Mobility effects on the evolution of co-operation in emotional robotic agents

Joe Collenette, Katie Atkinson, Daniel Bloembergen, Karl Tuyls
Department Of Computer Science
University Of Liverpool
j.m.collenette@liverpool.ac.uk

ABSTRACT

Simulating emotions within a group of agents has shown to support co-operation, in the prisoner's dilemma game [8]. Recent work on simulating these emotions in agents has focused on environments where the agents do not move, that is, they are static and their neighbours are fixed. However it has also been shown that when an agent is given the ability to move, then the type of the environment affects how co-operation between agents evolves in the group of agents [11]. In this paper, we will explore the effects on co-operation when emotional agents are given the ability to move within relatively small and structured environments.

We conclude that once mobility is introduced, different strategies are successful than in static agents. The successful strategies, regardless of environment type, respond quickly to defection, while not immediately reciprocating co-operation. The higher the density of the agents, the lower payoff all agents achieve. The further an agent travels, the higher its total payoff. The slower an agent is to copy another agent by imitating its strategy, increases its average payoff.

1. INTRODUCTION

It is well known in psychology that emotions in humans affect decision making [13]. By simulating these emotions within agents we can then show the evolution of co-operation between agents within the prisoner's dilemma game [7, 8]. The recent work in emotional agents and their co-operation has focused on agents which do not have mobility.

Whilst we recognise that emotions have both psychological and physiological grounds [5], we consider only the former in this paper. We will simulate the functional aspect of emotions, to the effect that emotions can change the current behaviour of the agents, such as anger driving a pacifist to fight [6].

When an agent is given mobility, initial work has started to explore the affect the environment type has on decision making when playing the prisoner's dilemma game [11]. By adding mobility to emotional agents, it will allow us to examine whether the environment structure has the same effects as it does on non-emotional agents. By placing our agents in a mobile environment we are hoping to have a more accurate description of the evolution of co-operation within simulated emotional agents, and to see if we can observe similar effects that the environment type has on decision making.

Our study addresses the following questions: Do these simulated emotions effect how the environment affects decision making? Does the added mobility affect the simulated emotions and the decision making in these agents? By answering these questions we can understand how the evolution of co-operation is affected when these agents with simulated emotions are placed in different types of environment. We can also see the effects that the addition of mobility has on the agents and so gain insight into the use of emotions in situated robots.

To achieve this we will be using two different types of environments, a regular environment and a small world environment. A regular environment is where the agents can only move within a small range of the other agents around them, so to play against agents on the other side of the map would require moving a long distance to reach them. It is regular as at all intersections there is the same number of exits. The small world environment is similar to the regular but it contains shortcuts across for these agents to move over to different parts of the map quickly.

In our environments we will be simulating e-puck robots, which are small disc shaped robots [9]. To simulate the e-pucks and the environment we will be using the player/stage simulator [4]. This allows us to simulate the e-pucks' movement and sensors, the environment type, in addition to letting us record the positions of each e-puck at any given time.

Isolating the effects that the environment has on decision making in our agents, we can observe the differences between mobile and fixed agents. This enables us to see what the effects this has on the evolution of co-operation in societies of agents.

2. BACKGROUND

The agents will be playing the prisoner's dilemma game. The prisoner's dilemma is a game where two players have the choice of either defecting or co-operating; choices are made simultaneously. They then get a payoff depending on the choices of both agents, the payoff matrix is shown in Table 1. When our agents are playing the game they have no knowledge of the payoff matrix or how many iterations of the game they will be playing. We are using this particular game as it has been shown that it can be used to explore the evolution of co-operation [2, 12, 3].

When looking at the prisoner's dilemma outcomes, it is in the best interest of both players to both play co-operatively since this would lead to the largest total payoff. However there is an incentive to defect as this can lead to higher individual payoffs. This then leads to a Nash equilibrium of *(DEFECT, DEFECT)*, which gives the worst outcome for the group as a whole. This outcome shows the dilemma of the game and it allows us to see if co-operation between

Table 1: Prisoner's Dilemma Payoff Matrix

	CO-OP	DEFECT
CO-OP	3,3	0,5
DEFECT	5,0	1,1

Table 2: Emotional Characteristics

Anger Threshold	Gratitude Threshold	Character
1	1	Responsive
1	2	Active
1	3	Distrustful
2	1	Accepting
2	2	Impartial
2	3	Non-Accepting
3	1	Trustful
3	2	Passive
3	3	Stubborn

agents can flourish.

