.

***Keywords*—**Artificial Immune Systems, Computational Intelligence, Negative Selection, Pattern Discovery.

I. Introduction

NEGATIVE selection is a branch of Artificial Immune Systems (AIS), which are inspired by the observation of the behaviors and the interaction of antibodies and antigens in a biological system [1-3]. Negative selection mimics the way a human body detects and destroys harmful antigens. It has found a wide range of applications, especially in abnormal detection [4-12].

However, despite of many successful reports, on almost every reflection on the achievements of AIS, including negative selection, the questions of the uniqueness and usefulness of AIS have always been asked, yet not fully answered. In Hart and Timmis [13], the question of "*was it worth it*" was discussed. In [14], one of the challenges put forward by Timmis is to find suitable applications ("*killer application*", as described) which are unique to AIS. In [15], Garrett tried to answer "*is AIS a useful paradigm*", with the conclusion that "*It seems extremely like that AIS will become increasingly useful over next few years*". Again, in [16], Dasgupta suggested that "*the uniqueness and usability of each model* (of AIS) *needs to be determined"*.

Negative selection has very strong temporal nature, yet not enough attention has been paid to fully utilize it. In a negative selection system, a detector starts with the maturing phase for a period of time and then waits to be activated by a certain number of successful matches against incoming antigens. If it cannot find enough matches within a period of time, it ages and dies. The incoming antigens is just a sequence of incoming data items, each of which is tested by matching against the detectors. The matching serves 2 purposes. If there is a successful match against an active detector, the incoming data item is regarded as a harmful non-self cell. If there is a successful match against a yet-to-be-activated detector, the matching count of that detector is increased by 1. The detector will become a fully activated one when the matching count exceeds the pre-defined antibody activation threshold. Therefore, the order of the antigens in the sequence, i.e., the order in a timeline, is critical to the detection performance. Changing the order results a completely different set of dynamics of detector's activation and ageing and, therefore, a different performance.

On the other hand, a negative selection process is trained only with "normal" data samples (self cells), and it learns, discovers, and detects abnormal patterns (non-self cells) on the fly by itself. Negative selection is neither a fully supervised learning process nor a completely unsupervised one. It can be regarded as a semi-supervised learning process. There are many real life applications where abnormal data samples are not available, hard to obtain, not accurate, or impossible for human beings to identify and label them as abnormal. Spam email is a perfect example. While the research community is improving the effectiveness of different mathematical models, for example, Bayesian classifiers, support vector machines, and neural networks etc., which all rely on the knowledge of previous spam emails, to detect spam emails, the spammers are also arduously inventing new techniques to hoodwink these models. They creatively construct new strains of spam emails, which could be very different from these in the past. Hence, no knowledge about them is available. In other words, the abnormal data samples are not actually available. The other possible application areas include the detection of faulty parts, credit fraud, and money laundry etc.

In this paper, we study the potential of using negative selection to discover unknown temporal patterns. We

Wanli Ma is with the Faculty of Information Sciences and Engineering, University of Canberra, Australia, Australia (phone: +61-2-62012838; fax: +61.2.62015231; e-mail: Wanli.Ma@canberra.edu.au).

Dat Tran is with the Faculty of Information Sciences and Engineering, University of Canberra, Australia, Australia (e-mail: Dat.Tran@canberra.edu.au).

Dharmendra Sharma is with the Faculty of Information Sciences and Engineering, University of Canberra, Australia, Australia (e-mail: DharmendraSharma@Canberra.edu.au).

experiment the idea on 2 different types of datasets. One dataset is on 2D synthetic geometry data. With the dataset, the issues of distance calculation, data value normalisation, and the accuracy of pattern occurrences etc. can be easily and also accurately defined, and we can thus concentrate mainly on the patterns and pattern discovery. The other dataset is TREC07 spam email dataset. We primarily focus on the spam email temporal patterns in this paper.

The rest of the paper is as follows. In Section II, we briefly review negative selection. Section III studies the temporal nature of negative selection. We report our experiments on the 2 different datasets in Section IV and V respectively. We summarize the paper with future work in Section VI.

