Emotional Robotics: Tug of War

David Grant Cooper DCOOPER@CS.UMASS.EDU
Dov Katz DUBIK@CS.UMASS.EDU
Hava T. Siegelmann HAVA@CS.UMASS.EDU
Computer Science Building, 140 Governors Drive, University of Massachusetts, Amherst, MA 01003

Abstract

Emotional communication skills are dominant in biological systems. Although the rules that govern creating and broadcasting emotional cues are inherently complex, their effectiveness makes them attractive for biological systems. Emotional communication requires very low bandwidth and is generally easy to interpret. Despite the advantages of emotional communication, little or no research has explored which emotional cues are the most effective when used by a robot. To study this question, we introduce an interactive environment in which a person can learn the robot's emotional responses through interaction. We then present a one player game in which a person attempts to attract the robot's attention, make it move towards and stay close to the person. We further develop this concept into a two player version, in which the players engage in a *Tug of War* game, competing for the robot's heart. We propose our system as a potential test bed for human-robot interaction, both for engineers, and clinical psychologists.

1. Introduction

Emotional communication is a complex interactive process. It may involve multiple agents with different desires and goals. Each agent communicates its emotional state and can deliberately try to manipulate the other's emotional state through interaction. This complex communication process becomes even more intricate when considered in the context of a multi-agent, dynamic and complex environment. Despite its inherent complexity, emotional communication is an efficient way to communicate goals and desires. In this paper, we hypothesize that humans can acquire and adapt to the emotional mechanism that governs the behavior of a robot. Moreover, we evaluate how effective using facial expressions as secondary feedback in facilitating the acquisition of the robot's emotional behavior.

To explore our hypothesis, we have developed an emotional robot with which human participants can interact. Through interactions with the robots, the participants are able to develop a model of its emotional behavior. A reliable model can predict the outcome of future interactions, thus enabling the participant to manipulate the robot to a desired emotional state. The participant's ultimate goal in our experiments is to make the robot happy. Much like preverbal communication with infants, this can be achieved through motion and voice.

Our main contribution is a simulated robot experimentation platform that can interact using visual and auditory sensors and a video representation of the robot's motion and facial expression. The robot must provide cues that can be understood by the human participants. Moreover, the robot must possess enough sensor capabilities to observe cues generated by the human participants. Based on these cues, both the robot and the human participant can interact. The robot demonstrates its desires, and human participants learn what pleases the robot.

Successful implementation of this robot-human emotional interface would create a platform for testing how well humans can model the emotional state of the robot given different feedback from the robot. This platform will enable us to investigate whether people can successfully attribute cause and effect relations to the behavior of the robot, and then use these relations to manipulate the robot into a state of happiness.

The testing platform is in the form of a game. The goal of the game is to get Danny, the robot, to come and stay close to the player. This is done by performing actions that Danny likes. The score in the game allows us to measure whether Danny is able to accurately and effectively communicate its desires. Danny communicates implicitly by moving back and forth. In addition, Danny can communi-

Presented at *North East Student Colloquium on Artificial Intelligence (NESCAI)*, 2008.

cate with facial expression, and the system as a whole can communicate by accumulating a score for the player.

We consider two versions of this game for our experiment. The first is a one-player version in which a human tries to convince the robot to move towards it. This can be achieved through emotional communication — by making the robot happy, the robot will be enticed to move towards the human. This version of the game enables us to explore the effectiveness and ease of use of our emotional interface. The second version is a two-player game. This version is a type of *Tug of War* (also known as rope pulling). In *Tug of War*, two teams are competing against each other. A team wins the competition by pulling the other team towards it. Successful teams utilize physical strength, mental strength, and coordination. In the two player version, competitors try to gain the robot's trust and affection. This is more intricate than two simultaneous one-player games, as the effect of direct competition between humans adds a new dimension of emotional communication. *Tug of War* will enable us to explore the reliability of emotional communication when the primary expression of approach and withdraw does not always correlate with the robots feelings.

Our implementation relies on cues that are very common in human communication. Running away represents fear. Getting closer signifies trust. Smiling, or putting on an angry face are strong ways to communicate an emotional state. In much the same way, a robot that plays *Tug of War* can elicit and express emotions by moving closer or running away, smiling or frowning.