The simulated emotions we will be implementing are based on the Ortony, Clore and Collins model of emotions, known as the OCC model [10]. This was developed from psychology research and has been used within the AI community [1, 7] to simulate emotions within agents. The emotions we will be modelling are *anger, gratitude* and *admiration.*

The OCC model provides 22 emotions that can be modelled; they take the view that each action is a response from the emotional makeup and that each emotion gives a different action to take. Since the OCC model describes the actions that an emotion can lead to rather than how that emotion is processed internally, this gives us a good platform to implementing this in a computational setting.

Our implementation of these emotions is similar to previous work into the emotional agents [7]; this allows us to compare the differences in mobility and environment structure rather than implementation. Each emotion has a threshold, when that threshold is reached it triggers a change in the agents behaviour. When the anger threshold is reached the agent changes its behaviour to defection, and when the gratitude threshold is reached the agent changes its behaviour to co-operation. Admiration, when triggered will cause the agent to take on the emotional characteristics of the agent that triggered the admiration threshold.

There are a number of emotional characters which have differing thresholds for these emotions. The full set of characters is shown in Table 2. Admiration thresholds can be rated as high (3), medium (2) or low (1). When any threshold is reached the value of that emotion is then reset back to 0.

An agent's anger increases by one when its opponent defects, gratitude increases when the opponent cooperates. Admiration increases when the agent believes that its opponent is performing better than itself. The exact implementation details of the admiration emotion is discussed further in Section 4.2.

Take for example the Active characteristic whose anger and gratitude values are currently zero and is set to initially co-operate. When this character receives a *DEFECT* from its opponent then the anger value will increase by one. The anger value is now at the anger threshold of one, the character will then change from its initial co-operation to defection, that is, in the next game against that same opponent, the character will choose to defect against that opponent.

3. METHOD

The agents will be simulated in an environment and given a random walk behaviour with some basic obstacle avoidance procedures. The prisoner's dilemma game will be initiated when two agents are within close proximity of each other and both have line of sight of each other. They will then continue their random walk behaviour. The two environments that we will be placing the mobile agents into include a basic regular environment and a small world environment as shown in Figure 1, with the black areas being the walls and the white areas the floor. The arena size is 5 metres by 5 metres and the e-puck has a diameter of 7 centimetres. The agents will be placed in a random location initially.

The random walk behaviour is simplistic. When the agent gets information about the world from its sensors it will first check for obstacles. If the sensors on the left detect anything they will stop and then turn to the right, and the reverse for the sensors on the right. The right sensors are located at 15°, 45°and 90°from the direction the e-puck is facing, and the reverse for the left sensors.

To place each agent into the environment, we do the following:

1. Calculate how many of each agent types we need from the given percentage.
2. Create each agent and place into the list.
3. Shuffle the list of agents.
4. Calculate the number of defectors from the given percentage, $D =$ the number of initial defectors.
5. Assign the top D of the list to defect, and the rest to co-operation.
6. Shuffle the list of agents.
7. Calculate the number of high admiration threshold agents, $H =$ the number of high threshold agents.
8. Calculate the number of medium admiration threshold agents, $M =$ the number of medium threshold agents.
9. Set the top H agents in the list to have the high admiration threshold.
10. Set from $H+1$ to $H+M$ in the list to have the medium admiration threshold.
11. Set the rest the agents to low admiration threshold.
12. Shuffle the list.

This ensures that we have the correct number of agent types and that initial moves and admiration thresholds are distributed randomly to each agent.

To ensure that we have placed each agent randomly, we use the list as created above, pull each agent off the top of the list and assign it a place in the environment though randomly generated X and Y positions. The randomly generated location is checked to make sure that it is not on any walls or a previously allocated positions, it if it is, then we generate a new random location.

When observing the environment and no obstacle is detected then the agent will move forwards with a random turn speed between no turning and the maximum turn speed for left or right, while moving forward. Since the rate in which the simulated e-puck receives data is around every second then this gives a random movement around the environment.

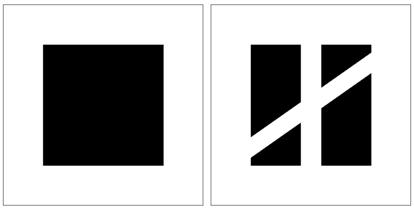

Figure 1: Environments to be used. The regular environment is on the left and the small world environment is on the right.

Each agent placement will be randomized to prevent pockets of identical agents, which cannot be broken down as they will use each other to prevent replications happening in their group.

In regards to the iterated prisoner's dilemma game that they will be playing, the payoffs can be seen in Table 1. The agents have no knowledge of the payoff matrix or the number of games to be played, they will purely use their own strategies to decide whether to co-operate or defect.

