

# Developing Co-operation Through Simulated Emotional Behaviour

Martyn Lloyd-Kelly, Katie Atkinson, and Trevor Bench-Capon

The University of Liverpool, Department of Computer Science, Ashton Street,  
Liverpool, L69 3BX, United Kingdom  
`mlk5060@liverpool.ac.uk`

**Abstract.** The emergence and maintenance of co-operation in a society of agents is an important issue. Hitherto it has largely been investigated in the framework of game theory, but we have previously explored the role that can be played by a functional model of emotions [11]. For example, we have shown that the emotions of *gratitude* and *anger* can be used to produce co-operative behaviour from agents acting solely in accordance with their current emotional state. The effectiveness of these emotions in producing co-operation, however, depends on the emotional volatility of the agents, which determines the strength of these emotions required to alter behaviour. In this paper we explore the role of the emotion of *admiration* in spreading the characteristics required to further co-operation throughout a society of agents. The research is conducted using simulations of a range of scenarios on a test bed based on the *Prisoner's Dilemma* game. Our results show how *tolerance* and *responsiveness* affect the utility of the system as a whole and how this varies according to the proportion of initially co-operative agents.

## 1 Introduction

In this paper we describe some experiments we have carried out to investigate the emergence of co-operation in populations of agents with varying emotional characteristics. While many agent simulations use agents that reason explicitly about their goals and the effects of their actions, we wish to explore the possibility of describing behaviours in terms of an emotional response to the agents' situation. We therefore characterise our agents in terms of a limited set of emotions, capturing certain functional aspects of emotions in so far as they affect decisions: we are not creating emotional agents as such, and pay no attention to issues such as the physiological effects of emotions.

In our simulations we have agents that interact with one another. Inspired by Axelrod's investigation of the emergence of co-operation [2], these interactions take the form of an *iterated Prisoner's Dilemma game*. Agents play against sets of their neighbours, but their choice between co-operation and defection will take no account of payoffs or history. Rather, (after using a default for the first round) their emotional attitude towards the other agents involved will determine their choice. Our model of emotions is based on the Ortony, Collins

and Clore model [14] (henceforth referred to as the *OCC* model), as interpreted by Steunebrink et al. [15], [16] and, from the twenty-two emotions available, we model two: *gratitude* and *anger*. In short, *defection* from an opponent excites anger towards that agent and sufficient anger will cause a co-operating agent to defect. Similarly, *co-operation* will excite gratitude and cause a defecting agent to co-operate when sufficiently grateful. The intensity of anger and the intensity of gratitude required to affect behaviour can be varied, and we say that an agent is more *tolerant* if a greater degree of anger is needed before defection occurs, and more *responsive* if less gratitude is required to induce co-operation. We use three degrees of tolerance/responsiveness and these are independent, giving rise to nine different emotional characters for our agents (a description of each character will be provided later in the paper). The effects of gratitude and anger when played against notable strategies submitted to Axelrod's tournament in the context of an iterated Prisoner's Dilemma game were reported in [11]. However, in this paper we extend our work by investigating the effects of the emotion *admiration*, as defined by the OCC model and [15].

After a round of interactions, agents compare their performance with their neighbouring peer group. At this point admiration can be evoked with respect to one of these peers, which may be sufficient to cause the admiring agent to modify its emotional character to imitate the admired agent. In this way, successful strategies will spread through the population so that we have a mechanism by which an initial population can develop particular characteristics as a result of agents comparing their performance with that of others. The experiments reported here will focus on: what kind of emotional character, in terms of tolerance and responsiveness, is conducive to individual success; what characters are conducive to system success, measured in terms of several criteria; and, how character affects the emergence of co-operation.

Our contribution is threefold: firstly, we model the emotions of anger, gratitude and admiration so that they are able to play a functional role in determining how agents behave in a simulated environment. Secondly, we propose that admiration is a functionally valid mechanism, capable of driving the spread of the most successful character traits (which are, we will show, high tolerance and high responsiveness) through a population. Thirdly, we illustrate the effects of employing different intensities of admiration and initial co-operators/defectors on both individuals and the system and on the emergence of the most prevalent character. To achieve this, we have implemented a novel test-bed inspired by [2].

Section 2 will provide some background support for the idea that emotions are an important factor in determining behaviour and provide more details of the OCC model and the use we make of it. Section 3 will describe our test environment, the research questions we address in this paper and the scenarios used in our experiments. Section 4 reports the results of these experiments and offers some discussion. Finally, section 5 gives a summary and describes our current and future work.

