Artificial Immune System for Discovering Heuristics in Othello

Milad Lagevardi
San Diego State University
5500 Campanile Dr.
San Diego, CA 92182
01-619-594-2014
milad@lagevardi.com

Joseph Lewis
San Diego State University
5500 Campanile Dr.
San Diego, CA 92182
01-619-594-2014
lewis@cs.sdsu.edu

ABSTRACT
The natural immune system found in mammals offers interesting insights for models with features such as decentralized control, pattern recognition and learning. The work presented here experiments with the effectiveness of using the natural immune system model as a method for discovering heuristics in the game of Othello. The features of the AIS that offer a novel approach to decision making in this context are compared to the classic heuristics of "lookahead", where an algorithm considers all moves in the game to find the best move. The game of Othello is used since it offers a compelling search space and simple rules. Claims as to the appropriateness of using the natural immune system paradigm for discovering heuristics in Othello, and what the results may imply for other games both like and unlike Othello, will be discussed.

Track: Artificial Immune System

Categories and Subject Descriptors
I.2.11 [**Artificial Intelligence**]: Distributed Artificial Intelligence – *coherence and coordination, intelligent agents, multi-agent systems.*

General Terms
Design, Experimentation, Measurement.

Keywords
Artificial immune systems, Othello, heuristic discovery

1. INTRODUCTION
Artificial intelligence has come a long way in offering novel approaches to problem solving during a time when traditional techniques seem to be approaching a plateau. One new approach which came about in response to the challenges of traditional techniques is from Rodney Brooks [2]. Brooks introduced the notion of Agent-Based systems or Multi-Agent systems. These systems are comprised of many locally acting agents which are responsible for their own individual components and processes.

These agents are able to focus on attaining their own goals, while at the same time collectively contributing to a global goal of which they possess very little information. Another interesting and important AI design is a complex adaptive system, where the individual actions of many locally acting agents results in a global behavior, such as in the natural immune system in mammals. The natural immune system consists of many cells whose purpose is to seek out invaders and alert one another to take action to rid the body of them. Using the natural immune system as a model for computational innovations led to the creation of Artificial Immune Systems, or AIS.

The field of AIS is still in its early stages compared to other biologically inspired models such as evolutionary algorithms and neural networks. With an increasing body of knowledge and research into the field, AIS is slowly being applied to domains outside the more popular computer security applications. The AIS framework offers compelling features such as autonomy and multilayered design and learning mechanisms, which offer building blocks to future innovative designs and methods. We are interested in a new application of this paradigm: the discovery of heuristics for move choice in Othello.

2. BIOLOGY OF THE IMMUNE SYSTEM
The natural immune system is composed of many interacting parts that contribute to the overall goal of removing foreign invaders, or antigens. If an antigen is able to bypass the physical defenses (skin, hair, mucus) and bio-chemical defenses (saliva, enzymes, stomach acids), it is met with two last lines of defense known as the innate immune response and the adaptive immune response. The two main contributors in these responses are cells known as T-cells and B-cells.

T-cells have two main purposes, to help regulate immune responses or to directly attack antigens. The T-cell can regulate the immune response by stimulating B-cells or activating more T-cells. On the other hand, B-cells play an important role during the adaptive immune response. Although some antigens are removed by the first two defenses, some progress and require a more aggressive eradication scheme.

A B-cell will release an antibody into the bloodstream that can seek out and attach itself to an antigen, marking it for destruction by a T-cell. Both T-cells and B-cells attach to antigens using receptors stationed on their surfaces. It is not necessary for the receptors to be of the same size on both the antigen and the B-Cell or T-Cell, rather, the degree of match between the two is what is most important. When a B-cell attaches itself to an antigen, the relationship is similar to a lock and a key. Some keys are a perfect match for the lock, while others serve as skeleton keys, matching many locks to varying degrees. The degree to which a B-cell

GECCO'06, July 8–12, 2004, Seattle, Washington, USA.
ACM 1-59593-186-4/06/0007.

receptor is similar to, or the complement of, an antigen receptor is referred to as the affinity. The higher the affinity, the greater the match. The more valuable the B-Cell is, namely the number of times it has attached to various antigens, the more likely it is to remain apart of the immune system as a memory cell, which are long-lived cells that accumulate to battle future invaders. The adaptive immune response gets its name from its ability to create clones and initiate mutation. After attaching to an antigen, a B-cell will begin to make copies of itself in hopes of producing another B-cell with a higher affinity to the given antigen. This is done using a process known as somatic hypermutation, where the B-cell will rapidly clone itself, while each clone undergoes a certain level of mutation. The amount of mutation is related to the affinity of the B-cell to the antigen. The greater the affinity, the less mutation that will occur. Those B-cells that have a higher affinity than the original B-cell will become memory cells, used to fight off future invaders.