In addition the agent has no knowledge of the strategies of its neighbours, but the emotions it has apply specifically to the agent it is playing against. That is the agent can differentiate between players, but has no knowledge of them. The agents also do not know about the environment they are placed in. They will only use their random walk behaviour to navigate around the environment.

4. EXPERIMENTS

4.1 Validation

The aim of this experiment is to show that our agents which move have the same emotional response and outcomes as the agents which do not move as in [8]. That is for, each emotional character they will choose the same response after receiving the same input. In this experiment we will only be using the emotions *gratitude* and *anger*, as these were the emotions used in original experiment [7]. In these environments we will be using two types of agent, one which is an emotional agent and the other being a set strategy which doesn't use emotions. In addition our emotional agents will be set to co-operate initially.

These non-emotional agents have the same knowledge of the world as the emotional agents. They have the same random walk behaviour and the same limited knowledge about their neighbours.

The strategies that the emotional agent will be played against are traditional ones from Axelrod's tournament and include:

Mendacious Always defects
Veracious Always cooperates
Random Equal chance of defection or co-operation
Tit-for-tat Initially co-operate then play the opponent's last move
Joss Tit-for-tat with a 10% chance of defection
Tester Defect on round n, if the opponent defects play tit-for-tat until the end of the game otherwise cooperate until round $n + 2$ then repeat from $n + 3$

For each emotional character as shown in Table 2, we will perform 10 runs against each strategy in turn. A run consists of simulating the mobile agents until 200 rounds of the prisoner's dilemma game have been completed.

In this experiment there are only two agents. By letting the agents run until the same amount of rounds have been completed as in the previous experiment, this then means that they have played the same amount of games against the same opponent as the agents which do not move. This should make the results identical, allowing for some slight variation with the strategies that use randomness. If the results are the same then it shows that our emotional agents behave in the same manner as emotional agents which do not move in [7].

4.2 Main Experiment

This experiment aims to highlight the differences and similarities between emotional agents that move and ones that don't, as well as showing what differences the environment type has on the outcomes. In addition to the *anger* and *gratitude* emotions we will be including the *admiration* emotion.

The admiration threshold in [7] increases when an agent compares its total payoff against each of its neighbours every five games. For our agents the neighbours are not as well defined because they will be moving constantly which changes who they are near to at a particular time. We will instead use the following to determine if admiration of an opponent has been triggered.

A mobile agent will complete five games of the prisoner's dilemma. After this the mobile agent will request the average payoff per game of its next opponent, before the game has started and compare the value to its own average payoff. Whoever has a higher average will gain the admiration point.

We are using average payoff rather than total payoff which was used in the original experiments because we cannot be sure that each mobile agent has engaged in the same number of games as its opponent. When the admiration threshold has been reached the agent will then take on the emotional characteristics of the agent that triggered the threshold which may be itself. Then the admiration of all agents is cleared, finally the agent plays the game with that opponent.

As per [7], there will be 14 scenarios that will be carried out. In each scenario there will be a number of initial defectors and cooperators, and number of agents with high, medium or low admiration thresholds. The first 5 have identical admiration threshold distributions, but have varying percentages of initial actions. This is to show how the makeup of initial actions can affect the evolution of co-operation. The remaining scenarios have varying admiration thresholds but identical distributions of initial actions, this will show us how differing distributions of admiration can affect co-operation. For a break down of each scenario see Table 3.

For each of these scenarios there will be a number of sub-scenarios which relate to the number of mobile agents there will be. The number of simulated robots will range from 9 to 144, with each emotional character being represented equally in each sub-scenario. They are represented equally so that when looking to see which emotional characteristic is dominant, we can say that the reason for the dominance is not because of a larger representation of the characteristic but due to the effects we are exploring. For the breakdown of the sub-scenarios that are used in combination with the scenarios see Table 4. We will then run each scenario and

Table 3: Experiment 2 scenarios

Scenario	Initial Defector %	Initial Co-operator %	Admiration % High	Admiration % Medium	Admiration % Low
1	90	10	34	34	32
2	70	30	34	34	32
3	50	50	34	34	32
4	30	70	34	34	32
5	10	90	34	34	32
6	50	50	50	25	25
7	50	50	70	15	15
8	50	50	90	5	5
9	50	50	25	50	25
10	50	50	15	70	15
11	50	50	5	90	5
12	50	50	25	25	50
13	50	50	15	15	70
14	50	50	5	5	90

Table 4: Sub-scenarios

Sub-scenario	No. of agents	No. of individual emotional characteristics
1 - Very low density	9	1
2 - Low density	36	4
3 - Medium density	72	8
4 - High density	144	16

sub-scenario combination ten times to gather a strong set of data to compare to the static agents. Each run will last ten minutes so that sufficient replication can take place.

