How Does It Feel? Emotional Interaction with a Humanoid LEGO Robot

Lola D. Cañamero[*]
Inference and Learning Group
LRI, Bldg. 490
University of Paris-XI
F–91405 Orsay Cedex, France
lola@lri.fr

Jakob Fredslund
LEGO-Lab
Dept. of Computer Science
University of Aarhus
Aabogade 34, DK–8200 Århus N, Denmark
chili@daimi.au.dk

Abstract

We report work on a LEGO robot capable of displaying several emotional expressions in response to physical contact. Our motivation has been to explore believable emotional exchanges to achieve plausible interaction with a simple robot. We have worked toward this goal in two ways. First, acknowledging the importance of physical manipulation in children's interactions, interaction with the robot is through tactile stimulation; the various kinds of stimulation that can elicit the robot's emotions are grounded in a model of emotion activation based on different stimulation patterns. Second, emotional states need to be clearly conveyed. We have drawn inspiration from theories of human basic emotions with associated universal facial expressions, which we have implemented in a caricaturized face. We have conducted experiments on children and adults to assess the recognizability of these expressions, and observed how people spontaneously interacting with Feelix respond to its emotional displays.

Introduction

Emotional exchanges constitute an important element in human interaction and communication. Social robots interacting with humans must incorporate some capabilities to express and elicit emotions in order to achieve natural, believable interactions. Bearing these ideas in mind, we have built Feelix[1] (Figure 1), a 70cm-tall "humanoid" LEGO robot that displays different facial emotional expressions in response to tactile stimulation. Feelix is a descendant of Elektra (Fredslund, 1998), a mobile, also "humanoid" LEGO robot first exhibited at the FIRA Robot World Cup in Paris in 1998. People, in particular children, found it very natural to interpret the happy and angry expressions of Elektra's smiley-like face; however, the interaction that humans could have with the robot to elicit these expressions was not so natural, as it consisted in inserting different color LEGO parts into its chest. Our motivation for building Feelix—Elektra's body with a new head and feet—was twofold. First, we aimed at a more plausible interaction with the robot. We wanted to focus on the interaction itself, and therefore we did not want it to be influenced by the robot performing a particular task. For this, we decided to exploit the potential that robots, unlike computer simulations, offer for physical manipulation, as this plays an important role in children's development and in human interaction in general. Interaction with Feelix is therefore through tactile stimulation rather than through other sensory modalities that do not require physical contact such as vision, used for example in (Breazeal, 1998; Thrun, 1999).

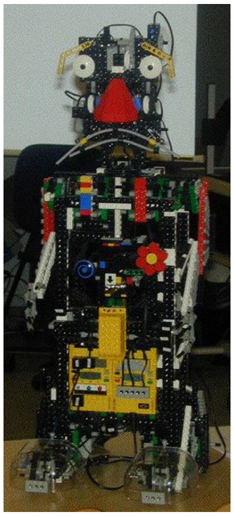

Figure 1: Feelix.

Our second motivation was to achieve a richer interaction so that a wider range of interaction patterns giving rise to a wider range of emotional responses (both in the robot and, by empathy, in the human) was possible. These emotional responses must be clearly recognizable, and therefore we limited ourselves to implementing the ones known in the emotion literature as "basic emotions", which supposedly have clear associated expressions. However, while we aimed at a richer interaction, we wanted to keep the robot very simple, both its architecture and its emotional expressions. In doing this, our purpose was to come up with a minimal set of features that make the emotional displays of the robot (and the interaction itself) believable, and to assess to what extent we could rely on the tendency humans have to anthropomorphize in their interactions with objects presenting human-like features.

Modeling Emotions in Feelix

The emotions we have implemented in Feelix correspond to the ones known as "basic emotions". Although the use of this term is still highly controversial among students of human emotions, as researchers do not agree neither in the number and subset of emotions that can be considered basic (classifications range from two to nine), nor in what sense they are so, we can characterize basic emotions in a general way as follows: they seem to be universally found across

---

[*] The work reported here was carried out while this author was a visiting researcher at the LEGO-Lab.

