

The Coevolution of Antibodies for Concept Learning^{*}

Mitchell A. Potter¹ and Kenneth A. De Jong²

¹ Navy Center for Applied Research in Artificial Intelligence,
Naval Research Laboratory, Code 5510,
Washington, DC 20375 USA

² Computer Science Department, George Mason University,
Fairfax, VA 22030 USA

Abstract. We present a novel approach to concept learning in which a coevolutionary genetic algorithm is applied to the construction of an immune system whose antibodies can discriminate between examples and counter-examples of a given concept. This approach is more general than traditional symbolic approaches to concept learning and can be applied in situations where preclassified training examples are not necessarily available. An experimental study is described in which a coevolutionary immune system adapts itself to one of the standard machine learning data sets. The resulting immune system concept description and a description produced by a traditional symbolic concept learner are compared and contrasted.

1 Introduction

Concept learning is a task that has been extensively studied by researchers in the field of machine learning. Much of this work has been in the area of inductive learning from examples using symbolic representation languages such as predicate calculus [5] and decision trees [9]. In most of the previous efforts to apply evolutionary computation to concept learning, binary-string representations have been evolved with a genetic algorithm and mapped into some form of symbolic representation for evaluation, such as propositional logic. For some examples of this approach, see [4, 1, 3]. In the work described here, we take a different approach by experimenting with a biologically inspired representation in which concept descriptions are evolved using a model of the immune system. For other approaches to evolving models of the immune system, see the pioneering work of Forrest et al. [2].

The motivation behind applying a model of the immune system to concept learning lies in its highly developed ability to discriminate between self and non-self. In biological immune systems, this consists of the discrimination between

^{*} To appear in Proceedings of The Fifth International Conference on Parallel Problem Solving From Nature, pp. 530–539, September 1998, Amsterdam, The Netherlands. Springer-Verlag (<http://www.springer.de/comp/lncs>).

the vast array of molecules that are an integral part of the body of an organism and foreign molecules that left unchecked could result in disease or death. In the research described here, we apply the immune system’s power of discrimination to the problem of differentiating between examples and counter-examples of a given concept. An advantage of this approach over the traditional symbolic approaches to concept learning is its generality. We believe that evolved computational models of the immune system could be successfully applied to a wide variety of discrimination problems that do not necessarily lend themselves to the supervised learning methodology typically used by symbolic concept learning systems. A controller for an autonomous vehicle, for example, may need to learn to discriminate between navigable terrain and a variety of hazards based on input from a noisy sensor array. Constructing a set of preclassified training examples that adequately covers the modalities of this task would probably not be practical. An evolutionary immune system could learn the necessary concepts by adapting instead to a simple reinforcement signal that captures the ability of the autonomous vehicle to move safely through its environment.

2 A Brief Overview of the Immune System

The purpose of an organism’s immune system is to protect it against infection. This is accomplished by recognizing the molecular signature of microbes or viruses that attack its body, and once identified, eliminating the foreign molecules in a variety of ways. The immune system consists of two interrelated components: an innate defense component and an adaptive component. Here we will focus on the adaptive component, which is responsible for acquired immunity.

Molecules capable of stimulating an acquired immune response are called *antigens*. When the immune system is working properly, only foreign molecules will produce a response. There are a number of ways antigens are recognized, depending on whether the foreign molecule is inside or outside a cell boundary. It is the job of *antibodies*—protein molecules displayed on the surface of a type of white blood cell produced in the bone marrow called a *B*-lymphocyte or *B*-cell for short—to recognize antigens that are located outside a cell boundary. Recognition by a *B*-cell occurs when one of its antibodies comes into contact with an antigen of complementary shape. Although all the antibodies on an individual *B*-cell have the same three-dimensional shape, the human body, for example, has about 10 trillion of these cells and they collectively have the potential of recognizing about 100 million distinct antigens at any one time.

One should realize that the immune system is quite complex and is the focus of much current research. Although we have only provided a brief and somewhat simplistic overview of one of its processes here, this description should be sufficient for an understanding of the rest of this paper. For more details concerning the workings of the immune system, see, for example, [10].

