Modelling Mood in Co-operative Emotional Agents

Joe Collenette, Katie Atkinson, Daan Bloembergen and Karl Tuyls

Abstract Simulating emotions based on psychological models has been a topic where work has focused on social dilemmas using simulated emotions to inform decision making within artificial agents. However human decision making is affected not only by emotions but also by other aspects of people's temperament: the mood of the person also affects their decision making, in conjunction with other factors such as inequity aversion. We propose a simulated model of mood, which is formed and validated through psychological research. We use this to inform decision making in conjunction with simulated emotions to improve the decision making within agents compared to emotions alone. We empirically evaluate our simulated model of mood in addition to emotions. We show that our mood model can be implemented in a robotic setting which can clarify aspects of multi-agent systems, such as cooperation within an agent society.

1 Introduction

We have developed an understanding of how simulated emotions and mood can be used to inform decision making in agents so as to avoid expensive computation. We propose a functional model of mood that can be used independently or in conjunction with the current work on simulated emotions. We show that our model of

Joe Collenette
University of Liverpool, UK, e-mail: j.m.collenette@liverpool.ac.uk

Katie Atkinson
University of Liverpool, UK

Daan Bloembergen
University of Liverpool, UK

Karl Tuyls
University of Liverpool, UK

simulated mood can be used to allow cooperation in a social dilemma to be achieved through choice in a multi-agent setting. With the addition of simulated mood, there is an improvement when compared to using simulated emotions exclusively. Our mood model is grounded in psychology research, using aspects of human emotions and mood to inform decision making [10, 13].

We use this developed mood model in practise to explore how cooperation flourishes within a society of agents. The resilience of cooperation growth is tested by the addition of defectors to the society, indicating the stability of the cooperation strategy that uses our model.

Psychology research has shown emotions affect human decision making [21]. Recent work has shown that simulating these emotions within artificial agents affects the evolution of cooperation within the prisoner's dilemma game [15]. Similarly psychology shows that mood affects decision making in humans [10]. There is a clear distinction between mood and emotion, emotions are short-term feelings that are directed towards a particular object or person [13]. Mood, in contrast, is a long term feeling which does not have this focus on a particular object or person [8].

Previous research has focused on simulating emotions within agents without regards for the effects that mood will have on the decision making process. [22] gives an overview of the different methods of integrating emotions into a computational model, however there has been no previous attempt to model mood. Our model for mood is integrated into previous research using a psychological background to justify the model. Whilst we recognise that emotions and mood both have physiological affects [11], we will only be considering the functional aspect where mood and emotions change the behaviour of the agents [13, 17].

We aim to provide a generic framework of mood that can be integrated into existing emotional models, which in turn provides a deeper level within the decision making captured. Within this framework interactions between agents can occur yet the agents do not need to know each others' strategies in order for cooperation to flourish. Our model of mood is grounded in psychology research. We have shown how this mood model reacts to an unknown strategy, which in our experiments is pure defection.

2 Background

We will first introduce the emotional model we will be using as part of our mood model and our experiments. Then we continue with the prisoner's dilemma game which is the setting for our experiment.

Table 1 Emotional characters used.

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

2.1 Emotional Characteristics

The simulated emotions that will be implemented in our agents are based on the Ortony, Clore and Collins model of emotions, known as the OCC model [18]. The model was developed through psychology research and has been used throughout the AI community [1, 15, 19, 4]. The OCC model takes a functional view of emotions, in which emotions influence changes in behaviour. The action taken is a result of the emotional makeup of the person. The emotional makeup is a result of previous outcomes. This functional view lends itself to being a good platform for implementing emotions as the descriptions are of the outward effects of the emotions rather than how emotions are processed internally. Of the 22 emotions defined in the OCC model we will be modelling *anger*, *gratitude* and *admiration*, so we can compare to previous work [4].

As in [4] each emotion has a threshold and a value. When that value increases past the threshold for *anger* or *gratitude* the action of the agent will change. When *admiration* reaches the threshold then that agent will imitate the emotional characteristic of the agent that triggered the admiration. This is how replication is implemented in our experiment. In this paper we will be using 9 different types of emotional characters who have differing emotional makeups but they all have admiration thresholds of 3. We have chosen that value based on previous work as it gives the highest payoff in [4]. The different characteristics can been seen in Table 1.