[1] FEELIX: FEEL, Interact, eXpress.

societies in humans (and in other mammals), have particular manifestations associated with them (facial expressions, behavioral tendencies, physiological patterns), have adaptive, survival-related roles that evolved through their value in dealing with situations which are or were fundamental in life, can be seen as prototypical emotional states of a number of emotion families (e.g., rage and anger belong in the same family, anger being the more standard or prototypical case, while rage corresponds to highly intense anger), and can be taken as building blocks out of which other, more complex emotions form. In our case, the main reason for using discrete basic emotions has been the potential they offer to be easily elicited and recognized from their facial expressions. The particular subset of basic emotions that we have adopted is the one proposed by Ekman (Ekman, 1992)—anger, disgust, fear, happiness, sadness, and surprise—with the exception of disgust. The reason for choosing this classification is that its main criterion to define emotions as basic is their having distinctive (universal) facial expressions.

Facial Expressions

Feelix makes different facial expressions (see Figure 2) by means of two eyebrows and two lips. To define the "primitives" (lips and eyebrow positions) for each emotion, we have largely adopted the features concerning eyebrows and lips described in (Geppert *et al.*, 1997), a coding system inspired by Ekman's Facial Action Coding System. Concerning an observer's perception of emotional expressions, we have adopted the hypothesis proposed by (De Bonis *et al.*, 1999) that the upper and lower parts of the face function as the building blocks at the basis of emotion perception, rather than the building blocks of finer granularity postulated by other authors. We have taken the most telling feature of each part of the face—eyebrows and lips—to express emotions in Feelix, making the hypothesis that these two features should allow humans to recognize its emotional expressions. Feelix' face is thus closer to a caricature than to a realistic model of a human face.

Feelix' eyebrows are two slightly bent LEGO parts, resembling the bent shape of human eyebrows, that are attached at their long end to a shaft around which they rotate symmetrically. Lips are flexible rubber tubes that can independently curve both ways. The mouth can be made narrow or wide by symmetrically moving its corners inwards or outwards. Five motors (arranged in two pairs plus a switcher) are used to achieve the four degrees of freedom (DoF) of the face, controlled by a LEGO Mindstorms RCX™ computer (www.legomindstorms.com).

Each of Feelix' emotional states has an associated distinctive prototypical facial expression (shown in Figure 2), characterized as follows:

- Anger: raised eyebrows, moderately open wide mouth with upper lip curved downwards and straight lower lip.
- Sadness: maximally lowered eyebrows, closed mouth curved downwards.
- Fear: lowered eyebrows, moderately open wide mouth.
- Happiness: straight eyebrows, closed wide mouth curved upwards.

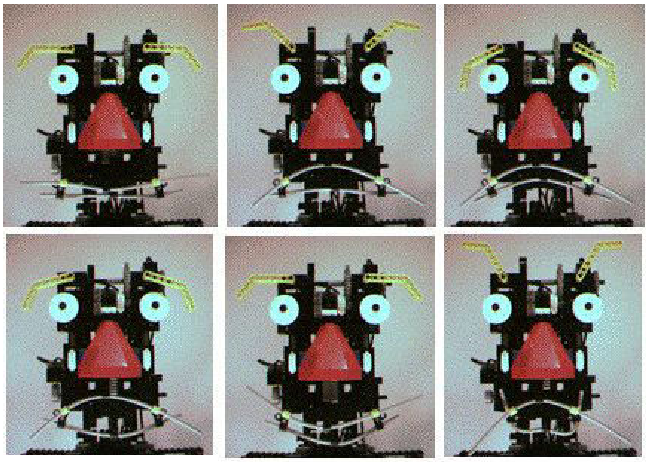

Figure 2: Emotional expressions displayed by Feelix. From left to right and top to bottom: neutral, anger, sadness, fear, happiness, and surprise.

- Surprise: highly raised eyebrows, very open narrow mouth.

Expressions are slightly modified (by moving the expressive feature in the dominant half of the face, eyebrows or mouth) when emotion intensity is very high.