Artificial Immune Systems, as defined in [3], "are adaptive systems, inspired by theoretical immunology and observed immune functions, principles and models, which are applied to problem solving." Using this definition as a basis for the AIS used in this paper, it is important to clarify what an AIS is and is not. Simply creating a model where members are given the names of immune components such as T-cell and B-cell is not sufficient criteria to be called an AIS. Rather, to be labeled an AIS, the model must at a minimum contain immune components such as cells, be designed using theoretical/experimental immunology, and lastly be intended for problem solving [3]. The representation that is used might contain components such as B-Cells, T-Cells, bone marrow or the thymus, using their respective functions from the immune model to contribute to the goal of the system. Some mechanism for affinity must be present that offers a way to compare, or match, two components of the system. Lastly, the design needs to embody an immune inspired algorithm such as negative, positive or clonal selection.

3. HOW TO PLAY OTHELLO

It is important to take a moment and describe the game that is used for these experiments. Othello is credited to Goro Hasegawa who wrote the book entitled, "How to win at Othello" in 1971 [4]. The board consists of an 8 x 8 grid and is played with two players, one represented with white pieces and one with black. The initial game board begins with two white and two black pieces in the middle of the grid. Alternating turns, a player must place their piece on the board in such a way as that one or more of the opponents pieces lie in between this and another one of their own pieces. The opponent's pieces can be considered "trapped" or "sandwiched" in. Once this occurs, all of the opponent's pieces that are captured turn to the color of the player capturing the pieces. If they are black, the pieces will turn white and vice versa. After this the players alternate turns and the game continues. If a player is unable to make a move, the opponent is allowed to move again until a move for the other player becomes available. There are three ways to end the game: The board becomes full, all pieces on the board are of one color, or neither player can make a move. If any of these three are met, then the pieces are counted and the player with the most pieces on the board wins.

4. ARTIFICIAL IMMUNE SYSTEM

The AIS designed in this study relies primarily on the natural immune system's model of clonal selection. Once a B-cell has coupled with an antigen, it proliferates by making clones of itself, with each clone being able to produce only one antibody (monospecificity). Some of these clones will undergo a rapid mutation in hopes of reattaching to the antigen with a higher affinity. Those cells with a high affinity to the antigen are then selected to be memory cells, helping to fight off future antigens of similar makeup. With the AIS, the game board will serve as the antigen where the entire eight by eight grid is the antigen's receptor. The B-Cells in the AIS are three by three squares with a given configuration of black, white and empty squares along with a direction and affinity associated with them. Figure 1 shows an example of what a possible B-Cell could look like. Figure 2 illustrates the relationship between the antigen board and the B-Cell receptors.

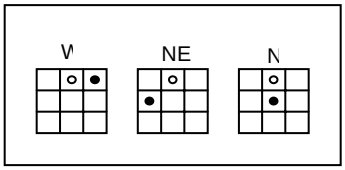

Figure 1: Three of many possible configurations of B-Cells with an associated direction based on Black's turn to move.

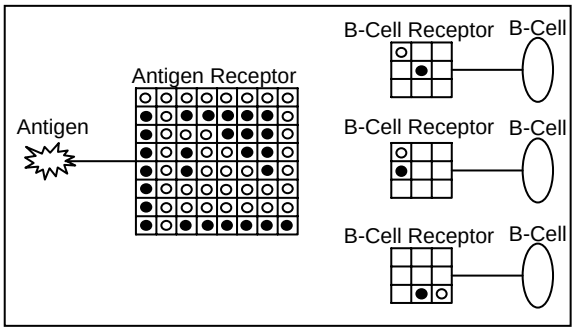

Figure 2: The Othello board serves as the antigen, while three by three squares with a given configuration are the B-Cell receptors. The B-Cell receptors will attempt to match different portions of the antigen board.

