Using NEAT to Evolve Cooperation and Intelligence in Game Theory Problems

Angela Mastrianni and Lucas Franca do Nascimento Farias

December 18th, 2015

Abstract

The goal of our study was to see if neural networks would cooperate with each other in two different game theory problems. We replicated a study entitled “Cooperation and the Evolution of Intelligence” by Luke McNally, Sam P. Brown and Andrew L. Jackson, however we evolved the neural networks with NEAT instead of a genetic algorithm. The neural networks would play every other individual in the population in the two game theory problems, choosing to either cooperate or defect and receiving a payoff based on their action and the action of their opponent. We evolved 10 runs of each problem to 10,000 generations, measuring the frequency of cooperation in the each generation. One of the game theory problems, known as the Iterative Snowdrift Dilemma, by the virtue of its structure, encouraged cooperation more than the other game theory problem, the Iterative Prisoners’ Dilemma. Our hypothesis was the neural networks evolved with the Snowdrift Dilemma would have higher levels of cooperation than neural networks evolved with the Prisoners’ Dilemma. Our results supported the hypothesis and led to some insights into the behavioral strategies of the neural networks and the different number of neurons in the networks in the different problems.

1 Introduction

Cooperation between artificial neural networks is an expanding area of research in cognitive science and artificial intelligence. Chern Han Yong and Risto Miikkulainen [6] researched if neural networks could work together to capture prey. They found that the neural networks would take on different roles in order to successfully capture the prey. Mitchell Potter, Lisa Meeden and Alan Schultz [3] found that neural networks would cooperate in order to herd sheep and protect them from the foxes that were trying to harm them. In both of these scenarios, all participants benefit from cooperation as it allows them to achieve their maximum goal and payoff.

Researchers Luke McNally, Sam P. Brown and Andrew L. Jackson wanted to discover what would happen to cooperation between individuals when they participate in a task where cooperation may not always achieve the maximum payoff. To study this, they evolved neural networks under a genetic algorithm and had these networks participate in two different social dilemmas, which were represented as game theory problems. One of the social dilemmas, through its structure, encouraged cooperation more than the other. The researchers looked at the frequency of cooperation and the different behavioral strategies that emerged as the intelligence of the neural networks increased as they participated in the social dilemmas. The researchers defined intelligence as the number of neurons in the neural network saying that a neural network with more neurons would be considered more intelligent. For example, a neural network with 8 nodes would be more intelligent than a neural network with 5 nodes. In the social dilemma that encouraged cooperation less, the Iterative Prisoners' Dilemma, they found that as intelligence increased, cooperation decreased as the neural networks entered in "arms races" with one another, trying to defeat the other individuals in the population. In the social dilemma that encouraged cooperation more, the Iterative Snowdrift Dilemma, cooperation slightly increased as intelligence increased [2]. The goal of our project is to replicate the results of this study using the NeuroEvolution of Augmenting Topologies (NEAT) to evolve the neural networks. Our hypothesis is that we will achieve similar results when evolving with NEAT– a lower level of cooperation in the Iterative Prisoners' Dilemma and a higher level of cooperation in the Iterative Snowdrift Dilemmaa.

1.1 Game Theory Problems

There are two different game theory problems used in the study, the Iterative Prisoners' Dilemma (IPD) and the Iterative Snowdrift Dilemma (ISD). Both problems are two player games, where the players can choose to cooperate or defect. In the Iterative Prisoners' Dilemma, two individuals can choose to either remain silent (cooperate) or rat out the other individual for the crime (defect). In the Snowdrift dilemma, the individuals can either help to shovel out the snowbank (cooperate) or do nothing (defect). In both scenarios, the best case for a player is that they defect and their opponent cooperates, as this will result in the best payoff. The second best scenario is that they both cooperate. However, the worst case scenarios differ in the two games. In IPD, the worst case for a player occurs when they cooperate and their opponent defects. On the other hand, the worst case scenario in ISD is that both players defect. Due to these differences, ISD encourages cooperation slightly more than IPD since mutual defection is the worst case scenario for the players in ISD [2]. To make the differences in payoff clear, figures 1 and 2 show the payoff grids for IPD and ISD respectively.