An agent's anger increases by one when its opponent defects; gratitude increases when the opponent cooperates. For example, take the two characteristics Responsive and Active. If Responsive chooses to cooperate, Active's gratitude increases to one, if Active chose to Defect then Responsive's anger increases to one. Responsive's anger level is at the anger threshold, so in the next game with that agent, Responsive will choose to defect and the anger level will return to 0.

Admiration increases when the agent believes that its opponent is performing better than itself. When a threshold is reached, the agent's behaviour changes to the emotional character that triggered the admiration emotion and the admiration value is then reset back to 0. When a mobile agent completes five games of the prisoner's

Table 2 Payoff matrix of the prisoner's dilemma.

	$COOP_i$	$DEFECT_i$
$COOP_j$	$3_j, 3_i$	$0_j, 5_i$
$DEFECT_j$	$5_j, 0_i$	$1_j, 1_i$

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 [14], 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 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.

2.2 Prisoner's Dilemma

The prisoner's dilemma is a social dilemma where two players are given the choice of cooperation or defection. This choice is made simultaneously with no communication prior to the decision made. Each player then will get a payoff according to the choices made by both players. The payoff matrix is shown in Table 2.

When looking at the prisoner's dilemma outcomes, it seems in the best interest of both players to both play cooperatively since this would lead to the largest total payoff for the group as a whole. However, 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 incentivised to cooperate in the prisoner's dilemma has been an active area of research in the past decades, with a particular focus on the evolution of cooperation within groups of agents [2, 20, 3]. It is for this reason that we adopt this model of interaction in the current work as well.

3 Mood Model

Here we define our model of mood, with justifications for each mood state from psychological research and how this mood will affect decision making. We split the mood into three parts: negative, neutral, and positive. Our mood model only affects

the decision made as we interested in what decisions are made rather than simulating how mood can affect the agent physically.

It was shown in [21, 10] that negative moods can lead to a more "rational" outcome in general as people tend to think more thoroughly about the action they will take. In our experiments we use low moods to lead to defection, as this is the Nash equilibrium and can be considered the more rational decision. Very low mood levels will lead to defection regardless of the emotional state of the agents.

Positive moods tend towards an ideal outcome even if that negatively affects that person [10]. In our experiment the riskiest behaviour is cooperation as it can lead to the worst outcome for the individual agent. Cooperation is the most ideal outcome as it gives the highest payoff for the group as a whole.

For neutral moods the mood model will not affect the agents' decision making. The mood will affect how agents react to unknown opponents since they do not have any emotional attachment to them. When the mood levels are extreme they will override the current emotional decision. We have done this to represent that mood levels in humans do not necessarily reflect cooperation as a whole, but affect the choice made [16].

We define the representation of each mood state as follows: a mood of below 10 is characterised as extremely low, below 30 as low, higher than 70 as high and above 90 as extremely high, and between 30 and 70 as neutral. Equation 1 shows how the agent chooses an action based on our mood model with the simulated emotions. We define an initial action as the action an agent would take if the mood model is unable to provide an action. This is often the first interaction the agent makes. The simulated emotions in our model are defined as one of the emotional characters as described in Section 2.1. How and when an interaction occurs in our experiment is given in Section 4.

Definition 1. Let Ag be the set of all agents, with i and $j \in Ag$. Let t denote time. Let m_i^t return the mood of agent i at time t, in the range $]0, 100[$. Let $\eta_{i,j}$ return the number of interactions agent i as with agent j. Let I_i return the initial action of agent i. Let $E_{i,j}^t$ return the action that agent i would take against agent j based on i's simulated emotions, at time t.

