The Effect of Mobility and Emotion on Interactions in Multi-Agent Systems

Joe COLLENETTE [a,1] Katie ATKINSON [a] Daan BLOEMBERGEN [a] Karl TUYLS [a]
[a] *Department of Computer Science, University of Liverpool, UK*

Abstract. Simulating emotions within a group of agents has been shown to support co-operation in the prisoner's dilemma game. Most work on simulating these emotions 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 in other work that when an agent is given the ability to move, then the type of the environment affects how co-operation evolves in the group of agents. In this paper, we investigate the combination of these two ideas in an experimental study that explores the effects on co-operation when autonomous agents that can show emotions are given the ability to move within structured environments. We observe that once mobility is introduced, different strategies become successful. Successful strategies respond quickly to defection, while not immediately reciprocating co-operation, regardless of the environment type. The further an agent travels, the higher its average payoff in a small world environment. The slower an agent is to copy another agent by imitating its strategy, the higher its increase in average payoff.

Keywords. Emotion, Multi Agent, Co-operation, Mobility, Prisoner's Dilemma

1. Introduction

It is well known in psychology that emotions affect human decision making [1]. Recently, it has been shown that simulating emotions within autonomous agents (which we refer to as emotional agents throughout the paper) can similarly influence the evolution of co-operation within the prisoner's dilemma game [2,3]. So far, the work on emotional agents and their co-operation has focused on static agents which do not have mobility. In related work, Ranjbar-Sahraei et al. have shown that when agents (without emotions) are given mobility, the environment type has a clear influence on the evolution of co-operation in the prisoner's dilemma game [4].

In this work, we combine previous efforts by giving simulated emotional agents the opportunity to move around in the environment, and therefore allowing them to interact with many other agents over time. We examine whether the environment structure has the same effect on emotional agents as it does on non-emotional agents. By giving our agents mobility we aim to give a more accurate description of the evolution of co-operation than within simulated emotional agents, in a multi-agent setting.

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

[1] Corresponding author: email: j.m.collenette@liverpool.ac.uk

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]. It is important here to make clear the distinction made in the psychology literature between mood and emotion since they are closely related. Emotions are short-term feelings that are directed towards a particular object or person [6]. Mood in contrast is a long term feeling which does not have this focus on a particular object or person [7]. In this paper we will focus purely on this short-term and directed characterisation of emotions, which defines the scope that an emotion has within our agents.

Our study addresses two main questions. Firstly, does the addition of emotions influence the way in which the environment affects the evolution of co-operation? To answer this question we compare our findings with those obtained by Ranjbar-Sahraei et al. [4] for mobile but non-emotional agents. Secondly, how do mobility and the environment type affect the simulated emotions and the resulting evolution of co-operation? This second question directly contrasts our work to that of Lloyd-Kelly et al. [3] who looked at emotional but static agents. As such, our work highlights the interplay between emotions and mobility, as well as their joint influence on co-operation within a society of self-interested decision makers. The main purpose of this work is to better understand how emotions influence interactions and decision making.

We will consider two different types of environments: a regular environment, and a small-world environment. The regular environment is grid-like, as all intersections have the same number of exits. This means that agents need to move long distances to interact with agents on the other side of the map. The small-world environment is similar to the regular environment, but contains shortcuts for agents to move over to different parts of the map quickly. This reduces the average distance between any pair of agents, in similar fashion to the small-world networks of Watts and Strogatz [8].

In our environments we will be simulating small disc shaped robots, and we will be using the player/stage simulator [9]. We are simulating robots rather than mathematical models of graph-based interactions, as this naturally allows to emulate a number of interesting properties such as asynchronous interactions, dynamic neighbourhoods, and differing rates of interaction between agents. We simulate these properties to better understand how mobility affects co-operation in a realistic setting. We anticipate these results will inform a broader study on designing multi-agent systems with desirable properties related to agent co-operation.

