university of groningen / faculty of science and engineering / artificial intelligence

OBSERVING THE INFLUENCE OF NETWORK STRUCTURE ON THEORY OF MIND EVOLUTION DYNAMICS IN A POPULATION: AN AGENT-BASED STUDY

Bachelor's Project Thesis

Ivan Vasnin, i.vasnin@student.rug.nl,
Supervisor: Dr H.A. de Weerd

Abstract: While the advantages of Theory of Mind reasoning have been extensively researched, evolution of Theory of Mind in a network requires additional attention. Differences in network structure may affect the Theory of Mind dynamics and this study aims to determine whether this results significant differences in Theory of Mind evolution patterns. The setup includes experimental and control conditions, which differ by the initialization type: the experimental network is initialized using Preferential Attachment and the control network uses Random initialization. Members of each network play Rock-Paper-Scissors-Spock-Lizard for multiple rounds in each run. After the final round, counts of each order of Theory of Mind strategies are collected as the results of the run. The results show a significant difference in evolution patterns of Theory of Mind across networks. They also suggest that presence or absence of nodes with a relatively higher amount of connections also affects the network dynamics and that they may increase or lower the chance of lower order of Theory of Mind outcompeting a higher order Theory of Mind strategy.

1 Introduction

Theory of Mind is a well-known psychological concept defined by the ability to attribute mental states, knowledge, beliefs and emotions to another person and themselves and reason about these mental states (Premack & Woodruff, 1978). Humans have Theory of Mind and it is one of the fundamental factors required for successful interpersonal interaction. Consider the following example of Theory of Mind use in an everyday situation: Alex is organizing a surprise party for Martha. While knowing about the party herself, Alice can also infer that Martha does not know about the party and Alice will act accordingly to not spoil the surprise. In this social context, Theory of Mind allows the person to recursively estimate and attribute a mental state and be aware of it. The previous example can be expanded to showcase a deeper recursion level: Alice invited Peter to come to the party with her and told him that it is a surprise party for Martha. Based on the situation, Peter knows that Alice knows that Martha does not not about the party. The recursive nature means that Theory of Mind has theoretically infinite amount of orders, where possessing a $k+1$ level allows to iterate over $0, 1...k$ range of levels which produces awareness about kth order of Theory of Mind. First and second order reasoning have been shown in humans using multiple tests, including Sally-Anne test (Baron-Cohen et al., 1985) and matrix game (Hedden & Zhang, 2002). However, other tests, e.g. director game (Keysar et al., 2003), have shown that even in adults, higher order of Theory of Mind reasoning may occasionally not occur automatically.

While humans can reason using the higher orders of Theory of Mind, other species do not have such capacity, which suggests that higher orders of Theory of Mind may convey a competitive advantage (Byrne & Whiten, 1988). Extensive research has been conducted into Theory of Mind in the past 40 years. This includes research conducted by de Weerd et al. (2013), which investigated the advantage that higher order Theory of Mind reasoning have over lower orders in a competitive setting. The setting included multiple single-shot and extensive form games, including Rock-Paper-Scissors, Rock-Paper-Scissors-Spock-Lizard, and Limited Bidding. This study showcases that higher orders of Theory of Mind have direct advantage over lower orders when agents play against each other and concludes that this may explain the evolution of higher order Theory of Mind. However, the study focuses on interactions between only two opponents during a single test run. What was not investigated and may have a significant impact on the evolutionary success of different orders of Theory of Mind is the situation, where each participant has to play against multiple opponents in a network of connections.

