Int. Workshop on Emotion Representation and Modelling for HCI System (ERM4HCI), Sydney, Australia, Dec. 2013.

An Exploratory Study of the Role of Mood and Social Relationship in Collaboration

Mohammad Shayganfar, Charles Rich, and Candace L. Sidner

Collaborative Cognitive Agents (COCOA) Lab.
Computer Science Department,
100 Institute Road, Worcester, Massachusetts 01609 USA
{mshayganfar,rich,sidner}@wpi.edu
http://web.cs.wpi.edu/~rich/always

Abstract. We investigate how different moods and social relationships influence the performance of a dyad collaboration in a simple desktop game. Would the participants rather use their own resources to achieve the shared goal as quickly as possible or collaborate with their partner? Both the appraisal theory of emotions and collaboration theory are foundations of this research. First, we conducted a human study demonstrating the effects of mood and relationship on game play. Second, we implemented a computational model for two virtual agents and compared it to the human study. In the future, we plan to extend the computational model and evaluate it in a human-robot collaboration.

Keywords: Collaboration, Appraisal Theory, Relationship Types, Mood, Emotion, Social Robotics, Human-Robot Interaction

1 Introduction

Humans collaborate with others in different environments to achieve their short or long term goals. Sometimes the reason is to learn and adapt quickly and effectively to an environment that is exceptionally complex and turbulent, as well as to do tasks one cannot do alone. Because collaborators cannot read their partner's mind, they rely on surface behaviors including emotional expressions to comprehend what is important and appropriate to their partner during the collaboration. A collaborator who cannot read the partner's emotional expression, reason about the partner's emotions, and comprehend what is important to the partner will not act appropriately and intelligently during such interactions [1].

Emotions are short-lived psychological-physiological phenomena that represent efficient modes of adaptation to changing demands of the environment [2]. Moods, are affective states that last longer than emotions, usually for hours or days [3]. Emotions give meaning to people's evaluation of their social environment. Indeed, the necessity of a social approach to understanding the affective aspects of human cognition is apparent in artificial intelligence, psychology and the social sciences. There are numerous ways that emotions can be social. They can be conceptualized as responsive to social events and entities, and regulated by social constraints and affordances, and perhaps most provocatively, we can conceptualize emotions as socially constituted.

In this research, our goal was to investigate the role of mood in human collaboration in two relationship types, friendship (communal) and dominant-submissive. Collaboration is a special type of coordinated activity in which the participants work jointly with each other, together performing a task or carrying out the activities needed to satisfy a shared goal [4]. Cognitive appraisal theories of emotion specify the interpretation of events in our surrounding environment and how they are related to emotional experiences [5]. Hence, emotions and affect in general play an important role in this social context.

Our focus is to study how the collaborative and emotional aspects of cognition impact the performance of a short-term dyad collaboration. The study reported here contributes to a better understanding of human collaboration in a simple desktop game. Exploration of the data from the collaboration guided us to implement a synthetic form of the collaborative desktop game running between two virtual agents. After implementation, we "tuned" the architecture to perform like the humans, which we discuss in Section 4.

2 Related Work

To our knowledge, no prior work has examined the relationship between appraisal theory and collaboration theory. However, numerous studies in related areas of science, including psychology, neuro-science, sociology, computer science and artificial intelligence have tried to shed light on the concept and function of individual emotions and affect in general, from several different perspectives. In computer science and artificial intelligence, some researchers are interested in cognitive models of generating emotions, while others focus on nonverbal behaviors and emotional facial expressions. Furthermore, better recognition of emotions through voice, gesture and facial expressions has provided better human-computer and human-robot interaction.

