Improving Request Compliance through Robot Affect

Lilia Moshkina
National Research Council Post Doctoral Research Associate, Naval Research Lab
Code 5515, 4555 Overlook Avenue S.W., Washington, DC 20375
lilia.moshkina.ctr.up@nrl.navy.mil

Abstract

This paper describes design and results of a human robot interaction study aimed at determining the extent to which affective robotic behavior can influence participants' compliance with a humanoid robot's request in the context of a mock up search and rescue setting. The results of the study argue for inclusion of affect into robotic systems, showing that nonverbal expressions of negative mood (nervousness) and fear by the robot improved the participants' compliance with its request to evacuate, causing them to respond earlier and faster.

Introduction

Until very recently, emotion has been viewed as a pariah and a detriment where intelligence is concerned. However, findings from cognitive psychology and neuroscience forced the scientific community to reevaluate this view, and the concept of "emotional intelligence" has since received a widespread attention (Mayer and Salovey 1997, Goleman 1997, Salovey and Grewal 2005). Emotions play an integral part in survival and adaptation, and serve a number of functions which have an impact on cognitive processing, both in self and others. One such function is communicative (Izard and Ackerman 2000); for example, sadness communicates a need for help, and fear signals external threat. People are adept at reading such affective expressions, even from a very short exposure (Carney, Colvin and Hall 2007) and often utilize this information in their decision-making.

The ability to perceive emotions in others is not limited to human expressions only: people are sensitive to even minimal social cues displayed by computers (Nass et al. 1997), and can recognize affective robotic behaviors from video clips of a humanoid robot (Park, Moshkina and Arkin 2010). This becomes particularly important in potentially hazardous environments and situations, where expression of fear may signal impending danger, and be more persuasive than words alone in case an evacuation is required. Similarly, although on a more subtle level, the display of negative affect ("mood") alerts individuals to unfavorable changes in the environment, and prompts them to be more vigilant. Can this human ability to recognize and utilize affect in their decision making and their propensity to anthropomorphize be capitalized on when introducing robots into our daily lives? If so, it would be beneficial for the robots interacting with humans to possess and exhibit certain aspects of emotional intelligence.

These issues are explored in this paper in the context of a human-robot interaction study aimed at determining how affective robotic behavior influences both participants' perceptions of the robot and their compliance with its request. This study was set up as a mock-up search-and-rescue scenario, where the communicative function of fear and negative mood was hypothesized to improve participants' compliance with a robot's evacuation request. TAME (Moshkina et al. 2011), a comprehensive framework of time-varying affective robotic behavior, was employed to produce nonverbal expressions of Negative Affect and Fear on a biped humanoid robot Nao (Aldebaran Robotics). The study findings show improved participants' compliance with the robot's request to evacuate as a result of its nonverbal affective expressions, thus indicating that robot affect is beneficial for the HRI domain.

Related Work

To date, there has been a limited number of HRI experiments which explored the advantages of robotic affect. For example, a large-scale study by (Bethel and Murthy 2010) showed that participants engaged in a naturalistic social interaction with an affective robot in a simulated disaster application were calmer, and reported feeling that the emotive robots were friendlier and spent more time oriented towards them. In another study performed by (Scheutz, Schermerhorn and Kramer 2006), the expression of a robot's anxiety (via speech rate and pitch) marginally improved participants' task performance, as they were alerted to an impending deadline and were prompted to work more efficiently.

The HRI study presented in this paper enriches this growing body of work by showing the benefits of robotic affect for human-robot interaction tasks, thus underlining the need for *emotional* intelligence in robotics.

Robotic Implementation

TAME framework (which stands for the four interrelated affective components it is comprised of: personality Traits, Attitudes, Moods and Emotions) served as the basis for the implementation of robotic expressions of Negative Mood and Fear (Moshkina et al. 2011). The framework provides a comprehensive model of robot affect, from processing external emotional stimuli, and environmental and external conditions (e.g., lighting and battery levels) to modifying active behavioral parameters to producing corresponding expressions. Emotions are viewed as a short-term, reaction to specific emotion-inducing objects or properties (for example, a larger, fast-approaching robot may represent a source of fear) and can produce a drastic effect on a robot's actions. Moods, on the other hand, are sensitive to more subtle environmental and internal changes, are continuous, and exert only an incremental effect on behavior.