$$Ac_{i,j}^t = \begin{cases} COOP, & \text{If } m_i^t > 90 \text{ or } (m_i^t > 70 \text{ and } \eta_{i,j} = 0) \\ DEFECT, & \text{If } m_i^t < 10 \text{ or } (m_i^t < 30 \text{ and } \eta_{i,j} = 0) \\ E_{i,j}^t, & \text{If } 30 <= m_i^t >= 70 \text{ and } \eta_{i,j} \neq 0 \\ I_i, & \text{Otherwise} \end{cases} \quad (1)$$

Our representation of positive mood values comes from psychology literature showing how people take riskier behaviour to achieve a more ideal outcome [10]. However if the mood is too positive, as it is when a person has mania, then the behaviour becomes extremely likely to hurt that person [12]. In [21, 10] it is shown that negative moods can be more likely to lead people to make a more considered choice. Research into human patients with depression shows that these people are more likely to choose defection. The research also showed that depressed patients

were more critical of themselves [9]. This provides us with grounding for our choice of defection as part of our implementation of the mood model in the prisoner's dilemma, and validates how the mood values are more greatly affected when the mood is low.

The agent's mood value will go up or down based on the difference between the payoff received and its average payoff, as this represents how well the agent thinks it has done in that game [5]. Then additionally the mood value will go up or down based on how the agent feels towards inequity between the average payoffs. We will be using the inequity aversion model *Homo Egualis* to represent inequity as a value [5]. In this model we need to find an α and β, where α represents how much an agent cares when it is doing badly and β represents how much an agent cares when its opponent is doing badly. Since we want to represent an ideal solution we will take $\alpha = \beta$. This represents that an agent cares about an opponent as much as it cares about itself.

The amount the agent cares is represented by applying the mood to our α value, such that higher moods give a lower α. This results in mood changes being larger when the mood is low. If the mood is low then the agent "thinks" it is are doing poorly in the environment when compared to other agents. We do this to represent the property that humans can care more about equality when doing poorly in society [5].

Definition 2. Let Ag be the set of all agents, with i and $j \in Ag$. Let t denote time. Let p_i^t return the payoff of agent i at time t. Let m_i^t return the mood of agent i at time t, in the range $]0, 100[$. Let μ_i^t denote the average payoff for agent i up to time t. Let F_i^t return the opponent of agent i at time t.

$$\alpha_i^t = (100 - m_i^{t-1})/100 \tag{2}$$

$$\Omega_{i,j}^t = \mu_i^t - \alpha_i^t \cdot \max\left(\mu_j^t - \mu_i^t, 0\right) - \alpha_i^t \cdot \max\left(\mu_i^t - \mu_j^t, 0\right) \tag{3}$$

$$m_i^t = m_i^{t-1} + (p_i^t - \mu_i^{t-1}) + \Omega_{i,j}^{t-1} \text{ where } j = F_i^t \tag{4}$$

In equation 2 we show how we get our α value from the current mood of an agent; this places the mood value in the range of $]0, 1[$ so it can be used as the α. For example a mood value of 75 will return an α of 0.25. Equation 3 is the simplified version of the Homo Egualis function [7], as we have only two agents in a single interaction and $\alpha = \beta$. The equation gives us a numerical representation of inequity that the agent has for that interaction. Equation 4 shows the overall implementation of mood using the previous mood value, the average payoff, the received payoff, and the Homo Egualis function to update the mood value after an interaction with another agent. This equation gives us the current mood value of an agent. The mood will increase or decrease depending on the difference in the received payoff and the average payoff, meaning that the mood will increase when this agent is doing better than expected and decrease when it is doing worse than expected. With the inclusion of Ω the amount that the mood moves can change based on how fair the agent thinks the result was for both agents.

4 Method

In this work, the agents are simulated mobile robots that drive around in an environment. The agents are given a random walk behaviour with some basic obstacle avoidance procedures. The prisoner's dilemma game is initiated whenever two agents are within close proximity, and have line of sight of each other. The game is played once, after which the agents will then continue their random walk behaviour.

The agents are placed in a random location in the environment with emotional characteristics and moods distributed randomly, uniformly and independently among the agents, given the specific proportions of each experiment. The details for the experiments conducted are given in Section 5.