The research set out in the following papers contributed to the design and concept of this study. Emotion, as one of the underlying concepts of our research, is discussed in [6], [7]. As is evident in these sources, researchers have not reached a consensus on these concepts. For this reason, we would like to focus more on the functionality of these psychological or sociological concepts, rather than on their underlying definitions. In [8] Keltner and Gross outline the history, elements and variations of functional accounts of emotions, and in [9] they integrate claims and findings concerning the social functions of emotions at different levels of analysis, including the dyadic level, in which we are interested because our human study and our computational model both investigate dyadic collaboration. Study [10] takes a psychological approach and presents much previous research to conclude that emotions are constantly changing as part of a dynamic social context. As such, artificial agents require dynamic models that are continuously adapting to the interaction and the environment. We have used this concept in our study as our desktop game provides a dynamic social context and our agents' emotions can change depending on the status of the game. However, we have restricted and fixed our agents' overall moods as happy or angry. In [11] the authors focus on the cognitive theory of emotions and its related details to answer the questions "what are emotions?" and "what is it about persons and situations that determine how situations are evaluated?" which relates to our evaluation processes of the situations in our game.

In [12] Scheutz addressed the possibility of designing affective artificial agents, discussed "how affective states could usefully interact with rational processes" and concluded that affect mechanisms should be systematically evaluated in a wide variety of situations to determine the general usefulness of affect for artificial agents. Our research is specifically evaluating how affective states affect humans in collaboration, and whether we can extend this effect to artificial agents. In [13] the researchers integrated an emotion model into a reinforcement-learning architecture to influence perception, provide reinforcement value, and determine when to re-evaluate decisions. This architecture was tested in the context of a solitary learning robot performing a survival task, and the researchers concluded that artificial emotions are useful in behavior-based autonomous agents in this type of environment.

Study [14] constitutes a conceptual comparison with artificial intelligence. Artificial agents were created and placed in an environment where they needed to compete for resources to survive. The agents could be social or non-social, adaptive or non-adaptive, and have a fixed or a variable conflict tendency. The emotional agents were considered as adaptive agents with variable conflict tendencies. These agents out-performed all others except social adaptive agents, revealing the utility of emotions, both in biological and artificial agents. We are taking a similar general approach with our computational model, in that our agents are placed in an environment where they must compete for resources to win, or survive. Our agents also have a similar characteristic as the conflict tendency, in that they have a tendency to keep higher value resources and propose that their partner play its higher value resources. A key difference between the environments is that, while our agents compete for resources, they are also working towards a shared goal.

In [15] researchers reviewed behavioral research to study the significant role of mood and emotion on cognition. They revealed that positive affect leads to "relational, global or category-focused processing," while negative affect leads to more "perceptual, local or item-level processing." For this reason, the researchers recommend that artificial agents should limit the applicability of negative affective information, and unconstrained or spread the implications of positive affect. Because of this difference in processing due to the type of affective state, our game includes both positive and negative affective states which allowed us to evaluate their effects in our specific collaboration context. [16] is a biologically inspired study that uses social referencing through robotic imitation of human expressions in an anthropomorphic robot. The robot can then evaluate its own imitated expression to determine the affective state of its human companion, and use this information to learn and remember affective appraisals for objects. Although similar in nature to our study, the concept of imitating a partner's expression to determine their affective state was not utilized in our study, due to its exploratory nature.

Collaboration, including basic concepts such as goals, beliefs and intentions, as another foundational theory for our research has been explored by studies such as [4] in which the researchers updated SharedPlans, a model of collaborative planning, to handle complex actions and allow for plans to be partial, and [17] in which the researchers establish "basic principles governing the rational balance among an agent's beliefs, actions, and intentions" to manage an agent's persistence and commitment to a task. These concepts of collaboration are the foundation of our research. They were used in design of the human study and will be used in our computational model. In [18] a computational model is developed for the interplay of emotions, beliefs and intentions in a group decision-making context based on insights from social neuroscience. The realization of the modeling of theory of mind is discussed in [19], as the researchers' appraisal model models other agents' goals, states and beliefs. Our goal is to apply theory of mind in future works, and both [18] and [19] will be useful in our project as we investigate the influence of affect in the collaboration structure including beliefs, intentions and goals.