2. Related Work

There are various ways in which researchers have introduced emotions into a computational framework. These frameworks vary from implementations of emotions in a logic [10] to more practical applications in human-computer interaction [11]. Most works use the OCC (Ortony, Clore and Collins) model of emotions [12] as the base for their implementation, although this is not the only model of emotions [13]. We use the OCC model in this work due to its accepted use in agent-based systems as well as its flexibility in implementation. Lloyd-Kelly et al. explore how these emotions influence co-operation in multi-agent interactions [2,3]. However, their work is limited to static agents, in the sense that agents always interact with the same opponent. We build on the work of Lloyd-Kelly et al. by considering a system of mobile agents, equipped with emotions, which

Table 1. Payoff matrix of the prisoner's dilemma.

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

interact with different partners over time, and study how mobility affects the evolution of co-operation in different types of environments.

In recent work, Ranjbar-Sahraei et al. show how co-operation evolves within a society of mobile agents [4]. The authors simulate robots in two types of environments, regular and small-world. However, in their work they do not consider the effect of emotions. We base our simulation model on the work of Ranjbar-Sahraei et al., while incorporating the emotional characters of Lloyd-Kelly et al., allowing to compare our results directly to theirs while simultaneously being able to isolate the effect of both emotions and mobility.

A large body of related work deals with the evolution of co-operation in (social) networks, in particular in those scenarios where co-operation is costly but ultimately beneficial for all agents, often modelled as the classical Prisoner's Dilemma [14]. A large body of work focuses for example on how structural network properties and interaction mechanisms determine whether co-operation is sustainable [15,16]. Others have focused on developing strategies from the ground up that will support co-operation [17,18]. Closely related is the work of Van Veelen et al., who extensively study a class of "Tit-for-Tat"-based strategies [19]. Although the emotional characters we consider in the work are highly related, Van Veelen et al. do not link their strategies to psychological character traits as we do here, nor do they consider mobility.

We are building on previous works from psychology, game theory, and graph theory. By adding emotions to social dilemma strategies we can take the first steps into understanding how simulated emotions and their differing characteristics interact in establishing co-operation and what outside effects can occur on this co-operation.

3. Method

In the following we will describe how our experiments will be conducted including a brief introduction to the prisoner's dilemma game, which is our model of interaction. We also introduce the different emotional characters used and how agents interact with the environment and each other.

3.1. Prisoner's Dilemma

In the prisoner's dilemma two players have the choice of either defecting or co-operating; choices are made simultaneously and without prior communication. Both players then get a payoff depending on the choices made. The payoffs for the game are 3 for each agent when they both co-operate, 1 for each agent when they both defecting and 5 for the agent which defects in a non-mutual outcome and 0 for the co-operative agent.

The game matrix is shown in Table 1, with player one choosing a row, player two choosing a column, and both players receiving the payoff indicated in each cell.

When looking at the prisoner's dilemma outcomes, it seems in the best interest of both players to play co-operatively since this would lead to the largest total payoff. How-

Table 2. Emotional characters, as used in this work and previous work [2].

Anger Threshold	Gratitude Threshold	Character	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			

ever, there is a temptation to defect as this can lead to a higher individual payoff. When both players reason this way, this then leads to the Nash equilibrium of *(DEFECT, DEFECT)*, which gives the worst outcome for the group as a whole. This highlights the dilemma of the game. Investigating methods by which self-interested agents can be made to co-operate in the prisoner's dilemma has been an active area of research in the past decades, with a particular focus on the evolution of co-operation within groups of agents [14,20]. It is for this reason that we adopt this model of interaction as well.

3.2. Emotional Characters

The simulated emotions we implement are based on the Ortony, Clore and Collins model of emotions, known as the OCC model [12]. This was developed from psychology research and has been used within the AI community [11,2] to simulate emotions within agents. The emotions we will be modelling are *anger, gratitude* and *admiration*. Our initial focus is on these emotions as they were the focus of previous work [2,3] and the most relevant as a starting point in our study.

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 the emotions in a computational setting.

Our implementation of these emotions is similar to previous work by [2]. This allows us to compare the differences caused by mobility and environment structure rather than implementation. Each emotion has a threshold, and when that threshold is reached it triggers a change in the agent's behaviour. Specifically, when the anger threshold is reached the agent changes to defection, and when the gratitude threshold is reached the agent changes 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, and are intended to show a range of characteristics that could reflect a simple simulation of personality differences.