To give an example, some of the contexts where agents frequently use Theory of Mind on a network level during competition are chess and professional e-sports. Members of a network frequently play with each other, however there are many ways in which it is decided who plays with whom. A network may be isolated into subpopulations (representing smaller local scenes) or agents can play against every other opponent in random order (in case of a league). This means that there are numerous ways in which the same number of players can be playing between each other in a realistic scenario. This, in turn, can be represented by different types of network connections. Some types of network initialization include Random, Preferential Attachment (Barabási & Albert, 1999), Small-World Initialization (Watts & Strogatz, 1998) and Hierarchical Initialization (Ravasz & Barabási, 2003). When competitors play against each other for a significant amount of time, they might adapt and change their strategies, based on what they observe from their opponents. This paper will try to investigate the effects that different networks of opponents may have on agents' strategies by focusing on differences between evolution of strategy distributions in two types of networks, which differ by initialization method and distribution of connections between the nodes.

1.1 The Study

Difference in the network structures may produce different patterns of evolution for the populations of the same size. In particular, the presence or absence of highly-connected nodes represent two different real world populations. Continuing example from the previous section, a presence of a few highly-connected nodes among a majority of nodes with few connections symbolizes a population where there are many relatively-isolated local tournaments and only a few players can travel freely between them. Identical number of node connections represents one big hub competitive scene, where players randomly find someone to play with. Difference in dynamics between these two conditions would be interesting to observe, as it can be extrapolated to the real sports and how they are played.

In this study, two networks of agents will play Rock-Paper-Scissors-Spock-Lizard and the evolution of Theory of Mind strategies distributions will be observed and compared. Rock-Paper-Scissors-Spock-Lizard was successfully used previously in de Weerd et al. (2013). It is a relatively simple zero-sum game, patterns of which can be accurately formalized in an algorithm that uses Theory of Mind to make decisions about moves. Moreover, while being simple, it is more complicated compared to the standard Rock-Paper-Scissors, which justifies use of Theory of Mind in this less trivial context.

The size of the networks in the experimental and the control conditions is 104 nodes, with each node representing an agent. In each network, there is exactly 208 connections between nodes, meaning that during each time step 208 unique games are played.

Preferential attachment was chosen for experimental condition, since it naturally produces nodes with significantly higher amount of connections than the rest. It also uses an intuitive probability-based algorithm to create new connections, which increases replicability and validity, due to its consistency. Random initialization with a guarantee of the same number of connections for each node is used as a control condition, as it creates a homogeneous population that is a good counterpart to a population with sub-clusters in experimental condition.

The research question is the following: does a presence of a few nodes with higher amount of connections than the rest affect the evolution of Theory of Minds-based strategy distribution in a population that plays Rock-Paper-Scissors-Spock-Lizard? This study will observe the populations, in which agents play with each other games of Rock-Paper-Scissors-Spock-Lizard for a predefined number of time steps. The agents will use Theory of Mind based strategies to try to predict opponents' moves and may switch to another strategy if they see that the other strategy yields better performance. There are two conditions that separated by the type of network initialization and presence of highly-connected nodes: the experimental condition is initialized using preferential attachment and has highly-connected nodes, while the control condition uses random initialization and all of the nodes have the same amount of connections.

The experimental hypothesis states that there is a significant difference in the distributions of strategies between the experimental and control conditions. The null hypothesis states that the distributions of strategies are similar and there is no statistical difference between them.

2 Methods

2.1 Environment

In order to test the influence that presence of highly-connected nodes has on the distribution of Theory of Mind-based strategies in the population, an experimental environment was created using Python code language. The environment conducts an agent-based simulation, where agents are playing Rock-Paper-Scissors-Spock-Lizard (Kass, 1995) between each other based on generated networks of connections. Rock-Paper-Scissors-Spock-Lizard is an expanded version of Rock-Paper-Scissors, with two additional moves: Lizard and Spock. Each of the five moves wins against two moves, loses against two other moves and draws against itself. The table and graph showcase the standard in which the moves interact with each other.

In the agent simulation, agents play for a defined number of time steps (T = 50), using Theory of Mind-based strategies, which will be described in section 2.2.2. There are three orders of Theory of Mind that the agents may use: zero order, first order and second order.