3 Coevolving Antibodies for Concept Learning

As in previous evolutionary computation models of the vertebrate immune system (cf. [2]), our model is limited to the interaction between *B*-cells and antigens. This model is applied to concept learning from preclassified positive and negative examples by equating the positive examples to foreign molecules and the negative examples to self. Once the fitness of the immune system evolves to a point where all the foreign molecules and none of the self molecules are recognized, the antibodies represent a description of the concept³. This model can easily be generalized to learn more than one concept by simply evolving a separate family of *B*-cells for each.

The *B*-cells in our model consist of antibody and a real-valued activation threshold that represents the binding strength required to initiate an immune response. Rather than represent antibodies and antigens as three-dimensional shapes, we use a binary *schema* representation for the antibodies and represent the antigens as simple binary strings. A linear matching function that returns the percentage of matching bits in the antibodies and antigens is used to compute the binding strength. The locations at which the antibody schema contains a “don’t care” are ignored. This abstraction greatly simplifies the matching process between antibodies and antigens while still exhibiting some of the key properties of the biological entities. One important property that is captured by this representation is the ability of some antibody shapes to match a wider range of antigens than others. This enables us to model a continuum of antibodies from specialists, that can only bind to a specific antigen, to more general antibodies that can bind to whole families of antigens sharing common characteristics.

The genotype to phenotype mapping of *B*-cells in our model is shown in Figure 1. We produce an antibody schema from binary pattern and mask genes. A mask bit of one generates a schema value equal to the corresponding pattern bit, while a mask bit of zero produces a “don’t care” schema value. This many-to-one mapping is an abstraction of another property of real proteins—dissimilar chains of amino acids may fold into the same basic three-dimensional shape. The length of the pattern and mask genes depends on the complexity of the antigens the antibody must recognize. The real-valued activation threshold of the *B*-cell in the range $[0, 1]$ is produced from an 8-bit threshold gene.

If we were to evolve a population of *B*-cells with a standard genetic algorithm, the population would converge into a collection of very similar cells. However, it is important to maintain enough diversity in the *B*-cells to adequately recognize many different types of foreign molecules. Earlier evolutionary immune system models solved this dilemma through the use of a diversity preserving algorithm called *emergent fitness sharing*. We take a different approach by evolving *B*-cells with a coevolutionary genetic algorithm in which individuals from multiple non-interbreeding species collaborate to solve the target problem [8, 7]. Each species represents only a partial solution—in this case, a collection of *B*-cells

³ Given noisy examples, the immune system would be evolved until *most* of the foreign molecules and *few* of the self molecules were recognized.

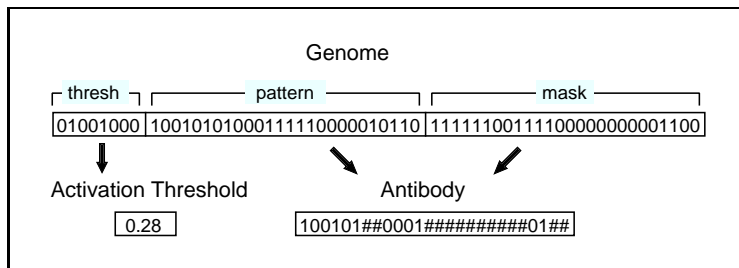


Fig. 1. Mapping from *B*-cell genome to activation threshold and antibody

having similar antibodies. The fitness of a *B*-cell is computed by adding it to a “serum” consisting of the current best *B*-cells from each of the other species in the ecosystem. Foreign and self molecules are then presented to the serum. A particular *B*-cell is considered to have recognized an antigen if the binding strength between its antibody and the antigen exceeds its activation threshold and the antigen binds to its antibody more strongly than to any other antibody in the serum. The fitness of the *B*-cell is defined to be the number of foreign molecules recognized by all the antibodies in the serum, minus the number of false-positives, that is, self molecules flagged as foreign. Therefore, each *B*-cell is rewarded based on how well it collaborates with *B*-cells from each of the other species to cover only the collection of foreign molecules. The final solution consists of the best *B*-cell from each species.

4 The AQ Approach to Concept Learning

We will compare the solutions produced by our coevolutionary immune system with those produced by AQ15, a symbolic inductive learning system developed by Michalski et al. [6]. This system is one of the latest in a series of AQ systems that constructs conjunctive descriptions from preclassified examples using an enhanced propositional calculus representation language. Each AQ concept description consists of a disjunction of conjunctive descriptions. Once a concept description has been constructed for each class of examples the system has been presented with, the system uses a conflict resolution procedure to discriminate between unclassified examples of one concept or another based on the strength of the match with the learned descriptions and the prior probability of the concepts. See [6] for more details concerning AQ conflict resolution and its method for constructing concept descriptions.

