Modeling Affect in Socially Interactive Robots

Rachel Gockley, Reid Simmons, and Jodi Forlizzi

***Abstract*— Humans use expressions of emotion in a very social manner, to convey messages such as "I'm happy to see you" or "I want to be comforted," and people's long-term relationships depend heavily on shared emotional experiences. We believe that for robots to interact naturally with humans in social situations they should also be able to express emotions in both short-term and long-term relationships. To this end, we have developed an affective model for social robots. This generative model attempts to create natural, human-like affect and includes distinctions between immediate emotional responses, the overall mood of the robot, and long-term attitudes toward each visitor to the robot. This paper presents the general affect model as well as particular details of our implementation of the model on one robot, the Roboceptionist.**

I. Introduction

As robots become more prevalent in social spaces, such as healthcare institutions and museums, people need to be able to interact with these robots in a smooth, natural way. We believe that one way to improve these interactions is to have robots display changing moods and emotions, just as humans do. This paper describes a generative model of affect that attempts to strongly mimic how people emote in order to produce as natural-seeming a system as possible. The model is designed particularly for robots that interact with people over long periods of time. As such, our focus is on modeling the long term aspects of, and interactions between, emotions, moods, and attitudes, rather than on how emotions should be triggered or displayed. We have implemented our affective model on the Roboceptionist (Fig. 1), a robot that interacts with people on a daily basis.

A. Human interaction

Emotions play a major part in human interaction. Quite often, emotional reactions are caused by social interactions, influenced by societal and cultural norms, or used to communicate desires to other people [1]. Emotions carry conversational content, allowing conversational partners to form common ground and communicate more effectively [2]. For instance, one might feel upset after a conflict with another or respond to a friend's sad expression with comfort. Furthermore, what mood a person is in has a strong impact on how that person interacts with others [3]; for example, people who are interacting may "catch" each other's moods and emotions, unconsciously matching their own emotional states to their conversational partners' [4]. Frijda argues that the primary reason for social interaction is, in fact, to experience emotions, which serve to form a "sense of coherence with others" [5]. Suppression of emotions can be highly detrimental to relationship forming and is disruptive to conversations [6].

Fig. 1. The Roboceptionist, displaying a sad expression after someone typed an insult.

A well-studied effect of human–computer interaction is that people tend to treat talking computers in the same way that they treat other people, forming social relationships with them [7]. We believe that this tendency to form social relationships with computers will also apply to robots, perhaps even more so. If that is the case, then people will respond to a robot's emotions as though the robot were human, and will expect that the robot's emotional responses will be consistent across multiple interactions.

B. Human–robot social interaction

In recent years, the robotics community has seen a gradual increase in social robots, that is, robots that exist primarily to interact with people. Museum tour-guide robots [8] and robots that interact with the elderly [9] demonstrate not only the benefits of having robots interact with people, but also the need for the interactions to be smooth and natural. Many of these robots have incorporated at least some rudimentary emotional behaviors, but such behaviors are usually ad-hoc and not generalizable to other robots. Robots with infant-like abilities of interaction, such as Kismet [10], have been used to demonstrate the ability of people to interpret and react appropriately to a robot's displays of emotions. Experiments with the robot Vikia [11] demonstrated the effectiveness of an emotionally expressive graphical face for encouraging

This work was funded in part by an NSF Graduate Research Fellowship to the first author and by NSF grants #IIS-0329014 and #IIS-0121426.

R. Gockley and R. Simmons are with the Robotics Institute, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213, USA {rachelg, reids}@cs.cmu.edu

J. Forlizzi is with the Human Computer Interaction Institute and the Design Department, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213, USA forlizzi@cs.cmu.edu

interactions with a robot. Furthermore, experiments with the Roboceptionist [12] showed that people interact differently with a robot depending on its apparent mood. All of these, however, have focused on short-term interactions with the robot. We believe that a richer model of affect is necessary for forming long-term human–robot relationships.