During the initialization of the environment, the first step is to create a set of agents. All agents are

	Rock	Paper	Scissors	Lizard	Spock
Rock	0	−1	1	1	−1
Paper	1	0	−1	−1	1
Scissors	−1	1	0	1	−1
Lizard	−1	1	−1	0	1
Spock	1	−1	1	−1	0

Figure 2.1: Payoff table for all move interactions in Rock-Paper-Scissors-Spock-Lizard. Entries show payoffs for the row player. Retrieved from de Weerd et al. (2013)

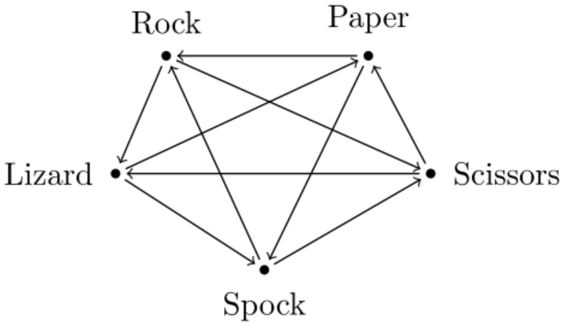

Figure 2.2: A visual representation of Rock-Paper-Scissors-Spock-Lizard rules. Each arrow represents a defeat relation, so that the action at the origin of the arrow defeats the action at the destination. Retrieved from de Weerd et al. (2013)

initialized with a random Theory of Mind level-based strategy. The strategies are adapted from de Weerd et al. (2013). Afterwards, based of the set of agents (the set size N = 104), two separate networks of connections between the agents are created, which determine which agents will play with each other for the rest of the test run. The two networks represent an experimental condition and a control condition. They are initialized based on the same agent distribution to ensure that the data sets obtained from both conditions are paired, which control for the effects of random initialization of beliefs and strategies and isolates the influence of the network connections on the final distributions of strategies.

2.1.1 Control network

The network for control condition is initialized to guarantee that each node has the same amount of unique bilateral connections. A choice was made to set the number of connections for each node to 4. All of the bilateral connections are added at random, which means that occasionally the network is initialized incorrectly, with some of the nodes having fewer connections, while all of the other nodes either already have 4 connections or are connected to the ones with insufficient amount of links. For that reason, after the initialization a check is conducted to guarantee that each nodes has 4 connections and if that is not the case, then the network is initialized again, until it passed the check. Next, due to the random generation of the network connections, there is a small chance for the network to have isolated subsets, that are not connected with the rest of the network. For that reason a second check is conducted, to verify whether every node in the network can be reached from every other node. Dijkstra's algorithm (Dijkstra, 1959) was used to calculate distances between all of the nodes and if any of the values was equal to infinity (meaning that it is impossible to reach one node from another), the network was generated again. This procedure for network initialization results in a network with no clustering, a constant number of connections per node and total number of connections in a network, and a relatively high average distance between nodes.

2.1.2 Experimental network

The network for experimental conditions is initialized using a modified preferential attachment algorithm (Barabási & Albert, 1999). This algorithm starts by initializing a graph, adding two nodes to it and creating a connection between them. Afterwards, the remaining nodes are added to the graph using preferential attachment to determine with which node it will be connected via a single edge. The preferential attachment formula calculates for each existing node in graph, the probability that the new connection will have the node at one of the ends of the edge. For each node, the formula for the probability is $P = C/N$, where C is the number of connection that the nodes already has and N is the total number of connections in the network. Next, in order to make sure that each agent has more than one other agent to play each round with, the algorithm iterates over all nodes and adds a random second connection if the node did not have it before. The connection is decided at random rather than based on preferential attachment, which is done to make the population slightly more interconnected and lower average distance between nodes. Finally, preferential attachment is used to select both endpoints of a new connection and create additional edges between nodes until the total number of edges in the network is equal to exactly double of the amount of nodes. This ensures that the total amount of connections in a network is constant between the experimental and control conditions. This procedure for network initialization results in the structure, where majority of nodes have two-three connections, while there are a few nodes that have a significantly higher amount of links. Average distance between nodes is less than in control condition. Moreover, it is possible to create an intuitive and easily-observable split of nodes into highly and low-connected in the experimental condition. For this thesis the four most connected nodes are counted as highly-connected and the remaining 100 are referred to as low-connected. Table 2.1 shows an example of distribution of connections per nodes in a single experimental run.