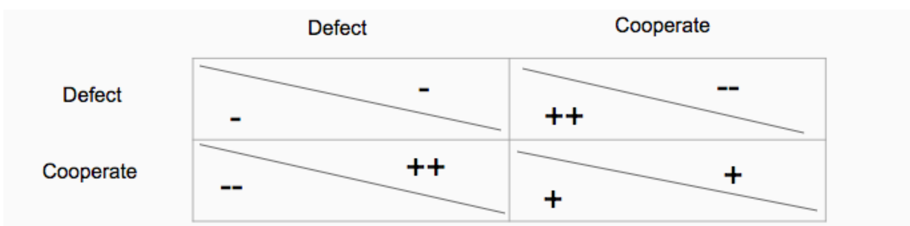

Figure 1: The payoff grid for the Iterative Prisoners' Dilemma

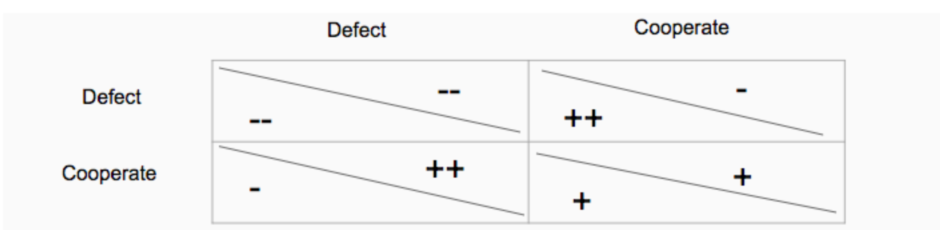

Figure 2: The payoff grid for the Iterative Snowdrift Dilemma.

1.2 Behavioral Strategies

The researchers divided the behaviors of the neural networks in the social dilemmas into four different strategies– Always Defect, Always Cooperate, Tit-for-Tat and Pavlovian. In Always Defect, the neural networks always chooses to defect, no matter what their opponent has chosen to do in the previous round. Similarly, in Always Cooperate, the neural network always chooses to cooperate. In Tit-for-Tat, the network does the same move that its opponent made in the previous round. For example, if its opponent defected in the previous round, the Tit-for-Tat neural network would defect in the next round. In Pavlovian behavioral strategies, networks repeat an action if it achieves a score within a certain threshold. A Pavlovian neural network will respond to mutual defection or mutual cooperation with cooperation and to other combinations of moves with defection [2].

1.3 Evolving Neural Networks with NEAT

Artificial neural networks have been used to represent neural networks in the human brain. According to Thomas Schultz, an artificial network is a "set of units and connection weights organized in a particular topology" [4, pg.15]. The neural networks have a layer of input node(s), a variable number of layers of hidden node(s) and a layer of output node(s). Input nodes take in information from the environment while output nodes return values back to the environment. Hidden nodes do not interact with the outside environment, only communicating with other hidden nodes and the input and output nodes. The connection weights between the nodes are not set but instead can change as the networks evolves [4].

One method of evolving neural networks is through the NeuroEvolution of Augment Topologies (NEAT). NEAT implements complexification, allowing the genome to start simply and complexify as needed. Neural networks evolved with NEAT, which was created by Kenneth Stanley and Risto Miikkulainen, become more complex over time as hidden nodes and connections are added to the existing topology [5]. Our project evolves neural networks using the MultiNeat library, which was created by Peter Chervenski and Shane Ryan as a generic NEAT implementation [1].

2 Experiments

To test our hypothesis, we ran 10 experiments of the Iterative Prisoner's Dilemma and 10 experiments of Iterative Snowdrift Dilemma, evolving 10,000 generations of neural networks. Each game would have a population of fifty neural networks that would play every other individual for fifty rounds of the dilemma. In each round, the individual would get a payoff score depending on their action and their opponent's action. We used the same payoff scores as the study, which can be found in Table 1 for IPD and Table 2 for ISD, normalizing them to one for NEAT. The individual's fitness score would be their mean payoff score over all rounds against all other individuals in the population. The neural networks were evolved using MultiNEAT and the parameter settings used in MultiNEAT are located in Table 3. The neural networks had three input nodes: the individual's payoff score from the previous round, their opponent's payoff score from the previous round, and a bias. They also had one output node which was the probability that the individual would cooperate in that round. Each network would begin with one hidden node and neural networks could be recurrent.