## 2 Background

In [8], emotions are terms ascribed to a collection of bodily responses felt in reaction to some stimuli. This stance is expanded by [6] and [10] where it is argued that emotions are functions that serve to directly mediate behaviour. However, the authors of [4] argue that emotions are key components in *feedback loops* whereby emotions are used to determine behaviour in a reoccurring situation based on past experience. Both views of emotions described are taken into account by [3], which argues that both interpretations are likely to co-exist. We adopt the view that posits emotions as functional mediators of behaviour. Whilst we agree that physiological factors and the notion of feedback loops are important in any comprehensive account, we do not consider such aspects here.

We use Axelrod’s tournament [2] as a way to both test the effectiveness of using emotions functionally as behavioural modifiers to enable and spread co-operation throughout a population and as a way of distinguishing between *rational* and *emotional* agents. Axelrod pitted rational strategies against each other in a number of iterated Prisoner’s Dilemma games to determine which is the most successful i.e. strategies which maximise their own score whilst maximising the total score of the system. The tournament led to the identification of four general rules that should be adhered to in order to create a successful strategy; one of these rules states that strategies should not be overly complex. This rule provides us with a basis for distinguishing between emotional and rational agents. The agents in Axelrod’s tournament rationally condition their behaviour by taking into account their past, present and future payoffs. In contrast, our emotional agents make no use of the concept of payoffs: they are simply reactive agents inspired by the notions outlined in [5]. Essentially, the emotional agent’s behaviour is a product of its character and its current emotional state (which is determined by its past experience and character) with the layer of rationality associated with consideration of future and past payoffs stripped away. Thus, while history determines the current state, the details of that history are discarded and behaviour is a simple response.

The decision to model and implement admiration in a multi-agent system in order to spread successful characteristics was inspired by the occurrence of *admiration* in both the OCC model and [15] and the observation that in both human and non-human societies there are individuals who possess *status* or *prestige* and act as role-models. As suggested by [7], prestigious individuals are the subject of *infocopying*, defined as a collection of direct social-learning by others. Such social-learning mechanisms are investigated further by [12] and [13] who conclude that familiar prestigious individuals with valued skills elicit *inspiration* in those evaluating the individual. The elicitation of inspiration therefore appears to be the product of comparison with a model example; however, how can such inspiration be elicited in agents?

In humans, [1] proposes that inspiration is the motivational output of admiration that results in the individual copying the valued traits possessed by the prestigious individual. The study makes a number of important findings that provide us with the motivation for using admiration in the way that we

do: firstly, participants that experienced admiration were more likely to try to emulate the prestigious individuals’ actions and secondly, participants that experienced admiration reported a unique feeling of wanting to achieve success themselves. Therefore, we have used such empirical evidence of the effects of admiration as the basis of our functionality.

Taking inspiration from [9], we decided to use Netlogo [17] to create an implementation of an iterated Prisoner’s Dilemma game. The next section details the *Emotional Population* simulation test-bed, the agent architecture used, our research questions and the experiments constructed and run to provide answers to these questions.

### 3 Test-Bed Implementation

Our *Emotional Population* test-bed is a multi-agent simulation created using the Netlogo programming language and inspired by the *Prisoner’s Dilemma* game. The test-bed consists of 338 agents; all capable of emotional responses. A typical round in an *Emotional Population* game proceeds as described below. A complete game entails 500 of these rounds being played:

1. Agents and a mediator are placed in the environment.
2. Agents consult current emotional state before co-operating/defecting with/against opponents (their decisions are also sent to the mediator).
3. Agent’s emotional state is updated by reacting to the opponent’s play in the last round and payoffs are distributed by the mediator as per payoff distribution in *Prisoner’s Dilemma* (see Table 1).

The payoff matrix is used by the simulation system to determine the effect of the agent’s actions. It is not used by the agents themselves, and we need not assume any awareness of this matrix on the part of the agents.