II. NEGATIVE SELECTION

Artificial Immune Systems (AIS) are based on the observation of the behaviors and the interaction of antibodies and antigens in a biological system. Negative selection, which is a branch of AIS, mimics the way a human body detects and destroys harmful antigens. A human body constantly produces lymphocytes, with randomly mutated surface peptides, from bone marrow. All newly generated lymphocytes are sent to the thymus to mature. The thymus has almost all types and shapes of self cells. During this period of maturing time, if a lymphocyte matches any cell in the thymus, the lymphocyte is just a copy of a self cell and is then destroyed. Only these which do not match any self cell in the thymus are sent to the body to match (or detect) antigens, which are invasion cells. The lymphocytes keep matching all the cells in the body. If a match happens, it means that a non-self cell (an antigen) is just detected. An alarm might be raised, and immune reactions may follow. The lymphocyte which matches the antigen may become a memory lymphocyte and stays in the body to quickly respond to the same antigen in the future. If for a period of time, a lymphocyte does not make any match, it will age and die. For a detailed explanation on how the immune system works, under the context of AIS, we refer the readers to [3, Chapter 2]. Negative selection, due to its ability of discriminating self and non-self, fits naturally into the area of anomaly detection. It has found wide applications, especially in abnormal detection. Ji and Dasgupta have a comprehensive survey paper [11] on the latest development of negative selection.

The terms used in AIS literature are yet to be standardized. In this paper, we use the term *antibody* and *detector* interchangeably, and we view the data to be verified, i.e., to be matched by the antibodies, as a sequence of *cells*. A cell can be either a self cell or an antigen cell. For simplicity reason, we call these cells just as antigens before they being verified, if there is no confusion based on the context.

From the mathematic model point of view, a self cell, an antibody, or an antigen is represented in a string or a vector format. If the distance (aka affinity) between 2 cells is within a pre-defined threshold, the 2 cells are regarded the same. It is also said that a match happens.

A basic negative selection algorithm has 4 concurrently working modules and 2 repositories, Fig 1. The module generating random detectors is responsible for generating (random) detectors. The newly generated detectors are sent to another module, detectors maturing module, to mature, where a newly generated detector is matched against the self cells from the selves repository and is culled if there is a match. For these detectors which survive the maturing module, they are kept in the detectors repository. The antigen matching module keeps matching detectors from the detectors repository against the incoming antigens. If a match happens, depending on the matching count of the detector, one of 2 actions will follow. If the count is still under the threshold, the matching count of the detector is increased by 1, and the incoming antigen is not regarded as being detected. If the count is over the activation threshold, the detector is an activated one, which means that the incoming antigen is detected. Finally, the detectors ageing module checks the ages of all detectors. If a detector cannot be activated within a pre-defined time window, the detector ages, dies, and is then removed from the detectors repository. The detectors ageing module then notifies the generating random detectors module to generate more detectors to replace these dead ones.

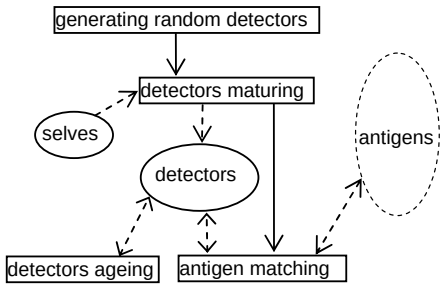

Fig. 1: A Negative Selection System

In summary, the basic issues for a negative selection algorithm are: the data representing format, the affinity calculation method, the affinity threshold, the detector activation threshold, the detector ageing threshold, the detectors generating method, and the size of the detector repository. Different settings result in different negative selection algorithms.

III. THE TEMPORAL NATURE OF NEGATIVE SELECTION

Negative selection has very strong temporal nature. A newly generated detector starts with a maturing phase, where it is matched against all self cells. If there is a match, the detector is just a copy of a self cell and cannot become an antibody to be used in the future for antigen detection purpose. If this newly generated detector survives the maturing phase, it is still not a ready-to-be-used antibody yet. It has to wait to be activated by the incoming antigens after a certain number of matches. In other words, before the

matching count exceeds the pre-defined antibody activation threshold, the newly generated detector is just a candidate with the potential to become a one. After it is fully activated with enough number of matches against the incoming antigens, the newly generated detector becomes an active one. Any subsequent match against this detector means successful detection. Finally, if a newly generated detector cannot reach the antibody activation threshold within a limited period of time (detector activation threshold), it will age and die, giving the way to other even newer detectors. Likewise, a fully activated detector may also age and die if there is no single match in a pre-defined period of time.

