
The Clonal Selection Algorithm with Engineering Applications¹

Leandro Nunes de Castro

lnunes@dca.fee.unicamp.br
School of Electrical and Computer Engineering (FEEC)
State University of Campinas (UNICAMP)
Campinas-SP – Brazil

Fernando J. Von Zuben

vonzuben@dca.fee.unicamp.br
School of Electrical and Computer Engineering (FEEC)
State University of Campinas (UNICAMP)
Campinas-SP – Brazil

Abstract

The clonal selection algorithm is used by the natural immune system to define the basic features of an immune response to an antigenic stimulus. It establishes the idea that only those cells that recognize the antigens are selected to proliferate. The selected cells are subject to an affinity maturation process, which improves their affinity to the selective antigens. In this paper, we propose a powerful computational implementation of the clonal selection principle that explicitly takes into account the affinity maturation of the immune response. The algorithm is shown to be an evolutionary strategy capable of solving complex machine-learning tasks, like pattern recognition and multi-modal optimization.

1 INTRODUCTION

Over the last few years, there has been an ever increasing interest in the area of *artificial immune systems* (AIS) and their applications. Among the many works in this new field of research, we can detach those of Ishida (1996); Hunt & Cook (1996); Dasgupta (1999) and Hofmeyr & Forrest (1999). The AIS aim at using ideas gleaned from immunology in order to develop systems capable of performing different tasks in various areas of research.

In this work, we will review the clonal selection concept, together with the affinity maturation process, and demonstrate that these biological principles can lead to the development of powerful computational tools. The algorithm to be presented focus on a systemic view of the immune system and does not take into account cell-cell interactions. It is not our goal to model exactly any phenomenon, but to show that some basic immune principles can help us not only to better understand the immune system itself, but also to solve complex engineering tasks.

First, we are going to apply the clonal selection algorithm to binary character recognition to verify its ability to perform tasks such as learning and memory acquisition. Then it will be shown that the same algorithm is suitable for solving multi-modal and combinatorial optimization. This work is concluded with a brief discussion relating the proposed clonal selection algorithm with the well-known genetic algorithms introduced by Holland (1995).

2 THE CLONAL SELECTION THEORY

When an animal is exposed to an antigen, some subpopulation of its bone marrow derived cells (B lymphocytes) respond by producing antibodies (**Ab**). Each cell secretes only one kind of antibody, which is relatively specific for the antigen. By binding to these antibodies (receptors), and with a second signal from accessory cells, such as the T-helper cell, the antigen stimulates the B cell to proliferate (divide) and mature into terminal (non-dividing) antibody secreting cells, called plasma cells. The various cell divisions (mitosis) generate a clone, i.e., a set of cells that are the progeny of a single cell. While plasma cells are the most active antibody secretors, large B lymphocytes, which divide rapidly, also secrete **Ab**, albeit at a lower rate. While B cells secrete **Ab**, T cells play a central role in the regulation of the B cell response and are preminent in cell mediated immune responses.

Lymphocytes, in addition to proliferating and/or differentiating into plasma cells, can differentiate into long-lived B memory cells. Memory cells circulate through the blood, lymph and tissues, and when exposed to a second antigenic stimulus commence to differentiate into large lymphocytes capable of producing high affinity antibodies, pre-selected for the specific antigen that had stimulated the primary response. Figure 1 depicts the clonal selection principle.

The main features of the clonal selection theory, that will be explored in this paper, are (Burnet, 1978):

- generation of new random genetic changes, subsequently expressed as diverse antibody patterns by a form of accelerated somatic mutation;

¹ In Workshop Proceedings of GECCO, pp. 36-37, Workshop on Artificial Immune Systems and Their Applications, Las Vegas, USA, July 2000.




Figure 1: The clonal selection principle.

- phenotypic restriction and retention of one pattern to one differentiated cell (clone);
- proliferation and differentiation on contact of cells with antigens.