After both networks are initialized, it is necessary to isolate the highly-connected nodes. For each initialization of an experimental condition network, the indexes of the highly-connected

Table 2.1: ***Counts of nodes with each number of connections in an example experimental network of 104 nodes***

Number of connections	Node counts
2	37
3	24
4	15
5	10
6	6
7	4
8	2
9	2
10	1
15	1
18	1
25	1

nodes are obtained. In order to determine whether they will have an influence on the strategy distribution in the network of agents, highly-connected nodes are manually given the same Theory of Mind order strategy. It should be noted, that in case of fourth and fifth highest connected nodes having the same amount of connections, only one of them is chosen to be considered highly-connected, while the other will be counted as low-connected. In order to control for the influence of the strategy level, all of the highly-connected nodes start with zero order Theory of Mind strategy in one third of all runs, in another third they start with first order Theory of Mind strategy and in the rest of the runs they start with second order Theory of Mind strategy. Furthermore, in the control network the nodes with the same indexes as the highly-connected nodes in the experimental condition network are given the same strategy as their experimental condition counterparts. That concludes initialization of the networks, after which the test runs may begin.

2.2 Simulation run

2.2.1 Theory of Mind orders and game strategies

During every game each agent makes a move based on the strategy that it currently uses and sets of multiple Theory of Mind level beliefs. The sets of beliefs are initialized for each opponent separately using random values at the moment when a connection with the opponent is created during network initialization. For each opponent an agent has 3 sets of beliefs: zero order beliefs, first order beliefs and second order beliefs. Each set of beliefs has 5 values, which represent the agent's estimated probability for each of the moves that their opponent may use (rock, paper, scissors, Spock, lizard).

During the decision process for the move, strategies use the belief values to determine the final perceived probability of each move. The strategies are based on the ones used in de Weerd et al. (2013). Each higher level Theory of Mind strategy integrates the beliefs of the same order with the beliefs of lower orders.

All of the 3 sets of beliefs are updated at the round's end using the opponent's and agent's own moves during the round at a learning rate of 0.8 (de Weerd et al., 2013). The learning rate of 0.8 means that after each time step the previous beliefs will have 20% of the total weight, while the beliefs based on the new observations will in total have 80% of the total weight. Learning rate of 0.8 is a high learning rate, meaning that the rate of decay of older beliefs is also high. All of the beliefs add up to 1 and the algorithm for the update of each set of beliefs is different.

For zero order strategy, the zero order beliefs become the final probabilities assigned to the opponent's moves and the agent chooses between one of the two moves that win against the opponent's last move, using the values for the other moves to decide between the two options. For example, if the agent believes that the opponent will use Rock, the agent can use Paper or Spock (see Figure 2.1). However, each of these two moves can also be countered by two moves, namely Scissors and Lizard for Paper or Paper and Lizard for Spock. Since Lizard counters both of the moves, the important values are the beliefs for Paper and Scissors. The agent will choose Paper if they believe that the opponent is less likely to play Scissors.

Zero order beliefs use zero order theory of mind to represent the probability of each of the opponent's options. These beliefs are updated whenever the agent observes the opponent making a move. At a learning rate of 0.8, it results in the agent believing that during the upcoming round, the opponent will use the same move as they did in the previous round. This means that the agent will choose to counter that move using one of the two moves that beat the opponent's move. Mathematically, during the update all of the zero order beliefs values are multiplied by 0.2, after which 0.8 is added to the belief for the move used in the last round. For example, in case where the opponent used Rock in the last turn, during the beliefs update the probability of all moves will be multiplied by 0.2, after which 0.8 is added to the probability of Rock.

