Computer Science Applications of Artificial Immune Systems (AIS)

Patrik Sternudd

June 24, 2009

Abstract

This paper gives an introduction to how the natural immune system works, what the current theories are, and which models and algorithms that has risen from the research. It then discuss how artificial immune systems can be used for problem solving. The paper finally highlight that while there is certainly possible to solve a problem with a particular method, it might be debatable whether it is suitable to do so.

Contents

1 Introduction 3

2 The Natural Immune System 3
- 2.1 Terminology . . . 3
- 2.2 Different Theories . . . 4
  - 2.2.1 The Classical Theory . . . 4
  - 2.2.2 The Immune Network Theory . . . 4
  - 2.2.3 The Danger Theory . . . 4

3 Models and algorithms 4
- 3.1 Artificial Immune Networks (AIN) . . . 4
- 3.2 Clonal Selection . . . 5
- 3.3 Danger Theory Models . . . 5
- 3.4 Negative selection . . . 5

4 Applications of AIS - Where can it be used? 6

5 The future of AIS - Is it worth the effort? 6

6 Conclusion and discussion 7

References 8

1 Introduction

The perhaps most important function of the Natural Immune System (NIS) is to classify external organisms that are potentially harmful, and subsequently remove those before they can cause damage. The immune system is very complex, and there are competing theories as to how it actually works. From a computer science (CS) perspective, it is assumed that if the immune system can be modelled or otherwise described in a sufficiently detailed manner, these models (or insights found during the development of the models) could be used for solving computational problems not directly related to the immune system itself. Indeed, Cohen in [1], which is also referenced to by Timmis, *et al.* [2], mention two CS approaches, each trying to find solutions and methods for different problems:

- The *literal* approach

  Modelling the immune system and using it in a similar, but computer related, setting. For example, antivirus and intrusion detection (or intrusion prevention), where the main goal is to detect malicious objects and prevent those from doing harm.

- The *metaphorical* approach

  Using the immune system for inspiration to develop algorithms and models which can be used for solving various computational problems.

There is also a third reason to study Artificial Immune Systems (AIS), and that is the *computational immunology*[1] approach where the models are using for gaining new insights into how the NIS actually works.

In this paper, the focus will be on applicability of AIS for solving computational- and other engineering problems (i.e. the literal and metaphorical approaches). For the sake of completeness, a very brief and simplified introduction of the NIS is included. Readers familiar to the subject may choose to go directly to chapter 3. For a more detailed introduction of the NIS, the reader is referred to Section 3.1 in [2], and chapter 18 in [3][2]

2 The Natural Immune System

The immune system have two main components - the innate and adaptive ones, where the first is inborn and have a number of fixed functions (for example the inflammatory response and temperature control).

The second one, as can be seen by its name, is more dynamic, and is what the rest of this section will focus on. Before plunging into the different theories, some terminology must be established.

2.1 Terminology

Cells and proteins

Pathogens are external organisms that cause illness (viruses, bacteria, etc).

Antigens are chemical substances (proteins) that triggers the immune system. Note that not only pathogens have antigens; both cells belonging to self and external organisms (not necessarily pathogens) have antigens.

Antibodies are proteins that is produced by the immune system and which has the ability to match and bind to specific types of antigens.

Cytokines are "message proteins" for cell-to-cell communications.

Lymphocytes are white blood cells, which developes either into T- or B-cells.

B-cells develops either into (B) *memory cells*, or (B) *plasma cells.*

- The plasma cells, when activated, produce antibodies.

- The memory cells helps the system to quicker response by proliferation next time the same antigen is detected[3] into plasma cells.

T-cells develops either into *T-Helper-Cells (THC)* or *Natural-Killer-T-Cells (NKTC)* [4]

- The killer cells detects viral proteins in infected cells, and destroy such cells. The killer cell is also destroyed in the process.

[1] This term is not used in the cited papers, but appears to be in use elsewhere.

[2] For even more details, a book on immunology is recommended.