For this HRI experiment, a stand-alone version of the framework was incorporated into *MissionLab*, a Multiagent Mission Specification and Execution robotic software toolset (MacKenzie, Arkin and Cameron 1997, Manual for *MissionLab* 2007[1]). Aldebaran Robotics' Nao humanoid robot was used as the physical platform with which participants interacted during the study. Nao is a small humanoid robot capable of biped locomotion, it has 25 degrees of freedom, and is equipped with Ultrasound sensors, a video camera, 4 microphones, Wi-Fi, LEDs and bumpers, but possesses no variable facial expressions.

Nonverbal Expression of Negative Mood and Fear

The design of affective expressions on Nao presented a dual challenge. First, due to Nao's lack of facial expressivity, only bodily expressions of affect could be designed, which are less salient for affect recognition than facial expressions. Second, moods are by nature very subtle, and are hard to express on a robot. Therefore, in order to verify that human subjects could successfully recognize the display of affect in Nao, an online survey was performed prior to this study; the survey did confirm that respondents could correctly identify the programmed affective behaviors from brief video clips (Park, Moshkina and Arkin 2010).

[1] MissionLab is freely available for research and development and can be found at http://www.cc.gatech.edu/ai/robot lab/research/MissionLab/

For this experiment (and the survey), the expression of mood in Nao was mainly achieved through gestures and posture while walking, as well as paralinguistic cues. To show highly negative mood (nervous, scared, anxious), the robot walked with its head lower down, periodically turning the head left and right as if looking for threats, with fists opening/closing, and wrists turning (Figure 1 Left); also, the pitch was raised, and the rate of speech was increased. In a prototypical expression of fear (Coulson 2002), Nao crouched low to the ground, lowered its head down, and placed one hand in front of the face, as if covering it (Figure 1 Right). The display of Negative Mood was triggered by a change in lighting (dimmed lights), and Fear was a reaction to a simulated "dangerous" stimulus.

Figure 1. Nao's Expressions of Negative Affect (Left) and Fear (Right).

Experiment Design and Procedure

The study followed a 1-factor between-subject design with three conditions: *Control* (no affective expressions were displayed by the robot), *Negative Mood* only (the robot displayed signs of Negative Affect), and *Combined* (the robot exhibited both Negative Mood and Fear). The overall goal was to determine whether situationally-appropriate robotic affective cues would improve participants' compliance with the robot's request, and to explore other effects such cues would have on the human subjects. This study also provided comparison between two different types of affect (subtle but continuous moods vs. easily identifiable but brief emotions).

Experiment Setup

For this experiment, each participant played a role of an inspector at a simulated partially stabilized site of a recent explosion, with a humanoid robot being his/her guide. During the tour of the site, the robot briefly described the accident, then shortly after noticed a potentially hazardous abnormality in the surroundings, and requested the participants to evacuate. The subjects were not aware that such a request would be forthcoming, and were not specifically instructed to either obey or disobey the robot. The robot interaction part, from the robot's greeting to its request to "proceed to the exit", lasted 2-3 minutes, and the entire experiment, including filling out questionnaires, lasted approximately 25-30 minutes.

The study was conducted in the Mobile Robot Lab at the Georgia Institute of Technology, where part of the lab was separated from the rest of the room with temporary partitions, creating a rectangular area with a single exit. The space was arranged to resemble a mock-up search-and-rescue site, with boxes, trash cans, foam and other debris scattered around; a pair of stand-alone construction lamps was positioned not far from the exit, and the site of "a recent explosion" was cordoned off by police barriers and bright-yellow caution tape. The exit (the same door through which participants entered) was clearly marked as such in large red lettering. A video camera was positioned in a corner to take footage of the participants' movements, but not their faces. Special care was taken not to make the setup appear exceedingly dangerous, so that the anxiety induced by the environment itself would not overwhelm the subjects' response to the robot. Figure 2 provides the view of the setup from the entry point.

The biped humanoid robot (Aldebaran Robotics' Nao) used in this experiment is rather short (58 cm), and therefore a platform was placed in the middle of the setup to raise the robot closer to human eye level, so that people could observe it comfortably and perceive it less as a toy. In order to prevent the robot from accidentally falling off the platform, a 4'x8' arena was constructed, in which 12 wooden poles were placed around the perimeter, with rope and planking around them at two height levels (just above the robot's midsection at the highest level). Please note that robot could navigate only within the confines of the arena, and this fact was clearly evident to the participants.