We also wanted to compare the behavioral strategies of the neural networks in each generation to the closest pure strategy (Always Cooperate, Always Defect, Tit-For-Tat, Pavolvian). To accomplish this, we followed the same strategy for classification as the study. We had each pure strategy play against itself and all other pure strategies, coding each move with a 1 for cooperation and a 0 for defection. We also had each network play against all five strategies, coding the moves in the similar manner. We than compared that sequence of moves to the sequences of moves between each pure strategies with itself and the other pure strategies. The network would then be assigned to the behavioral strategy that minimized the sum of squares between the sequences.

Moves	**Payoff**
Mutual Cooperation	6
Mutual Defection	2
Player Cooperates & Opponent Defects	1
Player Defects & Opponent Cooperates	7

Table 1: Payoff Scores for IPD

Moves	**Payoff**
Mutual Cooperation	5
Mutual Defection	1
Player Cooperates & Opponent Defects	2
Player Defects & Opponent Cooperates	8

Table 2: Payoff Scores for ISD

Parameter	**Setting**
Generations	10,000
Rounds	50
Population size	50
Input nodes	3
Output nodes	1
Prob. for recurrent link	0.3
Prob. to remove link	0.3
Prob. to add link	0.1
Prob. to add node	0.05
Elitism	0

Table 3: MultiNEAT parameter settings used in the experiments.

3 Results

Tables 4 and 5 show the average frequency of cooperation across all 10,000 generations for 10 runs of IPD and ISD. The average frequency of cooperation over the ten runs was 2.58% in IPD while the average frequency of cooperation over the ten runs of ISD was 23.89%, as highlighted by the tables. Figure 3 shows graphs of the frequency of cooperation over the 10,000 generations for a run of IPD and a run of ISD. As shown by the tables and graphs, the frequency of cooperation among the neural networks was much higher in ISD than in IPD. The graph in Figure 3 also shows that cooperation levels remain constant in ISD, as they did in the original study. The average frequencies of each behavioral strategy over the generations for the ten runs of IPD is shown in Table 6 while Table 7 shows the the average behavioral frequencies for 10 runs of ISD. IPD had a higher average frequency of Always Defect (98.45%) while ISD had a higher average frequency of Always Cooperate (6.91%), Tit-For-Tat (2.04%) and Pavlovian (1.62%). Figure 4 shows the frequencies of the behavioral strategies over the generations of a run of IPD and the frequencies of the behavioral strategies in a run of ISD. These tables and figures highlight how neural networks in ISD displayed a greater diversity of behavioral strategies than neural networks in IPD.

Run	**Average Freq. of Cooperation**
1	2.06%
2	2.23%
3	2.36%
4	2.31%
5	3.26%
6	3.21%
7	3.60%
8	1.80%
9	2.36%
10	2.58%
Average	2.58%

Table 4: Average Frequencies of Cooperation in Runs of IPD

Run	**Average Freq. of Cooperation**
1	23.67%
2	23.64%
3	24.76%
4	23.83%
5	24.46%
6	23.45%
7	23.89%
8	22.99%
9	24.83%
10	23.41%
Average	23.89%

Table 5: Average Frequencies of Cooperation in Runs of ISD

Run	**Always Cooperate**	**Always Defect**	**Tit for Tat**	**Pavlovian**
1	0.69%	98.99%	0.17%	0.14%
2	0.91%	98.68%	0.23%	0.17%
3	0.86%	98.68%	0.31%	0.14%
4	0.84%	98.75%	0.23%	0.18%
5	1.64%	97.91%	0.24%	0.21%
6	1.62%	97.92%	0.24%	0.22%
7	2.10%	97.40%	0.27%	0.25%
8	0.54%	99.21%	0.14%	0.11%
9	1.03%	98.56%	0.22%	0.19%
10	1.06%	98.41%	0.28%	0.26%
Average	1.13%	98.45%	0.23%	0.19%