The proposed platform provides an opportunity to conduct clinical psychological research on human participants. The emotional response that the robot generates in humans as part of the competition, as well as the emotional response that the competitors induce on the robot, create a complex emotional interactive environment. Observing this interactive emotional communication will provide an interesting test bed for interdisciplinary research on human and robot emotional communication. We hope that with the development of this test bed, researchers in psychology will be able to provides new insights and develop new models of human-robot emotional interactions.

On the application side, we believe that adding an emotional aspect to existing Human-Machine interfaces will create a new layer of security. For instance, in the battlefield, robots could choose to cooperate only with people they consider reliable and trustworthy. Databases could be protected by providing information only to the person that convinces the robot it is the rightful owner of that information. Although these applications are very promising, they all can be emotionally manipulated. Our platform will provide a test bed for exploring what measures needs to be taken to overcome this difficulty and guarantee both effective and reliable emotional communication.

In the following sections, we discuss related work, the details of the hardware and software components of our platform, as well as the experimental setting and results.

2. Related Work

(Takeuchi & Naito, 1995) compare the usefulness of a situated 3-D animated face pointing compared to an arrow pointing. A person is shown to perform better from just the arrow, but the face is better at grabbing the person's attention. In this case there is a neutral facial expression the whole time.

Cynthia Breazeal pioneered the use of an emotional model and emotional expression using her robot, named Kismet,((Ferrell), 1998),(Breazeal & Scassellati, 1999),(Breazeal, 2002), and (Bar-Cohen & Breazeal, 2003). Kismet was shown to be able to regulate its internal state based on social interaction. Kismet also used facial expressions, sounds, head and eye motion to convey its emotional state. Though a persons interactions were influenced by these expressions, no work was done to show how much each individual expressive feature accounts for the influence. In addition no work was done on whether Kismet's emotion could help in learning.

(Kringelbach & Rolls, 2003) used neutral and angry faces as a reinforcement signal for humans to learn to change their selection from one face to the other. The response time was a little bit slower with the neutral reinforcement signal, but both faces were learned. In this case the faces presented were black and white photos from Paul Ekman's collection(Ekman & Friesen, 1971). This work shows that faces can be used as a reinforcer, but they were not used in conjunction with other feedback.

(Bruce et al., 2002) used a robot speaking robot with a screen as a head to show that facial expressions and head tracking each independently had an effect on a person stopping to listen to what a robot is saying and to answer a polling question by stepping up to a microphone. The facial expressions were on a 3-D animated face. In the no facial expression condition, the screen was blank. The experiment showed that facial expressions caused an increased probability of stopping, head tracking caused a slightly higher probability of stopping, and the combination of head tracking and facial expression caused a significantly higher probability of stopping than just the facial expression. The facial expressions of the robot were based on the robot's ability to get the person to follow the script. It was happy at first, and got less happy as the person did not participate as asked.

3. Emotional Communication Algorithm

3.1. States and Transitions

The emotional algorithm continuously evaluates and acts upon the robot's internal emotional state. The state is represented by an emotional state vector. In the current implementation, the state vector has three states: friend, foe, and self interest. That is $s = [s_{friend}, s_{foe}, s_{absorbed}]$. Each state is updated as follows:

$$\begin{aligned} s_{friend} \leftarrow s_{friend} &\quad + \quad w \cdot f_{friend}(Input) \\ s_{foe} \leftarrow s_{foe} &\quad + \quad w \cdot f_{foe}(Input) \\ s_{absorbed} \leftarrow s_{absorbed} &\quad + \quad w \cdot f_{absorbed}(Cycle) \end{aligned}$$

where w is a constant multiplier between 0 and 1 acting as a low pass filter and all states s_i are between 0 and 1, and all functions f_i return a value between 0 and 1. The state vector is always normalized by the $\mathcal{L}_1$ Norm after each update.

The expressed emotional state is then dependent on the values of the state vector. If s_{friend} or s_{foe} are above $\frac{1}{2}$, then the state will be non-neutral as shown in figure 1. Since the values of the vector are normalized after each update, at most one such value can exist. If neither the friend or the foe state are above $\frac{1}{2}$, then the expressed state is neutral. In order to alleviate fast switching between states, a state must be expressed for a minimum of 3 time steps. The expressed states have 3 values for friend and 3 values for foe, and one neutral state. A state can change by at most one degree each step. Our implementation is easily extendible to support more dimensions, such as surprise, fear, disgust, and sadness.