Therefore, time plays a significant role in negative selection. If we change any of these time related parameters, we will get completely different results. For example, the following incoming antigen sequence.

$$...\alpha_i\alpha_j\beta\alpha_k\beta\alpha_l\alpha_m\alpha_n\beta\alpha_p\alpha_q...$$

can activate the detector which matches the antigen β, if the ageing threshold is 4 and detector activation threshold is 2. To maintain the general applicability of the discussion, we use rough relative time in this paper. A time unit equals to the time required to test 1 incoming antigen against all current active detectors. After being activated, the third β antigen in the sequence can then be detected. However, if we change the order of these antigens in the sequence to

$$...\beta\alpha_i\alpha_j\alpha_k\alpha_l\beta\alpha_m\alpha_n\alpha_p\alpha_q\beta...$$

The detector which matches the antigen β may never be activated, because the detector ages and dies before being able to get a match, and the subsequent third β antigen won't be detected any more. This time, because β is so sparsely occurring in the sequence of the incoming antigens, it is beyond the limit of the interested observation window. Therefore, we do not really care if it exists or not. However, if we do care about its existence, we will have to prolong our observation window, by changing the parameter of detector ageing threshold.

From the 2 fictitious examples of antigen sequences, we can see that negative selection is sensitive to repeated patterns, i.e., temporal patterns. It can efficient learn, discover, and detect these patterns. The duration of the interested observation time window and the behaviors of detectors can be easily adjusted by setting the 2 associated thresholds: detector activation threshold and detector ageing threshold.

The ability for a negative selection system to age and retire inactive detectors is very important. In many real life applications, such as the detection of spam emails, computer viruses, and credit fraud etc., when a particular technique has been detected and studied, the detection to the technique is put into the detecting systems to detect any future offences. By this time, the perpetrators would move on to invent new techniques to avoid being detected. Very unlikely, the perpetrators keep performing the obsolete techniques, because they are not effective anymore. While on the other hand, from the detecting system's point of view, if a system never retires the devices of detecting obsolete and inactive patterns, as time goes by, the system just accumulates everything. Sooner or later, the system won't be able to cope with the load.

Negative selection, on the other hand, only keeps active detectors. If a fully activated detector cannot make any match against the incoming antigens in a period of time (detector ageing threshold - should we wish, we can set 2 different detector ageing thresholds for yet to-be-activated detectors and active detectors), it is basically not in need anymore, at least, for the time being. It will then age and die. This attribute of negative selection keeps the size of its detector repository limited within an acceptable range. Without the attribute, the size of the detector repository may increase unlimited. Eventually, it becomes infeasible to compare an incoming antigen against every possible detector, because there are so many.

To age and retire inactive detectors is also justifiable. If a detector becomes inactive, it must not contribute anything useful in detecting harmful antigens for some period of time; otherwise, it won't be a one. To keep the detector in the system only wastes the precious computational resources, both space and time. If there are a large portion of inactive detectors, the system wastes a large portion of its resources on these inactive detectors. As a consequence, it cannot detect harmful antigens in a timely manner, and the system becomes virtually useless. Ageing and retiring inactive detectors provide an appropriate solution to the problem. On the other hand, should a retired detector become in need again, the negative selection system can easily reestablish it via the antigen feedback mechanism. Understandably, there is some delay in reestablishing the detector, which may lower the system performance in 2 peculiar situations. However, either can be alleviated by adjusting the relevant system parameters. The first situation is in dealing with the antigens of very rare occurrences: the time from one occurrence to the next one is much longer than the antibody activation threshold. Therefore, no detector for these antigens can be activated. If this happens, it basically means, should there be no change to the antibody activation threshold, that due to its so limited occurrences, the system is not interested in detecting them. Otherwise, the antibody activation threshold can be changed to a smaller value, meaning shorter period of time, to allow the detectors to be activated to catch them. The other situation is in dealing with antigens of periodical occurrences, but with very long periods, which are longer than the antibody ageing threshold. In essence, the situation is the same as the previous one, but with regular occurrences, and therefore ignoring their existence may not be a choice anymore. Should this happen, there are 2 choices of solutions. The first choice, this type of antigens is still deemed as rare events and ignored. The second is to increase the antibody ageing threshold, making it longer than the periods of these antigens. As a consequence, the size of the antibody repository will increase, as it bears

more antibodies in a longer term. To the extreme, a negative selection system may never retire any antibodies. In either of these 2 peculiar situations, the decision on the values of antibody activation threshold or antibody ageing threshold, or both, is a balanced choice based on a particular application.