Admiration thresholds can similarly be rated as high (3), medium (2) or low (1) — these are not listed in the table as they are independent from the emotional character, and will be detailed later. An agent's anger increases by one when its opponent defects; gratitude increases when the opponent co-operates. Admiration increases when the agent believes that its opponent is performing better than itself (explained further in the next sub-section). When a threshold is reached, the agent's behaviour changes as described above and the value is then reset back to 0.

The agents are placed in a random location initially, while checking that no agents are placed at the same position. Moreover, the different emotional characters, admiration thresholds, and initial actions, are distributed randomly and independently among the agents, given the specific proportions of each for that experiment. Details of the different scenarios are given in the *Experiments* section.

3.3. Agent Interactions

The agents are given a random walk behaviour with some basic obstacle avoidance procedures. Each agent has proximity sensors to detect walls and obstacles, located at $\{-90°, -45°, -15°, 15°, 45°, 90°\}$ w.r.t. the robot's heading. If the sensors on the left detect anything, the agent will stop and then turn to the right, and the reverse for the right sensors. The robot's speed is set at 10 centimetres per second and can turn at 45 degrees per second. When no obstacles are detected the agent randomly selects a turn speed between $-45°$ and $45°$ per second while moving forward. Since a new heading is generated each time the robot receives sensor data this results in a random movement pattern.

The prisoner's dilemma game is initiated whenever two agents are in close proximity, and have line of sight of each other. The game is played once, after which they will then continue their random walk behaviour. The payoffs of the game are given in Section 3.1. The agents have no knowledge of these payoffs or the number of games to be played, and will purely use the strategy given by their emotional character to play. The agent has no knowledge of the strategies or emotional characters of its neighbours, but it can differentiate between them, and the emotions it has apply specifically to the agent it is playing against. The agents have no knowledge of the environment; they will only use the random walk behaviour driven by their sensor inputs.

In the work of [2], the admiration threshold 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 a modified version of the trigger for admiration.

The modified version is that when a mobile agent completes five games of the prisoner's dilemma. After that, the mobile agent will request the average payoff per game of its next opponent, before the game has started, and compares this value to its own average payoff. The agent will increase its admiration value towards whoever has the highest average, this will be either itself or its opponent. We are using average payoff, rather than total payoff which was used by [2], 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, then the agent takes on the emotional characteristics of the agent that triggered the threshold, which may be itself, so the agent will then respond to other opponents in the same way as the agent who triggered the admiration threshold. Then the admiration threshold is reset to zero. Finally, the agent plays the game with its opponent.

The average payoff is obtained directly from the opponent, since we study how effective these agents are in an ideal situation we force all agents to be truthful. Similarly the agent will not lie when communicating the emotional characteristics it is currently inhabiting. Exploring how lying can affect these emotional agents is an interesting topic but it is out of scope of this paper since we are most interested with isolating the effects of movement on a mixed group of emotional agents.

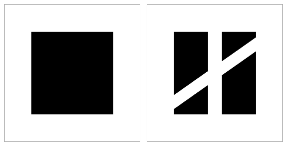

Figure 1. The two 5×5 meter environments used in this work.

4. Experimental Setup

We conduct two sets of experiments. First, a validation experiment is designed in order to check if our implementation reproduces the results obtained by [3] for static agents, which would allow for a direct comparison. Second, in our main experiment we extensively investigate the effects that mobility and the environment type have on the evolution of co-operation. In the following we describe the set-up of these two experiments.

4.1. Validation Experiment

The aim of this experiment is to show that our mobile agents have the same emotional response and outcomes as the static agents reported by [3]. In this experiment we will only be using the emotions *gratitude* and *anger*, as these were the emotions used in the original experiment [2]. The emotional agents will play the iterated prisoner's dilemma against a fixed-strategy agent that does not use emotions. The emotional agents will be set to co-operate initially. The 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 fixed strategies that the emotional agents will be tested against are the traditional ones from Axelrod's tournament [14] and are described in [2].

