

The effect of tags on non-local adaptation

by

Bradley James Powers

A thesis submitted to the graduate faculty
in partial fulfillment of the requirements for the degree of
MASTER OF SCIENCE

Major: Bioinformatics and Computational Biology

Program of Study Committee
Daniel A Ashlock, Co-major Professor
Kirk A Moloney, Co-major Professor
John E Mayfield

Iowa State University
Ames, Iowa
2004

Graduate College
Iowa State University

This is to certify that the master's thesis of

Bradley James Powers

has met the thesis requirements of Iowa State University

Committee Member

Co-major Professor

Co-major Professor

For the Major Program

As with all research and writing, a lot of work and encouragement has been required to complete my thesis. Without the help from friends, family, and others I could never have done this. Specifically, I would like to thank: Stacey, Dan, Mom, Dad, Jeff, Shannon, & Jess.

Special thanks to Sherlock & Bailey, who always kept my lap warm during the typing of this paper.

TABLE OF CONTENTS

Abstract	v
Chapter 1. Introduction	1
Current Evolutionary Theory	1
Game Theory	3
Iterated Prisoner’s Dilemma	5
Previous Work on Non-Local Adaptation	7
Further Insight into Non-Local Adaptation	11
Paintbots and non-local adaptation	11
Gridplants and non-local adaptation	12
Project Objectives	14
Chapter 2. Materials and Methods	16
Agent Specifications	16
Grid Play	18
Competition	18
Chapter 3. Experimental Design	20
A Higher Resolution Look at Non-Local Adaptation	20
A Survey of Different Tag and Tag Array Mutation Levels	21
A Look at Overall Mutation Rate	21
Chapter 4. Results	23
Chapter 5. Discussion	29
Non-Local Adaptation Over Time	29
A Survey of Tag and Tag Array Mutation Levels	30
Overall Mutation Rate	34
Chapter 6. Conclusions	35
Future Work	36
Acknowledgements	39
Works Cited	40

Abstract

This project investigates in greater depth the phenomenon of non-local adaptation previously observed in an evolutionary model based on the game iterated Prisoner's Dilemma. Non-local adaptation is the ability of an agent or population of agents to perform well against other agents that share no common history or ancestry with them. Populations of agents both with and without identifying tags are evolved to perform noisy iterated prisoner's dilemma on a toroidal grid. The agents consist of a finite state machine specialized for playing iterated prisoner's dilemma and simple tag recognition capability. The populations are allowed to evolve for 10,000 generations and the state of the world is stored every 500 generations. Populations from these samples are placed in competition with populations from generation 10,000. This procedure is repeated for varying levels of overall mutation rate, with and without tags, and varying frequencies of tag related mutations. Non-local adaptation is seen in these populations, however, tags seem to slow the acquisition of non-local adaptation. Although the concept of non-local adaptation is not a widely accepted phenomenon in biology, these results suggest that it may happen and that the effect is persistent in the face of changes in mutation rate and in the face of increased task complexity. Further analysis of the populations with tags shows interesting patterns of "tag space" usage. Specifically, the populations tend to have a predominant tag most of the time with punctuated periods of increased tag space usage that most likely correspond to invasion of the population by an opportunistic agent with a new tag identifier.

1. Introduction

Current Evolutionary Theory

After the introduction of the theory of evolution by Darwin that theory has grown to be the explanation for the development of the immense complexity of life. In the book *The Blind Watchmaker*, Dawkins outlines the current state of evolutionary theory and defends evolution against some opposing theories that still exist in the face of the overwhelming scientific support for evolution.[1] One point raised by detractors of evolution is that it is extremely unlikely that complex organs like eyes will arise through chance mutations. Dawkins reply comes in the form of an example. In this example, Dawkins looks at the probability of a computer randomly generating the phrase “METHINKS IT IS LIKE A WEASEL”. In this example, the alphabet is restricted to the 26 capital letters and an empty space. Even with these restrictions, the probability of generating the target phrase is 27^{28} , or 1 in 1.2×10^{40} . For perspective, the odds of winning the powerball lottery are 1 in 1.2×10^8 . [2] Given the extreme difficulty of generating this simple phrase by chance, people argue that it is near impossible that a complex organ, arguably much more complex than a simple string, to come about through chance. If these complex structures cannot come about by chance, they argue that evolution cannot explain the diversity of life currently found on the planet. As Dawkins points out, this example is missing some of the key features that are necessary for evolution to work. In evolution theory, it is not necessary for fate to spontaneously generate an eye, a foot, or a person. Instead, these complex features arise through incremental changes that are passed through the filter of natural selection and then inherited by future generations. Dawkins then alters the previous example so that an initial string is randomly generated, then several generations are allowed to pass. A generation consists of duplicating the initial string several times with a small chance of a mutation (changing a letter in the string), strings that more closely resemble the target phrase survive

and produce more offspring. Using this selection method, the string is consistently produced in a little over 40 generations of selection. [1]

This example illustrates the power of mutation combined with selection to bring about events that have low naive probability, however, it suggests a new question. In the previous example evolution is given a definite target, the target phrase. As a result there is a goal and progress can be measured. In evolution, natural selection is a blind force that cannot have a target or goal. As a result, it is widely believed that evolution will result in adaptations that solely benefit the organism in its current environment. For the rest of this paper we will refer to this environmentally constrained adaptation as local adaptation. Despite the support for local adaptation theory some biologists believe that evolution can also result in general adaptive features for the population. In other words, evolving populations can gain general abilities through the course of evolution. Dawkins himself alludes to this when he says:

“if ... predators from one era could meet prey from another era, the later, more ‘modern’ animals, whether predators or prey, would run rings around the earlier ones.

This is not an experiment that can ever be done ... [1]”

For the rest of this paper, evolution that results in the acquisition of general adaptive features will be referred to as non-local adaptation. The key to investigating the possibility of non-local adaptation is to obtain populations of the same type such that their only difference is position in evolutionary time and to place these different samples in competition. If the modern generation significantly out-competes the archaic generation, this would give evidence that non-localized adaptation can happen. Although this is a simple test to conceive of, lack of practical methods of time travel demand the use of some innovation to arrive at a suitable way to test this theory. Even in the examples where generation time is short, i.e. bacteria, the time that the organisms have had to evolve is so great that it is difficult to obtain samples far enough apart so that non-local adaptation can be seen.

Game Theory

The inability to perform direct experimentation is not unique to evolutionary biology. In the field of economics it is rarely possible to experiment directly with a country's economy and experimenting with untested theories could have serious consequences. As a result, economists have utilized game theory to model markets so that they can identify the possible effects of policies and strategies on an economy. This approach allows the economists to obtain a better understanding of the way that the economy works without endangering the economy, and the financial strength, of entire nations. As mentioned before, it is not possible to directly test the theory of non-local adaptation by using biological organisms, however, game theory can be used to conduct experiments.