The self absorbed state represents times when the robot is incapable of handling inputs. This behavior models periods in which the robot has other needs that have to be fulfilled. These needs include replenishing the battery, some non-interactive tasks, performing off-line learning, and other events that dictate an anti-social behavior. When the robot is in this self-absorbed state, it takes on the neutral expression. The self-absorbed cycle increases at a slow rate for one third of a period, and then decreases at a fast rate for the other two thirds of the period. As the self absorbed value gets larger, it becomes increasingly more difficult to keep the state away from neutral.

3.2. Input Features

The emotional algorithm uses the following 10 visual input features: amount of motion on the screen or face, brightness of screen or face, whether or not the person is facing the robot, the position of the person's face, the amount of motion on the persons face, jittering motion of the face, continuous motion of the face, no motion of the face, and the variance of the face motion. In the experiments described in this paper, we have focused on two of these visual features: whether the person is facing the robot, and the motion variance of the person's face.

State Value	Face	How Friendly	Steps Towards Person	Change in Score
friend > 0.80		+3	+3	+3
friend > 0.65		+2	+2	+2
friend > 0.55		+1	+1	+1
friend & foe <= 0.55		Neutral	0	0
foe > 0.55		-1	-1	0
foe > 0.65		-2	-2	0
foe > 0.80	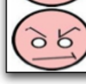	-3	-3	0

Figure 1. An example of our faces with the emotion, robot behavior, and game score output at each step. The threshold values on the left show when s_{friend} and s_{foe} will activate a particular expression.

The emotional algorithm uses the following 4 audio features: beats per minute (sampled from a second, or using a whole minute), average pitch in the last second (partitioned to high, medium, and low pitch), the mode pitch of the last second, and variance of the pitch. In the experiments described in this paper, we have only used one of these features: the variance of pitch.

We have explored several mappings from inputs to emotional state (the functions f_{friend} and f_{foe} described in section 3.1). In the experiments described in this paper, the mapping we use measures whether the person is facing the robot, and the variance in the person's pitch and motion. Figures 2 and 3 show the two Decision Trees that compute the function on the input. Facing the agent with which we communicate is an important feature of human communication. The motion variance and the pitch variance measure the jitteriness of the person.

The motion variance of the face is computed over a seven step period. Relatively low variance reflects little change in motion, which makes it easier for the robot to predict what the person is doing. Non-jittery behavior is indicative of comfortable and friendly communication. The pitch variance estimates the change in pitch that is typical of human speech. This was calibrated by the developer's speech

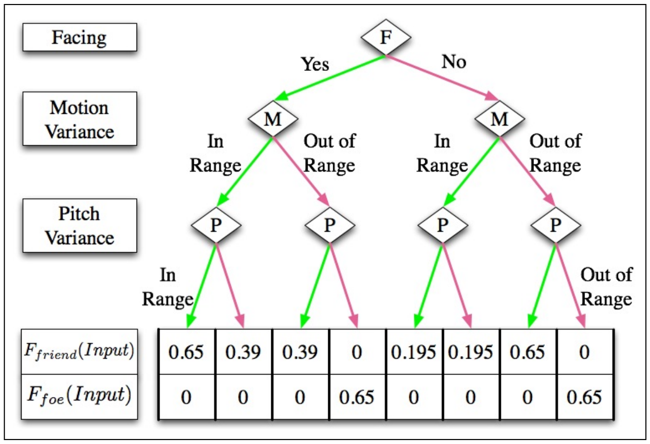

Figure 2. Emotional Input Decision Tree based on a variance threshold. Three thresholds were used: **low**: $1.0 < P < 1000.0$, $0.2 < M < 1.5$, **medium**: $1000.0 < P < 40000.0$, $1.5 < M < 3.5$, and **high**: $10000.0 < P < 40000.0$, $3.5 < M < 10.0$.