[3] i.e. multiplies by cell division.

[4] In [2], a different naming convention is used: T helper (T$_H$)- and T cytotoxic cells (T$_C$).

- The HTCs attach to B-cells that has already matched[5] an antigen. This cause the B-cell to go into proliferation, or alternatively (depending on the strength of the binding) suppress it from further action.

Organs

In addition to the cells, there are several organs that play an important role in the immune system:

The *bone marrow,* where all lymphocytes are created, and B-cells are matured.[6]

The *thymus*, where the T-cells mature.

The *lymph nodes*, (where immune cells may interact) and *lymph vessels* (which provides a way of transportation).

The *tonsils* and the *spleen*, where B-plasma-cells reside.

Related terms

Finally, there are a couple of related terms that are necessary to be aware of:

Affinity is the strength of the binding between an antibody and the antigen. The stronger the binding, the higher the affinity is said to be.

Negative selection is a process where any T-cells that match self-antigens are removed (killed in a controlled manner) before leaving the thymus. Having the immune system react to self would obviously be a very bad[7] idea.

Clonal selection is a process where a lymphocyte after activation starts to produce clones of itself, and where the average affinity for the antigen that triggered the cloning process increases over time. This makes the immune system much more effective against the cell with that particular antigen.

2.2 Different Theories

There is some contention between scientists who propose different theories that try to explain what factors that are responsible for certain actions. Three of those are presented below.

2.2.1 The Classical Theory

The classical view is based on the idea that the immune system operates by classifying the world into two parts, self and non-self, where the second part is unwanted and thus should be eliminated. The classification is done by B-cells together with Helper-T-Cells. The elimination is made by T-killer cells and phagocytes[8]. B-memory-cells remember earlier "attacks" on the self organism in order to respond quicker to hostile elements that has been seen before.

2.2.2 The Immune Network Theory

The Network Theory is an extension to the classical one. The difference is that B-cells are connected in a network, where a reaction (either activation or suppression) in one B-cell affects its neighbors, which in turn might affect their neighbors, etc.

2.2.3 The Danger Theory

The proponents of the Danger Theory argue that the self vs. non-self discrimination is not detailed enough; not all foreign cells should be reacted to. Instead, the immune reaction only occur when something dangerous is detected. The idea is that cells belonging to self are able to signal the manner of their death. A cell that dies a natural (apoptic) death is not given any attention[9], while one that signal unexpected (necrotic) death triggers the immune system. These signals are thought to be passed by cytokines.

3 Models and algorithms

3.1 Artificial Immune Networks (AIN)

Several different artificial immune network models are presented in the book *Computational Intelligence* by Engelbrecht [3]. One of those are aiNet [4], by de Castro and Von Zuben. As in NIS, the model is based on the assumption that B-lymphocytes are connected

[5] The matching is done by antibodies.

[6] Note that B- and T-cells have got their names from their respective place of maturation.

[7] There are immune-related illnesses where such thing happens (the term is *auto-immunity*).

[8] Phagocytes are cells that are able to devour other cells or particles.

[9] Except possibly suppression of the immune system.

in a network. In computation, this makes it possible to use such a system for classification and clustering problems. The authors of aiNet state that:

> The network approach is particularly interesting for the development of computer tools because it potentially provides a precise account of emergent properties such as learning and memory, self-tolerance, size control and diversity of cell populations.

It is likely that AINs could be used in many situations where SOFM (Kohonen maps) are used today. In the discussion, the authors also compare aiNet to ANNs and notice some similarities.

3.2 Clonal Selection

> Clonal selection in AIS is the selection of a set of ALCs[10] with the highest calculated affinity to a non-self pattern. The selected ALCs are then cloned and mutated in an attempt to have a higher binding affinity with the presented non-self pattern. [3]

One such algorithm is CLONALG by De Castro and Von Zuben [5]. CLONALG comes in two versions, depending on the problem to be solved; one for pattern recognition, the other for combinatorial or optimisation problems (e.g. the Travelling Salesman Problem). It is thus an example of the metaphorical approach to AIS. The authors do a comparison with Genetic Algorithm (GA) solutions, and claim that CLONALG gives better result for multimodal optimisation (i.e. finding multiple local or global optima). The motivation for the claim is that for non-uniform peaks, GA tend to favour *the* best candidate, while CLONALG find more of the local optima. For uniform optima, CLONALG did fewer misclassifications.