In this experiment there are only two agents in the environment: the emotional agent, and the fixed-strategy agent. For each emotional character of Table 2 we will perform 10 runs against each fixed strategy in turn. A run consists of simulating the mobile agents until 200 rounds of the prisoner's dilemma game have been completed, equal to the set-up used by [3]. This should make the results identical, up to a slight variation caused by chance in the Random and Joss strategies. The Joss strategy plays tit-for-tat with a 10% chance of defection.

4.2. Main Experiment

This experiment aims to highlight the differences and similarities between mobile and static emotional agents, as well as showing what influence the environment type has on the outcomes. In addition to the *anger* and *gratitude* emotions, here we will also include the *admiration* emotion. As in [2], there will be 14 scenarios that will be investigated. Each scenario is defined by the number of initial defectors and co-operators, and the number of agents with high, medium, or low admiration thresholds. The first 5 scenarios have identical admiration threshold distributions, but have varying percentages of initial actions. The remaining scenarios have varying admiration thresholds but identical distributions of initial actions. For a break-down of each scenario, see Table 3.

For each of these scenarios there will be a number of sub-scenarios using different numbers of agents. The number of simulated mobile agents will range from 9 to 144,

Table 3. Scenarios used in the main experiment. A_i is the admiration threshold.

	S1	S2	S3	S4	S5	S6	S7	S8	S9	S10	S11	S12	S13	S14
Initial Defect %	90	70	50	30	10	50	50	50	50	50	50	50	50	50
Initial Co-Op %	10	30	50	70	90	50	50	50	50	50	50	50	50	50
High A_i%	34	34	34	34	34	50	70	90	25	15	5	25	15	5
Med A_i%	34	34	34	34	34	25	15	5	50	70	90	25	15	5
Low A_i%	32	32	32	32	32	25	15	5	25	15	5	50	70	90

Table 4. Sub-scenarios.

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

with each emotional character being represented equally in each sub-scenario. For the break-down of sub-scenarios, see Table 4. We have included these sub-scenarios as we expect when the number of agents increases the density of the agents will increase, since the environment is still the same size. We predict that the increased density will prevent the agents from moving further afield in the environment as other agents will block their path. We then expect an agent will need to move less to initiate a game, and so give the agents a higher probability of playing against the same agents more often.

Having an equal distribution of emotional characters initially makes sure that we test character strength without being affected by characteristics having a initially higher represention. We will run each combination of scenario and sub-scenario 10 times, and compute average statistics for each. Each run will last for 10 minutes during which the agents move around and interact, which allows sufficient interactions and replication to take place. We expect Trustful to be the most dominant characteristic as in previous work [2] however as mobility is introduced we expect some changes in lower rankings due to agents playing against a larger range of characteristics. This will show us that while mobility does affect the characteristics, high performing characteristics should be affected to a minor extent.

5. Results

5.1. Validation Results

We investigate how our emotional characters perform against the static strategies discussed. To compare our results to those in [3], we focus on the Responsive and Trustful characters[2] as they are characters whose individual scores were reported. The results of this experiment show that our agents do indeed react in the same way. We observe that

[2]Characters Responsive and Trustful are referred to as E1 and E7 respectively in [3]

Table 5. The mutual outcomes that occur between two agents i and j with differing initial actions, where I_i is mutual co-operation or defection depending on the initial action of agent i. C is mutual co-operation, D is mutual defection and R is a repeated loop of $(COOP, DEFECT)$ then $(DEFECT, COOP)$.

Character	1	2	3	4	5	6	7	8	9
Responsive (1)	R	D	D	C	I_j	I_j	C	I_j	I_j
Active (2)	D	D	D	R	D	D	C	I_j	I_j
Distrustful (3)	D	D	D	D	D	D	R	D	D
Accepting (4)	C	R	D	C	C	I_j	C	C	I_j
Impartial (5)	I_i	D	D	C	R	D	C	C	I_j
Non-Accepting (6)	I_i	D	D	I_i	D	D	C	R	D
Trustful (7)	C	C	R	C	C	C	C	C	C
Passive (8)	I_i	I_i	D	C	C	R	C	C	C
Stubborn (9)	I_i	I_i	D	I_i	I_i	D	C	C	R

against agents which do not have randomness, our mobile agents perform identically to their static counterparts. Against agents which have randomness introduced (Random and Joss), 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 mobile agents react in the same way as their static counterparts, and that our results will be directly comparable.