2.1 REINFORCEMENT LEARNING AND MEMORY

Learning in the immune system involves raising the population size and affinity of those lymphocytes that have proven themselves to be valuable by having recognized any antigen. While doing technology, it's one's desire to solve any kind of problem using a minimal amount of resources. Hence, we need the engineering tools to seek high quality and parsimonious solutions. In our model, we do not intend to maintain a large clone for each candidate solution, but to keep the single best individual. A clone will be temporarily created, according to the clonal selection theory, and those progeny with low affinity will be discarded.

In the normal course of the immune system evolution, an organism would be expected to encounter a given antigen repeatedly during its life time. The initial exposure to an antigen that stimulates an adaptive immune response is handled by a spectrum of small clones of B cells each producing antibody of different affinity. The effectiveness of the immune response to secondary encounters is considerably enhanced by storing some high affinity antibody producing cells from the first infection (memory cells), so as to form a large initial improved clone for subsequent encounters. Rather than 'starting from scratch' every time, such a strategy ensures that both the speed and accuracy of the immune response becomes successively greater after each infection. This scheme is intrinsic of a *reinforcement learning strategy* (Sutton &

Barto, 1998), where the system is continuously improving its capability to perform its task.

One important characteristic of the immune memory is that it is associative: B cells adapted to a certain type of antigen A_1 presents a faster and more efficient *secondary response* not only to A_1 , but also to any structurally related antigen A_2 . This phenomenon is called *immunological cross-reaction*, or *cross-reactive response* (Smith *et al.*, 1997). This kind of associative memory is part of the process of vaccination and is called *generalization capability*, or simply *generalization*, in other artificial intelligence fields, like *neural networks*.

Some authors (Allen *et al.*, 1987; Coutinho, 1989) suggested that long-lived B memory cells are disconnected, at least functionally, from the other cells.

2.2 SOMATIC HYPERMUTATION, RECEPTOR EDITING AND REPERTOIRE DIVERSITY

In a T cell dependent immune response, the repertoire of antigen-activated B cells is diversified basically by two mechanisms: *hypermutation* and *receptor editing* (Tonegawa, 1983; Berek & Ziegner, 1993; Nussenzweig, 1998; George & Gray, 1999).

Antibodies present in a memory response have, on average, a higher affinity than those of the early primary response. This phenomenon, which is restricted to T-cell dependent responses, is referred to as the *maturation of the immune response*. This maturation requires the antigen-binding sites of the antibody molecules, in the matured response, to be structurally different from those present in the primary response.

Random changes are introduced into the genes responsible for the **Ag-Ab** interactions and occasionally one such change will lead to an increase in the affinity of the antibody. It is these high-affinity variants which are then selected to enter the pool of memory cells. Not only the repertoire is diversified through a hypermutation mechanism, but also mechanisms must exist such that rare B cells with high affinity mutant receptors can be selected to dominate the response. Those cells with low affinity receptors must be efficiently eliminated, become anergic or be edited, so that they do not significantly contribute to the pool of memory cells (Berek & Ziegner, 1993; Nussenzweig, 1998; George & Gray, 1999).

Recent results suggest that the immune system practices molecular selection of receptors in addition to clonal selection of lymphocytes. Instead of the expected clonal deletion of all self-reactive cells, occasionally B lymphocytes were found that had undergone receptor editing: these B cells had deleted their low affinity receptors and developed entirely new ones through V(D)J recombination (Nussenzweig, 1998).