5 Experimental Study

5.1 Congressional Voting Records Data Set

In this experimental study we evolve a political party classification system for members of the U.S. House of Representatives given their voting records. The

objective, therefore, is to learn to discriminate between the concepts *Republican* and *Democrat*. This is a supervised learning task in which we are given a number of preclassified training examples. The data set from which the training examples are drawn consists of 267 Democrat and 168 Republican voting records. Each record gives the vote cast by an individual on 16 different issues. Although the actual voting records are somewhat more complex, each vote in the compiled data set has been simplified to either a yea, nay, or abstain. For compatibility with the coevolutionary immune system, the symbolic voting records were converted into 32-bit strings (antigens) using the following two-bit codes: 00 for abstain, 01 for yea, and 10 for nay. Depending on one’s political orientation, the foreign molecules to be targeted by the immune system could represent either examples of Republicans or Democrats. The symbolic data set was originally used in a machine learning study by Schlimmer [11] and was compiled from actual voting records from the 98th Congress.

5.2 Experimental Setup

AQ15 is run with its default settings. Rather than simply learning a description for one of the two concepts, say Democrats, and interpreting non-conforming examples as members of the Republican party as suggested earlier in the section on coevolving antibodies, AQ learns a separate description for each concept. For example, it will first learn a description for Republicans using the Republican instances in the training set as positive examples and the Democrat instances as negative examples, and then learn a description for Democrats using the opposite orientation.

To provide a more fair comparison, our evolutionary immune system also uses this technique by coevolving two distinct classes of *B*-cells. One class recognizes Democrats and ignores Republicans, while the other class ignores Democrats and recognizes Republicans. As species are created, half are assigned to the Democrat class and the other half to the Republican class. Each species has a population size of 100, is initialized randomly, recombined with uniform crossover at a rate of 0.6, mutated by flipping bits at a rate of twice the reciprocal of the chromosome length, and evolved using scaled fitness-proportionate selection. Since AQ is strongly biased towards learning general solutions, the immune system is also given a generality bias by initializing approximately 90 percent of the alleles of each mask gene to zero. We begin the evolution of the system with only a single species of *B*-cells. New species are created and poor performing species are eliminated when evolutionary improvement stagnates.

5.3 Results

We first look at the quality of solutions produced by the immune system and AQ in terms of how well they are able to discriminate between Republicans and Democrats. Solution quality is compared using the *predictive accuracy* metric. Given the size of our data set, the tenfold cross-validation method is the recommended procedure for computing this metric [12]. One performs tenfold cross

Table 1. Final predictive accuracy of learning methods

<i>Learning method</i>	<i>Predictive accuracy</i>
Immune System	0.964 ± 0.018
AQ System	0.956 ± 0.023

validation by randomly dividing the complete set of positive and negative examples into ten partitions of approximately equal size. Ten runs are then performed, each using a different set of nine partitions as the *training set* and the remaining partition as the *testing set*. During each run, the concept learner will use the training set to construct a concept description. Once the run is complete, the concept description is applied to the testing set and the percentage of testing instances classified correctly is computed. The predictive accuracy is computed by averaging the percentage of correct classifications produced from the ten runs.

Each of the ten AQ runs was terminated when it produced a concept description capable of correctly classifying all the instances of Republican and Democrat voting records in the training set. Each run of the immune system was terminated after 100 generations of adaptation to the training set; however, learning flattened out after only a couple of generations. The predictive accuracy results are summarized in Table 1. The table includes 95-percent confidence intervals computed from the *t*-statistic. A *t*-test was also performed on these results and it was determined that there is not a statistically significant difference between the predictive accuracy of the methods.

In Table 2 we compare the number of elements in the concept descriptions produced by the two methods, specifically, the number of *B*-cells versus the number of conjunctive descriptions required to cover the voting record training examples. Over ten runs, the immune system consistently produced smaller descriptions than AQ. As shown in the table, on average the immune system evolved 7.0 *B*-cells while AQ generated 15.1 conjunctive descriptions. The table includes 95-percent confidence intervals on the mean computed from the *t*-statistic and a *t*-test was used to verify that there is a statistically significant difference between the number of cover elements produced by the methods.