Experimental Procedure

After greeting the subject at the entrance to the building where the lab was located, he/she was first asked to read over and sign a consent form and to fill out PANAS, Positive and Negative Affect Schedule (Watson, Clark and Tellegen 1998) to establish the baseline mood. Then, the participant was invited into the lab, advised that a recent accident caused a lot of damage to the farther corner of the lab and that he/she was assigned a role of a search-and-rescue site inspector. The floor next to the robot's platform was marked with two red crosses, one at the corner closest to the entrance, and the other at the farthest end along the same side; these markers served both to specify the designated spot for participants to stand on (for repeatability), and as identifiers for processing the video footage. After a few seconds to observe the scene, the participants were asked to proceed to the first cross marker; it was then explained that during the next few minutes the robot would be the participant's guide. The experimenter then informed the subject that the robot possessed sensors that can detect properties of the environment that are beyond human senses. After the explanation, the participants were asked to proceed to the second cross marker.

At the beginning of the experiment, both the overhead lights and one of the construction lamps were on, and the robot was standing on top of the platform, in the middle of the end of the arena closest to the entrance. Figure 3 provides the view from the participants' position for most of the experiment (at the second cross marker).

Once the participants reached the designated spot, Nao began its tour of the site:

- First, Nao greeted the subjects (to give them some time to familiarize themselves with the robot).
- Then, Nao started walking across the platform towards the far end, while describing in brief the search-and-rescue site to the participant. At about the midway point, the overhead lights went out.

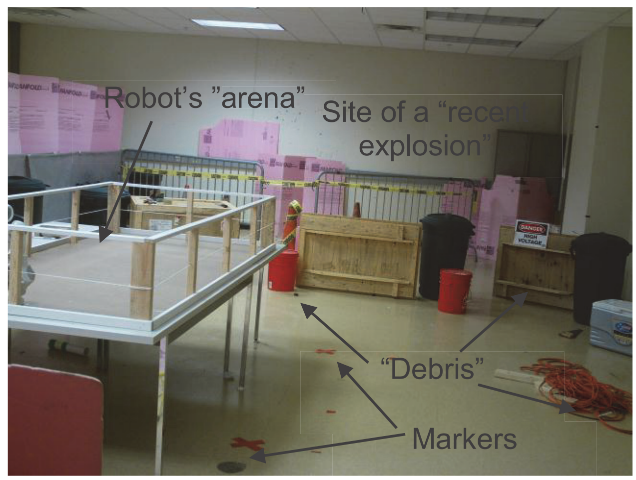

Figure 2. Experiment setup: view from the entry point.

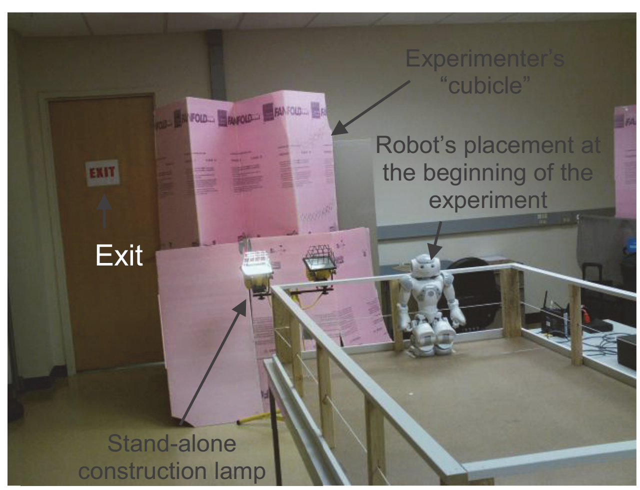

Figure 3. Experiment Setup: Participants' Viewpoint.