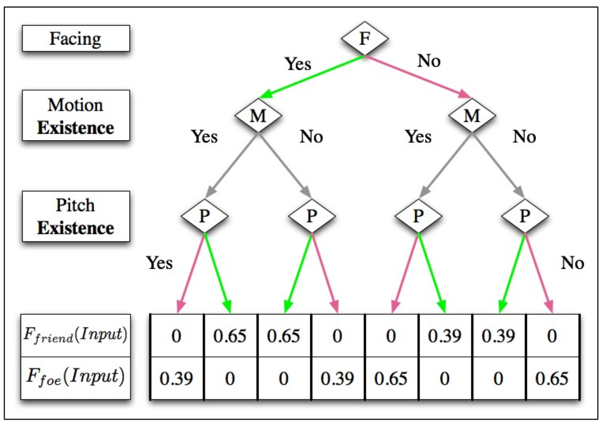

Figure 3. Emotional Input Decision Tree based on existence of motion or pitch.

and variance is computed over the last second. Regular speech, as opposed to yelling for example, is considered by the robot as a desired method for communication. Finally, we note that the system is designed to allow new input functions to include other inputs such as hand gesture detection or the human's facial expression.

3.3. Output Features

The emotional algorithm is composed of two features: the robot's facial expression, and the robot's motion towards/away from the person. The robot's facial expression and motion is updated after each processing step. Figure 1 shows an example of emotional output with transitions based on emotional state, and the faces we use. The system's design makes it trivial to add more output features in the future. Possible new outputs can range from adding more facial expressions. to verbal communication.

Currently, the output has 7 options for faces, and 7 options for motion, (approach 1-3, withdraw 1-3, and no motion). This creates a wealth of variety for how emotion is expressed in the physical space. Emotion can be expressed by how close the robot is to the person it is interacting with, how fast it is moving towards or away from the person, and whether it is oscillating back and forth. Currently, the gestures of motion are simple. However, extending the motion patterns to express particular emotions would be straight forward.

3.4. Implementation

Our model, which is inspired by (Breazeal, 2002; Bar-Cohen & Breazeal, 2003), is simple in order to focus purely on the emotional representation. Our model uses sensors, feature extraction, emotional transitions, and emotional expressions, while Breazeal's model has a visual attention system, a cognitive evaluation of the stimuli and drives, a set of higher level drives, an affective evaluation, an affective appraisal, emotion elicitors, emotion activation, behavior, and motor expression. This simplified model separates out the cognitive and reasoning aspects in order to get at the core concept of emotional representation. Rather than the three dimensional space in which emotions lie, we explicitly hold a value for each emotional state. This makes it so that more than one internal emotional state can change at the same time. In this way there can be strict breaks based on one state overwhelming the other states, or smooth transitions over a continuum as the states compete for expression.

Figure 4 shows examples of the emotional states being expressed from screenshots we took of a set of interactions with the simulated robot.

We now describe an example of an interaction between a human and the robot is as follows: The human faces the robot and rocks from left to right slowly for a short time with low variance. This behavior increases the friendliness state, and results in a friendly state (See Figure 1). The robot smiles and approaches the human. Next, the person starts to move faster, which displeases the robot as it makes predicting the person's behavior more difficult. In response, the robots changes its facial expression to less friendly values.

4. Experimental Validation

4.1. Experimental Platform

We use a simulated robot and a graphical environment (OGRE 3D) to display robot emotions. When the robot decides to approach the person with whom it interacts, it

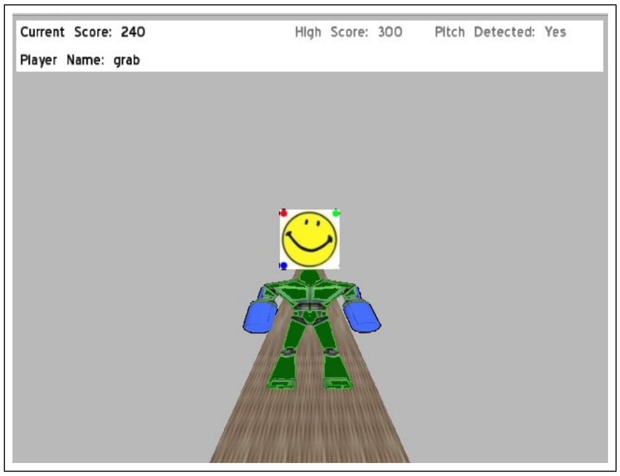

(a) Friendship being expressed by the robot: The robot approaches and shows a smile as it likes the participant