The exploration of affect in social robots currently lags behind similar research in software agents. For example, the Affective Reasoner [13] is an implementation of a virtual world populated by software agents, wherein the agents can detect and react to each other's emotions; however, the agents do not interact directly with humans. In contrast, Embodied Conversational Agents (ECAs) are designed explicitly to mimic human–human interaction [14]. Many ECAs are capable of expressing emotions (see [15] for an overview). One such system, called FearNot!, has been shown to invoke empathy in its users in response to the agents' emotional displays [16]. However, the emotional models used in these types of systems are typically ad-hoc, incomplete, or poorly documented, and thus difficult to extend to other systems.

The most developed emotional model for robots we are aware of at this time is the TAME architecture [17], which considers the four affective categories of personality traits, attitudes, moods, and emotions. Our model follows a similar breakdown of categories (namely, attitudes, moods, and emotions); however, the TAME model is designed primarily for behavior-based robotic systems and does not necessarily map directly to robots that interact socially with humans. Additionally, the TAME model does not currently specify how these phenomena influence human–robot interaction, nor does it explain how any of the affective phenomena interact with each other. In contrast, we have tried to provide a complete model with a strong psychological backing, designed specifically for long-term human–robot interaction.

II. Affect

Before presenting our affective model for robots, we provide a brief description of affect and how it relates to human interaction. **Affect** is a general term relating to emotions, moods, and other such states with varying degrees of positivity or negativity—that is, states with *valence*. While a great deal of psychological research has focused on affect, only a few researchers distinguish between terms such as "emotion" and "mood," and very few agree on the meanings of these (and similar) terms. Following the categorization suggested by Scherer [18], we consider the following phenomena to be of importance in developing an affective model: emotions, moods, and attitudes.

A. Emotion

An **emotion**, or "emotional response," is an immediate affective response to the evaluation of some event (or other stimuli) as being of major significance. Emotions are often assumed to have distinct facial responses [19], though some argue that this is not always the case [20]. For our purposes, however, emotions are meaningless unless they result in some outward change in the robot, including facial, vocal, or behavioral modifications.

In the psychological and cognitive science literature, there are two primary views on the representation of emotions: categorical and continuous. Ekman [21] and others argue for a set of "basic" emotions that are innate and universal across cultures. All other emotional categories are then built up from combinations of these basic emotions. Others, such as Russell [22], argue that all emotions lie in a continuous two-dimensional space, where the dimensions are typically taken to be *valence* (how positive or negative the emotion is) and *arousal* (the energy or excitation level associated with the emotion). Both representations have been used in various computing and robotic applications; for example, Sage [8] uses a categorical model, whereas Kismet [10] uses a continuous model. In both cases, people identified the robots' emotions and reacted appropriately; the underlying representation does not seem to have a major effect on people's understandings of a robot's emotions.

One of the commonly used emotional models in computing research is known as the OCC model [23]. In the OCC model, emotional events are evaluated as either consequences of events, actions of agents, or aspects of objects. The model provides a straightforward computational mapping of evaluations to emotions. It has been successfully applied (typically with some modifications) to both robotic and simulated behavior [13]. However, the OCC model does not specify salient issues such as how to compute the intensities of emotions or what happens if several emotions occur in rapid sequence [24]. Thus, while the OCC model may be suitable for rudimentary emotional responses, it does not provide the level of sophistication that we believe social robots need.

Exactly what events cause emotional responses, and the intensity of those responses, vary according to a person's personality traits. While traits may influence emotion processing [25], they are not affective states themselves, and so we have not directly incorporated them into the model.

B. Mood

Moods are more "diffuse" affective states [18] that typically do not have a single antecedent. They are typically of lower intensity than emotions and have fairly low variance over the course of a single day. Moods may be caused by any number of things, including changes in physiological state (such as lack of sleep or illness) [26], rapidly occurring emotional responses [26], or complex cognition regarding emotional life events [27].