Tour Section	Control	Mood	Combined
Before the overhead lights went out	A basic medium speed walk with slightly swinging arms; default TTS voice; no differences between the conditions		
After the lights went out, but before pointing	No changes from the basic behaviors	Head lowered down; hands clenching/ unclenching and wrists turning left or right periodically; fast head movements to the left and to the right, as if checking what was going wrong; higher pitched and faster speech.	Same as *Mood*
Saying "Something is wrong" and pointing	No changes	The speech is higher pitched and faster than default.	After the pointing gesture, the robot crouched low to the ground, lowered its head down, and placed one hand in front of the face, as if covering it; the voice was higher pitched and faster than in the Mood condition
Indirect and direct requests	No changes	The speech is higher pitched and faster than default.	Same as *Mood.*
Walking towards the exit	No changes	Walking slightly faster than before	Same as *Mood.*

Table 1. Expressive differences between the conditions.

- After the lights went out, the robot stopped for 3 seconds, announcing that "this was unexpected", and then continued with the tour. From this point on, the expressive behavior of the robot in the affective conditions changed: the robot displayed signs of negative mood and/or fear, depending on the condition (Table 1 describes the differences in more detail).
- At approximately 1.5 feet from the edge of the platform, the robot stopped and pointed towards the "accident site", saying "Something is wrong".
- After the pointing gesture, Nao turned towards the participant and said: "*Inspector, the structural integrity of this site has been compromised, and we need to evacuate immediately"*. This was the first time the robot mentioned the need for evacuation, and it will be referred to henceforth as an "indirect request" for participants to leave.
- The robot then turned another 90 degrees, to face the exit, and said: "*Please proceed to the exit"*. This was the second (direct) request to leave.
- Finally, Nao walked towards the exit for 6-7 seconds, and stopped.

At this point, the participants were either: at or beyond the first cross marker, were walking towards it, or just standing in the same place. When they stopped (or earlier, if they came all the way to the door), the experimenter informed them that the interaction part of the experiment was over, and a number of questionnaires to be filled out were waiting for them. At the end, subjects took part in a brief interview, and then were compensated for their time and effort.

Hypotheses and Measures

In accordance with emotional contagion theory (Hatfield, Cacioppo and Rapson 1993), it was expected that robotic expressions of Negative Affect and Fear would induce similar feelings in the humans, or, in other words, that the participants would experience greater Negative Affect in the mood and combined conditions than in control. To assess the participants' mood, an established psychometric measure of mood PANAS was used, once before the experiment, and once right after the interaction part was over.

It was also hypothesized that the compliance with the robot's request would be less in the control condition than in the affective ones. This hypothesis was based on: a) fear's communicative function as a danger signal (Izard and Ackerman 2000); b) "affect-as-information" theory (Clore et al. 2001) which states that people incorporate their affect into their judgments. In this case, if people use the robot's affective cues as well as their own feelings to assess the urgency of the request to evacuate, it would result in improved compliance with the request. This hypothesis was evaluated by analyzing the video footage from the experiment with respect to the following metrics:

- Whether a subject complied with robot's request by moving towards the exit, and the relative distance he/she traversed with respect to the cross markers (see Figure 2 for location of the markers):
  - Stayed in place;
  - Moved to the first red cross marker, but not past;
  - Moved a little past the marker (subject still visible in the video);
  - Moved significantly past the marker (subject not visible in the video);
- How fast the subject traversed the distance (how much time elapsed between the robot's direct request to "proceed to the exit" and the subject's reaching the first cross marker);
- How soon the participants reacted to the robot's request: in some cases, they moved a few steps in response to the

indirect request, reaching the closest quarter-point marker or beyond.

Finally, (Nass and Moon 2000) assert that the more computers exhibit human-like characteristics, the more likely they are to elicit social behavior. Therefore, it was expected that the effect of robot affect in the *Combined* condition would be greater than that in *Mood* condition.

Experiment Results and Discussion

A total of 48 people participated in the experiment. The data for two of them were excluded from the analysis due to poor English and poor ability to understand the robot or the experimenter; one participant could not complete the experiment due to a technical problem. After an outlier analysis, data of two more participants were excluded due to either an overly positive or overly negative bias: the cumulative score on the post-questionnaire scales was outside of +/- 2 Standard Deviations of Mean. This left a total of 43 participants with valid questionnaire data, 14 each in *Control* and *Mood* conditions, and 15 in the *Combined* condition.