The direction value of the B-Cell corresponds to the direction in which this particular B-Cell believes a possible good move might be located. The directions can be north, south, east, west, north-east, north-west, south-east and south-west. Given the configuration of the B-Cells, three by three squares on the antigen board will be the areas of the antigen receptor the B-Cells will try to match. Figure 3 shows three of many locations that a B-Cell could be dropped onto the board for affinity calculation and clonal selection.

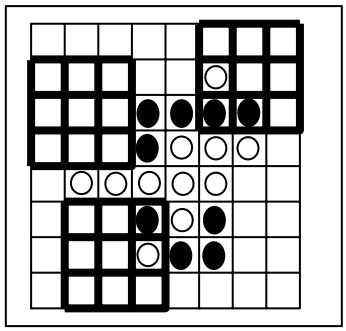

Figure 3: Three of many possible locations where B-Cells may be dropped onto the antigen board.

The natural immune system comes with a predefined repertoire of innate cells and in time builds a much more for specific repertoire of memory cells through interactions with antigens. This AIS will follow the same procedure and begin with innate B-Cells that have a configuration and direction. These cells will in turn give rise to possible memory cells. After memory cells have been created, a percentage of the memory cell repertoire will be included in the initial drop of B-Cells on the antigen board to be evaluated. There are 81 such cells that make up the entire initial repertoire.

After the AIS is initialized and the initial repertoire made active, a particular number of B-Cells from the initial repertoire and memory cell repertoire will be dropped onto the antigen board in random locations. There are 36 possible locations for the B-Cells to fall given an eight by eight board separated into three by three squares with overlap allowed. Once a B-Cell has been dropped on the board, the direction of the B-Cell will be followed to see if a valid move can be placed. This process can also be described as the B-Cell dropping and sliding across the antigen board in the direction it is associated with. If a valid move is found during the slide, the move is made using a copy of the board in order to capture the number of pieces flipped as a result of the move. This number is then multiplied by the number of times this particular B-Cell has been successful, which will in turn be the affinity of that B-Cell with respect to the antigen board. If there are more moves found during the slide of the B-Cell, the process is repeated and the move and affinity score that yielded the highest results are captured. Note that B-Cells can have only one direction associated with them.

After a particular number of initial B-Cells are dropped onto the antigen board and their affinities have been calculated, they are reinserted into the population of B-Cells and the highest ranking B-Cells will be chosen for clonal selection. These B-Cells will first be cloned at a rate proportional to their affinities, with the highest affinity B-Cells being cloned at a higher rate and vice versa. After initial cloning takes place, the B-Cells will undergo mutation at a rate proportional to their affinities, with a chance of their directions also being mutated, although with a much smaller probability. The higher the affinity of a B-Cell, the lower the chances of mutation. These newly cloned B-Cells will be randomly dropped onto the board and have their affinities calculated. The process of clonal selection will continue for a certain number of generations, with the highest ranking B-Cells at the end being chosen as memory cells.

When talking about the mutation of a B-Cell, this paper is referring to the rearrangement of the pieces on the three by three grid. The mutation function will choose one of the nine squares on the three by three grid to be mutated at random. It will never be the case that a B-Cell will mutate into a configuration it already has, assuring that some type of change is made with every mutation.

5. EXPERIMENTS

In order to be able to measure how well the AIS performs, it is necessary to have something reasonable to compare it to. In all of the experiments, the AIS was tested against an AI program that uses the square based evaluation approach to decide on moves in Othello. Part of the reason why such an AI program was chosen as the opponent is because it is usually the minimum heuristic required of an AI program that is able to beat its creator. This AI program uses no form of look ahead such as Minimax, thus it will never speculate on what could happen in the future, it simply looks at the current board and decides what the next best move is. If the AIS can offer even the slightest challenge to this AI, then perhaps some claim as to the effectiveness of the AIS model applied to decision making in Othello and possibly game play in general can be made. Figure 4 shows the values given to the squares as a part of the square based heuristic used by the opponent AI program. Corner pieces are given the most weight, while those squares either next to a side piece or corner pieces are given the lowest score. The AI program will hereafter be referred to simply as the AI.

10	8	8	8	8	8	8	10
8	1	1	1	1	1	1	8
8	1	4	4	4	4	1	8
8	1	4	4	4	4	1	8
8	1	4	4	4	4	1	8
8	1	4	4	4	4	1	8
8	1	1	1	1	1	1	8
10	8	8	8	8	8	8	10

Figure 4: The values for board spaces given to the square based heuristic of the opponent AI.