Emotional Interaction

Emotions are complex phenomena that involve a number of related subsystems and can be activated by any one (or by several) of them. Elicitors of emotions are for example grouped in (Izard, 1993) under the categories of neurochemical, sensorimotor, motivational, and cognitive. Some of these elicitors are emotion-specific, but emotions also show a certain degree of generality (Tomkins, 1984)—e.g., generality of object and time. This accounts for the fact that a person can experience the same emotion under different circumstances and with different objects. Unlike previous work that used specific stimuli to elicit emotions (Cañamero, 1997), we have adopted and extended the generic model postulated by Tomkins (Tomkins, 1984), which proposes three variants of a single principle:

- *Stimulation increase*. A sudden increase in the level of stimulation can activate both positive (e.g., interest) and negative (e.g., startle, fear) emotions. The degree of suddenness accounts for their differential activation.
- *Stimulation decrease*. A sudden decrease in the level of stimulation only activates positive emotions such as joy.
- *High stimulation level*. A sustained high level of stimulation activates negative emotions such as distress or anger.

The fact that this model is general makes it applicable to stimulation patterns of all kinds, and particularly suited for emotion activation in our robot, since we were aiming at physical interaction based on tactile stimulation. However, an analysis of the possible interactions that humans could have with Feelix revealed two cases that the model did not account for. First, the model did not propose any principle to activate negative emotions such as sadness or boredom. Second, happiness is only considered in the form of relief resulting from stopping a very high, annoying level of stimulation, i.e., as the cessation of a negative stimulus. Happiness, however, can also be produced by positive stimuli, such as gentle interaction. We have therefore refined the model by postulating two more principles:

- *Low stimulation level.* A low stimulation level sustained over time produces negative emotions such as sadness.
- *Moderately high stimulation level.* A moderate stimulation level produces positive emotions such as happiness.

We had two main reasons to adopt this model. On the one hand, the coarse sensory capabilities of the robot considerably limit the number of different types of stimuli it can recognize; the use of particular stimuli to elicit different emotions seems thus unrealistic. On the other hand, the generality of the model makes it particularly suited to experiment with different implementations using different sensory modalities, allowing to assess the significance and adequacy of each modality in human-robot interactions.

A second RCX controls the interaction with humans and communicates with the RCX controlling the face. We wanted the interaction to be as natural as possible, and since for this project we are not using Feelix as a mobile robot—the human is sitting in front of it so as to better observe the face—the feet seemed to be the best location for tactile stimulation, as they are protruding and easy to touch. We built two special feet for Feelix using touch-friendly (smooth, large, and rounded) LEGO parts. Underneath each foot is a binary touch sensor—pressed or not-pressed. To distinguish between different kinds of stimuli, we use *duration* and *frequency* of presses. For duration, it was enough to define three types of stimuli to implement the general stimulation patterns in our model: short (less than 0.4 seconds), long (up to five seconds), or very long (over five seconds). Frequency is calculated on the basis of a minimal time unit of 2 seconds that we call *chunk*. When a chunk ends, information about stimuli is analyzed and a message encoding Feelix' current emotional state and its intensity is sent to the RCX controlling facial expressions.

Although it is possible to combine two different expressions in Feelix' face, we have as for now adopted a winner-takes-all strategy[2] based on the level of emotion activation to select and display the emotional state of the robot. Emotions are assigned different intensities calculated on the grounds of stimulation patterns. At the end of each chunk, the emotion with the highest intensity determines the emotional state of Feelix. However, for this emotion to become active and get expressed, its intensity has to reach a certain threshold. By setting this threshold higher or lower, we can give Feelix different temperaments—make it more "extroverted" or "introverted". When a new emotion becomes active, it temporarily inhibits all the other emotions by resetting their intensities to 0. Emotion intensities are calculated by an *update function* that depends on time and reflects some of the distinctive features of basic emotions, namely quick onset and brief duration (Ekman, 1992). The intensity of the active emotion increases with appropriate stimulation depending on how long this emotion has been active. All emotions increase their intensities with stimulation except sadness, which is produced when Feelix gets no attention. A *time decay function* makes emotion intensities decrease when Feelix is not being stimulated. For sadness, this function applies only after a long period of inactivity, when its intensity has reached its highest level. When no emotion is active—i.e. when all emotions' intensities are below the activation threshold—Feelix displays a neutral face.

