Can a negative selection detect unique non-self cell in an infinitely large number of self cells?

Akira Imada

Brest State Technical University
Moskowskaja 267, 224017 Brest, Republic of Belarus
akira@bsty.by

Abstract: *Artificial immune system is said to be useful when we want to distinguish anomaly patterns from normal ones. For the purpose, an analogy of negative selection of T-cells in the real biological immune system is often exploited, which results in discrimination of non-self cells from self cells. Since biological immune system can detect so far unknown non-self cells, it is natural that we think in such a way that artificial immune system detects invasion of a virus in our computer network. Furthermore, negative selection has an advantage of using only self cells in training the system. Hence a plenty of results which assert to be successful have been reported so far. We are wondering, however, what if the number of non-self cells is only a few while we have infinitely large number of self cells surrounding them, which is of often the case both in real engineering and biological problems. In this paper, we argue an recognition or classification under an extreme case where exactly only one non-self cell exists in a huge amount of self cells.*

Keywords: *artificial immune system, negative selection, network intrusion detection, detection of self and non-self cells*

1 Introduction

> *A sultan has granted a commoner a chance to marry one of his n daughters. The commoner will be presented with the daughters one at a time with being shown the daughter's pre-fixed dowry. The commoner must immediately decide whether to accept or reject her and he is not allowed to return to a previously rejected daughter. The sultan will allow the marriage only if the commoner picks the daughter with the highest dowry. — "Sultan's Dowry Problem"* [1]

The above quotation might remind us of so-called *a-needle-in-a-haystack* problem once Hinton & Nowlan [1] proposed. That is, an exploration in the search space made up of 20 binary bits where only exactly one point is assigned fitness 1, while all the others are assigned fitness 0. The search space is made up of 2^{20} points and the point which is assigned fitness 1 is like a needle in a haystack, is it not? See Figure 1 bellow.

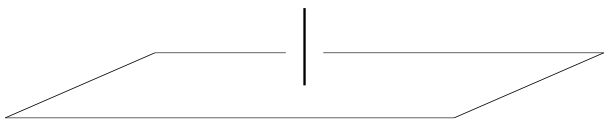

Figure 1: A fictitious sketch of fitness landscape of *a-needle-in-a-haystack*. The haystack here is drawn as a two-dimensional flat plane of fitness zero.

One thing which we have in mind is as follows. We create a huge amount of candidates of the needle and test them one by one, whether or not one candidate is the real needle we are searching for. If we

[1] According to the author(s) of the web-page of Cunningham & Cunningham, Inc. (http://c2.com) the problem was probably first stated in Martin Gardner's Mathematical Recreations column in the February 1960 issue of The Scientific American. To explore the problem more in detail, see, e.g., http://mathworld.wolfram.com. We thank Mariusz Rybnik at University Paris XII for suggesting that the problem is reminiscent of our context.

happen to find the real needle during this search procedure, then our goal is attained. This is a matter of *if-we-have-a-good-luck* or we should rather be *more-than-lucky*. Hence, in parallel with this procedure of testing those candidates of the needle one by one, we train the system using them as a set of training samples, assuming that we are more likely to fail to find the needle after all these training samples are tested exhaustively, but expecting the system will recognize the needle later as a successful result of this training procedure.

On the other hand, our implicit interest in this paper is if the needle can be detected as a non-self (anomaly) cell from enormous amount of self (normal) cells. This might be said to be an extreme case of pattern classification problem. When we talk about this issue, it is often recommended to exploit a so-called soft-computing technique such as *evolutionary computation*, *fuzzy rules*, *artificial immune system*, and so on, emphasizing among others negative selection in artificial immune system would be a good choice, especially if the context of *network intrusion detection* is under consideration. In fact, an illegal access to a network might be detected as an anomaly, and usually we know what do legal accesses look like, while we don't have any information as for illegal accesses in advance. Thus, negative selection in which the system can train itself using only the normal sample data has a great advantage, indeed.