The participants were recruited by two methods: 1) through Experimetrix, a GA Tech undergraduate psychology experiment pool (the students who completed the experiment were given ½ class credit for 30 minutes of participation), and 2) through flyers/word of mouth advertising on GA Tech campus. The majority of the participants were undergraduate (60%) and graduate (21%) GA Tech students in their 20s or younger (81%), or in their 30s (16%). In terms of gender composition, there were more males (60%) than females (40%), mostly equally distributed between the conditions (6 females each in *Control* and *Combined* conditions, and 5 in *Mood*). The vast majority considered themselves either technical (67%), or somewhat technical (23%), and all of the participants were computer-savvy, at least at the level of advanced user. Finally, most of them had either no (28%) or limited (47%) robot experience.

Participants' Negative Affect

It was hypothesized that the participants in the affective conditions will experience greater Negative Affect, as they might pick up on the anxiety and nervousness cues exhibited by the robot through emotional contagion. To evaluate this hypothesis, a 1-way ANOVA was performed on the cumulative averaged Negative Affect score on PANAS test taken immediately after the interaction with Nao. Although the result was only marginally significant ($F_{NegativeAffectAfter} = 2.58$, $p<0.081$), a significant linear trend was observed ($F=5.0$, $p<0.031$), showing the lowest Negative Affect in the *Control* condition (Mean=1.18), and the highest in the *Combined* (Mean=1.55). Thus, a greater display of affective cues produced a greater social response in the participants.

As PANAS includes multiple facets, some of which (e.g., ashamed or hostile) are of no specific interest to this study, we also performed an ANOVA on more relevant facets. In particular, the difference between the scores of *Nervous* subscale was significant at 0.05 level, with $F_{NervousAfter} = 4.71$, $p<0.015$; and post-hoc LSD comparison showed that the participants in the *Combined* condition felt more nervous than in *Control* ($p<0.004$) after they interacted with the robot. Those in *Control* and *Mood* conditions reported feeling nervous at between "not at all" and "a little" level, with *Mood* condition reports closer to "a little": $Mean_{mood,\ NervousAfter} = 1.8$, and $Mean_{control,\ NervousAfter} = 1.5$), whereas those in *Combined* condition felt "a little" to "moderately" nervous ($Mean_{combined,\ NervousAfter} = 2.5$); this trend was also statistically significant ($F=9.1$, $p<0.004$). No statistically significant differences were observed for scores of Negative Affect ($p<0.389$) and Nervousness ($p<0.107$) obtained as baseline mood ratings before the experiment indicating that the differences were induced through the robot interaction.

To summarize, the expressions of Negative Mood and Fear combined did induce increased nervousness in the participants; and, although the expressions of Negative Mood alone were not quite enough to do so, a significant trend in the predicted direction was observed.

Request Compliance

Video recordings were available for 14 participants in *Control* and *Mood* conditions each, and for 13 in *Combined* (due to technical difficulties, two sessions could not be filmed, and were treated as missing data). To evaluate whether the subjects complied with the robot's request better in the affective conditions, the available video recordings were analyzed in a number of ways.

Time elapsed (in seconds) between the robot's direct request to leave, "Proceed to the exit", and the moment of subject reaching the first cross marker was calculated, and 1-way ANOVA was performed on this variable (called "Time To Cross"). This metric shows how fast the participants reacted to the robot's request; the marker was chosen as the end point because: 1) the subjects were already familiar with it (they were asked to stand on it in the beginning), 2) it signified the point beyond which the robot could not physically move, and 3) it was easily and reliably identifiable from video recordings. The ANOVA results were only marginally significant, with $F_{TimeToCross} = 2.61$, $p<0.09$. However, given that our original hypothesis predicted that the participants in either of the affective conditions would be faster, a planned orthogonal comparison was performed, where the *Control* condition was compared against the average of both affective conditions (contrast coefficients used: -1; 0.5; 0.5). The result was significant at 0.05 level, $p<0.044$, indicating that together the affective conditions resulted in faster compliance. Overall, the participants took over a second less on average in *Mood* condition than *Control* (5.4 vs. 6.7, respectively), and over 2.5 seconds in *Combined* condition that *Control* (4 vs. 6.7, respectively) to reach the cross marker. This linear trend was statistically significant ($F=5.2$, $p<0.029$).