**Table 1.** Prisoner’s Dilemma payoff matrix.

	<b>Co-operate<sub>j</sub></b>	<b>Defect<sub>j</sub></b>
<b>Co-operate<sub>i</sub></b>	3 <sub>i</sub> , 3 <sub>j</sub>	0 <sub>i</sub> , 5 <sub>j</sub>
<b>Defect<sub>i</sub></b>	5 <sub>i</sub> , 0 <sub>j</sub>	1 <sub>i</sub> , 1 <sub>j</sub>

The determination of an agent’s opponents is achieved by setting the *player set* variable to one of the following values:

- Parallel: opponents are located east and west of the agent.
- Orthogonal: opponents are located north, south, east and west of the agent.
- Diagonal: opponents are located north-east, north-west, south-east and south-west of the agent.
- Octagonal: combination of orthogonal and diagonal sets.

Increasing the number of players in an agent’s player set places a character under more rigorous testing as it may have to compete with up to 7 different emotional characters. Therefore, by altering the player set used we may further ascertain how un/successful certain characters are as they have to compete against a broader variety of opponents. It should also be noted that an agent may co-operate with one agent in its player-set and defect towards another i.e. behaviour towards different opponents is distinct. As described in section 2, the decision to co-operate or defect is dependent upon the agent’s character and its current intensity of gratitude/anger towards its opponent. An agent’s character is set randomly but the proportion and distribution of character types within the initial population is kept equal, this ensures that no character has an unfair advantage with regards to propagation throughout the population; a character is only propagated if it is successful, not because it is more prevalent in the initial population or there is a small concentration of characters which can never be broken down. An agent’s character type can be set to 1-9; as previously explained, an agent’s character determines how quick or slow the agent is to reward (gratitude) or punish (anger). The slower an agent is to anger, the more tolerant it is, the quicker it is to show gratitude, the more responsive it is. Thus, an emotional agent of character 1 - our least tolerant and most responsive agent - only requires one defection/co-operation for its anger/gratitude to become sufficient to change its behaviour to defection/co-operation. Conversely, an emotional agent of character 9 - our most tolerant and least responsive character - requires three non-consecutive defections/co-operations before the agent feels sufficient gratitude/anger to change its behaviour. Table 2 details all 9 characters for reference. An initial disposition for each agent is set to either co-operate or defect; this value is used to determine the behaviour of an agent on the initial round of every game. The proportion of initial co-operators/defectors is user-defined but the determination of which agents are co-operators/defectors is left to the program as we are not interested in analysing how the proximity of co-operators/defectors influences the results obtained.

**Table 2.** Emotional agent character (Ch.#) descriptions.

		If def. #co-ops to co-op.		
		1	2	3
If co-op, #def. to def.	1	Ch.1	Ch.2	Ch.3
	2	Ch.4	Ch.5	Ch.6
	3	Ch.7	Ch.8	Ch.9

After 5 rounds have been played, each agent compares its current score against each agent in the agent’s *comparator* set. An agent’s comparator set

variable may be set to either *parallel*, *orthogonal*, *diagonal* or *octagonal* (as for the player set). Related to this is the agent's *admiration* level, which determines how successful a neighbour must be before the agent changes its character to match that neighbour. An agent considers its own success when comparing itself against others and, if an agent deems itself the most successful; its character will not change. An agent's level of admiration can be set to any of the following values:

- High: an agent's neighbour (or itself) needs to achieve the highest score once before this agent's character is copied.
- Moderate: an agent's neighbour (or itself) needs to achieve the highest score twice before this agent's character is copied.
- Low: an agent's neighbour (or itself) needs to achieve the highest score three times before this agent's character is copied.

### 3.1 Research Questions

To give focus to the question of how co-operative behaviour can be described in terms of emotional response, we have identified specific research questions that the implemented test-bed aims to answer. These questions are presented below along with information that we will extract from the system that we deem to be relevant to each question.

1. Do any initial conditions enable highly/moderately/less tolerant emotional characters to become prevalent and if so, why?
2. Do any initial conditions enable highly/moderately/less responsive emotional characters to become prevalent and if so, why?
3. Does the prevalence of a particular character have an effect on the total system score and if so, why?

With respect to questions 1 and 2, we aim to analyse the frequency at which each group of characters places first in each scenario; those that do can be considered more prevalent. The scenarios under which these characters become prevalent may then be analysed to determine why the prevalence of this character emerges. Question 3 may be answered by simply analysing the prevalence of characters and average total system scores of sub-scenarios. The total system score is the sum of all individual agents' scores on the final round of the current game. Average total system scores may then be calculated and ordered so that the highest sub-scenario scores for a scenario can be analysed. From this, we can see if there is any relation between high average total system scores and character prevalence.