What we are wondering about, however, is whether or not the negative selection can design detectors which can detect a very few anomaly patterns by learning only available normal samples even if the number is enormous, which is of often the case in real world problems.

In this paper we explore three different negative selection algorithms which Ayara et al. [2] picked up in their paper from a history of such approaches with the goal being to detect binary anomaly patterns among binary normal patterns.

2 Biological Immune System

> *"Immune systems evolved more than once."*
> *— Nature Magazine Web-site (7 July 2004) by Laura Nelson.* [2]

As can be paraphrased from Ayara et al. [2], The immune system is *"a natural defense mechanism which responds to foreign invaders called pathogens. The first line of defense is known as innate immunity which is a mechanism our bodies are born with If the innate immune system cannot remove the pathogen, then the adaptive (or acquired) immune system takes over. The adaptive immune system is made up of B and T-cells, which are capable of responding to certain antigenic patterns presented on the surface of pathogens. Receptors on B and T-cells match antigenic material and depending on the closeness of that match, T-cells stimulate B-cells into rapid proliferation and undergo affinity maturation by which stimulated B-cells are driven to become better tuned to the antigen responsible for initiating the immune response. During affinity maturation, stimulated antibodies undergo a somatic mutation with high rates, termed hyper-mutation. The amount of mutation that a B-cell will undergo is inversely proportional to how well it matches the antigenic pattern: the higher the affinity (match) the lower the mutation, and vice versa. Production of antibodies from these B-cells then ensues, which ultimately remove the antigenic material."*

Here we use a metaphor of B-cells which detect pathogens avoiding to match self cells of our body. In particular, binary detectors as an analogy of B-cells search for binary unique pattern as anomaly avoiding an infinite number of given binary normal patterns.

3 Algorithms

> *The strong appearance of design allows a disarmingly simple argument: if it looks, walks and quacks like a duck, then, absent compelling evidence to the contrary, we have warrant to conclude it's a duck. Design should not be overlooked simply because it's so obvious. — "Design for Living" from New York Times (7 February 2005) by Michael J. Behe.*

In this section we summarize three different negative selection algorithms for the detection of a subset of binary anomaly non-self patterns which Ayara et al. picked up in their paper [2]. The goal is to create such detectors by using negative selection mechanism.

Assuming our self and no-self cells, and thus detectors are all represented by a n-bit binary string, the simplest way to realize the idea of the negative selection is an exhaustive search.

[2] http://www.nature.com

3.1 Exhaustive Generation

Helman and Forrest (1994) proposed an Exhaustive Detector Generating Algorithm for comparison purpose for designing more effective detectors. The algorithm is as follows with the goal being to obtain D, the set of detectors.

Algorithm 1 (Exhaustive Generation) *After (i) setting D, the detector's repertoire, to be empty; and (ii) determining N_s and N_d, the total number of self samples and detectors, respectively:*

1. *Create N_s samples of self cells at random.*
2. *Create a population of candidate detectors at random;*
3. *For each of the candidate detectors, try to match the candidate detector with each of the self cells*
4. *If it matches with any self cell, then it is discarded, otherwise it is added to D*
5. *Repeat 2 - 5 until D has N_d valid detectors.*

3.2 Greedy Generation

Exhaustive Generation of Detectors in the previous subsection is not applicable in reality, if not totally impossible, as Ayara et al. [2] wrote, quoting Forrest, Perelson et al. [6] *"It was observed that the number of candidate detectors increases exponentially with the size of the self-set, at a fixed probability of not detecting non-self."*

In order to reduce the complexity, a so-called *partial matching*, or equivalently, *r-contiguous matching* based on the use of schemata proposed by Helman and Forrest [4]. In their paper, Ayara et al. [2] summarized *"The r-contiguous bits matching rule is a model of the affinity measure in the immune system. Assuming a binary representation of the self and detector strings, the r-contiguous bits matching rule compares a sequence of bits of length r in one string with a sequence of bits of the same length in the second string to see if they match."* Then the algorithm is:

Algorithm 2 (Exhaustive Generation) *In Step 3 in Algorithm 1 above "perfect matching" is replaced with "r-contiguous matching" with all the other procedure remaining the same.*

Assume here the 3-dimensional space where, e.g., four patterns (110), (111), (010), (011) is considered as self, while the other four patterns (000), (001), (100), (101) as non-self. Under this situation, in order to distinguish non-self from self, it is enough to test if any contiguous 2 bits in a same position match or not. Although it is easy to find a counter example such as the one Dasgupta et al. [5] showed: (110), (100), (011), (001) as self while (101), (111), (000), (010) as non-self, in which all detectors when $r = 2$ fails to detect

3.3 Negative Selection

As Ayara et al. [2] wrote *"The maturation of the T-cells is conceptually very simple. T-cells are exposed to self-proteins. If this binding activates the T-cell, then the T-cell is killed, otherwise it is allowed into the repertoire."* This is a principle of negative selection.

Ayara et al. [2] chose yet another algorithm among many other more effective algorithms than the greedy algorithm above. The bellow is one proposed by de Castro and Timmis in their book (pp. 279–281, 2002).

Algorithm 3 (Negative Selection) *The same procedures as above mentioned two algorithms described in this section except for a procedure after knowing the candidate detector matches a self or not, that is:*

- *If it matches, mutate the candidate detector, where the mutation flips some of the bits of the candidate detector that matches with the self data with the probability inversely proportional to the affinity of the closely matching self cells to a candidate detector.*

4 Experiment

> *"We haven't found any mathematics in here yet," he said, before blowing into the pipe. "But I'm sure it exists." If anyone can find the formalisms in an amorphous molten blob it must surely be he. — from the article in New York Times (15 February 2005) by M. Wertheim quoting Dr. Erik Demaine.*

Ayara et al. [2] created 8-bit binary random strings, as a set of training samples of self patterns assuming other comparable amount of 8-bit binary patterns as non-self. What is different of our experiment from the ones by Ayara et al. is the number of non-self.

Then our curiosity is how the system will behave when only one non-self versus infinitely large number of self cells. Our experiment was carried out

with a total number of N_s randomly created n-dimensional binary strings as training self samples, whilst only unique n-bit pattern was chosen as non-self pattern. Though we are not assumed to know what does the pattern look like principally, we chose the string whose entries are all zero, here.

5 Results and Discussion

> *"Gee, Bill, looks like we found the needle in the haystack." The red-shirted security officer grinned at his partner. — from Fairy Tale Ending by Joanne K. Seward.*

First of all, we show how many iterations were needed to create the detectors to be able to detect the needle, as a function of n — dimension of patterns. The result is shown in Figure 2. As Ayara et al. suggested, the complexity explodes exponentially as dimension increases.

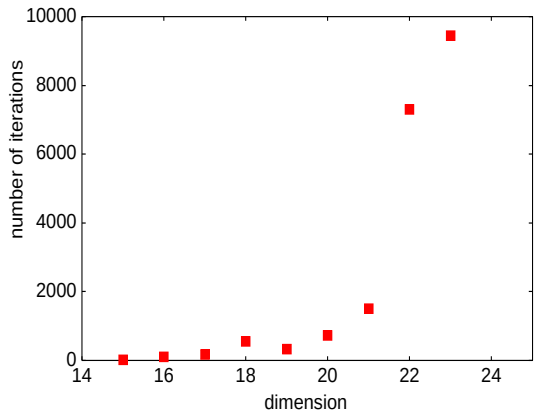

Figure 2: Number of iterations needed to find a detector to detect the needle, in this case the needle itself, as a function of dimension of the space where the needle exists.