The first order strategy integrates first and zero order beliefs to calculate the final probabilities, with first order beliefs having 80% of the final weight, while zero order beliefs have 20%. It then checks which two moves the opponent is most likely to use and chooses a move that counters both of them. In case of small chance, where the two most likely moves can't be beaten with a same move, the agent will operate like a zero order agent and counter the move with the highest likelihood.

First order beliefs represent the beliefs that an agent holds about probabilities of the opponent's moves, based on the assumption by the agent that their opponent will act using a zero order strategy. First order beliefs are updated based on the agent's own last move and are based on the idea that the opponent uses a zero order strategy and will try to counteract the move that the agent has chosen during the last round. For that reason, during update of first order beliefs, probability of the two moves that beat the agent's own last move is increased. At the learning rate of 0.8, it means that the during the update all of the beliefs values are multiplied by 0.2, after which 0.4 is added to the beliefs for both of the moves that counter agent's own last move. This calculation departs from de Weerd et al. (2013). For example, if the agent played Rock during the last game, it will assume that its opponent will try to counter it during this time step. After multiplying all beliefs by 0.2, since it believes in a higher chance of either Paper or Spock and 0.4 is added to chances of both of these moves.

Second order strategy integrates second order beliefs with the first and zero order beliefs by integrating second order strategy with first order strategy. 80% of the final probabilities are taken from the second order beliefs, while remaining 20% are taken from the first order strategy beliefs (meaning that first order beliefs have 16% weight and zero order beliefs have 4%). Then, same as in zero order strategy, second order strategy chooses between one of the two moves that beat the opponent's most likely move, using the values for the other moves to decide between the two options.

Second order beliefs assume that the opponent acts using first order strategy, which means that the opponent will act based on its last move. During the update of second order beliefs, opponent's last move is observed and the probability of the move, that beats the two moves that beat the opponent's last move is increased. The mathematical update is similar to the zero order beliefs update.

2.2.2 Environment update

During the test run, the simulation runs for a predetermined amount of time steps (T = 50), after which the execution is complete. During each time step every agent plays a game of Rock-Paper-Scissors-Spock-Lizard with every other agent that it has a connection with and records the results.

At the end of each round the agents are awarded with points based on the results of the played games. For a win, the agent gets 1 point, for a draw - 0.5, in case of a loss no points are given, which makes Rock Paper Scissors Lizard Spock a constant-sum game At the end of each time step a performance rate is calculated for each of the agents by dividing the number of obtained points by the number of agent's opponents (which is the same as the potential maximum amount of points the agent can get in one round). Each agent has access to its own performance rate and rates of the opponents with which it plays. After each round, every agent changes its strategy to match the one of the opponent with the highest performance rate if it is higher than its own.

During the first round of the run, all of the moves are chosen at random and the strategies do not change, while the beliefs are still updated. Starting from the second round, all moves are determined by the strategies and sets of beliefs and at the end of the round agents may change their strategy based on the points obtained in the round. After the predetermined number of time steps has passed, the simulation run is over and the set containing the strategy types of all of the agents is collected as the result of the run in order to conduct statistical test later.

3 Results

3.1 Data set

The simulation was run for 300 runs in control and experimental conditions. In total, 600 runs were conducted. Each run lasted for the 50 time steps and each of the 3 strategies was used in initialization of highly-connected nodes in 100 runs each. The total number of runs is 600. (see Table 3.1)

Table 3.1: ***Number of runs utilizing each level of Theory of Mind for initialization of highly connected nodes***

	Zero Order ToM	First Order ToM	Second Order ToM
Experimental	100	100	100
Control	100	100	100

Since each population had 104 agents, the results that were collected are two data sets with 31.200 data points, with each point having a values of either "0", "1" or "2" and signifying the strategy that an agent in a simulation had after the final time step of the simulation run.

3.2 Total strategy counts per conditions

The total count of the strategies per condition in the data sets is: in experimental condition there were 24676 accounts of second order strategy, 6072 accounts of first order strategy and 452 account of zero order strategy. In the control condition there were 30186 accounts of second order strategy, 702 accounts of first order strategy and 312 account of zero order strategy.