In game theory, researchers try to identify the essential elements of a situation and then use a simple game as a model for that situation. In simplifying the situation the hope is to identify general strategies from the simple games that can be used in a wide variety of models. To assist with the classification of games, researchers in the field of game theory have identified essential elements of a game. These include the number of players, the nature of the payoffs, the amount of knowledge each player has about the other, and the number of times that the players interact. [3] The number of players has a direct effect on the complexity of the game and the strategies that arise. In a game with only one player, the strategy adopted need only consider the current environment and the consequences of the player's actions. With two or more players, the strategies adopted by the opponents need to also be considered. The nature of the payoffs can either be zero sum or non-zero sum. In zero sum payoffs any gain by one player must come from a loss of the same amount by another player. A good example of a zero sum game is poker. The total money available does not change, instead one player increases their fortune by winning it from the other. A characteristic of zero sum games is that they place the players of the game in direct competition with one another. In non-zero sum games players are placed in competition,

however, the structure of the payoffs provide an incentive for cooperation as well. This cooperative element is the result of the lifting of the requirement that the total of losses in the game equal the gains. With the removal of this limitation the total payoff available (the sum of all of the payoffs given to all players) becomes a factor which, as we will see later, can have an interesting effect on game play. The amount of information that a player has is another important quality of a game. A game where the agents know their opponent's every move, even future moves, is said to have perfect knowledge. With perfect knowledge a player can use this information to develop a strategy that takes the opponent's future moves into account. A spectrum of games is formed as games can have decreasing amounts of information even to the point where the player only knows the players past moves. The number of interactions is another defining characteristic of a game. This is particularly important for non-zero sum games where there is a possibility of benefit from cooperation. If two players will only meet once, there is little incentive to cooperate as there will be no future rewards from this cooperative behavior. The result is that the player does best to take all they can get in that single interaction. If there will be multiple interactions, this greedy behavior can result in retribution from the opponent, under the right circumstances this retribution can reduce the total payoff to the point where the agents would do better to cooperate and maximize the total payoff that they can receive. As a result, a game that is iterated with the proper non-zero-sum payoffs can result in cooperation having direct and important benefits to the agents in the game.

In considering these qualities, many researchers choose a two player, iterated, non-zero sum game with imperfect information as a tool for studying evolution. This makes sense given the following observations: 1) Interactions in nature can usually be viewed as interactions between two entities. 2) The solutions that are arrived at are similar to non-zero sum game solutions. 3) The opponent's move is typically not known. 4) These interactions

usually happen more than once. One game that meets these requirements is iterated prisoner's dilemma.

Iterated Prisoner's Dilemma

Robert Axelrod's book *The Evolution of Cooperation* is an excellent introduction to the game iterated prisoner's dilemma (IPD). [4] The premise of this game is that two agents are given a choice of two moves, cooperate or defect. Each player makes a move knowing only its opponent's previous moves and both players announce their move simultaneously. If the player and its opponent both cooperate, both receive the payoff of C. If both defect, they receive a payoff of D. If one cooperates and the other defects, the cooperator receives a payoff of L and the defector receives a payoff of H. In IPD payoffs must follow these equations:

$$L \leq D \leq C \leq H \quad (1)$$

$$L + H \leq 2C \quad (2)$$

For Axelrod's studies the values $L = 0$, $D = 1$, $C = 3$, $H = 5$ were used. The use of these scores creates an interesting dynamic between the players. The naive solution is to simply defect, with this strategy you get a score of 5 every turn if you have an extremely nice opponent. The opponent, however, is receiving a score of 0 each time he cooperates. There is little incentive for the opponent to continue this behavior, as a result the opponent will likely retaliate by defecting as well. When this happens both players will receive a score of 1 per turn. If both players could cooperate, however, each would receive 3 points per turn. As a result there is an incentive for the players to cooperate, however, there is also an incentive to try to take advantage of the niceness of the other player. This forms the dilemma of prisoner's dilemma. The best time to "take advantage" of an opponent by sneaking in a defection is the last round of the game. If a player defects at the end of the iterations then there can be no retaliation. If a player can expect their opponent to defect in the last round, the best strategy is to defect in the round before that. This logic works recursively until the

best strategy is to always defect. To avoid this problem the number of iterations needs to be unpredictable to the agents.

The book considered two experiments using IPD and also contained a considerable amount of theory concerning the application of IPD to various types of competition. In the first experiment Axelrod asked fourteen experts in game theory to submit programs for playing IPD. All of the submissions played against each other as well as a totally random strategy. In the initial run a simple strategy called tit-for-tat stood out above the other submissions. The tit-for-tat strategy cooperates on the first move and then copies its opponent's last move. In further analysis of the competition Axelrod identified several qualities that seemed to separate the agents that performed well from the ones that did poorly. These characteristics were niceness, fast retaliation, quickness to forgive, and simplicity. Niceness was an important factor in the scoring. Niceness is defined as not being the first player to defect. The nice strategy favored cooperation which maximized total payoff (opponent's and player's payoff combined). Fast retaliation referred to the speed with which the agent responded to defection by an opponent. This was also important to avoid being taken advantage of by other agents. Although retaliation was important, forgiveness was also an important quality. Being quick to forgive allowed for cooperation to resume when the opponent stopped defection. Simplicity of strategy was also beneficial as it allowed opponents to quickly identify the strategy that they were playing against and adjust accordingly (in the case of playing against tit-for-tat, the adjustment is to play cooperatively). The best performing strategy, tit-for-tat exhibited all of these qualities.

In a follow up to the original experiment, Axelrod repeated the experiment with a larger number of participants. Unlike the previous experiment, entry into the competition was not restricted to game theory experts. Participants were given the analysis from the initial study that indicated the qualities that performed well in the initial experiment. Again tit-for-tat won the competition, and the strategies sorted themselves based on the four

qualities identified in the initial experiment. The conclusion from these experiments was that strategies like tit-for-tat are successful because they have the four qualities that encourage cooperation.

Axelrod further generalized his results by taking the entries from the second competition and putting them into an “ecological competition”. In this experiment a generation consisted of every strategy playing all of the other strategies (including itself). After the play was completed, the frequency of the different strategies was adjusted to correspond with their performance in the competition, better performing strategies would gain representation and poorer strategies would lose representation. Given that the agent strategies were deterministic, the same scores were used throughout all generations, only the weighting of the score resulting from a pairing was adjusted. This experiment was deemed an ecological rather than evolutionary competition as no mutation of strategy was allowed. The result was that the strategies that followed the four qualities increased in number and eventually weeded out the strategies that did not, again the best performance was tit-for-tat.

Axelrod spends the rest of the book applying these results to various subject areas including the live-and-let-live trench warfare situation in World War I, the cold war, and biological systems. This thorough study of IPD and its implications of biological species made it a strong choice when researchers wanted to study evolution using game theory.