Eight parameters were extracted that might have the greatest impact on the behavior of the system. The first official experiment was run with different values of these parameters in hopes of learning what effects they had on the system. The eight chosen parameters included: Generations, Drop Number, Low Cloning Number, Medium Cloning Number, High Cloning Number, Low Mutation Chance, Medium Mutation Chance, High Mutation Chance

The generations parameter is the number of times the clonal selection algorithm is allowed to be repeated. The drop number parameter is the number of initial B-Cells that is allowed to be dropped on the Othello board when a move is requested. Low/Medium/High cloning number is the amount clones that are made based on the affinity of the B-Cell. The higher the affinity, the higher the cloning number and vice versa. Low/Medium/High mutation chance is the chance that a B-Cell will be mutated one a

scale of 0 to 100. If the affinity of the B-Cell is low, it will have a high chance of being mutated and vice versa. The initial values experimented with for these parameters were as follows:

- Generations - 5, 6
- Drop Number - 10, 20
- Low/Medium/High Cloning Number - 2/4/15, 5/10/15, 8/10/11
- Low/Medium/High Mutation Chance - 15/40/65, 25/50/75, 35/60/85

Both the generation number and the drop number can have two values each, while the cloning number and mutation chance choose one of three configurations. The cloning number of 2/4/15 attempts to favor those B-Cells with a high affinity, while paying little attention to those with lower affinities. The 5/10/15 options give a more evenly distributed weight to all B-Cells, while the 8/10/11 option will accept and clone B-Cells much more frequently, cloning all at very close rates. Although this last method does not hold fast to the generic clonal selection model of cloning at a rate proportional to the affinity, it was important to experiment with it to observe the effects of different cloning rates on the system. Lastly, the mutation chance parameter uses an offset of -10, 0 and +10 from the base value of 25/50/75. Thus in the first grouping 10 is subtracted from all values, then the values are kept the same, then are increment by ten. To try all combinations of these 8 parameters, (2 * 2 * 3 * 3 =) 36 different tests needed to be run. Each test was run 1000 times for 36,000 total tests.

With respect to the win rate, the AIS performs anywhere between the 14 and 18 percent range. Although no claim that a particular configuration definitely yields better results was possible, the claim that the tests show a consistent playing ability and behavior is evident. Out of 1000 games, the chances of the AIS winning are 14-18 percent. According to the data, the higher the generation number, the greater average number of B-Cells are created. The reason for this lies in the design of the AIS. Since the AIS will choose the top 10 percent of the entire B-Cell repertoire as the cloning pool, one extra iteration of the clonal selection algorithm will only add to the pool. The larger the pool, the greater amount of B-Cells are taken when the top 10 percent are extracted. As added statistical analysis given the large data set, standard deviations and 95% confidence intervals were calculated for each of the 36 configurations with respect to the win rate. Using the number of wins from each configuration, a mean of 165 wins was calculated, with a standard deviation of 11.8 and a 95% confidence interval of 160.6 - 168.6.

The clone number and mutation number seemed to also produce varied behavior given their three different value sets. The cloning set of 5-10-15 resulted in the most wins, but not enough to claim that having this particular configuration was the best. Mutation rate, of all the parameters, seemed to offer the least information. In order to be thorough and put to rest the belief that the parameters did not effect the system in any major way, another set of experiments using more focused values for the clone number and mutation rate were used.

Instead of running a cloning number set with different configurations of values, two of the 36 configurations were run 1000 times each with a cloning set of 10-10-10 were run, giving equal chance for all of the B-Cells to be cloned, regardless of their affinities. Based on the prior experiments, it was speculated that this test would only show that an even range of clones would not produce better results against the opponent AI. The supposition was correct and the number of wins still resulted in the 14-18 percent range, just as the previous test run and verifying that the clone number is inconsequential.

To verify that the mutation chance also had little effect on the systems behavior, an experiment where the system would first try very little mutation and then a lot of mutation was tried. One configuration was taken and ran 1000 times with a mutation parameter set of 5-10-15 for low mutation occurrences and 75-85-95 for very high mutation occurrences. It was found that, as anticipated, the mutation chance was not a great factor in the success of the AIS.