### 3.2 Experiment Scenarios

In order to answer the research questions detailed in section 3.1 we have run a number of experiments. We tested 14 different scenarios as listed in Table 3.

**Table 3.** Scenario details - proportions of each variable.

Scenario	% Initial Defectors	% Initial Co-operators	% Admiration Level		
			High	Mod.	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

For each scenario there were 16 sub-scenarios run, each of which was repeated 10 times in order to harvest a data set large enough for us to construct a reliable analysis. These sub-scenarios are concerned with the player/comparator set configurations, details of these sub-scenarios can be found in Table 4.

**Table 4.** Sub-scenario details.

		Comparator Set			
		Parallel	Orthogonal	Diagonal	Octagonal
Player Set	Parallel	Sc.1	Sc.2	Sc.3	Sc.4
	Orthogonal	Sc.5	Sc.6	Sc.7	Sc.8
	Diagonal	Sc.9	Sc.10	Sc.11	Sc.12
	Octagonal	Sc.13	Sc.14	Sc.15	Sc.16

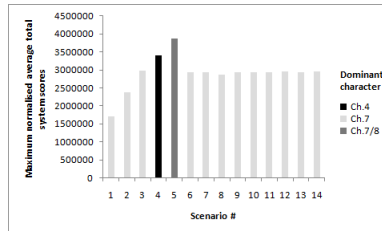
## 4 Results

Simulating the behaviour for 10 runs of the sixteen sub-scenarios of the 14 scenarios produced a wealth of detailed information. Here we provide some summaries that establish the key points with regard to the research questions posed in the previous section. Our nine characters were divided into three groups to reflect their tolerance and responsiveness. As an example, characters 1, 2, 3 have low tolerance and characters 3, 6 and 9 have low responsiveness. These groupings are shown in Table 5 below.

**Table 5.** Emotional characters and their group allocation according to degree of tolerance/responsiveness.

	Level		
	Low	Moderate	High
Tolerance	1, 2, 3	4, 5, 6	7, 8, 9
Responsiveness	3, 6, 9	2, 5, 8	1, 4, 7

To investigate system score we identified, for each scenario, the maximum average system score across the sixteen sub-scenarios (normalised to take account of the different number of encounters in the sub-scenarios), and the character that emerged as most prevalent in the sub-scenario in question i.e. the character that was represented by the most agents in the final population of the sub-scenario. As figure 1 shows, this was almost always character 7, the most tolerant and responsive of our characters. The exceptions are scenarios 4 and 5, in which the majority of the initial population are co-operators and where characters 4 (in scenario 4) and 8 (in scenario 5) become the prevalent characters. This is explained by the fact that since co-operation cycles are easier to achieve in these scenarios, a moderate degree of responsiveness is enough to establish co-operation. Moreover where, as in scenario 4, there is still a significant amount of initial defection, the moderately tolerant character succeeds more because it avoids exploitation. This also explains why scenario 5 has the highest overall system score and scenario 4, which has a majority of co-operators in the initial population, the second highest.

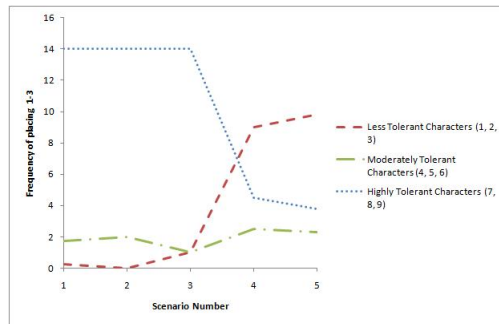


**Fig. 1.** Maximum normalised average total system scores and emerged prevalent character for each scenario.

Figure 2 shows the effect of different proportions of initial defectors on the character that emerges as most prevalent. Where two or more characters were equally prevalent in a sub-scenario, an appropriate fraction (e.g. 0.5 for two equally prevalent characters and 0.25 for four) was counted. This clearly indicates that a high degree of tolerance is more effective, until the initial proportion of co-operators becomes high. The reasoning is similar to that provided in the discussion of the total system score. Success depends on establishing co-operation and where co-operation is initially widespread, there is less of a requirement for

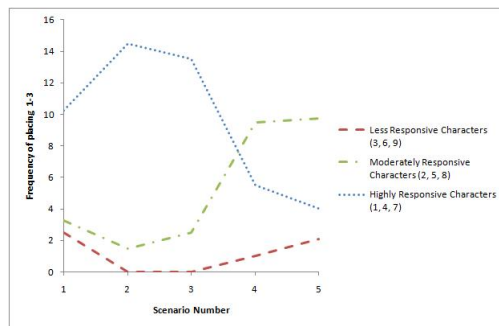


tolerance to establish co-operation with an unresponsive partner. Figure 2 also shows that intolerance has benefits when co-operation is widespread, since where defections are rare there is more value in not being exploited than in having the patience to win the defector over.