Upon inspection of the data sets, it can be observed that in both conditions a majority of agents use second order Theory of Mind, showcasing the advantage that second order strategy has in this environment. It can further be noted that in majority of runs all of the agents converge to a single strategy, which is in most cases either first or second order.

While the dominance of the second order strategy is consistent across the conditions, the prevalence of the other two strategies is not. The experimental condition has a noticeably higher proportion of first order strategy compared to control, while control condition has a higher occurrence of zero order strategy.

To verify whether the observed differences are statistically significant, a chi-squared test was conducted. The choice was based on the notion that both the independent variable (type of the network) and dependent variable (final strategies of the agents) are categorical variables. The test showed the relation between the variables is highly significant, X^2 (2, 62.400) = 4836.0421, $p < 0.00001$. This result suggests that the type of network significantly impacts the evolution of the Theory of Mind strategies in the population.

3.3 Strategy counts based on the Theory of Mind order used in initialization

The tables below shows the counts for each strategy level, based on the type of strategy, that the highly-connected nodes were initialized with.

Table 3.2: ***Number of agents with each strategy level after the 50 time steps, based on the strategy level given to the highly connected nodes at initialization in experimental condition***

		Initial Theory of Mind level of highly-connected nodes		
		Zero	**First**	**Second**
Counts	**Zero**	0	348	104
	First	2614	1099	2359
	Second	7786	8953	7937

Table 3.3: ***Number of agents with each strategy level after the final time step, based on the strategy level given to the highly-connected nodes at initialization in control condition***

		Initial Theory of Mind level of "highly-connected" nodes		
		Zero	**First**	**Second**
Counts	**Zero**	104	192	0
	First	207	208	303
	Second	10089	10000	10097

From the tables it can be observed that across all types of initialization strategies and both control and experimental conditions, the counts for Zero order strategy are minor.

While it was observed from the final counts that there is more cases of first order strategy in the experimental condition, it can now be specified, that the cases are also quite evenly distributed across the runs with different initialization strategies. However, it does appear that in cases when highly connected nodes initially started with first-order strategy, it was more likely by the final time step that the population will converge to second order strategy, compared to runs with a different type of initialization strategy.

In order to test the significance of this difference a proportion test was conducted between the counts for first order strategies in two conditions: first condition is initialization with first order strategy and experimental run, second condition is initialization with second order strategy and experimental run. The proportion test yields the following values: z-value is -23.4661, $p < 0.00001$.

The results of the proportion test suggest that first order strategy has different prevalence rate across initializations of strategies for runs in experimental condition.

3.4 Visualization of a typical simulation run

The figures 3.1 and 3.2 show the evolution of prevalence of different strategy levels in a population across all of the time steps. Figure 3.1 shows the dynamics of the environment that converge on the second order strategy, while figure 3.2 shows the less common case of the population converging on the first order strategy. Interestingly, this scenario mostly happens when the highly connected nodes were initialized with zero or second order Theory of Mind strategies.

It can be observed that when the population converges to the second order strategy, the first order strategy is eliminated first, followed by a decrease in the count of the zero order strategy. Meanwhile, when first order strategy prevails, zero order strategy is eliminated first, while first order strategy eventually outcompetes the second order strategy. The evolution pattern was similar across the runs in experimental and control conditions, aside from the environment converging sooner to a single strategy in the control condition.

Figure 3.1: Evolution of distribution of strategies in the population that converged on the second order Theory of Mind strategy

Figure 3.2: Evolution of distribution of strategies in the population that converged on the first order Theory of Mind strategy

4 Discussion

Firstly, it can be stated that the results and conducted statistical tests provide an answer to the research question: the presence of a few highly-connected nodes does affect the evolution of Theory of Mind of Theory of Minds-based strategy distribution in a population that plays Rock-Paper-Scissors-Spock-Lizard. The statistical significance of the difference between the results from the experimental and the control conditions leads to a rejection of the null hypothesis. The rest of the discussion will address potential confounds, alternative interpretation of results and provide suggestions for future research of Theory of Mind.