5.2. Main Results

When we look at the interactions between pairs of agents there are a number of patterns that emerge between them. When two agents start with identical initial actions the result of the game will be continued mutual co-operation or defection without deviation. When the initial actions are different then a number of different patterns emerge. The agents will play a series of $(COOP, DEFECT)$ cycles then after a number of interactions turn to mutual defection or co-operation and then continue this indefinitely. The agents may under certain conditions continue this $(COOP, DEFECT)$ cycle indefinitely without settling on a mutual action between them. The mutual action they choose is dependant on a number of conditions, namely their gratitude and anger thresholds, but it may also depend on the their opponent's thresholds. We will now define what the conditions are for each mutual outcome that can occur.

To show the mutual outcome they will produce indefinitely, we define $\Omega_{i,j}$ to return this mutual action where i and j are emotional agents, A_i is the anger threshold of agent i and G_i is the gratitude threshold of agent i. Ac_i returns the current action of agent i.

$$\Omega_i^j = \begin{cases} COOP, & \text{If } (Ac_i = Ac_j = COOP) \text{ or } (Ac_i = COOP \text{ and } G_j < A_i) \\ DEFECT, & \text{If } (Ac_i = Ac_j = DEFECT) \text{ or } (Ac_i = DEFECT \text{ and } A_j < G_i) \\ NotMutual, & \text{If } A_i = G_j \text{ and } G_i = A_j \\ \Omega_j^i & \text{Otherwise} \end{cases}$$

Table 5 shows what the mutual actions will be between our emotional characteristics when paired against each other. This table also shows us that when two agents are paired against each other and they have differing initial actions, they are more likely to settle into mutual defection.

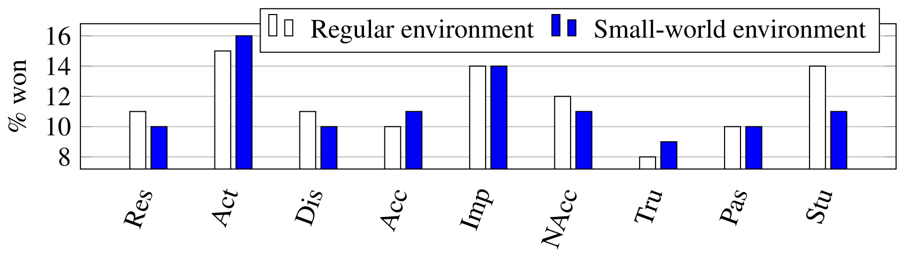

Figure 2. The percentage of games each characteristic was most dominant excluding draws, for all scenarios.

5.2.1. Dominant Characteristic

When we look at Table 5 we can see that the Trustful agent has almost all outcomes end in mutual co-operation, we then expect this characteristic to be the most dominant as this would give the highest average when compared to the other types of agent, which reflects our hypothesis and previous work [2]. However when we look at Figure 2, which shows which characteristics are dominant, we can see that Trustful is not only not the most dominant, but is consistently the worst performing agent in our experiments. This shows us that the environment does have an effect when compared to static agents, but the effect between environments is much smaller.

The reason behind Trustful's failure comes down to the fact that it takes a long time to settle into its mutual outcomes. Trustful will take up to three defections before it changes the opponent to co-operation and with the average number of interactions between two specific agents only being 3.37 (Standard Deviation 10.71) then the agent is unlikely to settle into mutual co-operation during the run. Active's success is due to its ability to be on the side that receives a payoff in $(COOP, DEFECT)$ outcomes. When it is on the receiving end of the sucker-punch, the agent punishes defection immediately, and so settles into mutual defection preventing the agent from receiving more than one zero payoff from its opponent.

Responsive is affected by responding to co-operation too quickly as it is unable to boost its score over time with the sucker-punches. While it may appear that Distrustful should therefore do well, the agent cannot get into mutual co-operation as the number of interactions is long enough for mutual co-operation to occur which brings its average score down. The Active characteristic differs from the Distrustful agent as it can get into these mutual co-operation cycles quickly raising the Active agent's average payoff. Impartial can take an advantage from other agents and can get into co-operation cycles quickly, however to achieve this it is open to being taken advantage of, lowering its average score.