Table 6: Average Frequencies of Each Behavioral Strategy in Runs of IPD

Run	**Always Cooperate**	**Always Defect**	**Tit for Tat**	**Pavlovian**
1	7.01%	89.04%	2.21%	1.75%
2	8.19%	87.68%	2.25%	1.86%
3	6.13%	90.22%	2.07%	1.57%
4	6.97%	89.64%	1.93%	1.56%
5	6.19%	89.83%	2.26%	1.71%
6	7.01%	89.31%	2.03%	1.58%
7	6.84%	90.27%	1.58%	1.31%
8	6.78%	89.78%	1.94%	1.52%
9	6.49%	89.73%	2.09%	1.68%
10	7.44%	88.74%	2.08%	1.69%
Average	6.91%	89.42%	2.04%	1.62%

Table 7: Average Frequencies of Each Behavioral Strategy in Runs of ISD

We also wanted to see how the topologies of the neural networks and their number of neurons, which under the original study's definition represent intelligence, would change as the networks evolved in IPD compared to ISD. The topology of the network evolved under ISD is much more complex, with many more hidden nodes and recurrent links, than the topology of the networked evolved under IPD. Figure 5 shows the average number of neurons in a network over the generations for IPD (on the left) and ISD (on the right). The average number of neurons in the graph for IPD remains fairly constant with networks ending with more hidden nodes than the initial ones that they began with. On the other hand, the average number of neurons increases in the graph of the 10,000 generations of ISD, ending at almost 450 neurons. We ran three different runs of ISD and IPD each, graphing the average number of neurons per network throughout the generations and the graphs all looked similar to the ones displayed in Figure 6.

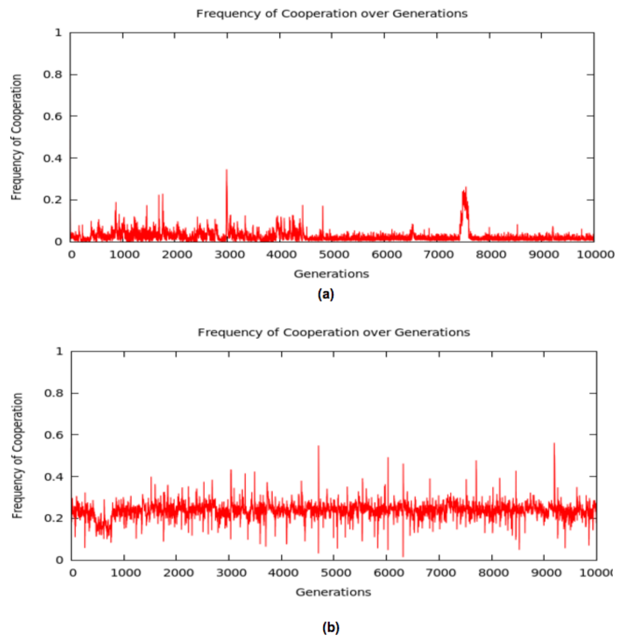

Figure 3: Frequency of Cooperation (Run 10) across 10,000 Generations of IPD (a) and Frequency of Cooperation (Run 2) across 10,000 Generations of ISD (b)

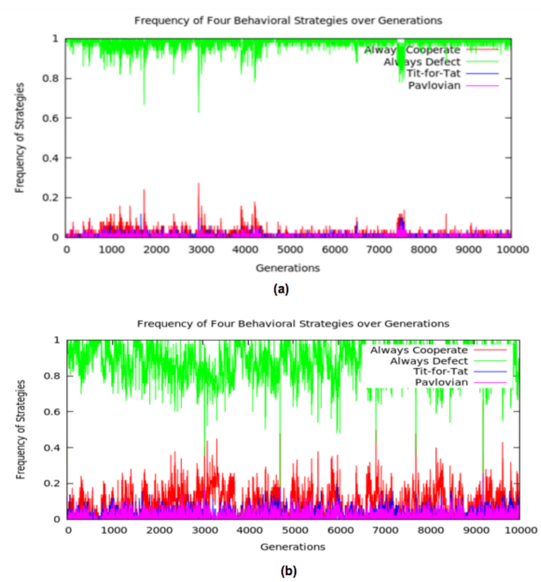

Figure 4: Frequency of Behavioral Strategies (Run 10) across 10,000 Generations of IPD (a) and Frequency of Behavioral Strategies (Run 2) across 10,000 Generations of ISD (b)