IV. The Experiments on the 2D synthetic dataset

The 2D synthetic dataset was created by Dipankar Dasgupta. We obtained it from Keogh [17]. In this dataset, there are 38 pairs of files which contain the data representing different 2D geometry shapes. Each pair has a file containing training data (named as ***_train.txt) and a file containing testing data (named as ***_test.txt).

Fig. 2: "Triangle small_test.txt" and "Triangle-small_train.txt"

In this paper, we use the pair "Triangle-small_test.txt" and "Triangle-small_train.txt" to illustrate our experiments on the discovery of temporal patterns.

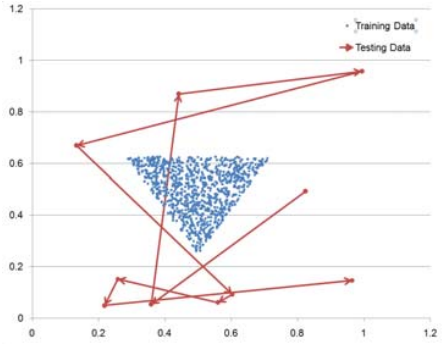

Fig. 3: The stream of incoming antigens

Each of the data items in either of the files is the coordination of a point on a 2D space of a 1×1 area. The data items in the training file form the blue triangle in Fig 2. The data which represent the points within the area are regarded as normal, and the data items in the file Triangle-small_train.txt are used to seed the negative selection experiments as self cells. The data items in the testing file are the red points scattered all over the 1×1 area, including some overlapping over the blue triangle area. The data which represent the dots away from the triangle area are regarded as abnormal. These which are overlapping on the triangle area are of course regarded as normal. Negative selection is used to detect these 2 types of data.

The original data items do not carry any time information. They are just 2D spatial data items. However, if we take these in the testing file as a stream of incoming antigens to be tested again the current active detectors, we can assign time order to these data items, Fig 3, the first 10 points in the order of an imaginary timeline. Therefore, when we see 2 or more points closing to each other, we can interpret it as the reoccurrence of the same antigens along the time line, and thus forming temporal patterns.

With the dataset, the distance between any 2 points is naturally calculated in Euclidean distance. If we set the affinity threshold at 0.16, and the antibody activation threshold at 1 match (a detector is activated after 1 successful match), we can achieve, on the original dataset, the detection rates of 100% for normal data items and 86% for abnormal data items. The detection rate on abnormal data does bear the failures of initial maturing and activating stage of each detector. If we repeat the original test data, the detection rate will go up, as the detectors are all active now. Our focus in this paper is to study the behaviors of detectors in dealing with temporal patterns. Therefore, we set the antibody activation threshold to 1 to simplify the illustration. With the help of the antigen feedback mechanism [18], the number of the detectors in use is 21. Among them, 14 detectors perform 90% of the detection.

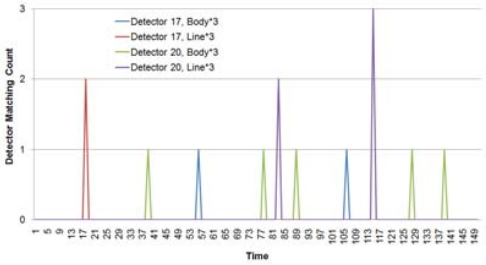

Fig. 4: The behaviors of 2 detectors in 2 different repeats of the original dataset

To create more and also predictable temporal patterns, we repeat the original testing data 3 times. We performed the repeats in 2 different ways, repeating the whole body of the testing data 3 times (Repeat A) and repeating each line 3 times on the spot (Repeat B). To clearly illustrate the discovery of temporal patterns, these detectors which have minimum matching count is of particular interest. With Repeat A, denoted as Body*3 in Fig 4., Detector 17 only made 1 match in a time window, and the behavior happened twice along the time line in the whole duration of the experiment. It indicates that in the first repeat, the detector is merely activated, and it then performed 2 matches in repeat 2 and 3. While with Repeat B, denoted Line*3 in Fig 4., after being activated, the detector matches the subsequent lines straightaway and continuously. The same observation can be

made by comparing the behavior of Detector 20: the 5 spikes of "Detector 20, Body*3" become the 2 spikes of "Detector 20, Body*3", Fig 4.