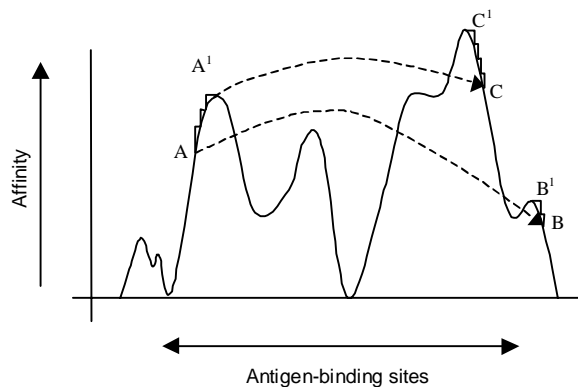


Figure 2: Schematic representation of shape-space for antigen-binding sites. Somatic mutations guide to local optima, while receptor editing introduce diversity, leading to possibly better candidate receptors.

Receptor editing offers the ability to escape from local optima on an affinity landscape. Figure 2 illustrates this idea by considering all possible antigen-binding sites depicted in the x -axis, with the most similar ones adjacent to each other. The **Ag-Ab** affinity is shown on the y -axis. If it is taken a particular antibody (A) selected during a primary response, then point mutations allow the immune system to explore local areas around A by making small steps towards an antibody with higher affinity, leading to a local optima (A^1). Because mutations with lower affinity are lost, the **Ab** can not go down the hill. Receptor editing allows an antibody to take large steps through the landscape, landing in a locale where the affinity might be lower (B). However, occasionally the leap will lead to an antibody on the side of a hill where the climbing region is more promising (C), reaching the global optimum. From this locale, point mutations can drive the **Ab** to the top of the hill (C^1). In conclusion, point mutations are good for exploring local regions, while editing may rescue immune responses stuck on unsatisfactory local optima.

In addition to somatic hypermutation and receptor editing, a fraction of newcomer cells from the bone marrow is added to the lymphocyte pool in order to maintain the diversity of the population. According to Jerne (1984), from 5-8% of the least stimulated lymphocytes is replaced per cell generation, joining the pool of available antigen recognizing cells. This may yield the same result as the process of receptor editing, i.e., a broader search for the global optimum of the antigen-binding site.

2.3 THE REGULATION OF THE HYPERMUTATION MECHANISM

A rapid accumulation of mutations is necessary for a fast maturation of the immune response, but the majority of the changes will lead to poorer or non-functional antibodies. If a cell that has just picked up a useful mutation continues to be mutated at the same rate during the next immune responses, then the accumulation of deleterious changes may cause the loss of the

advantageous mutation. The selection mechanism may provide a means by which the regulation of the hypermutation process is made dependent on receptor affinity. Cells with low affinity receptors may be further mutated and, as a rule, die if they do not become higher affinity cells. In cells with high-affinity antibody receptors however, hypermutation may be inactivated (Berek & Ziegner, 1993).

3 AN EVOLUTIONARY SYSTEM

The clonal selection functioning of the immune system can be interpreted as a remarkable microcosm of Charles Darwin's law of evolution, with the three major principles of diversity, variation and natural selection (Cziko, 1995).

The two central processes involved in the production of antibodies, genetic recombination and mutation, are the same ones responsible for the biological evolution of sexually reproducing species. In these species, the same two processes are involved in providing the variations on which natural selection can work to fit the organism to the environment (Holland, 1995). As a consequence, cumulative blind variation and natural selection, which over many millions of years resulted in the emergence of mammalian species, remain crucial in the day-by-day ceaseless battle to survival of these species.

Whereas adaptive biological evolution proceeds by cumulative natural selection among organisms, research on the immune system has now provided the first clear evidence that ontogenetic adaptive changes can be achieved by cumulative blind variation and selection within organisms. The clonal selection algorithm, to be described further in the text, aims at demonstrating that this cumulative blind variation can generate high quality solutions to complex problems.