The data we will be gathering during our experiments includes:

- Positional data of each agent, for every time it receives information about the world. This is usually every second.
- Each game that takes place with, who played the game, what time it occurred and what actions they chose. Including their total individual payoffs after the game takes place.
- The total number of games each agent played, the distance they have travelled and their final payoffs.
- How many of each emotional characteristic is represented at the end of the games.

Each scenario will be run first in the regular environment and then in the small world environment. This allows us to compare our results in each environment, noting if and how our emotional characters are affected by the change in environments. We can then show whether our simulated emotional agents are affected by environment types in a similar fashion to [11].

5. RESULTS

This section reports on the results of the above experiments. We will showing that our agents that move give the same results from the same games played as the agents that do not move. Then we will be showing the most successful characteristics, that is they were the most dominant characteristic by being the most prevalent characteristic after the ten minutes. We will then show what effects agent density, distance travelled and environment type has on an agent's average payoffs.

5.1 Validation

First to ensure that our emotional agents which move are reacting in the same way as the agents do not move, we compare how our emotional agents react to non-emotional agents. We compared out results to those in [8][1], Table 5 shows that our agents do react in the same way. From the table we can see that against agents which do not have randomness our mobile agents perform identically to their non-moving counterparts. Against agents which have randomness introduced, we can see that the average payoffs between the two types of agent are close, and that all of them have the same winners. This shows that our agents that move react in the same way as the agents that do not.

5.2 Main

5.2.1 *Effects of initial actions*

Figure 2 shows that the higher the percentage of defectors the lower the average score from each game an agent can expect to receive. This is an intuitive result as the more people that are co-operating the higher the chance of a *(CO-OP, CO-OP)* being achieved which raises the average. If the majority of games end in a *(DEFECT, DEFECT)* then this gives an average closer to one. The differences in environment type will be explained in Section 5.2.5

5.2.2 *Successful Characteristics*

We will compare which emotional characteristic is the most prevalent in our arenas and compare them to the prevalent character for the agents that do move which is character Trustful [7].

In Figures 3 and 4, we can see that in contrast to the static agents the most successful agent was the Non-Accepting agent, with Active and Distrustful not far behind. A similar contrast can be found in Figure 4 where Active is the

[1]Characters Responsive and Trustful are referred to as E1 and E7 respectively in [7].

Table 5: Comparison of average individual payoffs of initially co-operative emotional agents which move and those that do not move against non-emotional strategies

Character	Responsive Static	Responsive Mobile	Trustful Static	Trustful Mobile
Mendacious	204, 199	204, 199	212, 197	212, 197
Veracious	600, 600	600, 600	600, 600	600, 600
Random	451, 449	459.4, 457.4	630.4, 372.4	618.6, 367.4
Tit-for-Tat	600, 600	600, 600	600, 600	600, 600
Tester	533, 533	533, 533	668, 443	668, 443
Joss	233.4, 228.4	256.3, 251.3	523.4, 449.4	531.2, 467.2

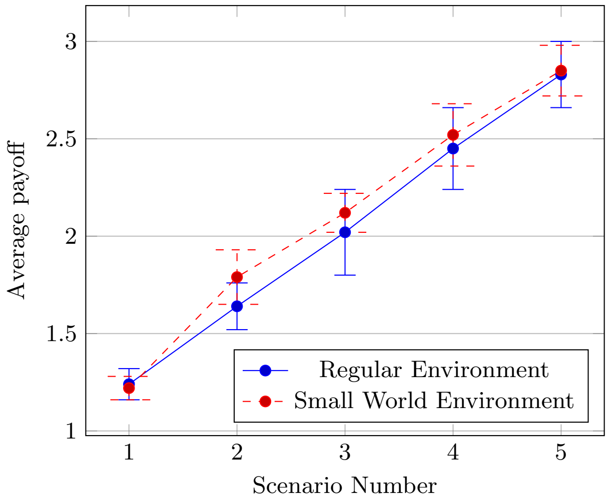

Figure 2: Average payoff per agent for scenarios with differing ratios of initial actions with standard deviations