Previous Work on Non-Local Adaptation

Ashlock and Mayfield used the IPD game in an experiment that provided evidence of non-local adaptation. [5] Unlike the Axelrod study, the interest was to identify if the players from more modern generations were better at playing IPD than their more archaic counterparts. The purpose being to test Dawkin’s assertion of the general superiority of modern generations over archaic generations. In keeping with that purpose, the study involved some important differences from the work of Axelrod. The game was changed by the addition of noise, in the format of a small probability ($\alpha = 0.01$) that an agent’s move would be replaced




Figure 1 Example diagram of FSM based agent used in the Ashlock Mayfield study. Transition labels are in the form of “opponent’s last move” / “agent’s next move”. For initialization arrow, only the agent’s first move is specified.

by the opposite move. The addition of noise served to eliminate the superiority of the tit-for-tat strategy. Such a change was desirable as the goal was to create a situation where there was no best strategy. If the simulation had a favored strategy, then the experiment would be much like Dawkin’s string example where there is a well defined target. If there is a defined target that selection is trying to achieve, it would be no surprise if steady progress toward that target were to appear. Another change that was required involved changing the experiment from an ecological study to an evolutionary study. For evolution to occur there must be heredity and mutation, in the Axelrod work there was a form of heredity, but no mechanism for mutation. The Ashlock and Mayfield study provided the ability for mutation through the use of agents that consisted of finite state machines (FSM’s) as shown in Figure 1. Each state in the FSM consists of two transitions originating from the state, one that is used if the opponent’s last move was defect, the other if the last move was cooperate. The transitions indicate the agent’s current move as well as the next state that the agent should go to. The agent’s keep track of their current state as play continues. Agent mutation was accomplished through changes to the FSM. Three types of mutation were allowed: Basic, Growth, and Contraction. A Basic mutation consisted of selecting a transition at random and changing either the agent’s action associated with the transition or the state that the transition pointed to. A Growth mutation consisted of selecting a state at random and duplicating it, the in-

bound transitions assigned to the duplicated state were then assigned randomly between the original state and its duplicate. The functional result of a growth mutation was no change in the current behavior of the agent, however, it did alter the impact of future mutations, much like gene duplications in biological organisms. A Contraction mutation consisted of randomly selecting a state, deleting it from the FSM, and randomly assigning the transitions that pointed to the deleted state to the remaining states in the FSM at random. A spatial aspect was given to the experiment by placing the agents on a 100 x 100 toroidal grid. Agents would play against their neighbors to the north, south, east, and west. In addition to playing against the neighbors, the agents would play against an “environmental” strategy that was the same for all places in the grid and did not change throughout generations. A generation consisted of every position on the grid playing against its neighbors and the environmental strategy. Identical agents were grouped into types and assigned an arbitrary type number. To calculate an agent’s score for a given grid position, a score matrix was utilized. In this score matrix, each agent type was given a unique identifier, the score that agent type i would obtain from play against agent type j was stored in position $x_{(i,j)}$ in the matrix. The score stored in the matrix was obtained by calculating the average score from 5 sessions of nIPD where each session consisted of 60 iterations of the prisoner’s dilemma game. The matrix was calculated on the fly and a new matrix was calculated for each generation. The score for a particular grid position consisted of the sum of the scores from the four neighbors plus the score against the environmental strategy. After each generation the grid was updated by copying the highest scoring neighbor into each grid position, if an agent outscored all of its neighbors, that position was not updated. When the grid was updated, there was a chance that a mutation would occur. Mutations happened on the average of one per generation. The relative frequency of the mutations was 2 Basic : 1 Contraction : 1 Growth. The grids were run for 10,000 generations and 30 grids were created for each of 7 environmental conditions.

In order to assess if non-local adaptation occurred, agents from generation 100 were placed in competition with agents from generation 10,000. Competition consisted of loading the right half of a grid with the right half of a grid saved from generation 100 and loading the left half of the grid with the left half of a grid saved from generation 10,000. Mutation was turned off and the grid was run for 100 generations. After the 100 generations, the relative numbers of agents from generation 100 and generation 10,000 were counted. If the number of agents from generation 10,000 was greater than those from generation 100 it was said that generation 10,000 dominated that competition. This competition was conducted for every pairing of the 30 archaic vs 30 modern generations yielding 900 trials. The result was a value indicating the proportion of the times that generation 10,000 dominated generation 100. The 99.5% confidence interval was calculated for this value using a normal approximation to the binomial distribution. If non-local adaptation did occur, it would be expected that the modern generation would consistently outcompete the archaic generation. If non-local adaptation did not occur there should be no significant difference between the performance of the modern and archaic generations. This would be seen as the 99.5% confidence interval overlapping the value 0.5.

The results gave strong support for the theory of non-local adaptation. For all of the environmental conditions the 99.5% confidence interval did not overlap 0.5. The evidence suggested that non-local adaptation does occur and that the phenomenon happens regardless of the environmental background. This study provides particularly strong support for non-local adaptation as most of the comparisons (all but 30 of the 900 competitions for each environmental condition) were between archaic and modern generations that had no common lineage. This means that these populations could not have evolved to specifically play against each other as they had never met until the competition.

Further Insight into Non-Local Adaptation

After the initial evidence supporting non-local adaptation, the next logical step was to see if the phenomenon appeared in other systems. Two studies looked at the phenomenon of non-local adaptation in different settings. The first study looked at a competitive robot painting task and the other looked at an artificial plant model called gridplants.

Paintbots and non-local adaptation

The first study looked at non-local adaptation using a competitive painting task.[6] The competitive painting task involved placing two virtual robots on a 12 x 12 toroidal grid. As each robot moved through the grid the spaces that they occupy are painted in that robot's color. The robots can paint over any space in the grid. A robot's score is the total number of spaces painted in the robot's color. Because a win consists of painting more spaces than its opponent, a robot gains 0 spaces over its opponent for painting over its own space, 1 space over its opponent for painting over a blank space, and 2 spaces over its opponent for painting over a space that is painted in its opponent's color. A fitness case consisted of placing two agents randomly on an unpainted grid, each agent can take 288 actions and the robots alternate actions. Each robot pairing consisted of running 5 fitness cases between the pairs.

The robots consisted of GP-Automata (essentially an augmented finite state machine). The input to the robots consisted of the color of the 8 grid squares immediately adjacent to the robot's position in the grid. The robot could perform one of four actions: left turn, right turn, forward, think. The think action is an immediate transition to the next state which is then executed, 8 consecutive think actions are allowed, and the fitness evaluation of the robot was terminated at the 9th consecutive think action. Each state had an associated "decider" function that returned a value. The decider consisted of a number of operators that were combined so that they would return an integer. Each state had two outbound transitions, one was used if the returned integer was odd and the other was used if the transition was even. Each transition had an associated action and an associated next state. There were 8 possible

types of mutation that were performed on the FSM: 10% changed the initial state, 10% changed the initial action, 20% changed the next state of a randomly selected transition with a new randomly selected state, 20% change the action of a randomly selected transition with a new randomly selected action, 10% replace a randomly selected decider with a randomly generated divider, 10% perform a crossover on two randomly selected deciders, 10% exchange two randomly selected deciders, and 10% copy one randomly selected decider over another randomly selected decider. The FSM is treated as a linear gene with each state being an atomic object, two point crossover is used in genetic recombination.