Lastly with regards to the drop number, this parameter was viewed as having the least effect on the systems behavior next to the mutation rate. The explanation for this again lies in the intent of the system's design. No matter what the number of random B-Cells chosen to be dropped on the antigen board, they are then re-inserted into the population, adding no new B-Cells to the already large pool, only altered ones. The drop number is really a way for the AIS to produce new affinity values for the repertoire, keeping the repertoire in an ever changing state. To help verify this, another experiment with the best performer, test number 29, was attempted 1000 times using a drop number of 5, 15, 25, and 30. Knowing the design of the AIS, it was expected to see a similar win rate as what was seen earlier. As expected, a varying value of drop numbers neither increases nor decreases the win rate.

Once it was shown with a few extra experiments that the parameters to the system are not offering increased performance outside of the consistent 14 to 18 percent success rate the system achieved, it was necessary to think of ways to modify the system design that might offer new behavior.

The first such design refinement was the directional value associated with each B-Cell. This is the value that determines the general area a B-Cell believes that a good move on the board might exist. The initial repertoire had predefined directional values that came from the authors' bias, and subsequent B-Cells received their directions via their parents. Testing the effects of the system when the directional value is also allowed to be mutated was the next step in the test plan. Due to the important role that the directional value plays in the system, allowing it to be mutated should either increase or decrease the win rate. Since this was not one of the original design decisions of the system, the necessary code to implement this new functionality was added and then put together a small test plan. It was learned from previous experiments that no new behavior resulted from the initial parameter set, so the most neutral configuration was chosen as the base configuration, with a generation value of 5, a drop number of 10, clone rate of 5-10-15 and mutation chance of 25-50-75. The new functionality of being able to mutate the directional value of the B-Cell uses a set value probability to decide whether or not the direction should be mutated. Three tests were compiled using the general configuration above where the effects of mutating the directional value 5, 50 and 85 percent of the time were tried. The speculations about the effects of mutating the directional value of the AIS were incorrect, since the win rate did indeed fell within the normal success rate for all three attempts at directional mutation.

The second attempt at producing new behavior in the system by changing the design of the AIS involved the nature of mutation. If a B-Cell is chosen for mutation, then one of the nine squares on the B-Cell will be changed to a piece other than the piece currently placed there. The effects of this choice for mutation were in question, so the functionality to mutate more than one of these squares was added. Again using a set value probability, when a B-Cell is chosen to undergo mutation, it will have a 25 percent chance of having 3 of its squares mutated, a 50 percent chance that 2 of its squares are mutated, and a 75 percent chance of just having one of its squares mutated, identical to the original design. By mutating more than one square, it was speculated that the added diversity in the configurations of the individual B-Cells would give rise to some new behavior. This extra diversity did nothing for the win rate as the AIS again performed within the normal success pattern.

With two failed attempts at producing new behavior when applying design changes to the AIS, one more functionality change was conjured up by the revisiting of the biological inspiration of the model itself. In the natural immune system, when memory cells are produced, they are retained to fight new and possibly returned invaders for years to come. In the design, every time a new game is initiated, a brand new memory cell repertoire is produced. The effects of keeping the memory cells' repertoire around from game to game was explored while allowing it to be manipulated as the progression of games continues. If the memory cells that carry on after many games are refined and successful contributors in game play, then it was speculated that the AIS should perform better than when it does not have these matured B-Cells to work with. It was surprising to see that the AIS did worse than before it had this extra information to work with. By using the memory cells, the AIS only won 11 percent of the time, a negative behavior not anticipated given the biological model.

At this point in the experiments, a change in the design of the AIS was not worth pursuing, seeing from the three different experiments above the lack of better behavior once these changed were incorporated. To try and better understand why these changes were not helping the system, an analysis of the opponent AI patterns of play was begun. The opponent AI uses the square base heuristic as illustrated in Figure 4 and has an advantage of knowing good and bad pieces on the board. The AIS has no concept of board location, as the B-Cells are only of size three by three and only know about the squares they land on and move to. Furthermore, the B-Cells in the AIS gain affinity by finding successful moves and recording the number of pieces that were flipped as a result of the move. So at the heart of the competition among these heuristics, the opponent AI incorporates a notion of good, bad and neutral positions while the AIS only cares about pieces flipped. This analysis gave a new way of thinking about the design of the AIS.