Agents move randomly throughout their environment, while avoiding collisions with the environment or other agents. Each agent has proximity sensors located at $\{-90°, -45°, -15°, 15°, 45°, 90°\}$ w.r.t. the robot's heading. When the left sensors detect something the robot will stop and turn to the right and the reverse for the right sensors. The agents move forward at up to 10 cm/s (the speed is constant except when accelerating from stationary as built into the simulator) and can turn 45 deg/s. When there are no obstacles detected the agent moves forward with a turn speed that is either $-45°$ or $45°$ per second. A new heading is generated when the robot receives data from the sensors, resulting in a random movement pattern.

In terms of the agents' knowledge of the world they are able to differentiate between agents, but have no knowledge of the strategies others will be using. They also have no knowledge of the environment apart from the sensor data they have at that moment in time. In addition the agents have no knowledge of the payoff matrix and will purely use their mood and emotion strategy to determine whether to cooperate or defect.

5 Experiment Outline

We will be describing the experiments that we have conducted to test our mood implementation. The first experiment will explore how mood affects the evolution of cooperation. Our second experiment introduces pure defectors as an invasion force to our environment and answers the question of how resilient the cooperation is to outside defectors.

Both of the experiments will be conducted in the environment which is four corridors in a square with each corridor having a length of 5 meters and a width of 1 meter. This is shown in Figure 1. The experiments will run for 10 minutes for each run. A run consists of a scenario and if applicable a sub-scenario. To ensure consistency of results we will be running each of the 10 minute runs 10 times. In our experiments each emotional characteristic is represented equally to prevent any characteristic becoming dominate due to them having a higher initial representation. In addition the initial actions of the emotional agents will be an equal split between cooperation and defection. The initial location of the agents will be generated ran-

Table 3 Mood experiment scenarios showing as a percentage the different distributions of starting mood levels for the agents

| Scenario	Low Mood	Medium Mood	High Mood
1	100	0	0
2	0	100	0
3	0	0	100
4	33	33	33
5	70	15	15
6	15	70	15
7	15	15	70

domly. Each of these aspects will be distributed randomly and independently of each other. For our experiments we will be simulating agents using the Player/Stage simulator [6].

5.1 Mood Experiments

The first experiment will explore how the evolution of cooperation is affected by differing initial mood levels. The initial level of mood will be categorised into three types, low, medium and high where low has a mood level of 30, medium is 50 and high is 70. There will be seven scenarios each with a different distribution of these levels among the agents which can be seen in Table 3. We refer to the outcomes given by neutral moods as medium as this better reflects the current mood value.

Each of these scenarios will be run against a number of sub-scenarios. The sub-scenarios define how many agents will be in the environment, with a range from 45 to 144 agents, the details of the scenarios can be seen in Table 4. We will be looking to see how different moods affect cooperation. We will explore if our mood model allows cooperation to increase in the society of agents over time.

Fig. 1 Environment used in this work. The environment has dimensions of 5 × 5 metres with the agents themselves having a radius of 7cm. The traversable areas are shown in white
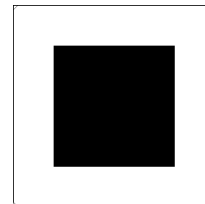
Table 4 Mood experiment sub-scenarios, showing the number of agents that will be simulated for each scenario

Sub-scenario	No. of agents
1 - Low density	45
2 - Medium density	81
3 - High density	117
4 - Very High density	144

Table 5 Resilience experiment scenarios showing the starting level of mood for the agents with our mood model and the number of pure defectors that will be added

| Scenario	Mood Level	No. of Defectors
1	Low	43
2	Low	63
3	Low	83
4	Medium	43
5	Medium	63
6	Medium	83
7	High	43
8	High	63
9	High	83

5.2 Resilience Experiments

Our next experiment is to test the resilience of the cooperation that evolves over time. To test this we will be introducing into our environment pure defectors at the beginning of the experiment ; they cannot replicate themselves but the emotional agents may take on the role of a pure defector due to their admiration emotion. Each scenario will have 63 agents whose initial mood is dictated by the scenario: the moods are categorised as high (70), medium (50) and low (30). The numbers of pure defectors are 43 (minority defectors), 63 (equal defectors and emotional agents) and 83 (majority defectors). The details of each scenario are shown in Table 5. This will show the resilience that our mood model has to these pure defectors.