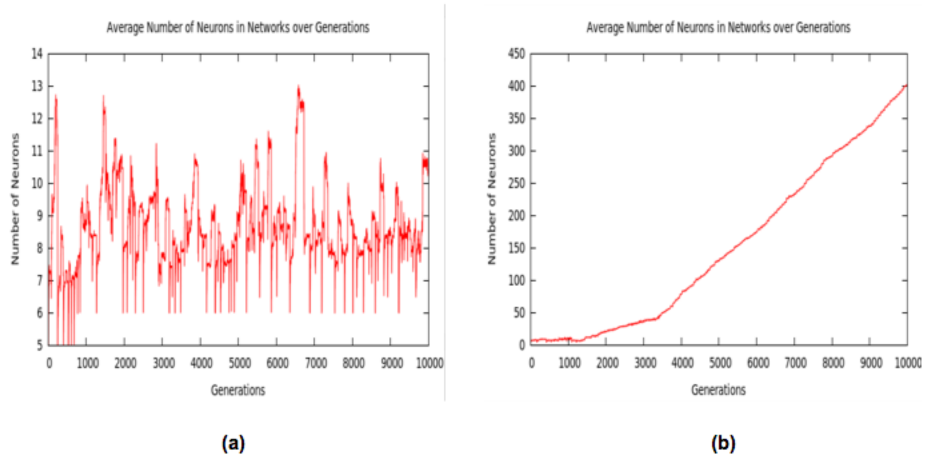

Figure 5: Average Number of Neurons in a Network over 10,000 generations of IPD (a) and ISD (b)

4 Discussion

Our results supported our hypothesis– that neural networks evolved under ISD would cooperate more than neural networks evolved under IPD as the structure of ISD encourages cooperation more than IPD. Also, networks evolved under ISD displayed a greater variety of behavioral strategies as compared to networks evolved under IPD. This work provides insights into what factors are effective at encouraging neural networks to cooperate. Also noteworthy was that fact that networks developed under ISD evolved to have topolgoies with many more hidden nodes than networks evolved with IPD. This could suggest that ISD is a more deceptive task than IPD, since in some situations defecting is the best option and in others it is the worst option, and therefore, requires more intelligence or hidden nodes.

There are several opportunities for future work with this research. The neural networks could be evolved playing against the four "pure" behavioral strategies instead of playing against each other. Also, the study could be redone using different types of behavioral strategies in addition to the four behavioral strategies in this study to see if the neural networks display strategies more similar to other behavioral strategies. Finally, it would be good to further investigate why the networks evolved in ISD evolve so many more neurons that those evolved in IPD. Also, many neural networks evolved to only have connections to one input– they would remove any existing links to the other input. It would be interesting to study if both inputs were really needed and if there were other inputs that would be helpful to the networks.

References

[1] Chervenski, P. 2012. MultiNEAT: Evolving Neural Networks. http://multineat.com/index.html.

[2] McNally, L., Brown, S. P. & Jackson, A.L. 2012 Cooperation and the Evolution of Intelligence. *Proc. R. Soc.* **279**, 3027-3034. (doi:10.1098/rspb.2012.0260)

[3] Potter, M.A., Meeden, L. A. & Schultz, A.C. 2001 Heterogeneity in the Coevolved Behaviors of Mobile Robots: The Emergence of Specialists. *Proceedings of the 17th international joint conference on Artificial intelligence* **2**, 1337-1343 .

[4] Schultz, T. Computational Developmental Psychology. Bradford Books, 1995.

[5] Stanley, K & Miikkulainen R. 2004 Competitive Coevolution through Evolutionary Complexification. *Journal of Artificial Intelligence Research* **21**, 63-100.

[6] Yong, C.H. & Miikkulainen R. 2009 Coevolution of Role-Based Cooperation in Multiagent Systems. *Transactions On Autonomous Mental Development,* **3**, 170-186.