4 THE SHAPE-SPACE MODEL

The shape-space model (S) aims at quantitatively describing the interactions among antigens and antibodies (**Ag-Ab**). The set of features that characterize a molecule is called its *generalized shape*. The **Ag-Ab** representation (binary or real-valued) determines a distance measure to be used to calculate the degree of interaction between these molecules. Mathematically, the generalized shape of a molecule (m), either an antibody or an antigen, can be represented by a set of coordinates $m = \langle m_1, m_2, \dots, m_L \rangle$, which can be regarded as a point in an L -dimensional real-valued shape-space ($m \in S^L$). Here, the precise physical meaning of each parameter is not relevant. In this work, we used binary (or integer) strings to represent the molecules. Antigens and antibodies were considered of the same length L . The length and cell representation depends upon the problem.

5 THE PROPOSED ALGORITHM

After discussing the clonal selection principle and the affinity maturation process, the development of the clonal selection algorithm (CSA) is straightforward. The main immune aspects taken into account were: maintenance of the memory cells functionally disconnected from the repertoire, selection and cloning of the most stimulated cells, death of non-stimulated cells, affinity maturation and re-selection of the clones with higher affinity, generation and maintenance of diversity, hypermutation proportional to the cell affinity.

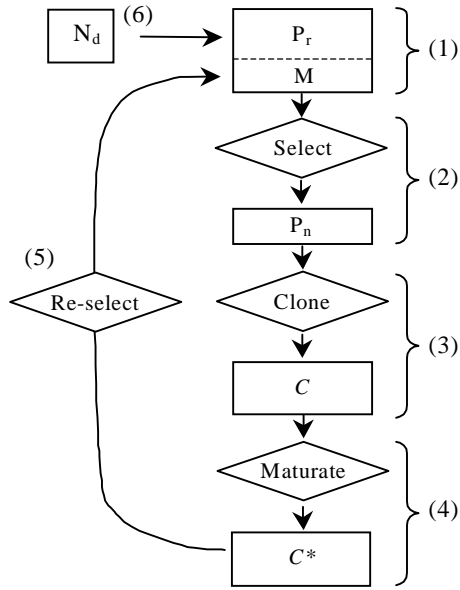


Figure 3: Block diagram of the clonal selection algorithm.

The algorithm works as in Figure 3 (after each six steps we have one cell generation):

- (1) Generate a set (P) of candidate solutions, composed of the subset of memory cells (M) added to the remaining (P_r) population ($P = P_r + M$);
- (2) Determine (Select) the n best individuals of the population (P_n), based on an affinity measure;

- (3) Reproduce (Clone) these n best individuals of the population, giving rise to a temporary population of clones (C). The clone size is an increasing function of the affinity with the antigen;
- (4) Submit the population of clones to a hypermutation scheme, where the hypermutation is proportional to the affinity of the antibody with the antigen. A matured antibody population is generated (C^*);
- (5) Re-select the improved individuals from C^* to compose the memory set M . Some members of P can be replaced by other improved members of C^* ;
- (6) Replace d antibodies by novel ones (diversity introduction). The lower affinity cells have higher probabilities of being replaced.


6 ENGINEERING APPLICATIONS

To evaluate the clonal selection algorithm (CSA), we considered three different problems:

1. a binary character recognition task will be used to test its learning and memory acquisition capabilities;
2. a multi-modal optimization task; and
3. a 30 cities instance of the Travelling Salesman Problem (TSP).

6.1 BINARY CHARACTER RECOGNITION

The learning and memory acquisition of the CSA is verified through its application to a binary character recognition problem. Our goal is to demonstrate that a cumulative blind variation together with selection can produce individuals with increasing affinities (maturation of the immune response). In this case, we assume that the antigen population is represented by a set of eight binary characters to be learned. These characters are the same ones originally proposed by Lippman (1987), in a different context. Each character is represented by a bitstring (Hamming shape-space, briefly discussed in Section 4) of length $L = 120$. The antibody repertoire is composed of 10 individuals, 8 of them in the memory set M .





(e) 200 generations

Figure 4: (a) Patterns to be learned, or input patterns (antigens). (b) Initial memory set. (c) Memory set after 50 cell generations. (d) Memory set after 100 cell generations. (e) Memory set after 200 cell generations.