Several studies have indicated that positive moods tend to reduce negative emotions (in frequency and intensity), and negative moods tend to reduce positive emotions [26], [28]. Additionally, evidence from psychology indicates that daily moods tend to have plateaus, during which repeated emotional events will not cause a large shift in mood [29]. In particular, mood due to life stresses is fairly stable; weak emotional reactions do not immediately cause a noticeable change in mood. However, many similar emotions in a row will eventually alter the mood significantly.

Interacting with other people can have a strong impact on one's mood, but the exact effects of such social interactions are still relatively unknown. Some evidence indicates an asymmetrical crossover model, wherein positive social exchanges will increase a positive mood but have little effect on a negative mood, while negative social exchanges will erode any mood [30].

C. Attitude

An **attitude** is an amalgamation of emotions experienced with a particular person (or thing), reflecting one's relationship with that person over time. Attitudes may change due to emotions inspired by the focus of the attitude, co-experienced emotional events, and the frequency and duration of the relationship. A person's attitude toward a conversational partner can influence many aspects of the conversation; for example, people are less likely to express emotions in the presence of strangers [31]. Attitudes are a key part of long-term relationships.

Emotions, moods, and attitudes all interact. In particular, both emotions and attitudes help to shape the overall mood, and at the same time, the mood affects how strongly emotional reactions occur. Additionally, changes in mood during social interactions result in changes of attitude. The exact interactions used in our affective model are described below.

III. Implementation

As discussed, many models of human affect have been postulated, and which model is "correct" remains undetermined. However, we require only that our robot's behavior is human-like; the underlying model need not precisely match human cognition. That is, we wish to design a generative model, rather than an explanatory one. Accordingly, we have taken the approach of selecting a simple, straightforward model of human affect to serve as the basis for our computational model.

Our current platform for social robotics research is the Roboceptionist [32]. The mechanical base of the robot is an RWI B21r with an LCD "head" that can rotate. The robot is housed in a custom-built booth near a high-traffic entrance to a computer science building at Carnegie Mellon University, and is available for use approximately 8 hours a day, 5 days a week. A keyboard and small monitor on the booth's desk allow for human input with visual feedback; speech recognition is not used do to high environmental noise. The LCD head displays a highly expressive, graphical face, as shown in Fig. 1. Because the face is so expressive, it can display a wide range of easily recognizable emotional expressions. In a collaboration with the Drama Department, the robot has been given a background and a complex, evolving storyline. Visitors to the booth can ask the robot questions about its life in order to hear the continuing story. For example, the major plot arc of the 2004-2005 academic year was that the robot, named Valerie, fell in love with a jukebox named Cal, but they were unable to marry legally because they were not human. These stories are intended to provide interest and to encourage the formation of long-term relationships with the robot.

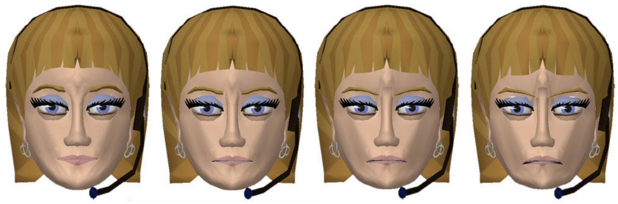

Fig. 2. Interpolation of emotional expressions between neutral (far left) and very angry (far right).

A. Emotions

We chose to implement a categorical model of emotions, primarily due to the relative ease of implementing several distinct emotional expressions versus defining a mapping between continuous multidimensional space and the robot's expressions. In particular, we implemented a subset of the "basic" emotions suggested by Ekman [21]. These emotions are: joy (happiness), sadness, disgust (frustration), and anger. While this list is in no way comprehensive, additional emotions can be added easily as desired. Each emotion has an associated intensity level, represented as a real number ranging from 0 (non-existent) to 1 (highest intensity), as well as a valence rating (positive or negative). For each emotion, we defined a series of expressions of differing intensities, which can be displayed on the robot's graphical face. Specific intensities are generated by a linear interpolation of the muscle positions between the two nearest defined intensity expressions (see example in Fig. 2). The emotional expressions used were based on Delsarte's code of facial expressions as implemented for the robot Vikia [11].