To compare and contrast the roles played by these cover elements, the results from both methods were converted into a rule-based representation using a straightforward mapping by which each antibody and conjunctive description is converted into a separate rule. Specifically, to produce a rule from an antibody, the first length-two antibody schema is mapped to a test of the first vote, the second length-two schema is mapped to a test of the second vote, and so on. Table 3 gives our schemata interpretation. The table reflects that half credit is given for partial matches. Rules are activated when their match strength exceeds a rule-specific threshold as described in Section 3. In contrast, AQ normally only activates perfectly matching rules. It will use a combination of the strength of the

Table 2. Required number of cover elements (conjunctive descriptions or antibodies)

<i>Learning Method</i>	<i>Cover Elements</i>		
	<i>Mean</i>	<i>Min</i>	<i>Max</i>
<u>Immune System</u>			
Democrat	3.80 ± 0.30	3	4
Republican	3.20 ± 0.74	1	4
Total	7.00 ± 0.67	5	8
<u>AQ System</u>			
Democrat	8.30 ± 0.68	6	9
Republican	6.80 ± 0.81	5	9
Total	15.10 ± 1.37	11	18

Table 3. Interpretation of antibody schema

<i>Schema</i>	<i>Interpretation</i>
00	abstain or yea (half credit) or nay (half credit)
01	yea or abstain (half credit)
10	nay or abstain (half credit)
11	yea (half credit) or nay (half credit)
0#	abstain or yea
1#	nay
#0	abstain or nay
#1	yea
##	ignore

match and the prior probability of the concepts to activate partially matching rules if no perfectly matching ones exist.

The rules for recognizing Democrats produced by the first run of the immune system and AQ are shown in Figures 2 and 3. To further visualize the roles played by these rules, the number of training set examples covered and classified by each is shown in Figures 4 and 5. By *covered*, we mean that the rule was activated but not actually chosen by the conflict resolution procedure.

The first observed difference between these rule sets is that, as previously noted, significantly fewer rules were produced by the immune system than by AQ. Furthermore, the total number of tests in the immune system rule set is smaller—the AQ rule set contains 29 tests while the immune system rule set contains only 20 tests. The second difference between the rule sets is that the

<p><u>Rule 1</u> THRESH: 0.08 IF vote4: abstain or nay vote9: yea THEN Democrat</p>	<p><u>Rule 3</u> THRESH: 0.85 IF vote2: abstain or yea vote3: yea vote11: abstain or yea vote16: abstain or nay THEN Democrat</p>
<p><u>Rule 2</u> THRESH: 0.78 IF vote1: abstain or nay vote5: abstain or nay vote7: nay vote8: abstain or nay vote9: abstain or yea vote10: nay vote12: abstain or nay vote15: yea vote16: yea THEN Democrat</p>	<p><u>Rule 4</u> THRESH: 0.77 IF vote3: abstain or yea vote4: abstain or yea vote5: yea (half credit) or nay (half credit) vote7: abstain or nay vote10: nay THEN Democrat</p>

Fig. 2. Rule-based interpretation of Democratic *B*-cells from immune system

<p><u>Rule 1</u> IF vote4: abstain or nay vote3: yea THEN Democrat</p>	<p><u>Rule 4</u> IF vote3: abstain or yea vote11: abstain or yea vote9: yea or nay vote7: abstain or nay THEN Democrat</p>	<p><u>Rule 7</u> IF vote13: nay vote2: yea vote3: nay THEN Democrat</p>
<p><u>Rule 2</u> IF vote4: nay vote12: yea or nay vote6: yea THEN Democrat</p>	<p><u>Rule 5</u> IF vote3: yea vote16: abstain vote13: yea THEN Democrat</p>	<p><u>Rule 8</u> IF vote12: nay vote11: abstain or yea vote16: abstain vote3: abstain or nay THEN Democrat</p>
<p><u>Rule 3</u> IF vote15: yea vote14: yea or nay vote2: abstain or yea THEN Democrat</p>	<p><u>Rule 6</u> IF vote5: nay vote15: yea vote3: nay THEN Democrat</p>	<p><u>Rule 9</u> IF vote11: yea vote2: nay vote1: nay vote16: nay THEN Democrat</p>

Fig. 3. Rule-based interpretation of Democratic AQ conjunctive descriptions




Fig. 4. Coverage and classification by immune system Democratic rules




Fig. 5. Coverage and classification by AQ Democratic rules

AQ rules are all at about the same level of generality⁴ while the immune system rules vary from very general to quite specific. This is a possible explanation for the smaller number of rules produced by the immune system. By being more flexible in constructing rules with a wide range of generality, the immune system is able to learn a more concise description of the concept.