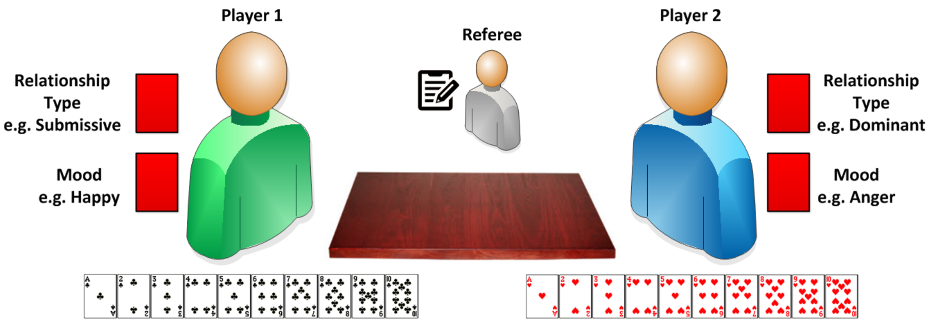

Fig. 1. Setting of the game for human study.

3 Human Study

This study examined the perceived effects of mood and social relationship type on three aspects of collaboration performance within a simple desktop game context using cards: number of steps, time to accomplish the shared goal, and the score difference between participants. In this game, players had two challenges: (1) to pursue their personal interests by playing lower value resources of their own and by proposing the other player use higher values; (2) to collaborate with their partner to achieve the shared goal and finish the game. It was hypothesized that, relative to individuals' given mood and relationship (a) participants assigned happy moods would collaborate better than participants assigned angry moods, (b) participants assigned to a friendship relationship would collaborate better than participants assigned to a dominant-submissive relationship. (See instructions in Table 1).

Thirty-six people in eighteen pairs were recruited as participants. We eliminated data from six participants due to human errors during the study. All analyses in this paper are based on the data gathered from thirty participants in fifteen random pairs (16 males and 14 females). Participants did not receive any reward for participation.

The study was conducted at Worcester Polytechnic Institute. The experiment took place in a room containing two seats for two participants and one for the referee (researcher) who provided the experimental instructions and logged the data and the results from each game, including the time when the players reached the predefined shared goal. There was also one table on which the desktop game was played (see Fig. 1). During all games, participants were seated upright in comfortable chairs. This study was designed to be within-participant in order to eliminate the effects of individuals' age and gender. Each pair played eight games, taking on all possible combinations of mood and relationship.

3.1 Procedure

On their arrival at the laboratory, participants were met by the experimenter, who introduced himself as the referee and provided information about the experiment. After participants read the one-page instruction sheet shown in Table 1, the referee asked whether they had any questions. The referee once again explained key points including the rules about winning and losing, the necessity of collaboration, the meaning of the two face-down cards, and the final goal of the game for each participant.

Two different moods were chosen for the study, happy and angry. These two moods were chosen in order to have one positive and one negative valence of related affective instances and to simplify the number of moods for social relationship types, i.e. dominant-submissive and friendship relationships. It is important to note that the participants were not placed in certain moods and relationships but assigned certain moods and relationships as roles to play. The participants were supposed to act out of actual happiness or anger but may have been attempting to recreate behaviors expected from happy or angry individuals. Likewise, the relationships were assigned and participants' behavior were based on their perceptions of dominant-submissive and friendship relationships, not the relationship types themselves. The purpose of the first goal (see Table 1) was to establish collaboration in the game context requiring the participants to cooperate to achieve the shared goal. The second goal (see Table 1) was designed to motivate participants to pursue their personal interests to win the game and thereby arouse participants' emotions.

Table 1. Instructions given to the participants

Each player will be given two cards face-down:

- One card will specify your **role** in the relationship.
- The other will specify your most prevalent emotion/**mood**.

You will also be given **10 playing cards** (resources) with values **1 through 10** which you will be required to play during the course of the game. Higher values represent higher resources.

There are **two goals** in this game:

1. The partners need to work together to build a **value of 60**.
2. Each partner wants to have as many resources as possible remaining at the end of the game.