Ayara et al. also suggested by quoting D'haeseleer, Forrest et al. [9] *"Greedy Detector Generating Algorithm ensures that generated detectors achieve as much coverage of non-self space as possible."* We observe this phenomenon including its dependency on the value r under the condition of only one non-self among a huge number of self patterns (not shown here).

Then we see how negative selection with mutation solve this difficulty or not (not shown here, either).

5.1 How good are created detectors?

We use a measure originally proposed by Lopes et al. [7] in which four quantities, i.e., (i) true-positive, (ii) true-negative, (iii) false positive, and (iv) false negative are used. Here we assume positive sample is non-self and negative sample is self, since detectors is designed to detect non-self cells. Hence, these four terms are defined in a sence that (i) t_p (true positive) — true declaration of positive sample, i.e., non-self declared as non-self (ii) f_p (false positive) — false declaration of positive sample, i.e., self declared as non-self (iii) t_n (true negative) — true declaration of negative sample, i.e., self declared as self (iv) f_n (false negative) — false declaration of negative sample, i.e., non-self declared as self. Under these definitions $d_r = t_p/(t_p + f_n)$ implies detection rate, and $f_a = f_p/(t_n + f_p)$ implies false alarm rate. Then we plot d_r versus f_a, and the resultant graph is called *Receiver Operating Characteristics (OCR)* [8] which reflects a tradeoff between false alarm rate and detection rate.

Though we now are analyzing the result using this OCR curve The fact we are facing so far is that our experiments have sometimes reversed our expectations.

6 Conclusion

In this paper we have reported a snapshot of our series of investigations into how already proposed methods work on a special situation of what we call *"a-needle-in-a-haystack"*. At this moment, we can just conclude a search with negative selection finds the needle a little more effectively than *if-we-have-a-good-luck*, and it is far from a universal way to find the needle. So far, we found the needle in 25-dimensional space but an explosion of complexity blocks us to further explore when the dimensionality becomes larger than 25. This report is not matured yet but still on-going and we try to find some real efficient and effective approaches, or at least to evoke interests in this problem. The challenge is awaiting us.

References

[1] G. E. Hinton and S. J. Nowlan (1987) *How Learning can Guide Evolution.* Complex Systems, 1, pp. 495–502.

[2] M. Ayara, J. Timmis, R. D. Lemos, L. N. D. Castro, and R. Duncan (2002) "Negative Selection: How to Generate Detectors."

Proceedings of 1st International Conference on Artificial Immune Systems pp. 89–98.

[3] Z. Ji and D. Dasgupata (2004) "Augmented Negative Selection Algorithm with Variable-Coverage Detectors." Proceedings of the Congress on Evolutionary Computation

[4] P. Helman, and S. Forrest (1994) "An Efficient Algorithm for Generating Random Antibody Strings." Technical Report 94-07, University of New Mexico, Albuquerque, NM.

[5] D. Dasugupta, S. Yu, and N. S. Majumdar (2003) MILA – Multilevel Immune Learning Algorithm GECCO 2003. LNCS 2723 pp. 183–194

[6] S. Forrest, A. S. Perelson, L. Allen, and R. Cherukuri (1994) "Self Nonself Discrimination in a Computer." Proceedings of IEEE Symposium on Research in Security and Privacy, pp. 202– 212.

[7] H. S. Lopes, M. S. Coutinho, W. C. Lima (1997) "An Evolutionary Approach to Simulate Cognitive Feedback Learning in Medical Domain." Genetic Algorithms and Fuzzy Logic Systems. World Scientific, pp. 193–207.

[8] F. Provost, T. Fawcett, and R. Kohavi (1989) "The case against accuracy estimation for comparing induction algorithms." Proceedings of international conference on machine learning, pp. 445–453.

[9] P. D'haeseleer, S. Forrest, and P. Helman (1996) "An Immunological Approach to Change Detection: Algorithms, Analysis and Implications." Proceedings of IEEE Symposium on Computer Security and Privacy, pp. 110–119.