Previously the focus was on design changes that had to do with the biological model and immune system components and how particular components and algorithms accomplished their goals. Now the focus moved to looking closer at the computational side of the system and how the heuristic components played a factor in the behavior of the AIS. Much like in the natural immune system where a receptor of a B-Cell might be more prone to attach to a particular region of an invader, those B-Cells that found a move that happened to be on particular locations of the antigen board, specifically the sides, were rewarded. At first the idea of rewarding the three by three B-Cells that happened to find a move that included the sides was considered; however, since the three by three blocks that include sides saturate almost the entire board, the desired effects would be lost because almost all B-Cells would be rewarded. So with some added code, a small boost of affinity was given to B-Cells that found moves laying on the outer edges of the antigen board, since the bias of the Othello game shows those are always good moves to be made. It was hypothesized that with this new positive reinforcement the AIS would perform better than average against the opponent AI.

When the first tests began running to verify the idea, it was quickly noticed that the AIS taking a very long time to play the game, recalling that a typical game should only last about 10 seconds. After some analysis, it was discovered that since the affinities of the B-Cells were growing larger due to the boost in finding moves on certain regions on the antigen board, many of the B-Cells were being cloned at the high end of the cloning algorithm, namely 15 clones each. This was causing the system to make many new clones which in turn resulted in the greater amount of time needed to make a move. To put the system back into a state of equilibrium, the affinity cut off points for which the system was determining how many clones to make was modified. Once this recalibration took place the system began to complete a game in about 30 seconds so the new run of 1000 games to test the hypothesis was begun. For the first time the hypothesis was confirmed and the AIS performed better than the normal success window. The AIS was able to win 22 percent of the time as opposed to the normal 14 - 18 percent.

The positive reinforcement test gave a new and inspiring perspective on the inner workings of the AIS and what an important role the heuristic in the system played on the behavior. One final test was run to see if negative reinforcement would have any added benefit to the system. We deduct points from the affinity of the B-Cell if it happened to match a move on the antigen board that was unfavorable, namely involving those pieces that were adjacent to all side pieces. If positive reinforcement caused an extra 3 - 8 percent increase in performance, then it was suspected that the notion of unfavorable locations on the receptors would produce even better results. Unfortunately an increase was not observed, but the system still had a 2.5 - 7.5 increase in performance over the usual results.

As a control group experiment, an AI program that finds all valid moves on the board and randomly chooses between them was implemented. The purpose of this AI program was to compare its win rate against the AIS win rate and determine at a minimum whether the AIS is at least better than random decision making. After 1000 runs of the AI program against the opponent AI that uses the square base heuristic, the AI program was able to win 9.5% of the time, 95 wins out of 1000. The standard deviation was .293 with a 95% confidence interval of .077 - 0.11. This value is less than the average 14-18% win rate of the AIS against the same opponent, offering evidence that the AIS is at least better than random game play.

The aims of the experiments presented here were to obtain a better understanding of the effects of the parameters on the behavior of the AIS with respect to game play and to explore the AIS design itself in this domain. It was shown that the parameters originally believed to have the most important roles in the system

were in fact not heavily responsible for the results received. Furthermore, design changes other than parameter manipulation were attempted in hopes of acquiring new behavior. This resulted in discovering the effects of the computational aspects of the game, specifically those associated with calculation of the affinity measure.

6. FUTURE WORK

Using an AIS for making decisions in game play is a new application in an approach still in its infancy. That being the case, all data collected so far has led to many more ideas and experiments to be explored. For any expansion on this work or for the particular application of an AIS to game play, the following ideas are offered as a initial outline.

Although it was determined from the experiments that the parameter set that was used was greatly affecting the behavior of the system, there are still some combinations of values that could have been tried if more time was available. It would be interesting to see the results of using the positive and negative reinforcement in conjunction with keeping the memory cell repertoire around from game to game. Perhaps the reinforcement technique would have further refined the memory cell set to one that better captured the notion of good directions of locations of moves on the board. The original idea of heuristic discovery was hoped to be realized in such a way; the ongoing refinement of a set of memory cells carried the information to support good choices. Patterns of play would be garnered over time.

In the AIS, three by three blocks were used to serve as the B-Cell, and the reason for the decision to use such as size was because two by two blocks were too small to justify a direction for a possible move to be located while four by four blocks offered too many combinations for the initial round of experiments. A worthwhile experiment would include producing four by four blocks for B-Cell and observing its affects on the system. To undertake such a task, an initial repertoire that is much larger than the one created will need to be produced. The benefit of having a larger block size is the increase view that each B-Cell receives to incorporate into deciding where good moves are located. A larger view means an increased opportunity for finding larger pattern sizes, which could be an integral component of the success of the AIS. Since the AIS was tailored to use three by three blocks, making a jump to four by four blocks is non-trivial. The initial design of the code did not warrant an abstract implementation aimed at growth since the scope of this project was short.