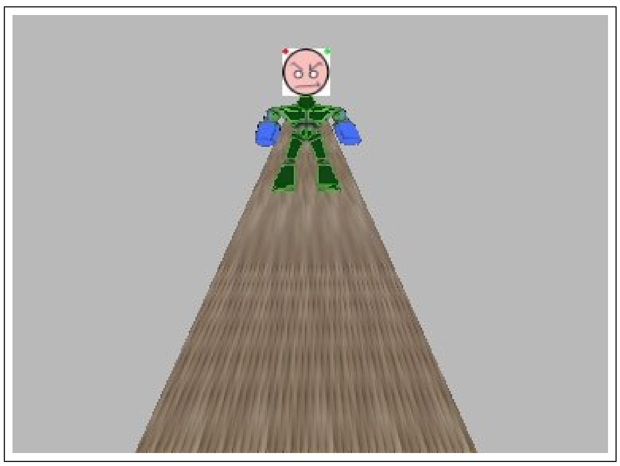

(b) Dislike being expressed by the robot: The robot withdraws with a scowl as it dislikes the participant.

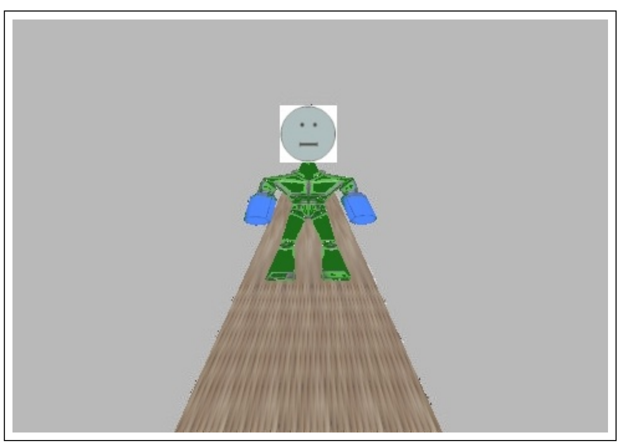

(c) Indifference being expressed by the robot. The robot stops and has a Neutral face

Figure 4. An example of 3 emotions elicited by a person and expressed by the robot. a) is Friendship, b) is Dislike c) is Indifference

moves forward in the simulated world. The robot accumulates sensory data from a camera and a microphone, and uses it to decide on transitions between its emotional states. In addition, the robot can express its emotional state using a facial expression.

The emotional architecture is composed of the emotional communication algorithm (described above), integrated with two sensor processing systems, and a 3D simulated robot environment. We use two open source libraries to processes sensor data: OpenCV for vision and MARSYAS for sound processing. With these libraries we were able to quickly produce some basic sensor processing. Consequently, we could focus on the development of the emotional algorithm.

OGRE is a 3 dimensional simulated environment which simulates the motion and cameras of the robot in its environment. In addition, participants can be simulated in the environment by projecting the video from a real camera onto the simulated environment. This allows the robot to move closer and further from the person while the real camera can stay stationary. The software is designed so that the drivers for the simulated cameras and robot control can be swapped for the drivers of the real cameras and robot.

The robot can be friend or foe: it can either trust or mistrust the human with whom it interacts. The emotional cues that the robot provides are of two types. The first is a facial expression. The robot has three levels of trust, and three levels of mistrust, and can be indifferent. The second is the distance between the robot and the user. The robot can choose to move closer or further away at three different rates, based on its emotional model.

The sensors we are using are a mono microphone and a web camera. The microphone records at 44.1 kHz. The WebCam has 640 x 480 pixel resolution with 24 bit color and a 10 fps maximum capture rate. The computer that runs the simulation, including sensor and emotional processing, is a 2.33 GHz Intel Core 2 Duo MacBook Pro with 2 gigabytes of RAM.

4.2. Experimental Setup

We conducted an experiment with 16 participants. A 4 x 4 mixed factorial design was used with Factor A as the Emotional Feedback, and Factor B as the Desired Behavior from the robot. Emotional Feedback types were, just motion (control), motion with facial expressions, motion with score, and motion with both facial expression and score. The goal of the participant was to determine what the robot likes given that the robot can detect the motion of their face, and the pitch of their voice.

Participants were randomly divided into 4 groups of 5. Each group received a different type of Emotional Feedback. Each participant interacted with the robot for 3 trials per desired behavior for a total of 12 one minute trials with a user specified break in between. The order of the desired behavior was randomized for each participant, and trials for a desired behavior were grouped together. One example order of trials is 3 low variance, 3 either pitch or motion, 3 high variance, and 3 medium variance trials.