How to play:

The judge will give each player his/her role and emotion/mood. He will also determine which player is to go first (Player 1). Player 1 will play one of the ten cards in his/her set, based on his/her relationship role and emotion. He will then recommend/suggest/propose a card for Player 2 to play from his set, again based on emotion and role. Player 2 faces a choice: to reject or accept the proposal. He can play any value, but the decision is to be based on the suggestion and his specified mood and role. After he plays, he makes a suggestion to Player 1, who will take his turn. This process repeats until the goal of 60 has been reached.

At this point, the game is ended. The judge will determine how many resources each player has remaining. He will then reset the game and repeat the previous steps 7 more times with different circumstances in each round.

3.2 Results

Box plots are used to show the results of the study. As shown in Fig. 2(a) and Fig. 2(b) and Fig. 3, participants playing in happy roles require fewer steps and less time to achieve the shared goal. The participants also acted such that those assigned happy moods had smaller score differences than those assigned angry moods. These three figures also show that participants believe that partners in a friendship relationship require fewer steps and less time than partners in a dominant-submissive relationship, although the time to completion does not show a substantial difference. Partners in a friendship relationship are also expected to have a smaller score difference based on the participants' behavior. Fig. 2(a) which depicts the number of steps to achieve the shared goal only has one outlier in the dominant-submissive bar.[1] Fig. 2(b) and Fig. 3 have at least one outlier on nearly every bar.[2]

We determined the p-values for the difference between means using paired comparisons. We used a one-tailed t-test to determine whether the data gathered supported our hypotheses, i.e. participants assigned to happy roles collaborate better than participants assigned to angry roles and participants acting as friends collaborate better than participants acting out a dominant-submissive relationship. The p-values that supported our hypotheses are marked in Fig. 2 and Fig. 3 with either one or two asterisks, depending on their level of significance.

[1] Outliers are represented by small circles.

[2] Mood conditions are aggregated over relationship types and vice versa.

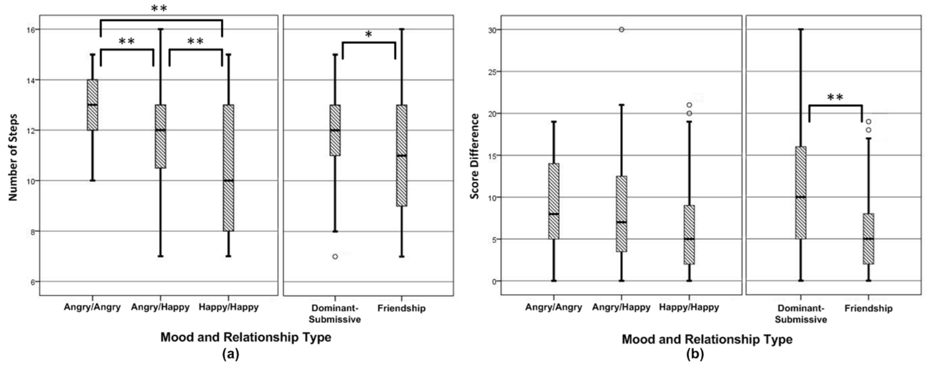

Fig. 2. (a) Number of steps to reach the shared goal, and (b) absolute value of score difference of the players versus players' assigned mood state and relationship type. ** p-values <0.01 (very significant), * p-values <0.05 (significant).

Significance is reported without the Benferroni procedure to enable evaluation of isolated comparisons. These plots show that the data generally supported our hypotheses.

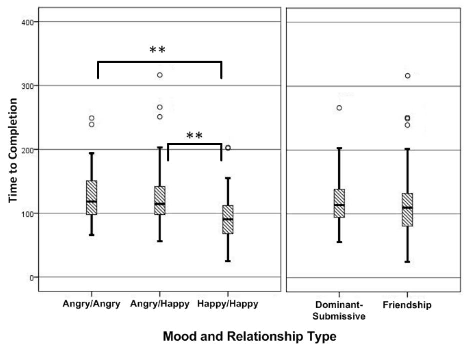

Fig. 3. Time to completion (in seconds) versus players' mood state and relatinship type. ** p-values <0.01 (very significant), * p-values <0.05 (significant).