A total of 400 evolutionary runs were performed and the best performing robots were saved from generation 500 and 5,000. The robots from generation 500 and 5,000 were then placed in competition with 5 fitness cases per competition. In these fitness cases, half of the time the modern robot moved first and half of the time the archaic robot moved first. The average score was calculated for the archaic and modern robots as well as the 95% confidence interval. The result was that the modern robots performance was better than that of the archaic robots by more than 7 standard deviations. Again these results occurred when the majority of the objects competing had a distinct evolutionary history. This provided further support for non-local adaptation. Additionally, it showed that the phenomenon of non-local adaptation is not restricted to IPD and that the effect still exists in a model that allows for genetic recombination between parents rather than the asexual reproduction method in the IPD example.

Gridplants and non-local adaptation

The gridplants study looked at non-local adaptation in a plant based model.[7] In this model, virtual plants are placed on a 100 x 100 toroidal grid. Each plant comes from a seed that contains a program for growing on the grid. Plant growth occurs in discrete steps and occurs at a growth tip for the plant. The growth tip of the plant moves as the plant grows and can perform one of 5 actions which are: move up, move down, move left, move right, plant

seed. The growth tip of the plant cannot move into a cell occupied by another plant and a seed cannot be planted in a grid space that already contains a seed. Before the growth tip can perform an action it must read the plant's genome, this also counts as a turn in the game. Each action has an associated "energy cost". The cost to read the genome is 1, the cost to move the growth tip is 3, and the cost to plant a seed is 10. A seed starts with 4 units of energy, additional energy units are provided at a rate of 1 unit per grid space occupied per turn. A generation consists of 40 moves which allows for the growth tip to perform 20 actions (the other 20 are used to read the genome). If an action is not allowed the growth tip performs no actions for that turn and no energy is expended by the plant. Plant growth is performed using a turn-based approach, and the order in which the plants take their turns is randomized to eliminate positional bias. After all 40 steps have been completed all of the plants on the grid die leaving the spaces on the grid unoccupied. The next generation of gridplants then sprouts and grows from the seeds that were left by the previous generation. Of the seeds left by a generation, 10% of the seeds do not survive (are deleted from the grid). Unless mutated, a seed is an exact copy of its parent's genome. There is one mutation type that consists of randomly selecting a position in the seed genome and replacing the action with a new randomly selected action. In a given generation, 10% of the seeds will have a mutation event. For the experiment 30 populations of gridplants were evolved for 100,000 generations and seeds were saved from generation 10, 100, 1,000, 10,000 and 100,000.

Competition consisted of seeding a grid such that 5% of the grid squares, chosen at random, contained seeds selected randomly from the archaic seedbank and another 5% of the grid squares, again chosen randomly, contained seeds selected from the modern seedbank. Mutation was turned off, the grid was run for 50 generations and the relative number of grid squares occupied by archaic and modern gridplants were calculated. Each pairing generated two matrices where one matrix corresponded to the grid space occupied by the archaic generation and the other corresponded to the grid space occupied by the modern generation.

For the archaic matrix, position (x,y) corresponded to the amount of grid space that archaic population x occupied when it was compared with modern generation y . In the modern matrix, position (x,y) corresponded to the amount of grid space occupied by modern generation x when it was compared with archaic generation y . The diagonal $x = y$ was left blank so that only competitions between populations with no shared evolutionary history were considered for analysis. A Wilcoxon signed rank test was performed on the two matrices. The results from the test indicated that there was a significant difference between the results from the archaic and modern populations, the differences indicated that the modern populations outcompete the archaic generations.

In order to isolate the effect of cooperation between different genomes in a given generation the experiment was repeated with the difference that only one seed type from the archaic and modern seedbanks were used. This meant that only one genome from the archaic and modern populations was represented which eliminated the possibility of cooperation across genomes. Again the results indicated that the modern population significantly dominated the archaic generation. These results suggest that non-local adaptation also occurs in plant models. Since this model did not make use of FSM's these results also suggest that non-local adaptation is not restricted to agents that use FSM's. This study also had no artificial selection criteria for plant evolution. As in nature, the genotypes that produced "organisms" that grew and reproduced the best were the ones that would go on to other generations. This characteristic of the study indicates that an artificial (human imposed) selection criteria is not responsible for the observed phenomenon of non-local adaptation.

Project Objectives

Despite the observation of non-local adaptation in multiple models, questions still remain about the nature and robustness of this phenomenon. The goal of this project is to obtain a deeper understanding of non-local adaptation. One of the specific goals is to obtain a better picture of non-local adaptation throughout the first 10,000 generations of an evolving

population. Another goal is to further generalize the phenomenon of non-local adaptation by adding a rudimentary ability to recognize other agents to the model. A third goal of the project is to observe the effect of overall mutation rate on non-local adaptation throughout the first 10,000 generations of a population. In pursuit of these goals, the IPD model of non-local adaptation used by Ashlock and Mayfield is used as a foundation.

2. Materials and Methods

This project makes use of the IPD game introduced earlier in this text. As previously mentioned, in this game each player can choose one of two moves, cooperate or defect. Both players make their move simultaneously. If both agents cooperate they both receive a score of C , if both defect they both receive a score of D . If one agent cooperates and the other defects, the cooperating agent receives a score of L and the defector receives a score of H . Throughout this project the common values $L = 0$, $D = 1$, $C = 3$, and $H = 5$ are used. In iterated prisoner's dilemma (IPD) the same two players face each other for multiple rounds of the game. These multiple iterations allow for complex play to develop. In IPD (without noise) a strategy known as tit-for-tat emerges as a particularly good solution. To eliminate this preference we introduce a small chance ($\alpha = 0.01$) that an agent's move will be replaced by the opposite move. With the addition of noise, tit-for-tat is transformed from an excellent strategy to one that is no better than totally random play. This is called noisy iterated prisoner's dilemma (nIPD).

Agent Specifications

An agent consists of an identifying tag, an array for handling opponent tag information, and a finite state machine (FSM) adapted to play nIPD (See Figure 2). The tag value is a number that serves as a label that is exchanged at the beginning of play, there is no explicit attempt to give meaning to the tags (i.e. associating them with a particular FSM arrangement) rather the tags are allowed to take whatever meaning that may arise through evolution in the population. Agent play begins by looking up the appropriate initial action and state in the tag array indexed by the opponent's tag number. The agents then make their initial move and set their current state. Each state has two outbound transitions, one for each of the opponent's possible moves. Each transition gives the action the agent should take for the next round of play and points to the next internal state. If, for example, two agents with the FSM pictured in Figure 2, one with a tag value of 1 (agent a) and the other with a tag