We mapped the general stimulation patterns from our model into tactile stimulation patterns as follows.

- *Stimulation increase* is achieved by frequent short presses on any of the feet. This pattern can give rise to two emotions, surprise and fear. *Surprise* is produced by a less intense increase—one or two short presses after a period of inactivity or low/moderate activity. An inhibition mechanism prevents the reoccurrence of surprise within a short period of time. *Fear* is produced when the increase is more intense, needing more frequent short presses to become active.
- A *sustained high stimulation level* overwhelms Feelix and produces anger. Very long presses, lasting three or more chunks, or many frequent short presses increase the intensity of anger.
- A *moderate level* of stimulation that neither overstimulates nor understimulates Feelix produces happiness. This level is achieved by gentle, long (but not too long) presses.
- Sadness is produced by a *sustained low level* of stimulation. As already mentioned, in Feelix this corresponds to lack of (or very little) interaction for a long period.

The amount of stimulation required to change Feelix' emotional state and its expression depends on the intensity of the currently active emotion—the more intense the emotion, the more stimulation is needed for a change to happen.

This model of emotion activation is implemented by means of a timed finite state machine—see (Cañamero & Fredslund, 2000) for details.

Experiments

We have investigated two aspects of the interaction with Feelix: the recognizability of the emotional facial expressions, and the suitability of the interaction patterns.

Experiments on emotion recognition

To evaluate this aspect we designed three tests. The first one is a free test—no list of emotion adjectives or any other cues are provided—in which subjects are asked to label a sequence of five expressions performed by Feelix: anger, sadness, fear, happiness, surprise. The second test is a multiple-choice one in which subjects are asked to label the same sequence of expressions, but this time they are given a list of

---

[2] We have also built some demos where Feelix shows chimerical expressions that combine an emotion in the upper part of the face—eyebrows—and a different one in the lower part—mouth.

nine emotion descriptors including four extra ones: disgust, anxiety, pride, worry. In addition, to test whether subjects can recognize the valence of the emotion, for each emotion they are asked whether they think the expression is elicited by something Feelix likes or by something it does not like. As a control, we designed a free test where subjects are asked to label emotional expressions from pictures of human faces—anger, sadness, happiness, fear, contempt, surprise, and disgust.

We have conducted experiments on 86 subjects—45 adults (ages 15-57) and 41 children (ages 9-10). Experiments were performed in four suites. Due to time constraints, only one group of 14 children could do the free test on human faces. Answers were considered to be correct when the subjects used the same descriptors we have employed or very close synonyms. The results we obtained were surprisingly close to those commonly reported in the emotion literature on recognition of facial expressions of basic emotions in cross-cultural studies (see e.g. (Cornelius, 1996) for an overview), in particular in the free tests. Results are summarized in Tables 1 (experiments with children) and 2 (experiments with adults). Average recognition of emotional expressions[3] was 58% for adults and 64% for children in the free test on Feelix' face, 55% for adults and 48% for children in the multiple-choice test on Feelix' face, and 82% for adults and 70% for children in the test on pictures of human faces. Children thus seem to be better than adults at recognizing emotional expressions in Feelix' caricaturized face when they can freely describe the emotion they observe, whereas they perform worse when they are given a list of descriptors to choose from. Contrary to our initial guess, providing a list of descriptors did not help recognize the observed emotion, but diminished performance in both adults and children. Results on recognition of emotional expressions from pictures of human faces were better than on the robot in both cases. Recall that all these results measure the ability to recognize emotions from the face alone—using some features in the case of the robot, the whole face in the case of human pictures—i.e. in the absence of any clues provided by body posture and contextual elements, which can indeed be crucial factors to assess observed emotion.