The pattern recognition version is described in [3], while both versions of the algorithm, including pseudo code is described in [5], where the main principles (for both) are said to be:

1. Maintenance of a specific memory set
2. Selection and cloning of the most stimulated Ab's[11]
3. Death of nonstimulated Ab's
4. Affinity maturation relation to antigen
5. Reselection of the clones proportionally to heir antigenic affinity, generation, and maintenance of diversity

> CLONALG is composed basically of two repertoires (populations) of strings: a set of antigens Ag and a set of antibodies Ab. The set Ab can be decomposed into several subsets according to to the application under study (pattern recognition or optimization).

3.3 Danger Theory Models

Danger theory is one of the newcomers among the theories. One example is an adaptive mailbox (described in [3]), where the goal is to separate interesting emails from uninteresting.

The concept is to use a danger signal, which is a the number of unread emails. If a the danger signal becomes too large (reaches a certain threshold), all unread emails are presented to a set of antibodies in order to classify and move uninteresting emails to a separate folder. Whenever an email is deleted, it is presented to the antibody set for re-learning (adaption)[12]. Unfortunately, the reference to the algorithm seems to be broken, so the publication date (or other details) could not be checked, but a more modern application would probably be "spam/not-spam".

Other examples for Danger Theory applications are intrusion detection systems. A paper [6] by Aickelin *et al.* give a good introduction to danger theory, and also express belief that this could be the future for AIS-based intrusion detection systems.

However, they do not present any practical research to prove this. Even so, it is very possible that they are right, because intrusion detection is a field that is very similar to detecting pathogens in nature, so any model that might work for natural immune systems should have a reasonable chance for success in an AIS-based computer system as well.

3.4 Negative selection

An algorithm for negative selection was originally proposed by Forrest *et al.* [7] 1994. The authors found inspiration from the natural immune system, specifically

[10] Artificial Lymphocyte.

[11] Ab: antibody; the authors use this term for both antibodies and B-cells.

[12] There are of course some practical aspects to this - either we must assume that the user never will delete an interesting mail, and thus would require large storage capacity, or there must be two delete functions, one for deleting uninteresting messages, and another for "housekeeping" deletions.

how the negative selection mechanism allowed T-cells to discriminate self from non-self, and thought to use this to detect changes to protected data in a computer system.

The idea of the algorithm is to, given a representation of self (e.g. a bit string), generate a set of detectors that do not match self. Then these detectors are continuously compared against self, and if a match is found, some untoward modification has been done. Note that symmetry of the algorithm - the definition of self also protect the detectors[13].

The authors primarily considered antivirus applications, but as they say:

> To date, we have only studied how the method can be applied to computer virus detection. However, we suspect that it is also applicable to a wide variety of network and operating system problems, and this is an area which we are currently investigating. Finally, our approach unifies a wide variety of computer and data security problems by treating them as the problem of distinguishing self from other.

This is a very accurate observation from an information security perspective - almost everything can be divided into a trusted self, which we want to detect changes to, and everything else. The problem is the representation of self, and especially when self changes during the lifetime (which is the case both in nature and in most computer systems). While Forrest *et al.* show successful application for antivirus systems in their experiments, Aickelin *et al.* [6] state that such solutions do not scale very well.