The first hypothesis that we evaluate is whether there is a main effect between the emotional feedback conditions. An analysis of variance shows an F-score of 1.92 and a p-value of 0.127 which leaves us with a failure to reject the null hypothesis that there is no difference between the mean scores of the participants based on the emotional feedback received. Given that there were only 16 participants, there may not be enough power to accept the null hypothesis.

The second hypothesis that we evaluate is whether there is a main effect between the desired behavior conditions. An analysis of variance shows an F-score of 8.86 and a p-value of 0.003 which allows us to reject the null hypothesis. This is further explored below.

4.3. Experiments

Figure 5 shows box plots of the scores in the different behavior conditions. The medium variance condition is clearly the hardest, and both the low variance and either pitch or motion conditions are the two easiest conditions.

Figure 6 shows how the different feedback conditions affect the score in each behavior condition. The medium variance condition is the only condition where there is a big difference. Figure 7 shows the individual differences for each participant in each trial of the medium variance condition separated by the emotional expression group that they were in. In the future, we will use more participants to test whether the reason for the drop can be explained by fatigue.

In addition to our quantitative results, we also performed a survey asking which was the hardest and which was the easiest condition. The survey also asked which type of feedback was the most useful, and what feedback would they like to see. All participants labeled the medium variance trials the hardest. The easier behaviors were also detected. The participants that had facial expression said that it was useful, and the participants that didn't have facial expression asked for it in the survey.

Despite the simplicity of our model, the robot expresses enough to let the user know how it feels, which allows the user to continue or change his behavior based on this feedback.

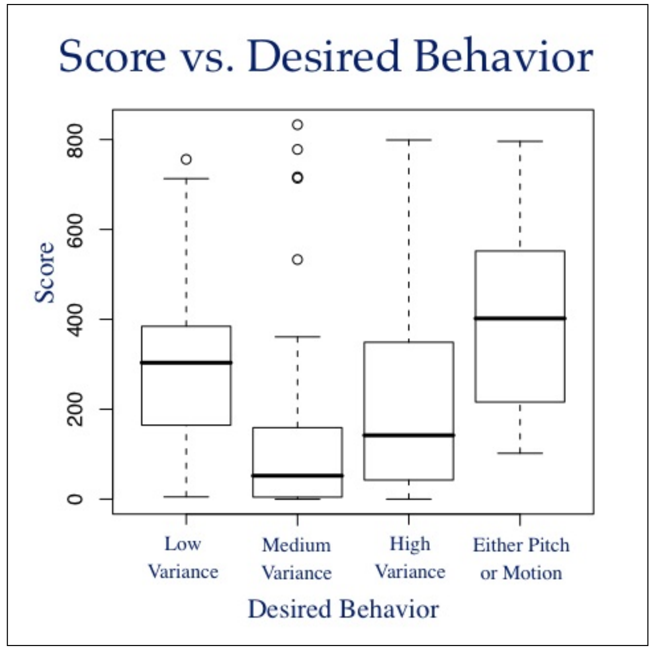

Figure 5. A box plot of the Score over the Desired Behavior factor. An analysis of variance shows an F-score of 8.86 and a p-value of 0.003, hence we reject the null hypothesis that there is no effect between desired behaviors. The box plot shows that the medium variance behavior was the hardest to learn, and both low variance and the either/or behavior were the easiest to learn.

5. Conclusion

The experimental results show that the framework we created has promise as an interactive test bed for emotional communication. The emotions elicited by the robot were clear enough for users to know whether or not their actions were pleasing or displeasing for the robot. The simulated environment also allows for repeated testing without having to utilize fragile robotic hardware.

More advanced sensor information could allow for many more options to make the robot happy. For instance if the robot is could do basic speech processing, then there may be words that the robot likes or dislikes. If the robot could do visual shape processing, then certain shapes may be potential inputs. With music processing, certain songs could effect the robot's emotional state.

Improved expressive output would also contribute to making the robot more believable. For instance, if the robot had arms, it could use them to express emotions by gestures. If it had speakers, it could play different sounds or music. And, with speech software, the robot could say different phrases or change its tone of voice.