4.1 Results interpretation

While the design of the environment was intended to control for as many extraneous variables as possible, the final results are still open for interpretation and different arguments can be made about which independent variables were most impactful.

The difference between the final counts for different levels of Theory of Mind strategies of control and experimental conditions was statistically significant. While an extensive size of the results data sets could be responsible for a higher statistical power and the low p-value may signify the great influence that the independent variable (influence of highly-connected nodes) has on the distribution of strategies, other arguments can be made. The observed difference between the conditions may be caused by the influence of multiple effects produced by a different network structures.

It can be observed that by using a different network initialization method, the direct result in the experimental condition is not only the appearance of highly-connected nodes, but also the emergence of clusters. In this case, clusters are a consequence of great inequality between the number of connections that highly-connected nodes have compared to the low-connected ones. The highly-connected nodes become centers of these clusters and their emergence affects the dynamics of the network. One of the effects of clustering on the network is that in experimental condition the average distance between the nodes is noticeably lower than in the control condition. In experimental condition the average distance varied between 1.5 and 3 nodes, while in control condition the range of averages was from 3.5 and 3.8 nodes.

The different method of network initialization in the experimental condition has extraneous effects, like the aforementioned difference in average distance between nodes. This means that it is impossible to separate the effect and influence of highly-connected nodes and their strategies from the additional differences in the networks by looking at the final counts between the experimental and control conditions alone. However, the significant statistical results are sufficient to suggest that these multiple effects do not contradict each other with the same power, which would result in no difference with the control condition as the multiple effects cancel each other. In turn, the effects of clustering and addition of highly-connected nodes are either complementary to each other or one of the effects is significantly stronger than others.

In order to isolate the effects that highly-connected nodes have on the population, the data for both conditions was split into groups, based on which strategy was given to the highly-connected nodes during the run initialization. Proportion test conducted between these groups in the experimental condition showed a significant difference between the final strategy counts. This additional test within the experimental condition controls for the clustering and average distance between nodes. Based on that a conclusion can be derived that highly-connected nodes have an impact on the network and strategy distributions as a sole isolated independent variable, which supports the experimental hypothesis.

4.2 Improvements and further research suggestions

Based on the interpretation of results it was possible to isolate the effect of highly connected nodes and determine that they have an effect on the network evolution in separation from the other effects. However, the next step would be determining the extent of effect that clustering and average distance between nodes may or may not have on the network dynamics. A possible setup for that could include two conditions, where all nodes have the same amount of connections, however in the control condition all of the connections are fully random, while in the experimental condition the population is split into subdivisions (clusters) and each node has exactly one connection with a node outside of the cluster, while the remaining connections have to be made with other members of the same subdivision.

Based on the results of the proposed study, it would also be useful to determine whether the effects of clustering complement or conflict with the effects of highly-connected nodes. There are many possibilities of how exactly these effects together affect the dynamics of the network and that deserves to be further looked into.

Another suggestion is related to the nature of Rock Paper Scissors Spock Lizard as a game and how Theory of Mind can be used in it. In this study, the population had strategies based only on the lowest three orders of Theory of Mind. While the three orders were sufficient at showing the dynamics that may occur in the population and provided both expected and unexpected results, it might be beneficial to include higher orders of Theory of Mind as well. When compared to the original Rock Paper Scissors, Rock Paper Scissors Lizard Spock may benefit more from higher orders of Theory of Mind due to a higher variety of interactions. In Rock Paper Scissors there are only three moves, which work in a circular manner. In a situation where an agent and its opponent both used Scissors in the last round, a zero order strategy chooses Rock as the optimal move, first order strategy would choose Paper and second order strategy would choose Scissors. Third order Theory of Mind strategy would once again use Rock to counter the second order strategy, which means that it is functionally identical to zero order strategy in terms of the most optimal move. This shows that the first three orders of Theory of Mind cover all possible move interactions and strategies based on orders of Theory of Mind higher than the second order would provide an identical optimal move to the one provided by one of the lower order strategies. However, this exhaustion of options using only the first three orders of Theory of Mind is not the case for Rock Paper Scissors Spock Lizard, which means that addition of higher order strategies until the exhaustion of options is met may significantly impact the dynamics of the population. The practical usefulness of that investigation may be limited, due to the fact that higher orders of Theory of Mind would require higher use of mental resources, making them less energy efficient. However, some important theoretical insight on interactions of different order of Theory of Mind might be obtained from such investigation.