As B-Cells were dropped onto the board, they follow a direction and see if a valid move exists. If a valid move does exist, the move is made and its effects calculated. The B-Cell, with its limited view of the environment, namely three by three blocks, will not know the effects of its decision across the board. In Othello, the placement of one piece can affect a string of pieces spanning from one side of the board to the other. One way of being able to capture the changes that B-Cells produce as a result of trying moves is to produce some mechanism of communication between B-Cells, much like in the natural immune system where the cell use chemicals to communicate and pass messages. By allowing the B-Cells of the AIS some form of communication, then they would be more informed of their surrounding environment, which might lead to better performance.

Keeping with the theme of alternate block size, the effects of using the entire board as one big B-Cell were contemplated. Now with a block size of three, four, or even five, there is a notion of travel direction, the direction that a B-Cell feels a good move might exist. If the B-Cell is one big board, then this notion of direction is lost and the model for the AIS is altered. The AIS now moves from an offensive strategy to a defensive one. In the offensive strategy, the AIS would use its clonal selection algorithm to find good moves and locations on the antigen board, a very proactive approach. In a defensive strategy, the AIS is more of an observer that captures the moves made by the opponent and saves them for future use in making decisions on moves that are worthwhile. These board configurations would turn out to be the memory cell repertoire of the AIS and could be used in future game play. The defensive strategy is similar to the notion of a board configuration database that is used by professional Othello playing programs. The notion of defensive playing need not be only applied when considering the whole board as one big B-Cell, but could also be applied when B-Cells are of smaller size.

In the strict implementation of the clonal selection algorithm described in [3], the B-Cells that attempt to match an antigen are the cells that are immediately cloned and undergo somatic hypermutation in order to find a B-Cell with a higher affinity. In the design described here, these B-Cell are first re-inserted into the cell population and those B-Cells with the highest affinity are chosen to undergo the clonal selection algorithm. This method was chosen to introduce diversity into the clonal cell pool, allowing the system to trying a range of options and not just those believed to be the best. It turns out that those B-Cells that ended up being cloned actually came from the set of those B-Cells that were chosen to encounter the antigen in the first place, allowing the design to stay close to the stricter implementation described above. An interesting experiment would be to follow the strict implementation and observe the affects on the system. From the knowledge about the AIS, using this approach would cause the drop number parameter to play a much more important role. The higher the drop number, the greater number of B-Cells that will undergo clonal selection. Subsequently, the memory cell repertoire would be smaller than what was apparent because far fewer B-Cells would be chosen to become memory cells.

Finally, the AIS was pinned against an opponent AI that used the square base heuristic, one that focuses on the value of each square and has no concept of look ahead in game play. One could try to put the AIS up against an AI that is far more informed than the opponent AI and verify its effects. Based on the win rate of the AIS, an opponent that is more informed would cause the consistent win rate of 14 -18 percent to drop dramatically. However, an AI that uses look ahead often assumes that its opponent will play that way it does, and since the AIS plays nothing like an AI that has look ahead, it would interesting to see how the opponent AI reacts.

7. CONCLUSION

The experiment set began as an ad hoc approach to understanding the boundaries of the AIS, it was quickly learned that there were resource limitations inherent in the design of the system. After this, a set of parameter values and alternatives were created to get a sense for the space of behaviors the AIS can exhibit, measured ultimately in the number of wins against the opponent AI. The

first round of experiments were promising but lacked enough accuracy to make strong claims as to the behavior of the system, so the necessary optimizations were made to run 1000 games for each of the 36 configurations of parameters sets, learning about and fixing discrepancies in the code as the implementation progressed. Once it was clear that the parameter set was not offering any radical impact on behavior and that the system performed consistently in the 14-18 percent range, the focus shifted to design changes that might lead to new behavior. Manipulating the number of times and ways mutations occur was attempted, only to find it did not increase the performance of the AIS. The biological model was mimicked more closely by allowing the memory cell repertoire to carry on from game to game, but that too failed to offer better results. Unable to find performance insights in the biological model, the focus was then altered to the computational aspects of the system, namely the affinity calculations happening behind the scenes. Once a positive reinforcement to the B-Cell repertoire was added for matching particular regions on the antigen board, the first real performance increase was attained, going from the 14-18 percent range to the 22 percent range. This effort gave rise to the idea that the affinity measure needed to receive more weight than previously thought, since it seemed to play such an important role in the behavior of the system. A final experiment to test the effects of negative reinforcement was added that reduced the affinity of a B-Cell when encountering an unfavorable region on the antigen. The results showed no increase in performance, leaving it necessary to research why the system was not performing any better even though affinity reinforcement seemed to be a step in the right direction. It is of note that the only changes with successful outcome undermined the very idea that this AIS would discover a heuristic by actually using a heuristic in the affinity calculation.