The original characters (antigens) are depicted in Figure 4(a). Figure 4(b) illustrates the initial memory set, and Figures 4(b) to 4(e) represent the maturation of the memory set (immune response) through cell generations. The affinity measure takes into account the Hamming distance (D) between antigens and antibodies, according to Equation (1):

$$D = \sum_{i=1}^L \delta_i, \text{ where } \delta_i = \begin{cases} 1 & \text{if } ab_i \neq ag_i \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

Notice that an exact matching is not necessary to obtain a successful character recognition. A partial matching is enough in most applications. The algorithm converged after 250 cell generations.

6.2 MULTI-MODAL OPTIMIZATION

The CSA reproduces those individuals with higher affinities and selects their improved matured progenies. This strategy suggests that the algorithm performs a greedy search, where single members will be locally optimized (exploitation of the surrounding space), and the newcomers yield a broader exploration of the search-space. This characteristic makes the CSA very suitable for solving multi-modal optimization tasks and, as an illustration, consider the case of maximizing the function $f(x,y) = x.\sin(4\pi x) - y.\sin(4\pi y + \pi) + 1$, depicted in Figure 5. Notice that this function is composed of many local optima and a single global optimum.

We employed the Hamming shape-space, with binary strings representing real values for x and y . The chosen bitstring length was $L = 22$, corresponding to a precision of six decimal places. The variables x and y are defined over the range $[-1, 2]$, and the mapping from a binary string $m = \langle m_L, \dots, m_2, m_1 \rangle$ into a real number z is completed in two steps:

- convert the binary string $m = \langle m_L, \dots, m_2, m_1 \rangle$ from base 2 to base 10:

$$\langle m_L, \dots, m_2, m_1 \rangle_2 = \left(\sum_{i=0}^{21} m_i \cdot 2^i \right)_{10} = z'$$




Figure 5: Function to be maximized by the CSA.




Figure 6: Optimized population after 100 cell generations.

- find the corresponding real value for z :

$$z = z_{\min} + z' \cdot \frac{z_{\max} - z_{\min}}{2^{22} - 1}, \text{ where } z_{\max} = 2 \text{ and } z_{\min} = -1.$$

The affinity measure corresponds to the evaluation of the function $f(x,y)$ after decoding x and y , as described above. Figure 6 presents the optimized population after 100 generations. Notice that the solutions (stars) covers most of the peaks, including the global optimum. The solution can be directly extracted by means of a sorting procedure.

6.3 TRAVELLING SALESMAN PROBLEM (TSP)

Simply stated, the travelling salesman must visit every city in his territory, exactly once, and then return to the starting city. The question is: given the cost of travel between all pairs of cities, which is the tour with the smallest cost? In this work, the cost of the tour is basically the length of the itinerary traveled by the salesman.

The TSP is a kind of combinatorial optimization problem and arises in numerous applications, from VLSI circuit design, to fast food delivery. In this case, the use of an Integer shape-space might be appropriate, where integer-valued vectors of length L , composed of permutations of elements in the set $C = \{1, 2, \dots, L\}$, represent the possible tours. Each component of the integer vector indexes a city. The total length of each tour gives the affinity measure of the corresponding vector.

Figure 7 presents the best solution determined by the CSA, which corresponds to the global optimum (Moscato & Fontanari, 1990). The population size is 300 individuals, with a rate of 20% of newcomers. In this case, low affinity individuals are allowed to be replaced after each 20 generations. This scheduling is supposed to leave a breathing time to allow the achievement of local optima, followed by the replacement of the poorer individuals.




Figure 7: Best tour determined by the CSA after 300 generations.

7 CONCLUSIONS

In this paper, we proposed a general-purpose algorithm inspired in the clonal selection principle and affinity maturation of the immune response. The algorithm was verified to be capable of performing learning and maintenance of high quality memory and, it was also capable of solving complex problems, like multi-modal and combinatorial optimization.