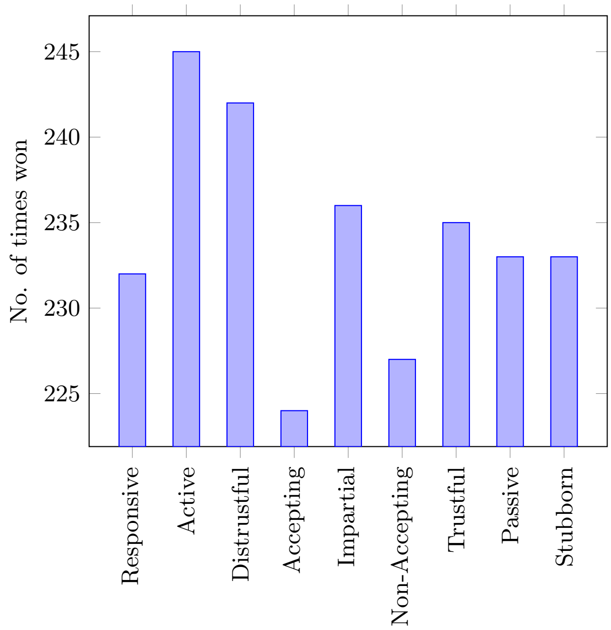

Figure 4: Dominant characteristic across all scenarios in a small world environment

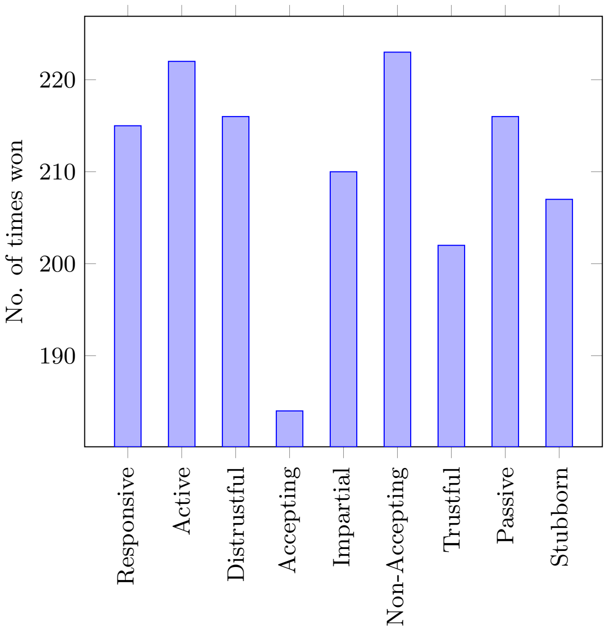

Figure 3: Dominant characteristic across all scenarios in a regular environment

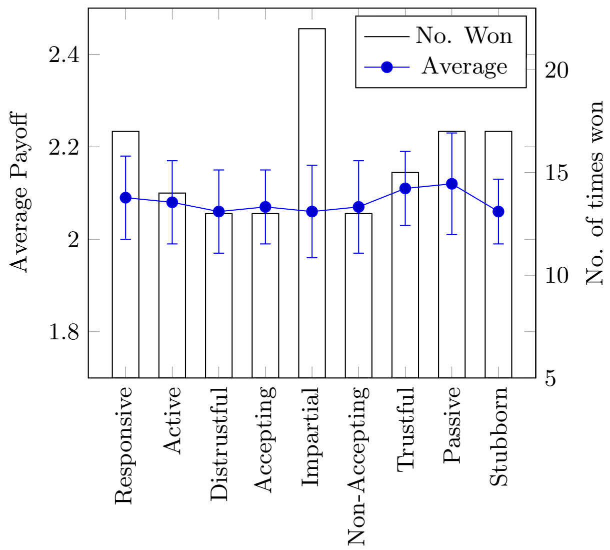

Figure 5: Average payoff with standard deviation against Dominant characteristic in a small world environment with a high density of agents

most successful characteristic with Distrustful again not far behind. The reasoning behind Trustfuls failure is that as it takes a long time to switch to defection, it does not end up punishing defectors since there is a chance that the opponent may not be played against again. Meaning that against agents that are constantly changing Trustful is taken advantage off to often without being able to punish that particular opponent. The differences in winning agents between the environment types will be discussed in Section 5.2.5.

We can also see that the Supportive characteristic, does not lend itself to being a dominant characteristic. The reasons for this is that with it reciprocating co-operation immediately it opens itself up to being taken advantage of, which the other characteristics do. The characteristic Supportive does not respond quickly to defection, it allows the advantage the other characters are taking to taken multiple times in a row. The reason this does not affect the Trustful character as harshly is that by waiting even longer to punish defection it allows co-operation to evolve between the two agents raising both their payoffs.