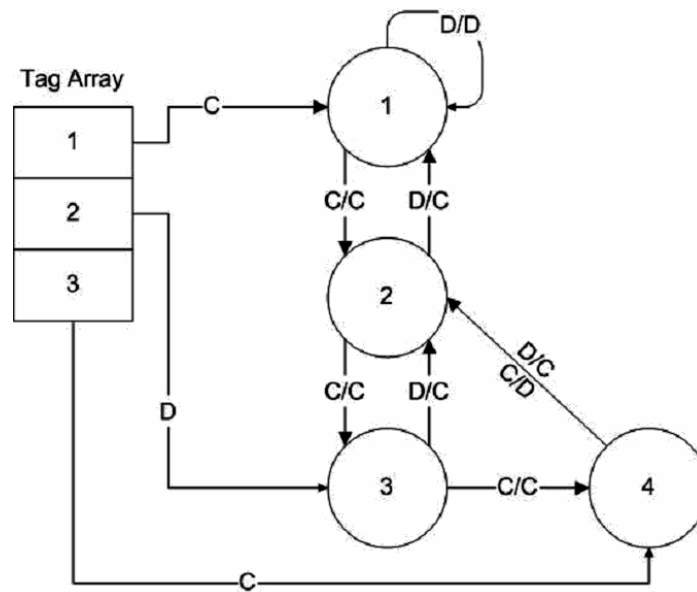


Figure 2 Example diagram of agent adapted for tag recognition. The array is indexed by opponent's tag value. The transition labeling follows the same convention as Figure 1

value of 2 (agent b), were to compete with each other play would go like this: The agents would read each others tags, set their initial state and make their first move. Play would be:

Agent a			Agent b		
Score	State	Move	Score	State	Move
0	3	D	0	1	C
5	4	C	0	1	D
5	2	C	5	2	C
8	3	C	8	3	C
11	4	C	11	4	C

Agent growth can happen through a number of mutations. A growth mutation consists of randomly selecting a state in the FSM, copying it, and randomly changing some of the in-bound transitions of the original state to the new state. A contraction mutation consists of selecting a state at random and deleting it. In-bound connections to the deleted state are assigned randomly to the remaining states in the FSM. A basic mutation consists of randomly selecting a transition and then changing either the state that it points to or the agent's action for that transition. A tag mutation consists of changing the agent's tag to a

new randomly selected value. A tag array mutation involves selecting one of the array positions and changing either the initial action or changing initial state to a state selected a random from the FSM. The rates of mutation used for individual experiments are given in the descriptions of the experiments later in this paper.

Grid Play

The world consists of a 100x100 toroid shaped grid. The grid is initially populated by randomly generating 16 agents and filling the grid with randomly selected copies of these agents. Agents play their neighbors to their North, South, East, and West. A generation consists of all grid positions playing their neighbors followed by a selection step. The score for an agent in a particular position in the grid consists of the sum of the scores obtained from playing all neighbors. Rather than playing all 1,200,000 pair wise games of IPD needed for a generation, the score for a pairing is computed using a dynamically generated score matrix where position $x_{(i,j)}$ is the average score obtained by agent(i) playing against agent(j) in 5 sessions of nIPD where each session consists of 60 rounds of prisoner's dilemma. A new matrix is generated for each generation and the matrix is indexed by the agents present in that given generation. After a generation of play each position in the grid is updated by copying the highest scoring agent for the neighborhood into the grid position or leaving the agent currently occupying the space in place if it outscored all of its neighbors. During the update there is a small chance (1 in 10,000 unless otherwise noted) that a mutation will occur. For each condition 30 populations are generated, each population is allowed to evolve for 10,000 generations, the grid is saved every 500 generations.

Competition

A competition grid is created by loading one half of a grid with the corresponding portion of the saved grid from the archaic generation and the other half with the corresponding portion of the saved grid from the modern generation. Mutation is turned off and the grid is run for 100 generations. The number of modern and archaic agents are

counted, if the number of modern agents is greater than the number of archaic agents it is counted as a win for the modern population. This competition is performed for every pairing of the saved modern and archaic populations, resulting in 900 trials. From this we can calculate win percentage of the modern generation. A 99% confidence interval is then calculated for the win percentage using a normal approximation to the binomial distribution.

[8]

As one of the major objectives of this project is to compare and contrast non-local adaptation with and without tags, there was a need to generate populations without tags. For this project, populations without tags were accomplished by setting the maximum label value to 1. With this setting the agents are functionally identical to the tagless agents in the paper by Ashlock and Mayfield[5]. This similarity is further confirmed by the similar results achieved for the tagless populations between the results for this project and the previous study.

3. Experimental Design

This project consisted of three different experiments. The first took a closer look at non-local adaptation by looking at more pairs of distinct generations than the previous nIPD study. This initial experiment also included a run where tags were introduced at low mutation levels to get an initial feel for the effect of tags on non-local adaptation. The next experiment varied the rate of mutation of both agent tags and their tag arrays to determine what effect this had on non-local adaptation. The final experiment varied the total mutation rate for populations with no tags to identify what effect the overall mutation rate had on non-local adaptation.

A Higher Resolution Look at Non-Local Adaptation

In order to obtain a closer look at non-local adaptation we repeated the “null environment” experiment from the previous study at a higher rate of sampling. In this run the only parameter that was different was the frequency with which the grid was saved. As mentioned previously, for this experiment every 500th generation was saved yielding a measurement of non-local adaptation for every 500th generation. For the agents without tags the relative mutation levels were 50 basic : 25 growth : 25 contraction.

A run with tags was performed in addition to the run without tags. For this run we selected an arbitrarily small tag mutation rate where the relative mutation rates were 40 basic : 25 growth : 25 contraction : 5 tag : 5 tag array. The maximum tag value was 20, meaning that agents could have a tag value from 1-20. Tag values and tag arrays were generated randomly when the agents were created. Again the populations were allowed to evolve for 10,000 generations and every 500th generation was saved and placed in competition against generation 10,000. For both data sets 99% confidence intervals were calculated using the previously described methods.

A Survey of Different Tag and Tag Array Mutation Levels

A better picture of the effect of tag mutation levels on non-local adaptation was generated by varying both tag mutation levels and tag array mutation levels over three levels. With this survey of tag related mutation rates, our goal was to identify what impact different combinations of these rates had different effects on non-local adaptation occurring in the population.

To accomplish this comparison we performed nine data runs that differed only in the tag and tag array mutation levels. The three different levels used were 10 (low), 25 (medium), 50 (high). These values were chosen as they represented a relatively broad range of frequencies from less frequent than the FSM mutations (growth, basic and contraction) to higher than the FSM mutations. As before, every 500th generation was placed in competition against generation 10,000, the percentage of modern “wins” was recorded, and the 99% confidence interval was calculated.

In order to get a better picture of how tags were being used by the populations, we calculated the number of agents displaying each particular tag for every 100th generation. This data shows if there is a large diversity of tags in a given grid or if most of the agents tend to have the same tag.

A Look at Overall Mutation Rate

As the relative frequency of tag mutations increases, the overall frequency of FSM mutation decreases. This occurs in our experiments because we change the relative frequencies of the mutation types without changing the overall mutation rate. For example, if the relative rate of tag mutations equals the relative rate of FSM mutations, the result is that the rate of FSM mutations equals half that of the overall mutation rate. We performed an experiment to determine what the effect of changing the overall mutation rate was. For this experiment two additional data runs were performed. These runs were identical to the “notags” data run with the exception of the overall mutation rate. The low mutation rate was

set at 0.1x the rate of the run without tags, this meant that there was roughly 1 mutation event every 10 generations. The high mutation rate was set at 10x the rate of the run without tags. This meant that there were roughly 10 mutations per generation. Again every 500th generation was placed in competition with generation 10,000, the percent modern generation wins were calculated, and the 99% confidence interval was generated.