5.2.2. Density Effects

When we look at the average scores of an agent in differing densities as shown in Table 6 we can see that the payoff is initially small with a large variance. Increasing the density lowers the standard deviation. The reasoning behind this is clearer when we look at how many interactions each agent expects to get in these densities, which is shown in Table 7. In very low densities the agents are only going to get one or two interactions which prevents them settling on an outcome and since there are so few games there is a large disparity in the results as very few agents can respond to their opponent's actions.

Table 6. Average Scores (Standard Deviation) for all agents across all scenarios for varying number of agents.

	Regular	**Small World**
9	1.64 (1.41)	1.24 (1.43)
36	2.04 (0.46)	2.01 (0.57)
72	2.03 (0.29)	2.11 (0.29)
144	2.02 (0.17)	2.08 (0.18)

Table 7. Average number of interactions (Standard Deviation) between specific agents for varying number of agents.

	Average Number Of Interactions
9	1.56 (0.83)
36	1.92 (1.93)
72	2.15 (1.73)
144	3.13 (15.08)

Table 8. Average payoff (Standard Deviation) for an agent per interaction based on distance travelled in a small-world environment.

	Sub-Scenario 3	**Sub-Scenario 4**
High	2.11 (0.06)	2.13 (0.09)
Medium	2.05 (0.06)	2.10 (0.09)
Low	1.97 (0.11)	2.02 (0.07)

When the density increases and the number of interactions increases as well we can see the variance in average scores becomes much closer, however it is slowly falling. When the number of interactions is very high the agents will settle into their mutual outcomes with the majority being mutual co-operation or defection, whereas in slightly lower densities they have not settled, leading to a majority of mixed outcomes. When half of the agents are in mutual co-operation and the other in mutual defection, the overall average will be 2, but the mixed outcome average will be around 2.5, showing the slight dip in the average scores in higher densities.

There is a large variance in the number of interactions in higher densities so we considered how this can affect the average score when we based this on distance travelled. We split the agents into three groups of roughly the same size and categorised the high movers as having moved 30 meters or more, medium as between 15 and 30, and low as 15 or less. The scores can be seen in Table 8. The table shows that the more an agent moves the higher the average score. When movement is high the number of repeated interactions decreases leading to these agents not settling on their mutual outcomes, which is the reverse for the low movers. When the number of interactions is low we have seen that the average score is slightly higher and this is again reflected by the distance moved.

5.2.3. Environment Effects

The environment has add a larger effect than expected, so we have examined more closely what effects it has. When we look at Figure 3, we can see that the more co-operators there are, the bigger the payoffs and that differing admiration thresholds do not have a significant effect on total payoffs. However we do see a difference in environment; this is due to the fact that small world environments have a larger surface area than our regular environment, and as we have seen previously there are less interactions between specific agents in lower densities. As the density is lowered we expect that the average score should also increase. Using scenario 3, as it has equal distributions in each category, the average score (Standard Deviation) of the regular environment is 2.02 (0.22) and the small world environment is 2.12 (0.10). We again see that repeated interactions are lowered the more average score increases.

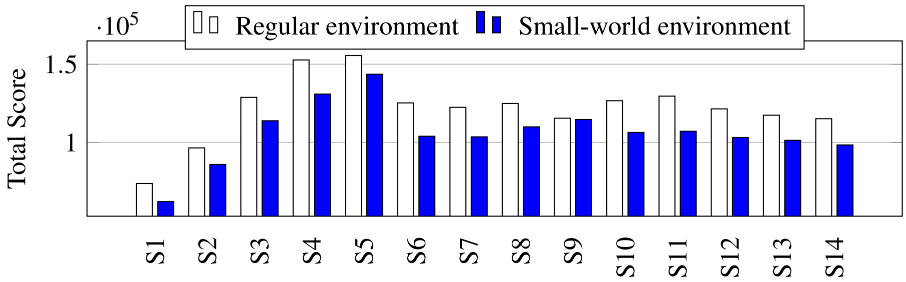

Figure 3. Total scores for each scenario

Table 9. Average payoff (Standard Deviation) per game for an agent based on distribution of admiration thresholds in both environments.