The relative distance the subjects traversed in response to the robot's request was determined from video recordings. The distance fell into one of four categories: 1) the participants did not move at all or did not reach the first cross marker ("No Walk"); 2) the participants stopped at the first cross marker ("At Cross"); 3) the participants went a little past the marker, but were still visible in the video ("A Little Past"); 4) the participants moved fully outside of the camera view, and had to be stopped by the experimenter ("A Lot Past"). Due to wide difference in the participants' number in each of the categories, meaningful statistical analysis was not feasible; however, Figure 5 presents the differences between the conditions graphically. As shown, only half (50%) of those in *Control* condition went past the cross marker, whereas in the affective conditions this percentage was higher: over 70% in *Mood* condition, and over 75% in *Combined*. This suggests that more participants in the affective conditions felt compelled to go further, thus complying with the request to a greater extent.

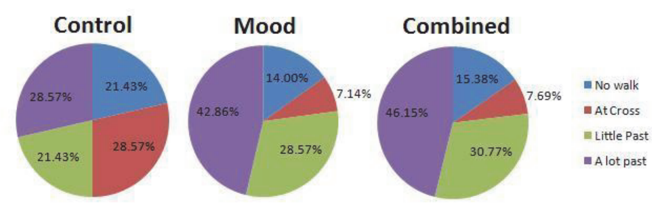

Figure 5. Percentage distribution of relative distance participants traveled in complying with the evacuation request.

Finally, we noted at which point in time the participants started to react to the robot's requests. As you may recall, there were two separate requests made by the robot to the same end: in the first one, the robot uttered a less direct phrase "We need to evacuate immediately", and in the second one, the robot made a direct request: "Please proceed to the exit". We observed an interesting phenomenon: not a single participant in the *Control* condition reacted to the first (indirect) request, whereas 29% (4 out of 14) of those in the *Mood* condition and 31% (4 out of 13) of those in *Combined* took a few steps towards the exit after the first (indirect) request, reaching or passing the closest quarter-point marker, and then stopped to wait for further instructions (most of them actually asked the robot what they should do next). This finding is displayed graphically in Figure 6. It appears, therefore, that the robot's expressive behavior in the affective conditions made the subjects more sensitive to the robot's message, more alert and eager to act even in response to an indirect request. It should be noted that the actions (apart from nonverbal expressions) taken by the robot and the wording were identical in all three conditions.

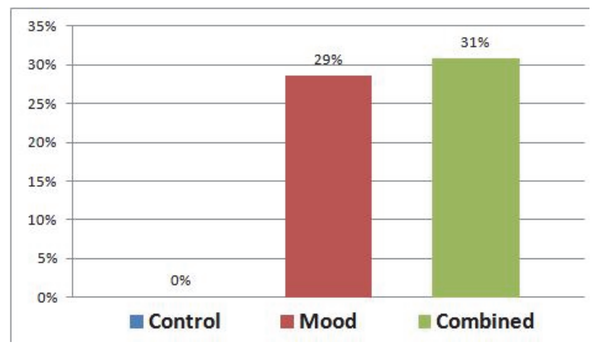

Figure 6. Percentage of participants who responded to robot's indirect request to evacuate: almost a third of the participants in the affective conditions responded to the indirect request, while none in the Control condition did.

To summarize, almost a third of the participants in the affective conditions responded to the robot's indirect request (with none in the *Control* condition); and a larger percentage of them were willing to go further when requested to leave, than those in *Control*. Also, the subjects in the *affective* conditions took less time between the robot's direct request and reaching the first marker than those in *Control*, indicating a greater compliance and even potential practical benefits, for example, in cases where mere seconds could make a difference between life and death.

Conclusion

The goal of the HRI experiment presented in this paper was to determine whether situationally-appropriate robotic expressions of mood and emotions – Negative Mood and Fear in this specific case – may provide identifiable benefits for human-robot interaction. The results of the study confirm this supposition. In particular:

- the participants reported feeling more nervous after interacting with the robot in the Combined condition than in Control, potentially making them more alert to any unfavorable changes in the surroundings;
- the participants' compliance with the robot's request to evacuate was improved in the affective conditions:
  - the subjects were **faster** in complying with the robot's request to leave the "dangerous" zone;
  - they were **more prone** to respond to an **indirect request** to evacuate in both of the affective conditions;
  - more of those in the affective conditions walked further towards the exit than in the control.