In Figure 5 we have taken average scores across scenarios 6 to 14 and the number of times each characteristic was most dominant in those scenarios. We have excluded scenarios 1 to 5 because as shown in Figure 2 the variation in the score is high which makes these figures less meaningful. The figure shows how having the highest average score makes that agent more likely to be dominant, but having a lower average score does not mean that the characteristic cannot dominate. This is because when a dominant character wins the majority of runs gets a higher score, such as the Impartial characteristic does in this particular instance. When Impartial is not dominant it performs particularly badly bringing its average score down. This is also shown by the higher standard deviation.

5.2.3 *Effects of Density*

The density of the robots, which was tested through the sub-scenarios can be seen to have an affect on the performance. Figures 6 and 7 show that the higher the density of agents the lower the average payoff per game for each agent. It has been shown that when the neighbours are fixed that cycles of defection occur within these emotional agents [7], with the higher densities the agents have less room to move. This lack of movement makes the agent play against the same group of agents as if they were fixed allowing these cycles of defection to occur, the higher the density the more these cycles appear in the environment. The differences in the environment will be discussed in Section 5.2.5

In low densities of agents we can see that the most successful agents are the ones which initially respond the same as the majority of the group and compete in the most games. This is because in low densities the number of games completed is very low, and by completing the most games you have the chance for the highest payoff. If the majority are defecting then there is not enough games for co-operation to evolve between two agents, if the majority are co-operating then there is not enough time for the advantage of defection to take affect, since the risk of getting a *(DEFECT, DEFECT)* reduces the payoff significantly.

Figures 6 and 7 also show how the density affects the range of potential average scores of an agent. When the density is very low, then the number of games completed between agents is also very low. For example an agent may only play

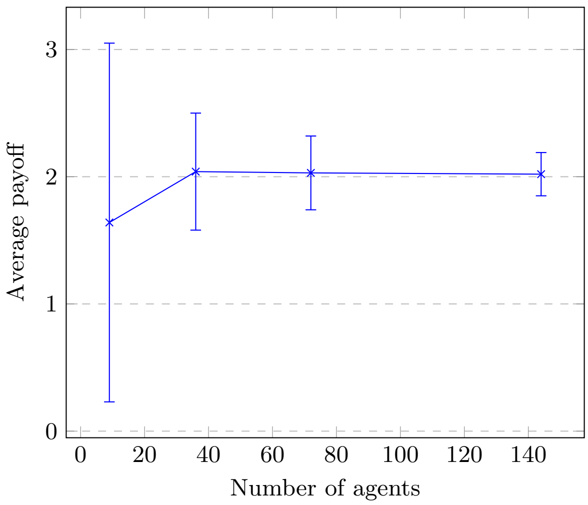

Figure 6: Average payoff per game for an agent in differing agent densities in a regular environment with standard deviation

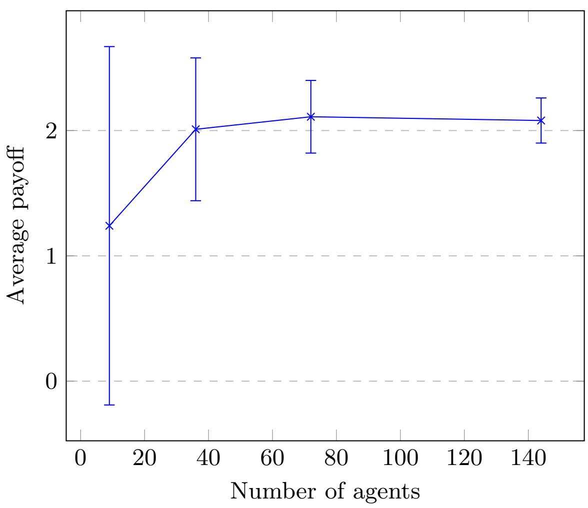

Figure 7: Average payoff per game for an agent in differing agent densities in a small world environment with standard deviation

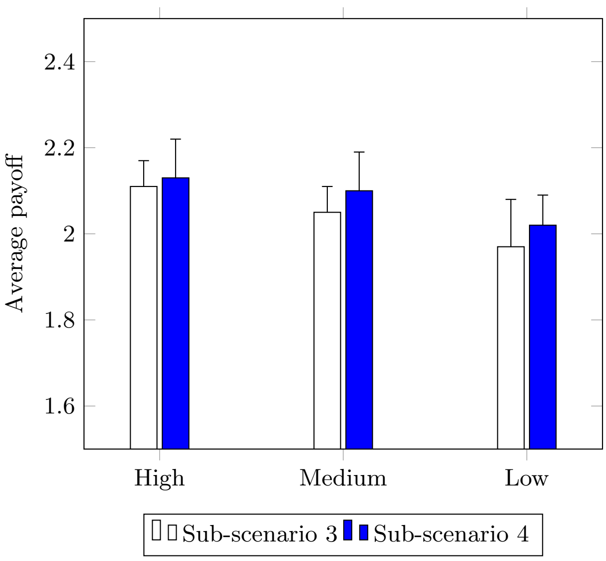

Figure 8: Average payoff per game for an agent based on distance travelled in a small world environment with standard deviation