6 Results

6.1 Mood Results

Figure 2 shows us the percentage of cooperation between each number of interactions for each scenario, with an extra scenario which excluded the mood model and only used the emotional strategy. The results given are quite intuitive, we see that
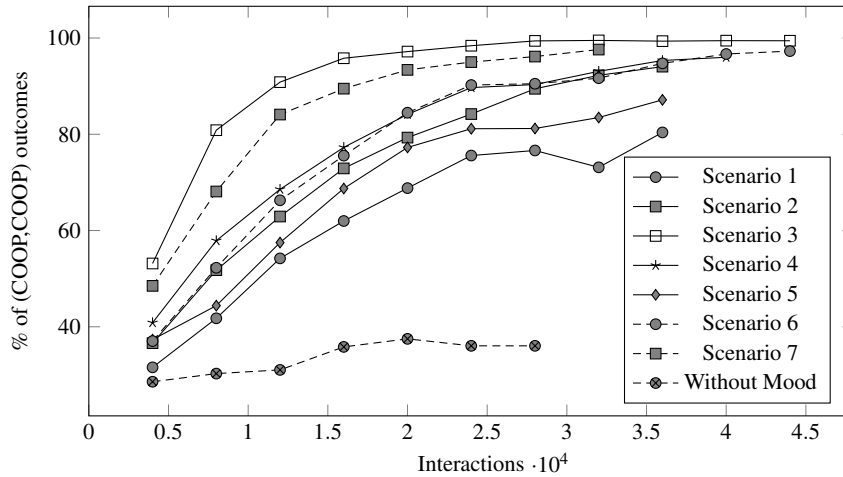
Fig. 2 Percentage (COOP, COOP) outcomes over all runs for each scenario in the mood experiment.

cooperation evolves throughout the agents, and the speed at which this is achieved is directly proportional to the average level of mood. The fastest is the scenario with 100% of agents starting with high mood levels and the slowest is the scenario with 100% of agents having low mood levels. To attribute this to the mood model we ran the same experiments but without our mood model, where the decision making was purely based on the agents' emotional decision making. We can attribute the rise in cooperation to the mood model as when it is removed the cooperation does not rise as quickly. The number of iterations between scenarios being uneven is due to the random nature of the agents movement.

The results show us that our mood model can support the evolution of cooperation over time and sustain cooperation; this was an expected result as when cooperation is high the mood moves very little. When two agents play the game, with one being in a high mood and one being in a low mood, the low mood will rise faster than the high mood can go down which is a property of our implementation of the egualis equation. This leads to more agents in a cooperative state raising cooperation. This effect is most apparent in the later stages of the simulation when the agents start with low moods, as the agents which are cooperating meet a group of agents which are not cooperating. This led to a dip in cooperation followed by the continuing rise of cooperation when a large amount of agents with opposing moods meet.

To justify our claim that the speed at which cooperation is achieved is proportional to the starting level of mood, we have plotted the average mood values against the number of (*COOP,COOP*) actions, as can be seen in Figure 3. We have shown this for scenario 1 as this is where the effect is most pronounced; we can see that cooperation between agents falls from 77% to 73% as agents who are cooperating meet larger groups of defecting agents. However the average mood level still rises,
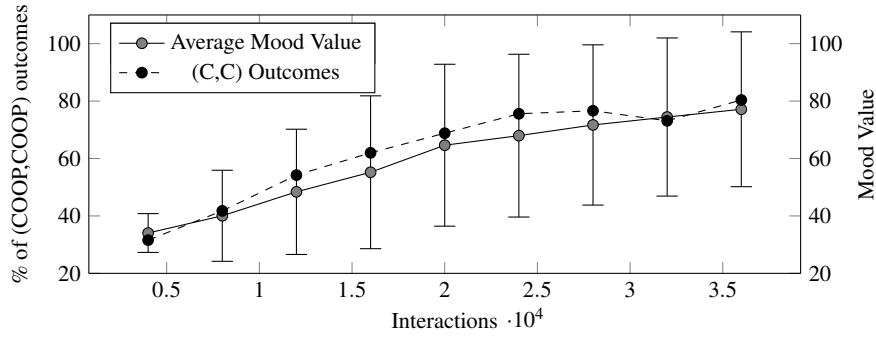
Fig. 3 Average mood value with standard deviation against percentage of (*COOP*,*COOP*) outcomes in scenario 1 of the mood experiment