	Regular	Small World
S8 (High)	2.09 (0.05)	2.26 (0.12)
S11 (Medium)	1.97 (0.16)	2.16 (0.03)
S14 (Low)	2.10 (0.11)	1.91 (0.02)

Table 10. Breakdown of interactions in both environments across all runs.

	Regular	Small World
Total Interactions	417372	361682
Unique Interactions	115261	115653
Percentage Unique	28	32

Table 9 shows the average scores for scenarios 8, 11 and 14 which have differing admiration thresholds. We see that the regular environment is relatively stable whereas the small world environment shows a drop in scores from high to low thresholds. The difference is due to the percentage of unique interactions in each environment, which is shown in Table 10. In small world environments we see that agents interact with individual agents less often. When they come to replicate using their admiration thresholds, the chance of them replicating into a characteristic which is not dominant is increased. This is due to the average scores not reflecting the performance of a characteristic accurately; this then prevents them achieving higher scores.

6. Conclusions & Future Work

We have investigated the evolution of co-operation in mobile emotional agents. Our 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 the agents. This is due to how these affect the number of unique interactions. The environment type affects which strategies are viable, with the Stubborn character being successful in a regular environment, but not as successful in a small world environment. However, strategies exist that do well regardless of the environment type, such as the Active character.

In answer to our questions posed in the introduction, we have shown that mobility and environment types do affect the simulated emotional agents, and as a result different emotional characters become more successful as compared to those of [3] which was an unexpected result as it shows that mobility has a large effect on the success of a characteristic. This answers our second question. In answer to the first question, we have seen that in regular environments payoffs increase with more co-operators, as also shown by [4]. In contrast, in the small-world environment, [4] found payoffs decreasing with the addition of co-operators, whereas we see an increase. While we also see effects of environment type further investigation will provide more evidence of these effects. An interesting prospect for future work is the addition of *moods* to our agents, to see whether this can improve co-operation. We are also interested in investigating the effect of state interventions on co-operation levels within a society of agents.

References

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

[2] 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.

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

[4] 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.

[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] Elizabeth K Gray, David Watson, R Payne, and C Cooper. Emotion, mood, and temperament: Similarities, differences, and a synthesis. *Emotions at work: Theory, research and applications for management*, pages 21–43, 2001.

[8] Duncan J. Watts and Steven H. Strogatz. Collective dynamics of 'small-world' networks. *Nature*, 393(6684):440–442, 1998.

[9] Brian Gerkey, Richard T Vaughan, and Andrew Howard. The player/stage project: Tools for multi-robot and distributed sensor systems. In *Proc. of ICAR'03*, pages 317–323, 2003.

[10] Bas R. Steunebrink, Mehdi Dastani, John-Jules C. Meyer, et al. A logic of emotions for intelligent agents. In *Proc of AAAI'07*, volume 22, page 142, 2007.

[11] Elisabeth Andr, 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.

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

[13] Stacy Marsella, Jonathan Gratch, and Paolo Petta. Computational models of emotion. *A Blueprint for Affective Computing-A sourcebook and manual*, pages 21–46, 2010.

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

[15] Lisa-Maria Hofmann, Nilanjan Chakraborty, and Katia Sycara. The Evolution of Cooperation in Self-Interested Agent Societies: A Critical Study. *Proc. of AAMAS*, pages 685–692, 2011.

[16] Bijan Ranjbar-Sahraei, Haitham Bou Ammar, Daan Bloembergen, Karl Tuyls, and Gerhard Weiss. Evolution of cooperation in arbitrary complex networks. In *Proc. of AAMAS*, pages 677–684. IFFAAMAS, 2014.

[17] Christian Hilbe, Arne Traulsen, and Karl Sigmund. Partners or rivals? strategies for the iterated prisoner's dilemma. *Games and Economic Behavior*, 2015.

[18] 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.

[19] Matthijs Van Veelen, Julián García, David G Rand, and Martin A Nowak. Direct reciprocity in structured populations. *Proceedings of the National Academy of Sciences*, 109(25):9929–9934, 2012.

[20] 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.