Since this model is intended for social robots, emotions are caused primarily by interactions that the robot has with people. For example, a new person interacting with the robot may cause happiness, and insults typed to the robot may generate sadness or anger. Currently, we have implemented a mechanism within the robot's language model in order to trigger specific emotions directly. That is, certain statements by visitors carry specific emotional content, such as compliments causing happiness or insults resulting in sadness.

Emotions are displayed immediately after an event, and last the duration of the robot's verbal response. In this way, emotions are short-lived, but are displayed long enough to be recognized. Emotions do not occur concurrently in our model, thus avoiding the question of how different base emotions might interact or "blend."

B. Moods

Our robot's moods are primarily caused by its personal history and "life" events. That is, because the robot has an ongoing life story, it can feel positive (or negative) about past (or future) events. Values for the moods are assigned to the storyline by the dramatic writers, as they see fit. The mood generated from these events is considered as the robot's "baseline" mood for each day. The robot's overall mood is influenced by the emotions it experiences throughout the day, as explained below. In different robots, moods may also be influenced by other circumstances, such as internal power levels or the ability to complete assigned tasks. Our robot's moods are indicated by posture, particularly the tilt of the head (e.g. a downcast face indicates a negative mood). A more sophisticated indication of mood might also include vocal and behavioral modifications.

After a life event has occurred, that event's contribution to the robot's mood fades over time. Events that cause stronger moods take longer to fade than lesser events. Similarly, anticipated events will increase in salience as the date of the event approaches, causing an increase in the mood contribution in the days preceding the event. Specifically, if the event will occur in d days and will cause a mood of strength m on the day it occurs, then that event's contribution c to the mood is computed according to the sigmoid curve given below:

$$c = m - \frac{m}{1 + e^{-|d| + |10m|}} \tag{1}$$

Events that have already occurred follow the same curve. Some examples can be seen in Fig. 3. Multiple events' contributions to the mood are additive, as shown in Fig. 4.

In our model, the robot's mood has one primary plateau [29], at the baseline mood as computed for each day. As discussed in the next sections, various events throughout the day may create an offset from the baseline, causing the mood to change dynamically. We constrain the mood to change smoothly and to approach the extremes of the facial expressions asymptotically, so that the model will never attempt to generate an expression outside the normal bounds of the facial muscles. To model this effect, the robot's displayed mood follows a generalized logistic (growth) curve, as given in the following equations:

$$m_d = \begin{cases} \frac{1+m_b}{1+e^{-B(m_o - M(1+m_b))}} - 1 & \text{if } m_o < 0 \\ m_b & \text{if } m_o = 0 \\ \frac{1-m_b}{1+e^{-B(m_o - M(1-m_b))}} + m_b & \text{if } m_o > 0 \end{cases} \tag{2}$$

where m_d is the displayed mood, m_o is the offset from the baseline mood, and m_b is the baseline mood. These equations follow a logistic curve with a growth rate of B; this influences how quickly the mood changes between the plateau and each extreme. Our current implementation uses $B = 15$, but this value can be modified as appropriate for specific robots. The time of maximum growth of each half-curve (that is, the inflection points) occurs at $M(1 \pm m_b)$, where in our implementation $M = 0.35$. This point is scaled according to the baseline to keep the rate of change of the mood similar across different baselines. These equations may be better understood by viewing some examples for different baseline moods, shown in Fig. 5. While the baseline mood and offset may be any real numbers, the curves given in Eq. 2 constrain the displayed mood to a range of -1 (extremely negative) to $+1$ (extremely positive), with 0 representing a neutral mood.