from 71.7% to 74.5%. When the cooperation rises again the standard deviation of mood levels is reduced to 26.9 from 27.5. This shows us that the mood reflects the level of cooperation, and the higher the starting level of mood the faster cooperation is achieved.

6.2 Resilience Results

Figure 4 shows that when the mood is low, the emotional agents as a group are more resilient to an invading population of pure defectors. In high moods cooperation between the emotional agents rises quickly, this in turn raises the mood of the agents as well. It therefore does not take long for the mood to increase to the point where the agents can be considered pure cooperators due to the mood level being very high. When this happens and agents are faced with the pure defectors the only outcome between an emotional agent and a pure defector can be (*COOP*,*DEFECT*). This causes the average score of the defectors to increase and the emotional agents' average to decrease. These changes in average payoffs will be affected rapidly because of the payoff difference. When replication occurs in the emotional agents they choose to become pure defectors because of this payoff difference, which leads to the collapse of cooperation as there are more pure defectors.

In contrast when the emotional agents are in a low mood it takes longer for them to get their moods to the level where they are indistinguishable from pure cooperators; this allows them to protect themselves from the pure defectors by using their emotional choice which switches their action to defection for that particular opponent. Actions driven by emotions rather than mood are bounded to a particular opponent which allows the agents to evolve cooperation with other emotional agents without replicating into pure defectors since the defectors have a low average as the number of (*DEFECT*,*DEFECT*) actions they receive increases over time.
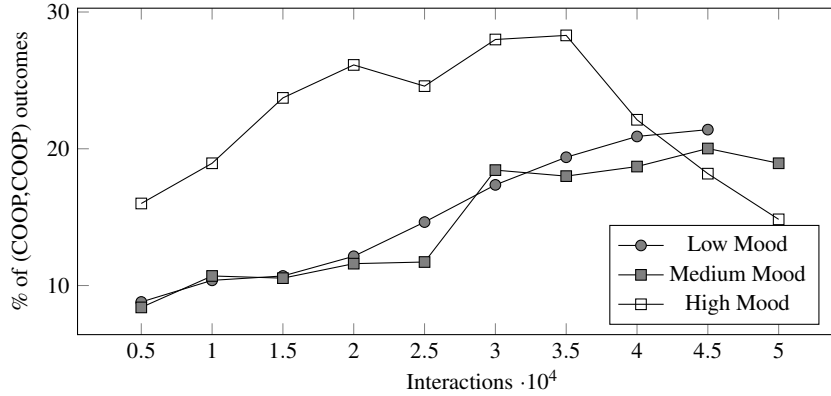
Fig. 4 Percentage of (COOP, COOP) outcomes for each initial mood level in the resilience experiment.

Table 6 Average scores with standard deviation of the defectors, for each mood level in the resilience experiment

Starting Mood Level	Average	Standard Deviation
Low	1.45	0.49
Medium	1.82	0.69
High	2.17	0.85

These results show both expected and unexpected results. We had expected that cooperation would continue to be stable over time as the simulated moods and emotions would adapt to the invasion force, as seen in the low and medium starting moods. However the collapse of the high mood was unexpected.

To justify why the high mood levels have collapsed, we have shown that high moods do not adapt quickly to the pure defectors and therefore are taken advantage of. The advantage taken then leads to the emotional agents becoming pure defectors as their average score is not high enough when compared to the pure defectors. We took the difference between average score of the defectors and the average score of the emotional agents for each starting level of mood. The results showed that the difference in average score between low starting moods and the pure defectors was 0.04, for medium moods was 0.22, finally for the high moods the difference was 0.5. We can see that the high mood difference is more than double the medium mood difference. The defectors are clearing taking advantage of the high moods the most.