The algorithm introduced constitutes a crude version of the clonal selection principle. Many heuristics could be

inserted in order to improve its performance in solving particular tasks, like the travelling salesman problem.

By comparing the proposed algorithm, called CSA, with the standard genetic algorithm (GA), we can notice that the CSA can reach a diverse set of local optima solutions, while the GA tends to polarize the whole population of individuals towards the best candidate solution. Essentially, their coding schemes and evaluation functions are not different, but their evolutionary search processes differ from the viewpoint of inspiration, vocabulary and sequence of steps. We do not advocate that the CSA performs better than the GA, on average, in all applications. Instead, we demonstrate that the proposed algorithm is also derived from a biologically inspired approach, which performs learning and multi-modal search. Like the GA, the clonal selection algorithm is highly parallel and presents a fine tractability in terms of computational cost.

Acknowledgments

Leandro Nunes de Castro would like to thank FAPESP (Proc. n. 98/11333-9) for the financial support. Fernando Von Zuben would like to thank FAPESP (Proc. n. 98/09939-6) and CNPq (Proc. n. 300910/96-7) for their financial support.

References

- Allen, D. *et al.* (1987), "Timing, Genetic Requirements and Functional Consequences of Somatic Hypermutation during B-cell Development", *Imm. Rev.*, **96**, 5-22.
- Berek, C. & Ziegner, M. (1993), "The Maturation of the Immune Response", *Imm. Today*, **14**(8): 400-402.
- Burnet, F. M. (1978), "Clonal Selection and After", In *Theoretical Immunology*, (Eds.) G. I. Bell, A. S. Perelson & G. H. Pimbley Jr., Marcel Dekker Inc., 63-85.
- Coutinho, A. (1989), "Beyond Clonal Selection and Network", *Immun. Rev.*, **110**, 63-87.
- Cziko, G. (1995), "The Immune System: Selection by the Enemy", In *Without Miracles*, G. Cziko, A Bradford Book, The MIT Press, 39-48.
- Dasgupta, D., (1999), *Artificial Immune Systems and Their Applications*, Ed., Springer-Verlag.
- George, A. J. T. & Gray, D. (1999), "Receptor Editing during Affinity Maturation", *Imm. Today*, **20**(4): 196.
- Hofmeyr S. A. & Forrest, S. (1999), "Immunity by Design: An Artificial Immune System", *Proc. of GECCO'99*, 1289-1296.
- Holland, J. H. (1995), *Adaptation in Natural and Artificial Systems*, 4th Ed., MIT Press.

- Hunt, J. E. & Cooke, D. E. (1996), "Learning Using an Artificial Immune System", *Journal of Network and Computer Applications*, **19**, 189-212.
- Ishida, Y. (1996), "The Immune System as a Self-Identification Process: A Survey and a Proposal", In *Proc. of the IMBS'96*.
- Jerne, N. K. (1984), "Idiotypic Networks and Other Preconceived Ideas", *Imm. Rev.*, **79**, 5-24.
- Lippmann, R. P. (1987), "An Introduction to Computing with Neural Nets", *IEEE ASSP Magazine*, 4-22.
- Moscato, P. & Fontanari, J. F. (1990), "Stochastic Versus Deterministic Update in Simulated Annealing", *Physics Letters A*, **146**(4): 204-208.
- Nussenzweig, M. C. (1998), "Immune Receptor Editing; Revise and Select", *Cell*, **95**, 875-878.
- Smith, D. J., Forrest, S., Hightower, R. R. & Perelson, S. A. (1997), "Deriving Shape Space Parameters from Immunological Data", *J. Theor. Biol.*, **189**, 141-150.
- Sutton, R. S. & Barto, A. G. (1998), *Reinforcement Learning: An Introduction*, A Bradford Book.
- Tonegawa, S. (1983), "Somatic Generation of Antibody Diversity", *Nature*, **302**, 575-581.