4. Results

The key test in this project was the competition between the archaic and modern generations. This competition was run for a number of populations with and without tags and with varying mutation rates. Figure 3 is a graph that shows the frequency with which the modern generation dominated the archaic generation for the initial tags/notags data runs. For all of these graphs, the x-axis represents the archaic generation that was placed in competition with generation 10,000. The y-axis represents the proportion of the 900 trials where the modern generation dominated the archaic generation. The error bars represent the 99% confidence interval for the values. Figure 4 holds the results for the 3x3 survey of tag mutation levels. The naming convention for the mutation levels is tag level + tag array level. For instance, the name medhig represents the population with the tag mutation level of 25 (medium) and the tag array mutation level of 50 (high). In order to increase readability, the results have been split into three graphs, labeled A-C. All three graphs follow the convention of Figure 3. Figure 4A contains the results from the populations where the tag mutation rate was 'low'. The results from the experiment without tags are included for reference. Figure 4B is similar to Figure 4A with the only difference being that the results are from the populations with the medium level of tag mutation, Figure 4C displays the results from the populations with the high level of tag mutation. Figure 5 is a graph of the results from the data runs where the overall mutation rate was changed. This graph follows the conventions of Figure 3 with the x-axis representing the archaic generation and the y-axis representing the proportion of the competitions that the modern generation dominated the archaic generation.




Figure 3 Graph showing the proportion of 900 trials where the modern generation (10,000) dominated the archaic generation (x-axis) for populations with tags and without tags. Error bars indicate 99% confidence interval.




Figure 4a Graphs of the data generated in the 3x3 tag mutation study. This graph represents the results for the three experiments where the tag mutation was at the lowest level. The results from the agents without tags are included as a reference.




Figure 4b Graphs of the data generated in the 3x3 tag mutation study. This graph is from the medium tag mutation level. The results from the agents without tags are included as a reference.




Figure 4c Graphs of the data generated in the 3x3 tag mutation study. This graph is from the high tag mutation level. The results from the agents without tags are included as a reference.




Figure 5 Graph showing the proportion of trials there the modern generation (10,000) dominated the archaic generation (x-axis) for three tagless populations at different overall mutation levels.

5. Discussion

In the competitions the vast majority of the comparisons are between archaic and modern generations that come from different lineages. This means that these populations have different evolutionary history so they have not evolved to play against each other. If the theory that evolution only yields specific adaptations to local challenges is correct, we would expect the modern generation to have no competitive advantage over the archaic generation. The null hypothesis, that non-local adaptation does not occur, is confirmed if the 99% confidence interval of the proportion of modern wins over the archaic generation includes the value 0.5. If the 99% confidence interval does not overlap 0.5, we can reject the null hypothesis in favor of the alternative hypothesis that non-local adaptation is occurring.

Non-Local Adaptation Over Time

The previous study tested generation 100 against generation 10,000. This test found that the modern generation dominated the archaic generation a significant proportion of the time. Although these results are significant, one still wonders if non-local adaptation can be seen at later points in evolution or if non-local adaptation is something restricted to a relatively “young” evolving population.

The data from this study suggest that non-local adaptation occurs throughout the entire 10,000 generations. In the run which replicates the previous study, the 99% confidence interval of the probability of modern victory never overlaps with 0.5. This implies that non-local adaptation is occurring throughout the entire data run. Most impressively, the competition between generation 9,500 and generation 10,000 still proves to be significantly different from 0.5 (0.5356 ± 0.0353). This implies that even at generation 9,500 the population is still developing its general ability to play IPD.

When tags are added to the mix, a small shift appears after generation 3,500 (See Figure 3). Throughout the entire run there are few significant differences between the run with tags and without, however, there is a trend in that the run with tags consistently shows

the modern generation dominating the archaic generation in a smaller proportion of the 900 comparisons. This trend becomes more pronounced after generation 3,500. Despite this difference, the majority of the values are not significantly different between the run with tags and the run without tags as determined by the overlapping confidence intervals. Based on this data, it appears that tags have little to no effect on non-local adaptation. This result is somewhat surprising as one would surmise that the addition of tags would make the task more complex. This added complexity should give the agents a larger strategy space to explore, which should lead for more opportunity for non-local adaptation.

A Survey of Tag and Tag Array Mutation Levels

One potential reason for the minimal effect of tags could have been their relatively low mutation level in the initial experiment. To address this possibility, we performed a 3x3 survey of tag and tag array mutation levels. With this type of an approach it is possible to see the effect of increasing tag related mutation levels as well as any interaction that may result from different relative tag and tag array mutation levels.

An initial look at the data in Figure 4 shows that for most combinations there was little effect on non-local adaptation. One condition, lowmed, showed a consistently higher proportion of dominance, however, many of the data points are not statistically significant. In terms of significance, one condition that jumps out is the medhig condition. For this set of mutation rates, the proportion of dominance by the modern generation is consistently and significantly lower than the results for the data set without tags. The medmed condition gives a similar performance, however, there is a stretch without significance around generation 4,500. Despite these individual results, there is no apparent trend based on relative combinations of tag and tag array mutation levels. Regardless of the combination, there is little difference between populations with tags and without tags.

In light of these results, we decided to take a look at tag usage as these populations evolved. Our initial expectation was that most of the tag space would be occupied by agents.

The data show the opposite. For all conditions the agents tend to all have the same tag. This means that there is only selective pressure for the portion of the FSM that deals with the current tag. Eventually, some agent's tag will mutate and the resultant mutant will successfully take over the world. It may well be able to do this because the other agents have no strategy for dealing with this new tag. The result is a population locked into a single tag value with punctuated changes in the tag value when a mutated agent is able to fill the open niche by mutating to a tag value that the rest of the population can not adequately defend against. This invasion and takeover may slow the rate of non-local adaptation to the nIPD task by creating periods of time when the selective pressure is reduced by a lack of strong "defensive" strategies to this new label.

When looking at the mutation levels that produced interesting results in the 3x3 survey, all of these conditions exhibited this tag cluster phenomenon. That is, most of the time the entire population had the same tag value. Occasionally, the population would exhibit multiple tag values during the time when an opportunistic agent was taking over the world, then the population would settle back down to predominantly one tag. Figure 6 shows some typical populations. In this graph, the x-axis represents the generation, the y-axis represents a simple measure of the amount of the tag space occupied. To calculate this metric, the total number of agents occupying each tag value is divided by number of agents that occupy the dominant tag value. These values are then summed to arrive at a single value for every 100th generation. For example, if there were 4 possible tag values and the distribution of 100 agents within the tag space was: 1:0, 2:50, 3:50, 4:0. The measure of tag space utilization would be $(0/50 + 50/50 + 50/50 + 0/50) = 2$. With this metric, values close to 1 represent a population with very little diversity of tags, a value of 20 would represent a population evenly occupying the entire tag space.