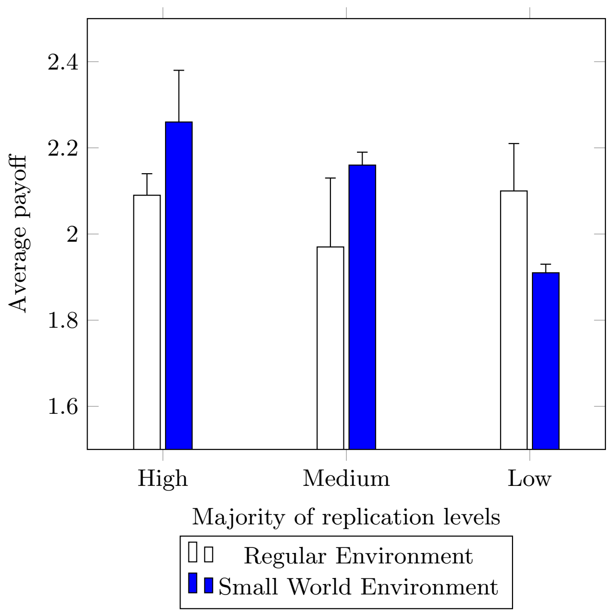

Figure 9: Average payoff per game for an agent based on distribution of admiration thresholds with average deviation.

two games over the run if the results are *(DEFECT, CO-OP)* and *(CO-OP, CO-OP)* then the average of one agent will be 4, but the other agent will have an average of 1.5. This occurs less often as the density increases as the number of games completed also increases. The average score will better reflects the performance of the agents at higher densities, this is shown by the decreasing standard deviations.

5.2.4 Effects of Distance Travelled

To see what effects movement distance has on the agents, we first defined a high mover as an agent that travels for more than 30 metres in a game, medium as over 15 but 30 or below and a low mover as 15 or below. Figure 8 show the average payoff per agent is affected by the distance travelled. We have excluded sub-scenario 1 and 2 due to the lack of low movers. We can see from the figure that the more an agent moves the higher its average payoff. The differences between each distance threshold is more pronounced the fewer agents there are in the environment. This is because agents that move more do not get stuck in cycles of defection as often because they are not stuck playing against the same agents.

5.2.5 Effects of Environment Type

After taking the fact that lower densities have large variations in them, we can see from Figures 2, 6 and 7 that the average payoff in the small world environment is slightly higher than the regular environment. The average payoff is reducing as cycles of defection are occurring as agents are playing against the same agent multiple times. The small world environment has shortcuts throughout the environment which allows agents to break these cycles by moving to another part of the environment, whereas in the regular environment the groups of agents are larger without these shortcuts which enable agents to move away from these defection cycles bring the average payoff down more quickly.

The success of a particular agent is related to the environment type, as shown in Figures 3 and 4 the success of the Non-Accepting agent is dependant on the type of environment. We see that the Non-Accepting agent is successful in a regular environment but not in a small world environment. The reason it is successful in the regular environment is that it takes advantage of the fact that agents can not move away easily and that it will play against the same agents multiple times. It is not unreasonable as it allows intermittent defections to take place, ensuring that co-operation is not broken. However it takes the advantage by not reciprocating co-operation until many games have been played.

When the Non-Accepting agent plays in the small world environment the advantage it tries to take from the co-operators in the group is reduced as the co-operators are free to move to other areas of the environment. This leaves the agent without the ability to create the co-operation cycles to increase its payoff.

We can also see in these figures that there are agents that do well in both types of environment, namely the Active agent and the Distrustful agent. They are both quick to switch to defection ensuring that they do not get taken advantage of. They both withhold reciprocating co-operation, the Active agent does better as cycles of co-operation are created more quickly.

The environment type also affects how quickly the agents should choose to adapt their characteristics, as seen in Figure 9. From this figure you can see that in a small world environment an agent should wait until more games have been played before it changes its characteristics. This is due to the fact that more games against different opponents can have a effect which characteristic is doing best, so waiting until a characteristic is clearly dominant is better for the agent.