We believe that competitive interaction with the robot, such as emotional *Tug of War*, is a useful framework for testing

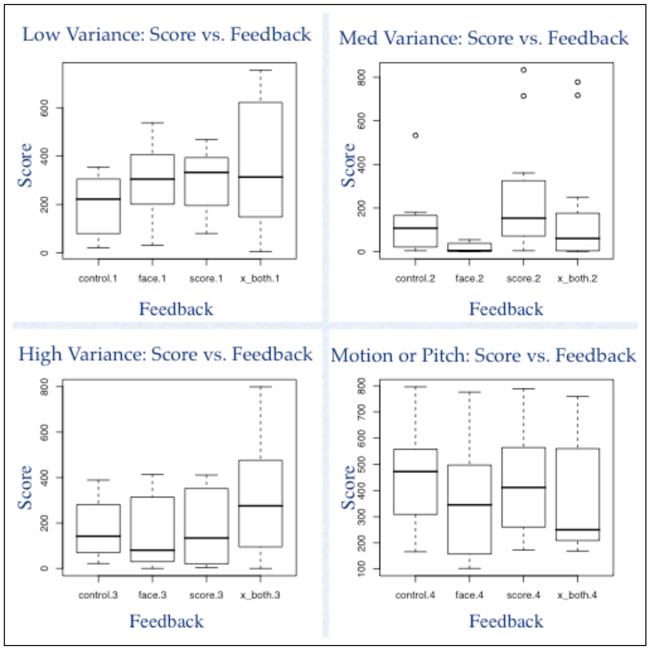

Figure 6. These box plots are a break down of the previous box plot. Each box plot represents a particular desired behavior, and each box represents a specific type of feedback. The feedback types are shown from left to right: just motion, motion with face, motion with score, and motion with both face and score.

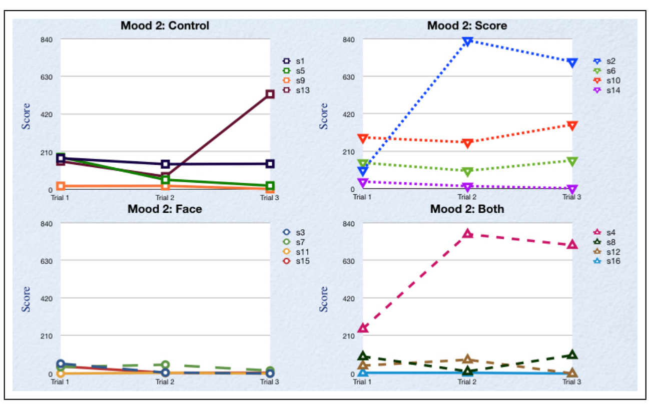

Figure 7. These plots detail the scores of each of the 16 participants for the **Medium Variance** desired behavior. Each plot shows 4 participants' individual scores on the 3 trials for one type of feedback. *top left* just motion, *top right*: score *bot left*: face, *bot right*: face and score.

other emotional interactions. With the addition of more advanced input and output capabilities, we believe that this platform could develop even further and provide a more interesting environment for research.

The next phase would be to extend our implementation to create a real *Tug-of-War* game. This will require using two computers, one associated with each contestant. It will also require addressing issues such as synchronizing behaviors based on multiple users, communication over a wireless medium, and doubling the number of sensors used.

Once the system has been tested in a simulated environment, the final step would be putting it onto an actual robot. There are several robotic platforms that could make use of our emotional software (Azad et al., 2007; Brock et al., 2005; Brooks et al., 2004; Deegan et al., 2006; Edsinger & Kemp, 2006; Katz & Brock, 2007; Khatib et al., 1999; Neo et al., 2006; Nishiwaki et al., 2007; Saxena et al., 2006; Wimboeck et al., 2007). *UMan*, for example, is a potential future platform. *UMan* is a mobile manipulator, a robot that is both mobile and capable of manipulating its environment. More importantly perhaps, *UMan* has the height of a human, and is able to create an impression on people. *UMan* can support multiple sensors, among which are multiple cameras, force sensors, and laser scanners.

Future work would include discovering ways to characterize the set of human emotions that are easy for the robot to perceive. Once this is done, the input space of the robot could be tuned to human emotion rather than arbitrary sounds or motions. One version of the robot could then have an affinity to happy and angry people. Another version could be attracted to sad and scared people, and put off by happy people.

Finally, we intend to add learning into the emotional agent. One notable characteristic of emotional behavior is the ability to adapt to new circumstances. We would like to create similar behavior in our robot. An agent should learn what is annoying for other people, as well as what is not pleasant for itself. Learning how to achieve goals using emotional reaction can be very beneficial. A robot that can take advantage of emotional communication may be able to communicate more efficiently, and change the state of the world in a mutually beneficial way.