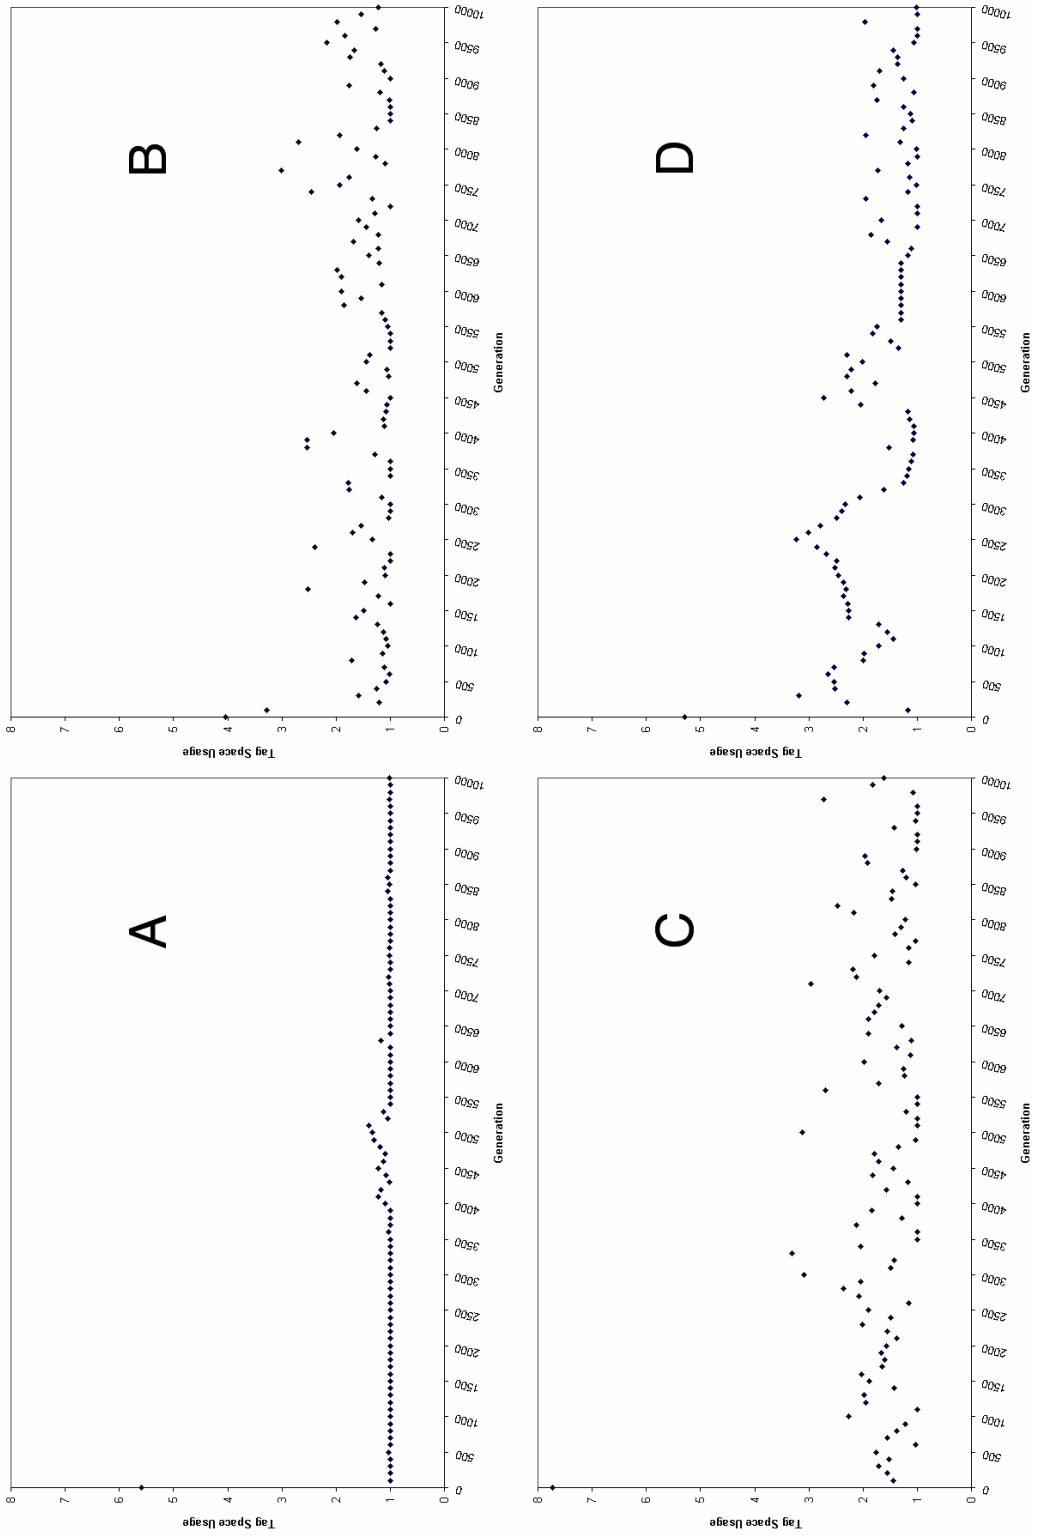


Figure 6 Graph showing tag space usage for four different populations in the “medmed” experiment. Tag space usage of 1 indicates that only one tag value is being used by the population. The maximum value for tag space usage is 20.

The examples in Figure 6 are taken from the medmed mutation level. In this figure, Population A demonstrates a population that consistently occupies only one tag value at a time. When a new tag value takes over, the transition is fast and the population quickly returns to one dominant tag value. Population B is a population that has a consistent, but extremely small level of other populations but again one tag is typically dominant. In this case opportunistic agents seem to occur more frequently. These opportunistic agents prevent the population from completely settling on a predominant tag level. Despite this, there is still one dominant tag throughout the evolutionary history of the population. Population C is a variation on the same theme as population B. The only difference is that there are a couple of periods when there are multiple opportunistic agents at the same time, these periods do not last for a long time and the population quickly settles down. Population D is an interesting example where there are fairly long periods of time where more than one tag is being used. This happens early in the population, but eventually a single tag prevails. Although interesting, the behavior of population D is quite rare and even when it does exist, it eventually comes back to the tag clustering phenomenon.

This kind of a pattern can be seen in nature. The immune system utilizes surface proteins as tags to identify and attack foreign organisms. Some organisms will change their surface proteins so that they can evade the immune response. One specific example of this can be seen in the malaria parasite. Strains of malaria, i.e. *Plasmodium falciparum*, have a set of genes called the *var* gene family that encode surface proteins. This particular parasite eludes the immune system by changing *var* gene expression which changes the proteins being displayed to the immune system, allowing it to stay ahead of the immune response. [9] In the population with tags it is easy to imagine that a similar phenomenon is occurring. When the entire population has the same tag, there is selective pressure to develop a response to the prevalent tag. Eventually a mutation develops where the tag of an agent is changed.

Suddenly the entire population has no defense against this new tag. The agent then takes over the world by virtue of the population having no adequate defense.

Overall Mutation Rate

As mentioned before, the introduction of tag related mutations with no corresponding increase in the overall mutation rate is similar to reducing the rate of FSM based mutations. We performed a several data runs to isolate the effect of this decreased FSM mutation rate on the populations. Figure 5 shows these results.