Fig. 5: The behaviors of all detectors in 2 different repeats of the original dataset

Fig. 5 demonstrates the detection activities of all detectors. From the diagram, "Body*3" curve generally has less fluctuation than "Line*3" line. To some extent, "Body*3" is just 2 more repeats of the matching activities of the original testing data, while "Line*3" reflects the concentrated patterns due to on the spot 3 repeats of each line. The wild fluctuation of the curve reflects the learning (the generating, maturing, and activating of detectors) and detecting activities of the negative selection experiment. On "Body*3" line, a detector may be activated by the first occurrence of the data item. When coming to the second and third repeats, the detector is mature and thus ready to detect similar data items. On the other hand, on "Line*3" line, after being activated, a detector continuously perform the detection on the adjacent incoming data items. The activities of a detector, maturing, being activated or ageing and dying, and detecting, happened within a very short period of time.

V. The Experiment on the TREC07 Spam email Corpus

The TREC07 Corpus dataset [19] isn't as neat as the 2D synthetic dataset. First, it is difficult to decide if an email is similar to the other one, and second the distance (how similar) between 2 emails is very hard to define. However, we still can make some interesting observations.

In this paper, we use Nilsimsa digests [20] to represent emails (The commonly cited web site for Nilsimsa http://lexx.shinn.net/cmeclax/nilsimsa.html is no longer available. In this paper, we use CPAN Perl module Digest-Nilsimsa-0.06 for the experiment. The original C implementation was credited to Chad Norwood chad@455scott.com). A Nilsimsa digest of a piece of text is a 256-bits string, which is written in a hexadecimal format as shorthand. Therefore, a self cell, a detector, or an antigen is coded as a 256-bits binary string. Nilsimsa digests have the ability to preserve the distance between 2 pieces of text. If the difference between the 2 pieces of text is big, then the difference between the 2 corresponding Nilsimsa digests is big. On the other hand, if the difference between the 2 pieces of text is small, the difference between the 2 Nilsimsa digests is small as well. Damiani et al [21] studied the robustness of Nilsimsa digests against text obfuscations, for spam email detection purpose, and demonstrated that Nilsimsa digests are resistant to text obfuscations up to 3 times bigger of the original text. Although the differences among Nilsimsa digests and the differences of their original pieces of text are not monotonically preserved in a mathematic sense, to a large extend, Nilsimsa digests are good enough for our experiments.

TREC07 Corpus contains 75,419 emails, among which 25,220 are ham, and 50,199 are spam. These emails were received by an email server between Sun, 8 Apr 2007 13:07:21 (GMT-0400) and Fri, 6 Jul 2007 07:04:53 (GMT-0400). The emails are ordered by their coming time. For our experiment purpose, we only extract the body part of an email, including the Base64 formatted attachment text. The extracted email body is then undergone the following steps of processing:

- The HTML tags are stripped off.
- The meta lines, such as Content-Type: text/plain and charset= etc. MIME attachment marks, are stripped off.
- All alphabets are converted into their lower cases.
- All blank spaces (white spaces, tabs, and new lines) are removed.

After these steps, we obtain the clean email body of each email and then calculate the Nilsimsa digest for the email based on its clean body. For this experiment purpose, we only take the first 20,000 emails. Of the 20,000 emails, there are 4,890 ham emails and 14,489 spam emails. There are also 621 blanks (620 from spam emails and 1 from ham email), which will be skipped in the experiments. We use the first 2500 ham emails (about 50% of the ham emails) as the seeds for selves. We set the affinity threshold to 80 different bits. If 2 Nilsimsa digests have less than 80 bits in difference, the 2 original emails are regarded the same; otherwise, different. The affinity threshold 80 yields the best detection rates, 94% detection rate on ham email and 78% detection rate on spam email [22]. Although the results are not as good in comparison with other approaches, they are justifiable under the assumption – no available prior knowledge about spam emails. The detection rates bear the failures at the beginning of each learning process.