4 Applications of AIS - Where can it be used?

There have been several successful applications of AIS-inspired algorithms. According to Hart and Timmis[8], they can be divided into three larger areas:

1. Learning
2. Anomaly detection
3. Optimisation

These in turn can be divided into a more fine-grained list of direct applications, such as pattern recognition, robotics, computer security, etc. A search of articles relating to AIS published by the IEEE yields several practical examples. Two of these are highlighted below:

- *Applying Artificial Immune System to Minimize Construction Cost of Water Distribution Networks* [9]
- *Artificial Immune System and its Applications in GPS Single Frequency Precise Point Positioning* [10]

There are also examples where AIS is combined with other machine learning concepts such as Genetic Algorithms (GA). In the example below, the AIS is deployed in order to make the system able to adapt to new problems as efficient as possible:

- *Multiprocessor Scheduling and Rescheduling with Use of Cellular Automata and Artificial Immune System Support* [11]

There are of course also other research that is not published by IEEE which also are very interesting. To conclude this section, it appears that AIS can be used for a multitude of different problems, and with good results.

5 The future of AIS - Is it worth the effort?

In [12], Garret propose that an algorithm, to be worth considering, must be both distinctly different (from others) and effective:

> An algorithm may be distinctly different from other algorithms but ineffective, or it may be highly effective but be reducible to other, existing paradigms, and therefore lacking in distinctiveness. However, if a method is both distinct and effective, then it offers a truly useful means of computation.

However, Hart and Timmis [8] do not think this is enough for proving the usefulness of AIS (and particularly for moving it from academic research into real world applications). They reason that benchmarking is problematic in itself, and conclude that:

> [..] in a climate in which there are any number of algorithms (biological and otherwise) inspired by an equally diverse range

[13] i.e. an attacker must change both self and the detectors at the same time.

> of processes, it seems critical that an algorithm must achieve something that cannot be achieved by any other means in order to earn its place in history.

Or, in other words, even if a problem can be solved using AIS, is it worth the effort of learning yet another tool as opposed to sticking to other algorithms (machine learning or more specific ones) that works equally (or almost equally) well?

For example, several algorithms in AIS are very similar to Genetic Algorithms or ANNs. Sometimes similar concepts are derived but with different inspiration. Perhaps it is best to stick with GA, SOFM and Reinforcement Learning for the most part, and only use AIS when we must? If this is the case, then we should focus on what makes AIS unique, and look for problems that have such properties so that they can be solved by AIS methods.

This reasoning would lead us back to to the *literal approach*, and particularly to the study of computer immune systems. On the other hand, as Somayaji, Hofmeyr and Forrest write in a study on such systems [13]:

> Not only might biological solutions not be directly applicable to our computer systems, we also risk ignoring non-biological solutions that are more appropriate. A more subtle risk, however, is that through imitation we might inherit inappropriate "assumptions" of the immune system.[14]

The two main problems, according to them, is that the NIS is a natural system built by chemicals and cells, and thus have several constraints that a computer system might not have.

The other problem is that in an information security settings, our goals are quite different from the NIS. The NIS goal is to ensure that the host survives, a concept which is quite vague in information security; on the other side, the NIS has no notion of security goals such as confidentiality.

So, should we leave AIS altogether? There are several reasons why we should not[15]:

- Even if all researchers were to agree that uniqueness is required, there is no telling whether the AIS field will outgrow other machine learning techniques in the future. AIS is still pretty young, and Danger Theory in AIS is younger than other theories in the field.
- At least one of the papers cited in this work use a combination of other machine learning techniques together with AIS to achieve better performance of the overall system. In order to exploit such things, more research on AIS is necessary. It is perhaps so that AIS should be seen as an integral part of machine learning, which is combined with other parts, instead of viewing it as a competing approach. This is no more different than using an ANN in Q-learning or Sarsa.
- Even though a computer immune system cannot be directly modelled on the NIS, there are many similarities. Even if all parts not are directly applicable, they are certainly good for inspiration.

6 Conclusion and discussion

Artificial Immune Systems seems to be both an interesting and vibrant field of research. There are a lot of publications on different techniques for solving problems, both in computer security as well as more general computational or engineering problems.