We also investigated the effect of gender on the results. Due to the fact that gender was not assigned to the participants, the effects recorded are possibly the direct result of gender, and not behavioral expectations of the participants. We compared the measured number of steps, score differences, and time to completion between male-male, male-female and female-female groups. In this study, gender only affected time to completion. The mixed-gender groups took longer to finish the task than single-gender groups. The p-value comparing female groups to mixed groups was 0.08, a weak trend, and p-value comparing male groups to mixed groups was 0.0001. The p-values for other comparisons including male-male to female-female time to completion, were greater than 0.25. In this comparison, the p-value showing a 0.05 level of significance, as calculated by the Benferroni procedure is 0.017. These gender effects do not confound the mood and relationship effects we observe above, because each pair performed all eight conditions. We believe the effect of gender on collaboration in this setting requires further study.

3.3 Discussion

Studying humans' affect, and more specifically mood, in dyad collaboration in different social relationship types was a challenge. We were curious whether our participants would be able to actually feel the assigned mood for each game as well as understand and play their own social role in the given relationship type. Participants' comments during each game revealed that most of the participants were (a) trying to express their assigned mood and related emotions in a perceivable way using verbal and nonverbal cues, and (b) trying to map the given mood to the given relationship type independent of their gender and age. For instance, one participant assigned to an angry submissive role said: "How come you never listen to anything I suggest!" while the same player as a happy friend negotiated with his partner to have equal resources remaining at the end of the game.

As mentioned previously, participants acting as happy players reached the shared goal in fewer steps as supported by the p-values. We believe this is because happy collaborators are perceived as being more likely to focus on the shared goal and collaboration instead of their personal desires to win the game. This leads the participants to use more resources to benefit the collaboration.

In Fig. 2(a), the wide range between the maximum and the minimum number of steps where the players' moods are different implies that when participants are acting out heterogeneous moods, their behavior creates a certain level of confusion which impacts the number of required steps to reach the shared goal. In Fig. 2(b), all the minimum values of the box plots are zero, indicating that at least one pair of players who finished the game had tied scores, independent of their assigned mood or relationship type. Also, the fact that the median values in most of the boxes are closer to the minimum rather than the maximum, means the players had a tendency to tie the game rather than winning with highest possible score. This fact indicates that the participants believed that collaborators view a shared goal as more important than personal advantage, except among pairs assigned dominant-submissive roles, where the dominant player drove the submissive one to achieve the goal quickly.

In Fig. 3, the outliers were the result of discourse between players who attempted to negotiate the resources they would apply during the game. Most appear when one or both of the players' assigned moods are angry. The maximum number of outliers occur when the players are assigned heterogeneous moods, implying that participants believed that collaborators in different moods are more likely to struggle. Alternatively, assigned friendship relationships have more outliers than dominant-submissive ones, meaning the participants expect friends to negotiate more and show tolerance during collaboration.

4 Computational Experiment

The computational experiment described here was our first step towards implementing a cognitive architecture inspired by the cognitive appraisal theory of emotions [5] and collaboration theory [4]. Our goal was to create a simple computational model that would produce behaviors matching those of our human study. In the future (see Section 5), we plan to run the same model in a human-robot collaboration setting. Hence, we built our own cognitive architecture with many manually tuned parameters, such as percentage of proposals rejected in various cases, that allowed us to match the human data. We ran our computational model as the cognitive part of two agents in a virtual environment simulating the same desktop game context. We tuned our model so that the agents' behavior matched the human data, and then compared our results.

4.1 Computational Architecture

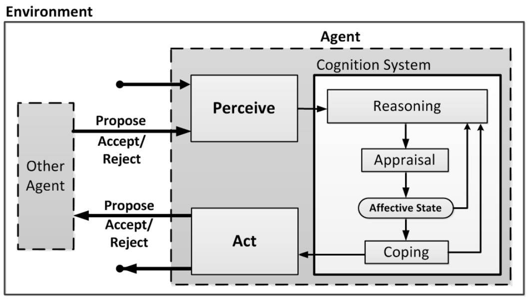

Fig. 4. Architecture of the computational model.