In the 20,000 emails tested, from email 5,361 to 7,938, there are 2,497 spam emails (out of a total of 2,578 emails in this segment) which are exactly the same. They all have the subject line "Avis Important et Personnel" (The system does not use the information from the subject line; nor any other header lines. We list the subject line here as a convenient indicator for the readers to quickly identify these emails). The system performed remarkably well in these intervals, with 99% to 100% detection rates. Fig. 6 displays the behavior of the detector which detects this spam email. This spam exemplifies a typical spam email pattern. Upon being created, it was sent out in a burst mode. Gradually, this

type of spam is dying off, and different types of spam email will emerge. From negative selection point of view, upon receiving a large number of the same spam emails, a detector for this type of spam emails can be quickly put in place and be activated. After the type of spam emails fades away, the detector may stay in the system for a whole, to quickly respond to similar spam emails if there are any. While as the time drags on, and no this type of spam emails coming anymore, the detector will age and die. In the future, should similar spam emails come again, negative selection will learn and detect them the same way as it does this time. This process resembles the immune system in a human body. Antibodies may become weak and gone without the stimuli for a period of long time – boast dose vaccine is a solution to the problem in human immunization case.

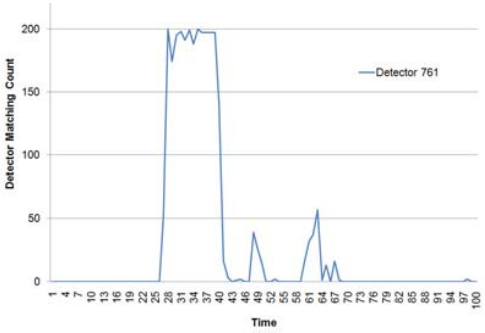

Fig. 6: The behaviors of Detector 761

Fig. 7 displays the behavior of another 2 detectors. The choices of these 2 detectors are just for the visual celerity of the diagram for the demonstration purpose.

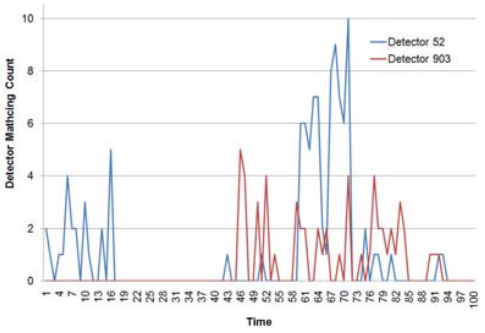

Fig. 7: The behaviors of Detector 52 and Detector 903

Finally, Fig. 8 demonstrates the behaviors of all detectors. The y axis of the diagram is the matching counts of all detectors combining together. Its value does not reflect the detection rate of the corresponding time period.

VI. SUMMARY AND FUTURE WORK

In this paper, we primarily concentrate on the temporal nature of negative selection through the behaviors of its individual detectors.

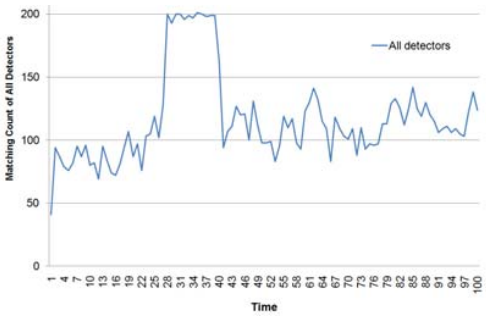

Fig. 8: The behaviors of all detectors

Being able to be generated through the antigen feedback mechanism, upon in need, a detector can be quickly put into action to detect subsequent similar incoming antigens. The maturing process will reduce the possibility of false detection and false findings. At a particular time window, a detector may find frequent matches against the incoming antigens and yield a high match count. After this period of time, it may not find a match at all. If it happens, the detector will age and die. This maintains a healthy collection of detectors. As time goes by, older detectors may just be not useful anymore. Should the same antigen comes back again in the future, the system can learn the changes and adapt itself to the new environment easily. In short, negative selection has strong temporal nature. The nature makes it a good candidate in detecting unknown temporal patterns.

We report our preliminary findings in using a negation selection system to discover unknown temporal patterns. We believe this capacity of the negation selection systems could define their uniqueness in some special applications. In the near future, we will broaden our experiments on more datasets and also compare negative selection systems with other machine learning systems. We anticipate the finding of a group of applications where negative selection can outperform other methods.

REFERENCES

[1] Hofmeyr, S.A. and S. Forrest. *Immunity by Design: An Artificial Immune System.* in *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO 1999).* 1999. Orlando, Florida, USA: Morgan Kaufmann.