From an intrusion detection standpoint, it is clear that the commercial systems available today are using old (static pattern matching) techniques that will fail to be viable in the future. The same problem exist in the antivirus market - the existing systems are reactive and centralised, and the virus definitions just increase in size for every new virus, while the old ones seldom are removed. The immune system differs, since it does not need to have antibodies for every possible pathogen; instead clonal selection develops what is needed on the fly.

Personally, I predict many possible applications in security field, although more research surely is needed. The biggest advantage with AIS (or at least NIS) as a whole are that a) it is adaptive and b) it has memory. Those factors alone gives it a big potential for anomaly detection and computer immunity. There are many promising ideas and concepts described in a multitude of papers regarding AIS.

When it comes to to solving computational problems, it appears to exist quite a future there as well. A problem is perhaps that the entry step for newcomers

[14] Just to clarify - the paper discuss how a computer immune system could be designed; they are not opposed to the concept - the quote is one of the risks they point of with such a design.

[15] It should also be noted that Hart and Timmis in no way implies that AIS is useless - their paper deal with different ways of modelling it, and is in itself a speculation on the future development of the field, just as is this section.

to the field is quite high, since the practitioner need to gain an understanding of the NIS before he or she is able to understand the motivation behind an algorithm. On the other hand, it is quite possible that the researchers will be able to create a "toolbox" with different algorithms that computer scientists might be able to use without understanding the background.

References

[1] Irun R. Cohen. Real and artificial immune systems: computing the state of the body. *Nature Reviews*, 7:569–574, July 2007.

[2] J. Timmis, P. Andrews, N. Owens, and E. Clark. An interdisciplinary perspective on artificial immune systems. *Evolutionary Intelligence*, 1(1):5–26, January 2008.

[3] Andries P. Engelbrecht. *Computational Intelligence: An Introduction.* Wiley, 2007.

[4] L. Nunes de Castro and F.J. Von Zuben. An evolutionary immune network for data clustering. In *Neural Networks, 2000. Proceedings. Sixth Brazilian Symposium on*, pages 84–89, 2000.

[5] L.N. de Castro and F.J. Von Zuben. Learning and optimization using the clonal selection principle. *Evolutionary Computation, IEEE Transactions on*, 6(3):239–251, Jun 2002.

[6] Uwe Aickelin, Peter Bentley, Steve Cayzer, Kim Jungwon, and Julie McLeod. Danger Theory: The Link between AIS and IDS? *Proceedings of the 2nd International Conference on Artificial Immune Systems*, 2787:147, 2003.

[7] S. Forrest, A.S. Perelson, L. Allen, and R. Cherukuri. Self-nonself discrimination in a computer. In *Research in Security and Privacy, 1994. Proceedings., 1994 IEEE Computer Society Symposium on*, pages 202–212, May 1994.

[8] Emma Hart and Jon Timmis. Application areas of AIS: The past, the present and the future. *Applied Soft Computing*, (8):191–201, February 2008.

[9] Min-Der Lin and Chien-Wei Chu. Applying artificial immune system to minimize construction cost of water distribution networks. In *Natural Computation, 2008. ICNC '08. Fourth International Conference on*, volume 6, pages 628–632, Oct. 2008.

[10] Chengquan Xu, Suxia Xu, and Wei Chen. Artificial immune system and its applications in gps single frequency precise point positioning. In *Intelligent System and Knowledge Engineering, 2008. ISKE 2008. 3rd International Conference on*, volume 1, pages 180–183, Nov. 2008.

[11] A. Swiecicka, F. Seredynski, and A.Y. Zomaya. Multiprocessor scheduling and rescheduling with use of cellular automata and artificial immune system support. *Parallel and Distributed Systems, IEEE Transactions on*, 17(3):253–262, March 2006.

[12] Simon M. Garrett. How do we evaluate artificial immune systems? *Evolutionary Computation*, 13(2):145–177, June 2005.

[13] Anil Somayaji, Steven Hofmeyr, and Stephanie Forrest. Principles of a computer immune system. In *In New Security Paradigms Workshop*, pages 75–82, 1997.