As mentioned above, the cognitive appraisal theory of emotions and collaboration theory are the underlying foundations of our work. This is reflected in our computational architecture shown in Fig. 4. From appraisal theory we have the concept of appraisal frames (containing appraisal variables) as an intermediate description of the agent-environment relationship and a mediator between the stimuli and responses. These variables characterize the significance of events from the agent's point of view [20]. In our model, we used *Desirability*, *Likelihood* and *Causal Attribution* as three appraisal variables to generate our two emotions, anger and happiness. For instance, a positive value for the desirability variable and a high probability for the likelihood variable elicits joy as the emotion instance. After appraising the significance of the events, the coping mechanism helps the agent determine how to respond. Our preliminary computational model does not contain the existing distribution of coping strategies in the literature [21], namely problem-focused and emotion-focused strategies. However, our model's response to the environment could be considered to be one of the problem-focused strategies.

Our model includes the concept of proposing a task as postulated in collaboration theory. Each agent proposes a primitive task to the other agent – which resource (card) the other agent should play – in its turn. The other agent either agrees to play the proposed card, or rejects and plays its own desired card. Rejection appears in two different ways. The agent could play a lower value or a higher value than the resource proposed by the other agent. In all cases, the agent's mood and the social relationship role in the game affects the agent's decisions.

Turn taking in the game was used to give the agent time to perceive the environment. This perception includes (a) the other agent's remaining resources, (b) the resources that have been played by both agents to reach the shared goal, (c) the value proposed by the other agent, (d) the resource that was most recently played by the other agent. The agent then applies the acquired data from the environment to make a decision about which value to play. The current self-mood state, the social relationship role of the agent, and the last emotion instance, as well as whether and how the other agent accepted or rejected the last proposal, are criteria that help the agent to make the decision about how to act on its turn.

4.2 Results

Fig. 5(a) shows the comparison of the mean number of steps for different pairs of moods and relationship types, between the synthetic data and the data we gathered from our human study. As the figure shows, the difference between the human and the synthetic means are less than two steps in all cases.

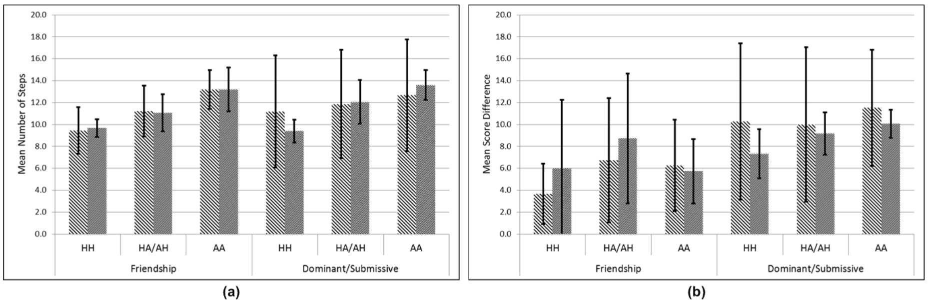

Fig. 5. Mean value and standard deviation of (a) number of steps, and (b) score differences versus mood for human (left bar) and synthetic data (right bar).

Fig. 5(b) shows that the mean absolute value of score differences of the synthetic data is similar to that of the human data in most mood and relationship type combinations. The only instances in which the difference is greater than two points is when both of the players are assigned happy moods, independent of the relationship type.

Fig. 6(a) illustrates the comparison of the mean number of steps of the synthetic data and the human data for different pairs of relationship types, according to the existing four pairs of moods. As the figure shows, the difference between the number of steps for the synthetic data and the human data is less than one step, except when both players are assigned happy moods in a dominant-submissive relationship. Error bars in each pair of mean values have less differences when compared with Fig. 5(a) and Fig. 5(b).