[2] Dasgupta, D., Z. Ji, and F. Gonzalez. *Artificial immune system (AIS) research in the last five years*. in *The 2003 Congress on Evolutionary Computation (CEC-03)*. 2003: IEEE Press.

[3] Castro, L.N.D. and J. Timmis, *Artificial Immune Systems: A New Computational Intelligence Approach*. 2002: Springer.

[4] Hofmeyr, S., *An Immunology Model of Distributed Detection and Its Application to Computer Security*, in *Department of Computer Science*. 1999, University of New Mexico, USA.

[5] Forrest, S., A.S. Perelson, L. Allen, and R. Cherukuri. *Self-Nonself Discrimination in a Computer*. in *Proceedings of the 1994 IEEE Symposium on Research in Security and Privacy*. 1994. Oakland, CA, USA: IEEE Computer Society Press.

[6] Hofmeyr, S.A., S. Forrest, and A. Somayaji, *Intrusion Detection Using Sequences of System Calls.* Journal of Computer Security, 1998. **6**: p. 151-180.

[7] Hofmeyr, S.A. and S. Forrest, *Architecture for an Artificial Immune System.* Evolutionary Computation, 2000. **8**(4): p. 443-473.

[8] Balthrop, J., S. Forrest, and M.R. Glickman. *Revisiting LISYS: Parameters and normal behavior*. in *Proceedings of the Congress on Evolutionary Computing (CEC-2002)*. 2002.
[9] Gabrielli, N. and M. Rigodanzo. *An Artificial Immune System for Network Intrusion. Detection on a Web Server: First Results*. in *Proceedings of the 2nd Italian Workshop on Evolutionary Computation (GSICE 2006)*. 2006.
[10] Gonzalez, F.A. and D. Dasgupta, *Anomaly Detection Using Real-Valued Negative Selection.* Genetic Programming and Evolvable Machines, 2003. **4**(4): p. 383-403.
[11] Ji, Z. and D. Dasgupta, *Revisiting Negative Selection Algorithms.* Evolutionary Computation, 2007. **15**(2): p. 223-251.
[12] Dasgupta, D., K. Kumar, D. Wong, and M. Berry. *Negative Selection Algorithm for Aircraft Fault Detection*. in *Proceedings of the Third International Conference on Artificial Immune Systems (ICARIS 2004)*. 2004
[13] Hart, E. and J. Timmis. *Application Areas of AIS: The Past, The Present and The Future*. in *Proceedings of Artificial Immune Systems: 4th International Conference, ICARIS 2005*. 2005. Banff, Alberta, Canada: Springer.
[14] Timmis, J., *Artificial immune systems - today and tomorrow*. Natural Computing: an international journal, 2007. **6**(1): p. 1-18.
[15] Garrett, S.M., *How Do We Evaluate Artificial Immune Systems?* Evolutionary Computation, 2005. **13**(2): p. 145 - 177.
[16] Dasgupta, D., *Advances in artificial immune systems.* IEEE Computational Intelligence Magazine, 2006. **1**(4): p. 40 - 49.
[17] Keogh, E, General Time Series Tutorial. [cited 20 December 2009]; Available from: http://www.cs.ucr.edu/~eamonn/Keogh_Time_Series_CDrom.zip.
[18] Ma, W., D. Tran, and D. Sharma. Negative Selection with Antigen Feedback in Intrusion Detection. in 7th International Conference on Artificial Immune Systems (ICARIS 2008).
[19] Cormack, G. and T. Lynam. 2007 TREC Public Spam Corpus, http://plg.uwaterloo.ca/~gvcormac/treccorpus07/. 2007 [cited 15 Janurary 2009].
[20] Prakash, V.V. Digest-Nilsimsa, http://search.cpan.org/dist/Digest-Nilsimsa/. 2002 [cited 20 February 2009].
[21] Damiani, E., S.D.C.d. Vimercati, S. Paraboschi, and P.S. Damiani. An Open Digest-based Technique for Spam Detection. in Proc. of the 2004 International Workshop on Security in Parallel and Distributed Systems. 2004.
[22] Ma, W., D. Tran, and D. Sharma. A Novel Spam Email Detection System Based on Negative Selection, in 4th ICCIT: 2009 International Conference on Computer Sciences and Convergence Information Technology. 2009.