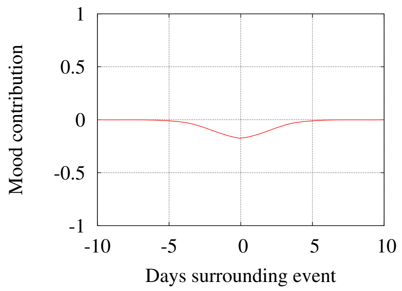

(a) Impending dentist visit ($m = -0.2$)

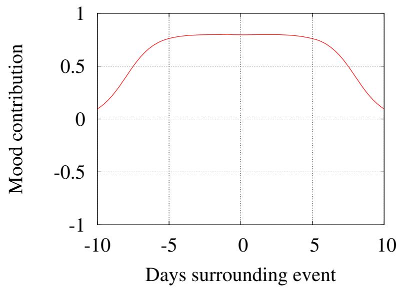

(b) Visit from a close friend ($m = 0.8$)

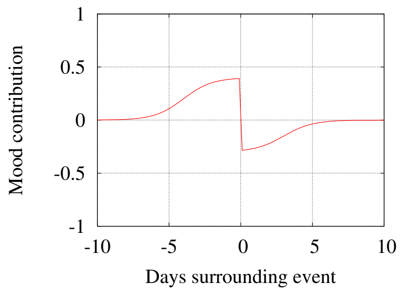

(c) Blind date that went badly (anticipated $m = 0.4$, actual $m = -0.3$)

Fig. 3. Examples of different events and how they contribute to the overall mood over several days. Each event occurs at day 0.

C. *Interaction between mood and emotion*

To model human behavior, the robot's mood level should modulate the intensities of its emotions [26], [28]. In our implementation, we model this effect by scaling the strength of the emotional response linearly, according to the current mood. Mood-congruent emotions (that is, emotions that have the same valence as the mood) are increased in intensity, while incongruous emotions are diminished. Specifically, if an emotion of nominal strength s ($0 \leq s \leq 1$) and valence v ($v = \pm 1$) occurs during a mood of m, then the emotion's strength is scaled according to Eq. 3:

$$s' = s\left(1 + \frac{1}{4}vm\right) \tag{3}$$

The exact scaling may be moderated by a particular robot's personality traits [25].

In addition to the mood influencing emotional responses, emotions in turn influence the mood. In our implementation, the robot's mood dynamically changes due to its interactions with people. For example, the robot may be in a negative

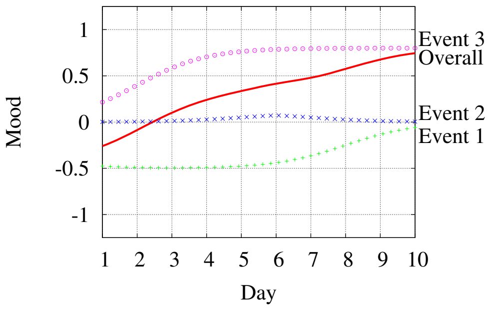

Fig. 4. Three hypothetical events and their contributions to the overall mood across ten days. Event 1, scheduled minor surgery, occurs on day 3 with contribution -0.5, event 2, an evening with friends, occurs on day 6 with contribution 0.1, and event 3, the start of a vacation, occurs on day 10 with contribution 0.8.

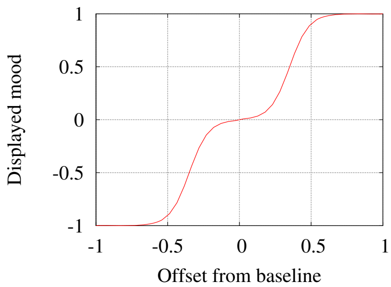

(a) Baseline = 0 (neutral)

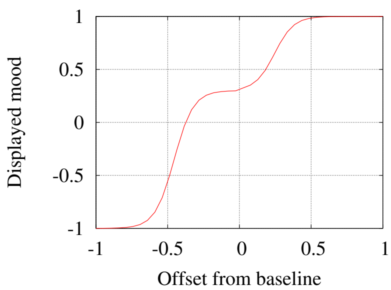

(b) Baseline = 0.3 (somewhat positive)

Fig. 5. How the displayed mood changes based on the underlying mood level, expressed as an offset from the "baseline" mood. Offsets result from changing emotions throughout the day.

mood due to some event in the storyline, but can be "cheered up" through repeated happy interactions during the day.