The rule sets also have a number of similar characteristics. First, the initial rules produced by both the immune system and AQ are very similar; specifically, they both consider an abstain or nay on issue number four to be strong evidence that the voting record belongs to a Democrat. This also happens to be the most general rule produced by both methods, and from Figures 4 and 5 one can see that this rule classifies most of the examples. In other words, both methods have discovered that the vote on issue number four is the most important discriminator. The second similarity is that the concept descriptions produced by both methods must rely on rules that match only a few examples to cover the training set adequately.

6 Summary and Conclusions

In summary, we have presented a novel approach to concept learning in which a coevolutionary genetic algorithm is applied to the construction of an artificial immune system whose antibodies can discriminate between examples and counter-examples of a given concept. The results from the adaptation of this immune system to one of the standard machine learning data sets is compared and contrasted with the results from AQ15—a sophisticated symbolic inductive learning system. The immune system approach produced a description of the concepts significantly more concise than that produced by AQ, while its predictive accuracy was just as high. These preliminary results suggest that the

⁴ Although the first AQ rule classifies most of the examples, all of its rules are fairly general.

immune system approach is able to produce such a concise description by being more flexible than AQ in constructing discriminator elements with a wide range of generality.

But perhaps the most significant advantage of the coevolutionary immune system model is that it can be applied to other machine learning problems, such as acquiring task-oriented behaviors. In such cases it is necessary to recognize appropriate situations, or concepts, required to perform a task well without being explicitly trained on each. Rather, feedback is in the form of a reinforcement signal from which important concepts must be indirectly learned.

What we have described in this paper is admittedly an extremely loose model of an actual vertebrate immune system. We emphasize that the focus of this paper is to explore a new method for concept learning inspired by the immune system—not to accurately model the complex biology of immunology. It is our belief, however, that there is a potential for building more biologically faithful coevolutionary models of the immune system that may lead, not only to better machine learning systems, but also to greater insight into the workings of actual biological systems.

Acknowledgments

This work was supported by the Office of Naval Research.

References

1. K. A. De Jong, W. M. Spears, and D. F. Gordon. Using genetic algorithms for concept learning. *Machine Learning*, 13(2/3):5–188, 1993.
2. S. Forrest, B. Javornik, R. E. Smith, and A. S. Perelson. Using genetic algorithms to explore pattern recognition in the immune system. *Evolutionary Computation*, 1(3):191–211, 1993.
3. A. Giordana and F. Neri. Search-intensive concept induction. *Evolutionary Computation*, 3(4):375–416, 1995.
4. C. Z. Janikow. A knowledge-intensive genetic algorithm for supervised learning. *Machine Learning*, 13(2/3):189–228, 1993.
5. R. S. Michalski. A theory and methodology of inductive learning. In R. S. Michalski, J. G. Carbonell, and T. M. Mitchell, editors, *Machine Learning*, pages 83–134. Morgan Kaufmann, 1983.
6. R. S. Michalski, I. Mozetic, J. Hong, and N. Lavrac. The AQ15 inductive learning system: An overview and experiments. Technical Report UIUCDCS-R-86-1260, University of Illinois, Urbana-Champaign, IL, 1986.
7. M. A. Potter. *The Design and Analysis of a Computational Model of Cooperative Coevolution*. PhD thesis, George Mason University, Fairfax, VA, 1997.
8. M. A. Potter, K. A. De Jong, and J. J. Grefenstette. A coevolutionary approach to learning sequential decision rules. In L. Eshelman, editor, *Proceedings of the Sixth International Conference on Genetic Algorithms*, pages 366–372. Morgan Kaufmann, 1995.
9. J. R. Quinlan. Induction of decision trees. *Machine Learning*, 1:81–106, 1986.

10. Ivan M. Roitt. *Essential Immunology*. Blackwell Scientific Publications, eighth edition, 1994.
11. J. C. Schlimmer. *Concept Acquisition through Representational Adjustment*. PhD thesis, University of California, Irvine, CA, 1987.
12. S. M. Weiss and C. A. Kulikowski. *Computer Systems that Learn*. Morgan Kaufmann, 1991.