As the high moods are being taken advantage of the most, we expect that the payoffs for the defectors should be the highest when faced with the highest mood. The average scores of the defectors are shown in Table 6 and clearly show that the defectors do the best when faced with high moods, meaning that they will replicate the fastest in the high moods. The medium and low moods do not collapse as they

adapt to the newly replicated defectors through the use of their directed emotion strategy. The high moods do not do this as when the mood is very high the emotional agents act as pure cooperators against the pure defectors.

7 Conclusions

We have proposed a model of mood that can be used either independently or in conjunction with emotions. We have constructed this model using psychological research, and have identified general cases of outcomes that can be expected. We have then applied this to an experiment with validation from psychology research for our choices made, which explores cooperation in a multi-agent setting.

For our experiments conducted we have shown that a combination of mood and emotion can support positive levels of cooperation within an agent society. We have also shown that mood levels in our agents are related to the level of cooperation that is achieved as a group. By adding an invasion force of pure defectors the cooperation between the emotional agents collapses over time when the mood levels of the emotional agents are high. In contrast when the mood is not high the cooperation over time is more stable since the agents do not give the benefit of doubt to the pure defectors, preventing the defectors from achieving a higher average. For future work we will looking into adapting our mood model to take into consideration mood fluctuations over time.

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. Joe Collenette, Katie Atkinson, Daan Bloembergen, and Karl Tuyls. Mobility effects on the evolution of co-operation in emotional robotic agents. In *Proc. of ALA Workshop*, 2016.
5. Ernst Fehr and Klaus M Schmidt. A theory of fairness, competition, and cooperation. *Quarterly journal of Economics*, pages 817–868, 1999.
6. 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.
7. Herbert Gintis. *Game theory evolving: A problem-centered introduction to modeling strategic behavior*. Princeton university press, 2000.
8. 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.
9. William E Haley and Bonnie R Strickland. Interpersonal betrayal and cooperation: Effects on self-evaluation in depression. *Journal of Personality and Social Psychology*, 50(2):386, 1986.

10. Guido Hertel, Jochen Neuhof, Thomas Theuer, and Norbert L Kerr. Mood effects on coop-eration in small groups: Does positive mood simply lead to more cooperation? *Cognition & emotion*, 14(4):441–472, 2000.
11. Dacher Keltner and James J. Gross. Functional accounts of emotions. *Cognition & Emotion*, 13(5):467–480, 1999.
12. Robert L. Leahy. Clinical implications in the treatment of mania: Reducing risk behavior in manic patients. *Cognitive and Behavioral Practice*, 12(1):89 – 98, 2005.
13. Robert W. Levenson. Human emotion: A functional view. *The nature of emotion: Fundamental questions*, 1:123–126, 1994.
14. 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.
15. Martyn Lloyd-Kelly, Katie Atkinson, and Trevor Bench-Capon. Fostering co-operative be-haviour through social intervention. In *Simulation and Modeling Methodologies, Technologies and Applications (SIMULTECH), 2014 International Conference on*, pages 578–585. IEEE, 2014.
16. Robert B. Jr. Lount. The impact of positive mood on trust in interpersonal and intergroup interactions. *Journal of Personality and Social Psychology*, 98(3):420 – 433, 2010.
17. William N Morris. A functional analysis of the role of mood in affective systems. *Emotion*, 1992.
18. Andrew Ortony, Gerald L Clore, and Allan Collins. *The cognitive structure of emotions*. Cambridge university press, 1990.
19. A. Popescu, J. Broekens, and M. van Someren. Gamygdala: An emotion engine for games. *IEEE Transactions on Affective Computing*, 5(1):32–44, Jan 2014.
20. 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.
21. Norbert Schwarz. Emotion, cognition, and decision making. *Cognition and Emotion*, 14(4):433–440, 2000.
22. Rodrigo Ventura. Emotions and empathy: A bridge between nature and society? *International Journal of Machine Consciousness*, 2(02):343–361, 2010.