As discussed above, positive social exchanges increase positive moods while negative exchanges erode any mood [30]. If the robot experiences a negative emotion during any mood or a positive emotion during a positive mood (e.g. $m_b + m_o > 0$), the mood offset (m_o) changes by one fifth the strength of the emotion—a change that is small enough to be unnoticeable after only a small number of exchanges, but may cause a significant change in the robot's expression (see Eq. 2). Positive exchanges during a negative mood have little effect, increasing the mood offset by only half what other exchanges cause (that is, one tenth the strength of the emotion).

While the overall mood changes dynamically due to interactions with people during the day, the "baseline" mood is designed to dominate the robot's mood. That is, after a person has stopped interacting with the robot, any effect that the person had on the robot's mood will begin to decay. Currently, we implement this decay by decreasing the mood offset (m_o) by small increments (currently, 10%) at regular intervals. Thus, if there are no further interactions over time, the mood will eventually return to the baseline level.

D. Mood and attitude

An attitude is essentially a long-term mood associated with some person (or thing). Each person who visits the robot may cause various emotional responses in the robot, thus changing the robot's mood, and those changes influence the robot's "opinion" of that person. How well the robot "knows" the person may influence how the robot's attitude changes toward that person. For the robot to maintain believability, its attitude toward someone should remain consistent across multiple interactions with that person.

In our model, attitudes consist of a mood level (A_m) and a familiarity rating (A_f). Familiarity is computed by combining how many hours a person has ever spent interacting with the robot and how many days since the person last visited the robot. Because we could find no psychological evidence for a computational model of familiarity, we have chosen a simple linear combination of these two values, as given in Eq. 4. A person has the highest possible familiarity rating (1) if she has interacted with the robot for at least 10 hours (including all past interactions with the robot), and the last time she interacted with the robot was that same day. Visitors are assigned familiarity ratings as follows:

$$A_f = \frac{1}{2}\left(1 + \frac{1}{10}\min(hours, 10) - \frac{1}{30}\min(days, 30)\right) \tag{4}$$

A stranger is assumed to have interacted infinitely many days ago, and so his familiarity rating is 0.

When a person begins an interaction with the robot, the robot's attitude toward that person has an immediate effect on its mood. Specifically, the mood offset changes to the average of the current offset (m_o) and the attitude's mood level (A_m). Strangers to the robot are given an attitude level of 0, which will pull the robot's mood closer to the baseline plateau—making the robot less likely to display emotions to strangers. When a person leaves the robot, the robot's mood becomes the average of its mood immediately before the person arrived and its mood when the person leaves. Furthermore, the robot's attitude toward that person changes by a fraction of the robot's mood change during the interaction (Eq. 5). Higher levels of familiarity result in a more stable attitude across interactions. An example of how a person might alter the robot's mood can be seen in Fig. 6.

$$A'_m = A_m + (\Delta m_o \text{ during interaction})(1 - A_f) \tag{5}$$

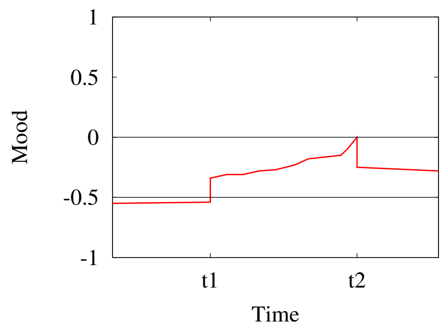

Fig. 6. An example of how attitude and interactions influence mood. Here, the baseline mood is -0.5, and a previous interaction has lowered the mood even further. A well-liked person ($A_m = 0.4$) approaches to begin interaction at time t1, which immediately improves the robot's mood. The person interacts from t1 to t2, largely trying to cheer up the robot, improving its mood considerably. When the person leaves, the mood drops once again, but remains at a higher level than before. After the interaction, the mood begins to decay toward the baseline.