When looking at the graph we can see that increasing the mutation rate by a factor of 10 had no significant effect on non-local adaptation in the population. Interestingly, lowering the mutation rate by a factor of ten only had an effect on the initial rate of non-local adaptation, however, it eventually merges with the other two data sets. This data seems to indicate that the overall mutation rate primarily has an effect in the early phase of an evolving population. After about 3,000 generations the effect of the overall mutation rate seems to fall off.

6. Conclusions

Non-local adaptation seems to be a fairly resilient quality of an evolving population. As one might expect, populations acquire non-local adaptive features rapidly while they are evolutionary young. This project suggests that populations are still developing non-local adaptive features at later evolutionary times, as late as generation 10,000. Despite altering a number of variables (tags/notags, tag mutation rates, overall mutation rates) we do not see any consistent and significant impact on non-local adaptation in these populations. The vast majority of significant differences occur in the initial generations and generally disappear after generation 3,500. One trend that can be seen in the data is that it becomes harder and harder to detect non-local adaptation as the populations evolve. Specifically, after generation 3,500 it becomes nearly impossible to detect the differences. This result is not unexpected as it is easier for new mutations to enhance IPD play in the young initial populations rather than in the more evolved populations. An analogy would be the acquisition of language by people. Infants learning to talk will rapidly improve, however, the progress continually slows. If one were to arrange a competition where individuals of differing ages were given vocabulary tests and a win consisted of whichever person had the largest vocabulary it would not be unreasonable to expect the same kind of curve that we see in the data. If these individuals had very similar backgrounds, we would expect there to be an extremely large difference between a one year old and a 20 year old, as a result the 20 year old population would totally dominate the 1 year old population. The gap would continue to close until we were comparing 19 year olds with 20 year olds. As people are continually expanding their vocabulary throughout their lives, we could expect there to be some difference between the 19 and 20 year old populations, however, we would need very sensitive tools (in this example vocabulary tests) to pick up this difference. Also, we would need increasingly larger samples to pick up the difference. We can expect a similar situation to occur with any individual or population learning a new task, be it dolphins learning to interact with people or

on the population level like humans learning how to build buildings or perform scientific experiments. When a population or an organism is new to task very few changes will decrease performance and many will result in some improvement, however slight it may be. In this project, the result is that a mutation is more likely to result in a better strategy early in evolution rather than later. Consistent with this observation, the data from this project indicate that changes in the overall rate of mutation exhibit most of their effect early in the evolutionary process. The increase in mutation rate gives the population the ability to explore the strategy space more quickly, this faster exploration seems to result in the population finding better strategies at a faster rate. As the population becomes more mature, the data tends to converge as mutations are more likely to be harmful than helpful than early in the population.

A surprise from the project is the use of tag space. Although we expected the tag space to be uniformly used, instead the space is used in a pathogen-like manner. This phenomenon seems to slow non-local adaptation. Although more study on the phenomenon is needed, it appears that this slowdown is the result of invasion by opportunistic agents. When the invasion occurs, it appears to act as a pause button on non-local adaptation as the opportunistic agent is taking advantage of an available “tag niche”.

Future Work

Although this project has provided some insight into the phenomenon of non-local adaptation, questions still remain. This project, in combination with the previous studies, has shown non-local adaptation to be a resilient and robust phenomenon. Although it is interesting to see all of the conditions where non-local adaptation occurs, it would be helpful to find situations where the phenomenon does not occur. With this information one could begin to identify the conditions necessary for non-local adaptation to occur. As the requirements for non-local adaptation become known, these requirements can be compared with what is known about biological evolution. If these conditions are met by biological

organisms, then it would be possible that living organisms exhibit non-local adaptation. The result would be a change in current evolutionary theory. Specific boundaries to explore include task complexity, agent complexity, and overall mutation rate. As task complexity decreases, the number of possible strategies will also decrease. With this reduced strategy space one would expect agents to hit a non-local adaptation ceiling where they had completely explored the strategy space and achieved maximum ability to play the game. Once the ceiling is hit, only local response to changes in population (local adaptation) would be left. Similarly, agent complexity should have an effect on non-local adaptation. If agents are restricted to only one state, the agents can only exhibit four strategies: always cooperate, always defect, tit-for-tat, and anti tit-for-tat. Although extremely limited, this example serves to show that reduced agent complexity reduces the strategy space available to the evolving population. Again, this should limit the amount of non-local adaptation that is possible within the evolving population.

This project made a close examination of the first 10,000 generations of an evolving population. In the first 10,000 generations non-local adaptation was observed, however, this study did not attempt to look at the duration of the phenomenon of non-local adaptation. By running the experiments for a longer time and altering various parameters in the simulation it would be possible to determine what factors have an effect on the length of time that non-local adaptation can be detected in an evolving population. Some factors that might increase the duration of the phenomenon may include task and agent complexity as well as overall mutation rate. Increased task and agent complexity should result in a larger strategy space. This larger strategy space may provide greater room for further non-local adaptation. A lower overall mutation rate may result in slower exploration of the strategy space. This slower exploration may prolong the phenomenon of non-local adaptation.

Another area of interest is the use of tag space by the evolving populations. In this project, the populations used a very small part of the available tag space. Adjusting

parameters to observe the effect on the use of tag space would be interesting. Changing the number of tags available may have an effect on the use of tag space in the population. At the very least, reducing the number of tags available would increase the proportion of the tag space used by the population as 1 is a greater proportion of 5 than 20. Despite this obvious change, reducing the overall tag space would also mean that the evolution of effective strategies against all possible tag values would be more likely.

Despite the length of time that has passed since Darwin first published *Origin of the Species* there are still questions about the specifics of evolution. This project took a closer look at one area of controversy, that of non-local adaptation. The evidence suggests that non-local adaptation is a real phenomenon and that it is rather robust.

Acknowledgements

I would like to thank NewLink Genetics for use of their LINUX cluster in support of this research. This work was supported by NSF IGERT grant DGE-9972653

Works Cited

- [1] Dawkins, R., *The Blind Watchmaker*, W.W. Norton & Company Ltd, 1986.
- [2] Iowa Lottery, <http://www.ialottery.com/Games/Online/Powerball.html>
- [3] Davis, Morton D., *Game Theory A Nontechnical Introduction*, Basic Books, New York, 1973.
- [4] Axelrod, R., *The Evolution of Cooperation*, Basic Books, New York, 1984.
- [5] Ashlock, D., Mayfield, J., “Acquisition of General Adaptive Features by Evolution”, *Conference on Evolutionary Programming*, March 1998.
- [6] Ashlock, D., Blankenship, E., Gandrud, J., “A Note on General Adaptation in Populations of Painting Robots”, *Congress on Evolutionary Computation*, 2003.
- [7] Doty, D., “Non-local Evolutionary Adaptation in Gridplants”, *Submitted to Congress on Evolutionary Computation*, 2004
- [8] Larsen, R., Marx, M., editors, *An Introduction to Mathematical Statistics and its Applications*, Prentice-Hall Inc., 1981.
- [9] Flick, K., Chen, Q., “var genes, PfEMP1 and the human host”, *Molecular and Biochemical Parasitology*, vol 134, pp. 3-9.