Fig. 6(b) compares the mean score differences of the synthetic data and the human data for different pairs of relationship types within four pairs of assigned moods. In most cases, the average values are close and less than three points. There are two exceptions, when both of the players are assigned to a dominant-submissive relationship and the second player is assigned to a happy mood (see circled labels in Fig. 6(b)). Again, the differences between error bars for each pair of mean values are similar to those in Fig. 5(a) and Fig. 5(b).

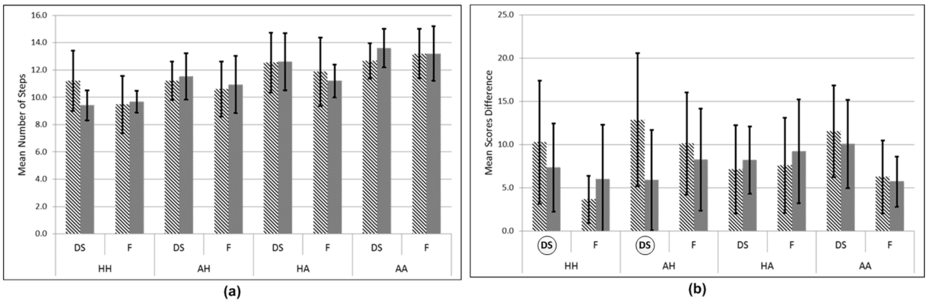

Fig. 6. Mean value and standard deviation of (a) number of steps, and (b) score differences versus relationship type for human (left bar) and synthetic data (right bar). Circles indicate large differences between synthetic and human data.

4.3 Discussion

The similarity between the means and the standard deviations of the human data and the synthetic data in Fig. 5 and Fig. 6 show that we were successful in tuning our computational model to the human data with reasonable accuracy. The patterns in both sets of data were similar. The figures show that we were able to tune our model to nearly match the number of steps to those of the human data. However, we were only able to produce collaborative behavior in our virtual agents that would match the mean score differences of the human data in some cases.

The inability to adequately mimic humans' collaborative behavior is due to the granularity of details in our current model. For instance, if our agents were able to elicit more than two emotions during collaboration regardless of the mood states, they would be able to evaluate internal and external events more accurately. Our current model supports having more emotions; however, we need to assign more appraisal variables to our appraisal system as other emotions require. These new emotions would allow the agents to evaluate the events with greater detail, leading to a better understanding of the environment.

As mentioned above, our current coping system only has one problem-focused strategy. A coping system possessing a series of emotion-focused strategies could help the agents to be more successful in mimicking humans' emotional behavior.

One final reason for the non-matching data might be that the agents currently have a lack of beliefs about the other agents' beliefs and intentions towards the game and the collaboration. If our model were to include the concept of theory of mind, our agents would be capable of forming beliefs about their collaborator's beliefs, and consequently infer intentions behind their actions. This inference mechanism also benefits the values given to *Causal Attribution* as one of the appraisal variables in our model. This could notably improve the collaboration performance which directly impacts the score difference and the other measures.

5 Conclusion and Future Work

In this paper, we have described a collaboration of two players in a desktop game context working based on collaboration theory and the appraisal theory of emotions. We also described the human study and our first step computational model which was tuned based on our observations from the human study. We recommend a similar human study in which participants are placed in certain moods, rather than assigned moods, so the effects of mood can be directly observed, rather than the expected effects of mood.

As mentioned before, our next step is to use the same computational model on an expressive robot, REETI [3] (Fig. 7). However, we will first improve our model as mentioned in Section 4.3 and check whether we can tune it to be a perfect match for the human data. We plan to run another study in which our robot and a human participant are involved in the same desktop game context and try to collaborate with each other to accomplish the same shared goal. We look forward to determining where the results are similar or different from what we observed in our previous human study.