IV. Conclusion

Humans increasingly need to communicate with robots, and we are working toward making human–robot interaction smooth and natural. This paper has presented a model of affect that covers the full range of emotions, moods, and attitudes, including the interactions between each. The model mimics human behavior, particularly with regard to long-term affective responses. Additionally, we have discussed our implementation of this model on the Roboceptionist.

We have had this model running on the Roboceptionist system for over six months. During this time, the mood has updated automatically to reflect the various events happening in the robot's storyline, which we have previously shown influences people's interactions with the robot [12]. While measuring short-lived emotional reactions is difficult, we have anecdotal evidence suggesting that people identify and respond to the robot's emotional expressions. We are currently working to improve the robot's person identification system in order to fully test the attitude portion of the model. We believe that this model will prove useful for a wide variety of social robots.

References

[1] B. Parkinson, "Emotions are social," *The British Psychological Society*, vol. 87, pp. 663–683, 1996.

[2] H. H. Clark and S. E. Brennan, "Grounding in communication," in *Perspectives on Socially Shared Cognition*. APA, 1991, pp. 127–149.

[3] J. P. Forgas, "Feeling and speaking: Mood effects on verbal communication strategies," *Personality and Social Psychology Bulletin*, vol. 25, no. 7, pp. 850–863, Jul 1999.

[4] B. Wild, M. Erb, and M. Bartels, "Are emotions contagious? Evoked emotions while viewing emotionally expressive faces: quality, quantity, time course and gender differences," *Psychiatry Research*, vol. 102, no. 2, pp. 109–124, 2001.

[5] N. H. Frijda, "Emotion experience," *Cognition and Emotion*, vol. 19, no. 4, pp. 473–497, 2005.

[6] E. A. Butler, B. Egloff, F. H. Wilhelm, N. C. Smith, E. A. Erickson, and J. J. Gross, "The social consequences of expressive suppression," *Emotion*, vol. 3, no. 1, pp. 48–67, 2003.

[7] B. Reeves and C. Nass, *The Media Equation*. Cambridge: CSLI Publications, 1996.

[8] I. Nourbakhsh, J. Bobenage, S. Grange, R. Lutz, R. Meyer, and A. Soto, "An affective mobile robot educator with a full-time job," *Artificial Intelligence*, vol. 114, no. 1-2, pp. 95–124, 1999.

[9] M. Montemerlo, J. Pineau, N. Roy, S. Thrun, and V. Verma, "Experiences with a mobile robotic guide for the elderly," in *Proceedings of the National Conference of Artificial Intelligence (AAAI 02)*, Edmonton,AB, July 2002, pp. 587–592.

[10] C. Breazeal, "Emotion and sociable humanoid robots," *International Journal of Human Computer Studies*, vol. 59, pp. 119–155, 2003.

[11] A. Bruce, I. Nourbakhsh, and R. Simmons, "The role of expressiveness and attention in human-robot interaction," in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, May 2002, pp. 4138–4142.

[12] R. Gockley, J. Forlizzi, and R. Simmons, "Interactions with a moody robot," in *Proceedings of Human-Robot Interaction*, Mar 2006, pp. 186–193.

[13] C. Elliott, "The affective reasoner: A process model of emotions in a multi-agent system," Ph.D. dissertation, Northwestern University, 1992.