**Fig. 2.** Prevalence frequency of tolerant character groups in scenarios focused upon initial numbers of co-operators/defectors.

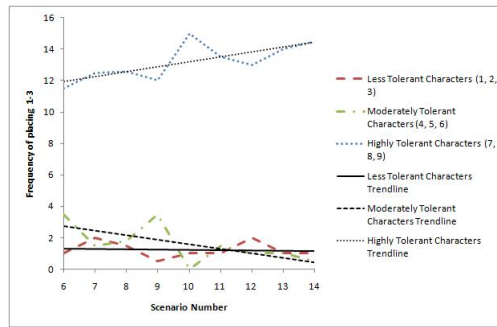
The effect of the proportion of defectors on responsiveness is shown in figure 3. Like tolerance, responsiveness is important when defection is frequent, but of less importance when defections are rare. Here, however, when co-operation is frequent, moderately responsive characters thrive. This is because such characters benefit from establishing co-operation without allowing themselves to be over-exploited by infrequent defections.



**Fig. 3.** Frequency of placing in positions 1-3 for responsive character groups in scenarios focused upon initial numbers of co-operators/defectors.

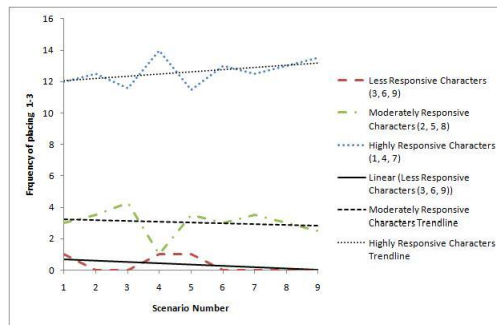
We now turn to questions about the speed at which admiration leads to a change of character, as explored in scenarios 6-14. Figure 4 shows that while

tolerance is generally a feature of success, the effect is less pronounced than in figure 2. In general we found that as admiration became less immediate in its effect, the effect of tolerance was more pronounced, with moderately tolerant characters thriving at high admiration levels. We explain this by the fact that where admiration acts quickly, characters change to allow co-operation to be easier to establish, and, as with the defectors, the harder co-operation is to establish, the greater the premium on tolerance.



**Fig. 4.** Frequency of placing in positions 1-3 for tolerant character groups in scenarios focused upon initial numbers of co-operators/defectors.

Figure 5 explores the effect of admiration on responsiveness. Although the results are broadly similar, with high responsiveness being increasingly favoured as the effect of admiration takes longer to elicit, we note that here moderately responsive characters do relatively well at all levels of admiration, whereas unresponsive characters rarely thrive at any level of admiration. This again shows that where co-operative characters spread quickly, it is more important to be responsive to co-operative behaviour than to punish defectors.



**Fig. 5.** Frequency of placing in positions 1-3 for responsive character groups in scenarios focused upon initial numbers of co-operators/defectors.

We also considered whether we could detect any significant effects or trends as a result of altering the numbers of players or comparators in the various sub-scenarios. We were unable to observe any striking patterns in this data, but may explore this issue further in future.

#### 4.1 Discussion

With regards to research questions 1 and 2, we found that when the initial proportion of defectors in the population was high, the most successful characters are those that embody high tolerance and high responsiveness. This remained true unless initial co-operators were in the majority. As co-operation becomes increasingly widespread, the benefits of being less tolerant and less responsive emerge. With a very large proportion of initial co-operators, the most successful characters were those that embodied intolerance, resulting in a reduced chance of exploitation by the rare defectors, and moderate responsiveness, a characteristic that enables co-operation to be established without being over-exploited. Similar, although less pronounced, observations were made when the speed with which admiration took effect was varied. The point here is that the quicker admiration is to be elicited, the easier co-operation is to establish, so moderately responsive characters do relatively well when a greater percentage of the population is easy to elicit admiration from. However, due to the fact that there is never a high proportion of *initial* co-operators in the scenario where admiration was varied, characters from the least tolerant group never have a chance to flourish. With regards to research question 3, we found that the highest system scores were always achieved when the most tolerant characters become prevalent. In addition the most prevalent character that emerged was also highly responsive i.e. character 7.