Fig. 7. REETI

Acknowledgments. This work is supported in part by the National Science Foundation under award IIS-1012083. Any opinions, findings, and conclusions expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

References

1. Picard, R.W.: Affective Computing. The MIT Press (2000)
2. Levenson, R.W.: The interpersonal functions of emotion. Cognition and Emotion **13**(5) (1999) 481–504
3. Frijda, N.H.: Varieties of affect: Emotions and episodes, moods and sentiments. In Ekman, P., Davidson, R.J., eds.: The Nature of Emotion, Fundamental Questions. Oxford University Press (1994) 59–67
4. Grosz, B.J., Kraus, S.: Collaborative plans for complex group action. Artificial Intelligence **86**(2) (1996) 269–357
5. Scherer, K.R., Schorr, A., Johnstone, T.: Appraisal Processes in Emotion : Theory, Methods, Research: Theory, Methods, Research. Oxford University Press (2001)
6. Cowie, R., Sussman, N., Ben-Ze'ev, A.: Emotion: Concepts and definitions. In: Emotion-Oriented Systems. Cognitive Technologies. Springer-Verlag Berlin Heidelberg, London, New York (March 2011) 9–30
7. Frijda, N.H., Manstead, A.S.R., Bem, S.: Emotions and Beliefs: How Feelings Influence Thoughts. Cambridge University Press (2000)
8. Keltner, D., Gross, J.J.: Functional accounuts of emotions. Cognition and Emotion **13**(5) (1999) 467–480
9. Keltner, D., Haidt, J.: Social functions of emotions at four levels of analysis. Cognition and Emotion **13**(5) (1999) 505–521

[3] REETI is an expressive and communicating robot built by Robopec. (http://www.robopec.com)

10. Claudia Marinetti, Penny Moore, P.L., Parkinson, B.: Emotions in social interactions: Unfolding emotional experience. In: Emotion-Oriented Systems. Cognitive Technologies. Springer-Verlag Berlin Heidelberg, London, New York (March 2011) 31–46
11. Ortony, A., Clore, G.L., Collins, A.: The Cognitive Structure of Emotions. Cambridge University Press (1988)
12. Scheutz, M.: Architectural roles of affect and how to evaluate them in artificial agents. International Journal of Synthetic Emotions **2**(2) (2011) 48–65
13. Gadanho, S.C., Hallam, J.: Robot learning driven by emotions. AI Magazine **9**(1) (2001) 42–64
14. Scheutz, M.: Useful roles of emotions in artificial agents: A case study from artificial life. In: Proceedings of AAAI, AAAI Press (2004) 31–40
15. Clore, G.L., Palmer, J.: Affective guidance of intelligent agents: How emotion controls cognition. Cognitive System Research **10**(1) (2009) 21–30
16. Thomaz, A.L., Berlin, M., Breazeal, C.: Robot science meets social science an embodied computational model of social referencing. In: Workshop Toward Social Mechanisms of Android Science (CogSci). (2005) 7–17
17. Cohen, P.R., Levesque, H.J.: Persistence, intention, and commitment. In Cohen, P.R., Morgan, J., Pollack, M.E., eds.: Intentions in Communication. MIT Press, Cambridge, MA (1990) 33–69
18. Hoogendoorn, M., Treur, J., van der Wal, C.N., van Wissen, A.: Modelling the interplay of emotions, beliefs and intentions within collective decision making based on insights from social neuroscience. In: Proceedings of the 17th international conference on Neural information processing: theory and algorithms - Volume Part I. Volume 6443 of Lecture Notes in Computer Science., Springer (2010) 196–206
19. Si, M., Marsella, S.C., Pynadath, D.V.: Modeling appraisal in theory of mind reasoning. In: Journal of Autonomous Agents and Multi-Agent Systems; Proceedings of the 8th International Conference on Intelligent Virtual Agents. Volume 20(1). (January 2010) 14–31
20. Gratch, J., Marsella, S.: Evaluating a computational model of emotion. Journal of Autonomous Agents and Multi-Agent Systems. Special Issue on the Best of AAMAS 2004 **11**(1) (July 2005) 23–43
21. Lazarus, R.S.: Emotion and Adaptation. OXFORD University Press (1991)