In a regular environment the effect is less pronounced but there is a difference, the agent should either change its characteristic as quickly as possible or wait until the dominant characteristic is known.

6. CONCLUSIONS

These experiments have shown that the distance travelled, the type of environment and the density of the agents all have an effect on the success of agents. By travelling more the agent can increase its payoff. An agent can also increase its payoff by waiting for a dominant characteristic to show and then copying that characteristic, rather than changing its characteristics more often.

The type of environment effects which strategies can be viable with the Non-Accepting agent being successful in a regular environment but not the small world environment. However there are strategies that are successful regardless of the environment type, namely the Active characteristic.

We can come up with a general set of rules on how to succeed when mobility is introduced regardless of the environment type, the rules are:

- Follow the group initially
- Keep moving
- Punish defection quickly
- Wait to reciprocate co-operation, but not too long.
- If there is a more dominant strategy wait before copying it.

7. FUTURE WORK

Now that these experiments have been completed, we can now further expand this body of work. We will show that this work completed is applicable to the real world by implementing this experiment on real world e-pucks. We will also be including the addition of mood to our robots to see how this can improve co-operation between agents. Mood is distinct from emotions but is influenced from the same input, we will be using a positive or negative mood which will effect how the agent responds to new agents. The main difference between emotions and mood is that emotions are short-term and mood is long term feelings.

In addition we will also be looking into how to improve co-operation through influencing the decision making of select agents. By selecting an agent through a given criteria we can force that particular agent to either co-operate or defect, and we will be able to see how this affects co-operation.

We will also be testing our strategy that we have developed through our results to validate this strategies in these scenarios against emotional agents.

REFERENCES

[1] Elisabeth André, Martin Klesen, Patrick Gebhard, Steve Allen, and Thomas Rist. Integrating models of personality and emotions into lifelike characters. In Ana Paiva, editor, *Affective Interactions*, volume 1814 of *LNCS*, pages 150–165. Springer, 2000.

[2] Robert Axelrod and William Donald Hamilton. The evolution of cooperation. *Science*, 211(4489):1390–1396, 1981.

[3] Daan Bloembergen, Bijan Ranjbar-Sahraei, Haitham Bou Ammar, Karl Tuyls, and Gerhard Weiss. Influencing social networks: An optimal control study. In *Proc of ECAI'14*, pages 105–110, 2014.

[4] Brian Gerkey, Richard T Vaughan, and Andrew Howard. The player/stage project: Tools for multi-robot and distributed sensor systems. In *Proc. of the International Conference on Advanced Robotics*, pages 317–323, 2003.

[5] Dacher Keltner and James J. Gross. Functional accounts of emotions. *Cognition & Emotion*, 13(5):467–480, 1999.

[6] Robert W. Levenson. Human emotion: A functional view. *The nature of emotion: Fundamental questions*, 1:123–126, 1994.

[7] Martyn Lloyd-Kelly, Katie Atkinson, and Trevor Bench-Capon. Developing co-operation through simulated emotional behaviour. In *13th International Workshop on Multi-Agent Based Simulation*, 2012.

[8] Martyn Lloyd-Kelly, Katie Atkinson, and Trevor Bench-Capon. Emotion as an enabler of co-operation. In *ICAART (2)*, pages 164–169, 2012.

[9] Francesco Mondada, Michael Bonani, Xavier Raemy, James Pugh, Christopher Cianci, Adam Klaptocz, Stephane Magnenat, Jean-Christophe Zufferey, Dario Floreano, and Alcherio Martinoli. The e-puck, a robot designed for education in engineering. In *Proceedings of the 9th conference on autonomous robot systems and competitions*, pages 59–65. IPCB: Instituto Politécnico de Castelo Branco, 2009.

[10] Andrew Ortony, Gerald L Clore, and Allan Collins. *The cognitive structure of emotions*. Cambridge university press, 1990.

[11] Bijan Ranjbar-Sahraei, Irme M. Groothuis, Karl Tuyls, and Gerhard Weiss. Valuation of cooperation and defection in small-world networks: A behavioral robotic approach. In *Proc of BNAIC 2014*, 2014.

[12] Francisco C. Santos, Marta D. Santos, and Jorge M. Pacheco. Social diversity promotes the emergence of cooperation in public goods games. *Nature*, 454(7201):213–216, 2008.

[13] Norbert Schwarz. Emotion, cognition, and decision making. *Cognition and Emotion*, 14(4):433–440, 2000.