### 5 Conclusion

In this paper we have presented a mechanism for agents that simulates the emotions of *gratitude* and *anger* to drive the emergence of co-operation. Furthermore, we have presented a functional implementation of the emotion *admiration* and investigated how this emotion spreads characteristics through a population that enable success under various conditions. We have also identified the conditions under which the most desirable emotional characteristics (high tolerance and high responsiveness) become prevalent, and how these characteristics have a beneficial effect on the total score of a system/population. A key point is that these behaviours, typically produced by agents using rational decision-making that depends on the expected payoffs of various actions, can equally be explained as a simple response to the emotional states produced by the preceding actions.

In future work, we will undertake further analysis of the effect of player and comparator sets upon character prevalence since we found no real trend to this in these experiments. We also wish to functionally model the emotion of *contempt*, which will act in the opposite way to admiration and we will analyse

how this affects both character prevalence and total system scores. We will then simulate the additional emotion *greed* by which agents will periodically defect when co-operation is established in an attempt to secure an advantage.

## References

1. S. B. Algoe and J. Haidt. Witnessing excellence in action: the “other-praising” emotions of elevation, gratitude and admiration. *The Journal of Positive Psychology*, 4(2):105–127, 2009.
2. R. Axelrod. *The Evolution Of Cooperation*. Basic Books, Inc., 1984.
3. R. F. Baumeister, C. N. DeWall, K. D. Vohs, and J. L. Alquist. *Does Emotion Cause Behavior (Apart from Making People Do Stupid, Destructive Things)?*, chapter 7, pages 119–136. Oxford University Press, 2009.
4. R. F. Baumeister, K. D. Vohs, C. N. DeWall, and L. Zhang. How emotion shapes behavior: Feedback, anticipation, and reflection, rather than direct causation. *Personality and Social Psychology Review*, 11:167–203, 2007.
5. R. A. Brooks. Intelligence without representation. *Artificial Intelligence*, 47:139–159, 1991.
6. N. H. Frijda. *The Emotions*. Cambridge University Press, 1987.
7. J. Heinrich and F. J. Gil-White. The evolution of prestige: Freely conferred deference as a mechanism for enhancing the benefits of cultural transmission. *Evolution & Human Behavior*, 22:165–196, 2001.
8. W. James. What is an emotion? *Mind*, 9:188–205, 1884.
9. H. Jiang, J. M. Vidal, and M. N. Huhns. EBDI: An architecture for emotional agents. In *Proceedings of the 6th International Joint Conference on Autonomous Agents and Multi-Agent Systems, AAMAS '07*, pages 38–40. ACM Press, 2007.
10. D. Keltner and J. J. Gross. Functional accounts of emotions. *Cognition and Emotion*, 13(5):467–480, 1999.
11. M. Lloyd-Kelly, K. Atkinson, and T. Bench-Capon. Emotion as an enabler of co-operation. In *4th International Conference on Agents and Artificial Intelligence, ICAART 2012*, volume 2, pages 164–169, 2012.
12. P. Lockwood and Z. Kunda. Superstars and me: predicting the impact of role models on the self. *Journal of Personality and Social Psychology*, 73(1):91–103, 1997.
13. P. Lockwood and Z. Kunda. Increasing the salience of one’s best selves can undermine inspiration by outstanding role models. *Journal of Personality and Social Psychology*, 76(2):214–228, 1999.
14. A. Ortony, G. L. Clore, and A. Collins. *The Cognitive Structure of Emotions*. Cambridge University Press, 1988.
15. B. R. Steunebrink, M. Dastani, and J.-J. C. Meyer. A logic of emotions for intelligent agents. In *22nd Conference on Artificial Intelligence*, pages 142–147. AAAI Press, 2007.
16. B. R. Steunebrink, M. Dastani, and J.-J. C. Meyer. A formal model of emotions: Integrating qualitative and quantitative aspects. In *European Conference on Artificial Intelligence*, volume 178, pages 256–260. IOS Press, 2008.
17. U. Wilensky. Netlogo. <http://ccl.northwestern.edu/netlogo>, 1999. Date Accessed: 23/6/2010.