[14] J. Cassell, T. Bickmore, L. Campbell, H. Vilhjálmsson, and H. Yan, "Human conversation as a system framework: Designing embodied conversational agents," in *Embodied Conversational Agents*, J. Cassell, J. Sullivan, S. Prevost, and E. Churchill, Eds. MIT Press, 2000, pp. 29–63.

[15] J. Cassell, J. Sullivan, S. Prevost, and E. Churchill, Eds., *Embodied Conversational Agents*. MIT Press, 2000.

[16] A. Paiva, J. Dias, D. Sobral, R. Aylett, S. Woods, L. Halle, and C. Zoll, "Learning by feeling: Evoking empathy with synthetic characters," *Applied Artificial Intelligence*, vol. 19, no. 3-4, pp. 235–266, Mar-Apr 2005.

[17] L. Moshkina and R. C. Arkin, "On TAMEing robots," in *IEEE International Conference on Systems, Man and Cybernetics*, Oct 2003.

[18] K. R. Scherer, "Psychological models of emotion," in *The Neuropsychology of Emotion*, J. C. Borod, Ed. Oxford: Oxford University Press, 2000, pp. 137–162.

[19] C. E. Izard, "Innate and universal facial expressions: Evidence from developmental and cross-cultural research," *Psychological Bulletin*, vol. 115, no. 2, pp. 288–299, 1994.

[20] J. A. Russell, J.-A. Bachorowski, and J.-M. Fernández-Dols, "Facial and vocal expressions of emotion," *Annual Review of Psychology*, vol. 54, pp. 329–349, 2003.

[21] P. Ekman, Ed., *Emotion in the Human Face*, 2nd ed. Cambridge, England: University of Cambridge Press, 1982.

[22] J. Russell, "Core affect and the psychological construction of emotion," *Psychological Review*, vol. 110, no. 1, pp. 145–172, 2003.

[23] A. Ortony, G. L. Clore, and A. Collins, *The Cognitive Structure of Emotions*. Cambridge University Press, 1988.

[24] C. Bartneck, "Integrating the OCC model of emotions in embodied characters," in *Workshop on Virtual Conversational Characters: Applications, Methods, and Research Challenge*, 2002.

[25] C. L. Rusting, "Personality, mood, and cognitive processing of emotional information: Three conceptual frameworks," *Psychological Bulletin*, vol. 124, no. 2, pp. 165–196, 1998.

[26] P. Ekman, "Moods, emotions, and traits," in *The Nature of Emotion: Fundamental Questions*, P. Ekman and R. J. Davidson, Eds. New York: Oxford University Press, 1994, pp. 56–58.

[27] R. Lazarus, "The stable and the unstable in emotion," in *The Nature of Emotion: Fundamental Questions*, P. Ekman and R. J. Davidson, Eds. New York: Oxford University Press, 1994, pp. 79–85.

[28] J. Neumann, B. Seibt, and F. Strack, "The influence of mood on the intensity of emotional responses: Disentangling feeling and knowing," *Cognition and Emotion*, vol. 15, no. 6, pp. 725–747, 2001.

[29] N. Bolger, A. DeLongis, R. C. Kessler, and E. A. Schilling, "Effects of daily stress on negative mood," *Journal of Personality and Social Psychology*, vol. 57, no. 5, pp. 808–818, 1989.

[30] K. S. Rook, "Emotional health and positive versus negative social exchanges: A daily diary analysis," *Applied Developmental Sciences*, vol. 5, no. 2, pp. 86–97, 2001.

[31] H. L. Wagner and J. Smith, "Facial expression in the presence of friends and strangers," *Journal of Nonverbal Behavior*, vol. 15, no. 4, pp. 201–214, Winter 1991.

[32] R. Gockley, A. Bruce, J. Forlizzi, M. Michalowski, A. Mundell, S. Rosenthal, B. Sellner, R. Simmons, K. Snipes, A. C. Schultz, and J. Wang, "Designing robots for long-term social interaction," in *Proceedings of IROS 2005*, Edmonton, Alberta, 2005, pp. 2199–2204.