Lastly, there was a surprising observation made while analyzing the evolution of network strategy distributions per time step. While the network dynamics followed an expected pattern in cases when the network converged on the second order strategy, that was not the case when the network converged on the first order strategy. By definition of the strategies, first order of Theory of Mind is effective against zero order Theory of Mind, while first order is in turn countered by the second order. Unexpectedly, in the test runs, when the network converged on the first order Theory of Mind, the strategy that was eliminated first was zero order and afterwards first order slowly directly outcompeted second order. In majority of test runs, the expected pattern took place and second order outcompeted the lower orders, nevertheless the proportion of runs where first order theory of mind was triumphant is not dismissible and the conditions for such outcome should be determined in further research. One of the potential reasons for first order Theory of Mind to be able to compete with second order Theory of Mind is the integration of beliefs and a learning rate of less than 1, which means that the move decision is not deterministic. Considering the currently available information it is possible for this outcome to be a result of unpredictable belief dynamics, but it is also possible that some other specific conditions have to be met for the network to converge on the first order strategy and that may be a topic of further research in the future.

5 Conclusion

The obtained results provided support for the experimental hypothesis, which stated that addition of highly-connected nodes to the network would influence evolution of strategy distribution in the population. However, this was not the sole observation extracted from the data and future research should be conducted to answer the newly posed questions. Specifically, concerning the effects of clustering on network evolution and the conditions under which first order Theory of Mind outcompetes the second order Theory of Mind.

References

Barabási, A.-L., & Albert, R. (1999). Emergence of scaling in random networks. *Science*, *286*(5439), 509-512. doi: 10.1126/science.286.5439.509

Baron-Cohen, S., Leslie, A. M., & Frith, U. (1985, October). Does the autistic child have a "theory of mind" ? *Cognition*, *21*(1), 37–46.

Byrne, R. W., & Whiten, A. (1988). Machiavellian intelligence: social expertise and the evolution of intellect in monkeys, apes, and humans.

de Weerd, H., Verbrugge, R., & Verheij, B. (2013). How much does it help to know what she knows you know? An agent-based simulation study. *Artificial Intelligence*, *199-200*, 67-92. doi: https://doi.org/10.1016/j.artint.2013.05.004

Dijkstra, E. W. (1959). A note on two problems in connexion with graphs. *Numerische Mathematik*, *1*(1), 269–271.

Hedden, T., & Zhang, J. (2002, August). What do you think I think you think?: Strategic reasoning in matrix games. *Cognition*, *85*(1), 1–36.

Kass, S. (1995). Retrieved from `http://www.samkass.com/theories/RPSSL.html`

Keysar, B., Lin, S., & Barr, D. J. (2003, August). Limits on theory of mind use in adults. *Cognition*, *89*(1), 25–41.

Premack, D., & Woodruff, G. (1978). Does the chimpanzee have a theory of mind? *Behavioral and Brain Sciences*, *1*(4), 515–526. doi: 10.1017/S0140525X00076512

Ravasz, E., & Barabási, A.-L. (2003, February). Hierarchical organization in complex networks. *Phys. Rev. E Stat. Nonlin. Soft Matter Phys.*, *67*(2 Pt 2), 026112.

Watts, D. J., & Strogatz, S. H. (1998, June). Collective dynamics of 'small-world' networks. *Nature*, *393*(6684), 440–442.