Only two experiments that touched on the importance of the affinity measure as an important component to the system were completed. One idea that was struggled with was adding even more information to the affinity component to make it more informed. Consider the heuristic technique of mobility value as described in [5]. Such a technique would undoubtedly increase the playing power of the AIS, but for two reasons was not going to be accomplished. The first reason is the amount of effort it would take to give the AIS such a broader view of the environment, much larger than each B-Cells view. Secondly, and more importantly, doing so would make the AIS no different than a regular mobility- based heuristic, losing its meaning of being an AIS altogether. This method would also go against the original intent of the uninformed AIS "learning" the moves made by the opponent and then mimicking them in its own game play. After all the experiments that have been run thus far, an explanation as to why the AIS does not play Othello the way it was desired was formulated.

The design of the AIS is myopic in nature, allowing a very limited view of the environment through the B-Cells of the system. The opponent AI looks at every piece of the board and then makes a decision based on its internal mechanisms. The AIS, however, could very well never encounter all squares on the board. No matter how many B-Cells are dropped on the antigen board, given all the different directions a B-Cell could travel, the board will not be completely saturated and thus not every best possible move found. Drawing from the understanding of games like Chess, making one bad move now can affect the outcome of the game later, and Othello is no different. The moment the AIS fails to make the best move it could have made, it has decreased its chances of winning, and if this occurs during every move cycle, its not surprising why the AIS (in its several variations) only won 14-22 percent of the time. Furthermore, the fact that a B-Cell cannot see the effects of its moves across the board limits the AIS even further. The AIS was designed to find local patterns on the board, when it could have benefited by knowing information from the environment outside of its field of view. Thus the combination of the chosen representation of the B-Cells and the method by which affinity is calculated was probably an incorrect choice for Othello. Given these limitations, the system as it stands has reached a plateau and will not be improved unless these limitations can be addressed.

Othello is not the only game where understanding global patterns play an important role in the outcome of the game. Games such as Chess and Checkers cannot simply be played by deciding on which moves to make just by looking at a few pieces, rather, the entire board and its pattern must also be taken into consideration when deciding on the best move to make. However, some games, such as Go might be better suited for the design of the AIS, because they can be played effectively by focusing on local patterns [1]. Future work might focus on such a game that would best exploit the AIS design strategy.

It was the goal of this work to use the artificial immune system model to build a program that could mimic the moves of a first time Othello player by using cell patterns trained and matured over many moves and many games. To mimic a novice, the AIS used no brute force look ahead such as Minimax and had no long term experience because it was always offered a fresh new memory cell repertoire to begin with. Given that we can only beat the opponent AI about 20 percent of the time too, the AIS does pretty well, all things considered. It was hoped that the memory cell set would in some way embody a heuristic that could then be transported to any other Othello playing program and compete like a beginner would. Being a type of complex adaptive system, it was inappropriate to hand the AIS any specific information about Othello, but rather to allow the notion of good moves and good patterns to emerge on their own. Based on the experiments, it is unclear that that actually happened. It is also unclear from these experiments whether Othello is the appropriate game to study AIS behavior for. Whatever the case, further progress must first address the limitations described above.

8. REFERENCES

[1] British Go Association - *How to play Go*. Last updated August 8, 2002, http://www.britgo.org/intro/intro2.html

[2] R. A. Brooks. *Intelligence Without Representation*, Artificial Intelligence Journal (47), 1991

[3] L. N. de Castro. and J. Timmis. *Artificial Immune Systems: A New Computational Intelligence Approach*. Springer, 1st Edition Chapter 3, 2002

[4] G. Hasegawa. *How to win at Othello*. Jove Publications, 1971

[5] Mandt, Marc. *Basic Othello Strategy*. Generation5 Articles, 2002