Acknowledgments

Thanks are given to the three anonymous reviewers who provided useful comments to improve the clarity of this work.

References

Azad, P., Asfour, T., & Dillmann, R. (2007). Toward an Unified Representation for Imitation of Human Motion on Humanoids. *ICRA*.

Bar-Cohen, Y., & Breazeal, C. (2003). *Biologically inspired intelligent robots*. SPIE Press.

Breazeal, C. (2002). *Designing sociable robots*. The MIT Press.

Breazeal, C., & Scassellati, B. (1999). How to build robots that make friends and influence people. *Intelligent Robots and Systems, 1999. IROS '99. Proceedings. 1999 IEEE/RSJ International Conference on*, *2*, 858–863 vol.2.

Brock, O., Fagg, A., Grupen, R., Platt, R., Rosenstein, M., & Sweeney, J. (2005). A Framework for Learning and Control in Intelligent Humanoid Robots. *International Journal of Humanoid Robotics*, *2*, 301–336.

Brooks, R., Aryananda, L., Edsinger, A., Fitzpatrick, P., Kemp, C., O'Reilly, U.-M., Torres-Jara, E., Varshavskaya, P., & Weber, J. (2004). Sensing and manipulating built-for-human environments. *International Journal of Humanoid Robotics*, *1*, 1–28.

Bruce, A., Nourbakhsh, I., & Simmons, R. (2002). The role of expressiveness and attention in human-robot interaction. *Robotics and Automation, 2002. Proceedings. ICRA '02. IEEE International Conference on*, *4*, 4138–4142 vol.4.

Deegan, P., Thibodeau, B., & Grupen, R. (2006). Designing a Self-Stabilizing Robot For Dynamic Mobile Manipulation. *Robotics: Science and Systems - Workshop on Manipulation for Human Environments*.

Edsinger, A., & Kemp, C. C. (2006). Manipulation in Human Environments. *IEEE/RSJ International Conference on Humanoid Robotics*.

Ekman, P., & Friesen, W. V. (1971). Constants across cultures in the face and emotion. *Journal of Personality and Social Psychology*, *17*, 124–129.

(Ferrell), C. B. (1998). A motivational system for regulating human-robot interaction. *AAAI '98/IAAI '98: Proceedings of the fifteenth national/tenth conference on Artificial intelligence/Innovative applications of artificial intelligence* (pp. 54–62). Menlo Park, CA, USA: American Association for Artificial Intelligence.

Katz, D., & Brock, O. (2007). Interactive Perception: Closing the Gap Between Action and Perception. *ICRA 2007 Workshop: From features to actions - Unifying perspectives in computational and robot vision*.

Khatib, O., Yokoi, K., Brock, O., Chang, K.-S., & Casal, A. (1999). Robots in Human Environments: Basic Autonomous Capabilities. *International Journal of Robotics Research*, *18*, 684–696.

Kringelbach, M. L., & Rolls, E. T. (2003). Neural correlates of rapid reversal learning in a simple model of human social interaction. *NeuroImage*, *20*, 1371–1383.

Neo, E. S., Sakaguchi, T., Yokoi, K., Kawai, Y., & Maruyama, K. (2006). Operating Humanoid Robots in Human Environments. *Workshop on Manipulation for Human Environments, Robotics: Science and Systems*.

Nishiwaki, K., Kuffner, J., Kagami, S., Inaba, M., & Inoue, H. (2007). The experimental humanoid robot H7: a research platform for autonomous behaviour. *Philosophical Transactions of the Royal Society*, *365*, 79–108.

Saxena, A., Driemeyer, J., Kearns, J., & Ng, A. Y. (2006). Robotic Grasping of Novel Objects. *Neural Information Processing Systems*.

Takeuchi, A., & Naito, T. (1995). Situated facial displays: towards social interaction. *CHI '95: Proceedings of the SIGCHI conference on Human factors in computing systems* (pp. 450–455). New York, NY, USA: ACM Press/Addison-Wesley Publishing Co.

Wimboeck, T., Ott, C., & Hirzinger, G. (2007). Impedance Behaviors for Two-Handed Manipulation: Design and Experiments. *ICRA*.