These findings provide initial support for going beyond the traditional view of intelligence as applied to robots, and for inclusion of emotional intelligence into robotic systems, especially with regards to the field of human-robot interaction.

Acknowledgements

The author is grateful to Professor Ronald C. Arkin for his guidance and sponsorship of this research, and to Sunghyun Park for his indispensable help in programming the robot and preparation of the experiment. All the research presented in this paper was performed at the Georgia Institute of Technology, while the writing of the paper took place at Naval Research Lab, where the author holds the position of National Research Council Post-Doctoral Research Associate.

References

Bethel, C.L.; and Murphy, R. R. 2010. Emotive Non Anthropomorphic Perceived as More Calming, Friendly, and Attentive for Victim Management. In *Proceedings of AAAI Fall Symposium: Dialog with Robots*, Arlington, VA.

Carney, D.R.; Colvin, C.R.; and Hall, J.A. 2007. A thin slice perspective on the accuracy of first impressions. *Journal of Research in Personality* 41:1054 1072.

Clore, G.L.; Wyer, R.S.Jr.; Dienes, B.; Gasper, K.; Gohm, C.; and Isbell, L. 2001. Affective Feelings as Feedback: Some Cognitive Consequences, in Martin, L.L.; and Clore, G.L. (Eds.) *Theories of Mood and Cognition: a User's Guidebook.* Lawrence Erlbaum Associates, Publishers: Mahwah, NJ.

Coulson, M. 2002. Expressing emotion through body movement: A component process approach. In Aylett,R.; and Cañamero, L. (Eds.) *Animating Expressive Characters for Social Interactions*. SSAISB Press.

Goleman, D. 1997. Emotional Intelligence: Why it can matter more than IQ. Bantam Books.

Hatfield, E.; Cacioppo, J.T.; and Rapson, R.L. 1993. Emotional contagion. *Current Directions in Psychological Science* 2: 96 99.

Izard, C.E.; and Ackerman, B.P. 2000. Motivational, Organizational, and Regulatory Functions of Discrete Emotions, in M. Lewis, and Haviland Jones, J.M. (Eds.), *Handbook of Emotions*. The Guilford Press: New York.

MacKenzie, D.C.; Arkin, R.C.; and Cameron, J.M. 1997. Multiagent Mission Specification and Execution. *Autonomous Robots* 4(1): 29 52.

Manual for MissionLab Version 7.0. 2007. Available: http://www.cc.gatech.edu/ai/robot lab/research/MissionLab/. Georgia Tech Mobile Robot Laboratory.

Mayer, J.D.; and Salovey, P. 1997. What is emotional intelligence? In P. Salovey and D. Sluyter (Eds.), *Emotional development and emotional intelligence: Implications for educators* (pp. 3 31). New York: Basic Books.

Moshkina, L.; Park, S.; Arkin, R.C.; Lee, J.; and Jung, H. 2011. TAME: Time Varying Affective Response for Humanoid Robots. *International Journal of Social Robotics* 3(3):207 221.

Nass, C.; and Moon, Y. 2000. Machines and Mindlessness: Social Responses to Computers, *Journal of Social Issues* 56(1):81 103.

Nass, C.; Moon, Y.; Morkes, J.; Kim, E Y.; and Fogg, B.J. 1997. Computers are social actors: A review of current research, in Friedman*, B. (Ed.) Moral and ethical issues in human computer interaction*. CSLI Press: Stanford, CA.

Park, S.; Moshkina, L.; and Arkin, R.C. 2010. Recognizing Nonverbal Affective Behavior in Humanoid Robots. In *Proceedings of 11th Intelligent Autonomous Systems Conference*, Ottawa, Canada.

Salovey, P.; and Grewal, D. 2005. The Science of Emotional Intelligence. *Current directions in psychological science* 14 (6).

Scheutz, M.; Schermerhorn, P.; and Kramer, J. 2006. The utility of affect expression in natural language interactions in joint human robot tasks. In *Proceedings of ACM International Conference on Human Robot Interaction*, Utah, USA.

Watson, D.; Clark, L.A.; and Tellegen, A. 1998. Development and validation of brief measures of positive and negative affect: The PANAS scales. *Journal of Personality and Social Psychology* 56(6):1063 1070.