Assessing the interaction

We have not (yet) performed formal tests to evaluate the plausibility of the interactions with Feelix. So far we have observed people spontaneously interacting with the robot, or trying to guess the stimulation patterns we used to elicit different emotions in it. These are thus informal observations that should not be taken as conclusive results.

Some interaction patterns (those of happiness and sadness) seem to be very natural and easy to understand, while others present more difficulty (e.g., it takes more time to learn to distinguish between the patterns that cause surprise and fear, and between those that produce fear and anger). We have observed some interesting mimicry and empathy phenomena, though. In people trying to elicit an emotion in Feelix, we have observed their mirroring—in their own faces and in the nature of the presses applied to the feet—the emotion they wanted to elicit (e.g., displaying an angry face and pressing the feet with much strength while trying to elicit anger). We have also observed people reproducing Feelix' facial expressions during emotion recognition, this time with the purpose of using proprioception of facial muscle position to assess the emotion observed. During recognition also, people very often mimicked Feelix' expression with vocal inflection and facial expression while commenting on the expression ("ooh, poor you!", "look, now it's happy!"). People thus seem to empathize with the robot quite naturally.

Children (ages 9-14)	Robot, free test	Robot, multiple choice test	Human, free test
Anger	64%	44%	100%
Sadness	83%	57%	79%
Fear	0%	22%	64%
Happiness	93%	57%	100%
Surprise	17%	37%	50%
Contempt	-	-	0%
Disgust	-	-	29%

Table 1: Results of emotion recognition by children.

Adults (ages 15-57)	Robot, free test	Robot, multiple choice test	Human, free test
Anger	57%	37%	71%
Sadness	81%	84%	91%
Fear	2%	9%	62%
Happiness	64%	62%	98%
Surprise	29%	36%	93%
Contempt	-	-	0%
Disgust	-	-	76%

Table 2: Results of emotion recognition by adults.

Related Work

Not much research has been conducted on the use of emotions and their facial expressions in the context of human-robot interaction, although the interest in this issue is fast growing. Perhaps the project that most closely relates to ours is the robot Kismet, developed by Cynthia Breazeal. Kismet is a testbed for learning social interactions in situations involving an infant (the robot) and her caretaker (a human). Kismet is a head with active stereo vision and configurable expressive features—controllable eyebrows, ears, eyeballs, eyelids, a mouth with an upper lip and a lower lip, and a neck that can pan and tilt—with many DoF. All these features allow Kismet to express a wide variety of emotional expressions that can be mapped onto a three-dimensional space with dimensions arousal, valence, and stance. Humans can interact with Kismet either by direct face-to-face exchange or by showing it a toy. (Breazeal, 1998) reports

---

[3] These figures exclude results for fear in the robot tests and for contempt in the human faces one, since all subjects agreed that these expressions were very bad (results were close to 0%). Their inclusion lowers figures by about 10 points.

on some early experiments where Kismet uses nine different facial expressions to manipulate its human caretaker into satisfying its internal drives—a social drive, a stimulation drive, and a fatigue drive. More recently, some web-based experiments have been conducted to assess how well humans can recognize Kismet's expressions (Breazeal & Forrest, 1999). We will further discuss this model and compare it with ours in the next section, since we have taken almost opposite working hypotheses concerning emotional expression—a sophisticated, configurable set of features in her case, a minimalist one in ours.

Minerva, developed by Sebastian Thrun, is a mobile robot that gives guided tours to visitors of the Smithsonian's Museum of American History (Thrun, 1999). It displays emotional expressions—neutral, happy, sad, and angry—using a caricaturized face and simple speech. Emotional states arise as a consequence of travel-related interaction (e.g., anger results from its inability to move due to the presence of people), and their expressions aim at affecting this interaction towards achieving the robot's goals—traveling from one exhibit to the next one, engaging people's attention when describing an exhibit, and attracting people to participate in a new tour. Although very successful interactions attributed to empathetic feelings in people are reported, emotions in Minerva are purely a means to an end and not an integral part of the robot's architecture and interface.

A surprising experiment, although using a computer instead of a robot, was conducted by Clark Elliott to test the computer's ability to express emotions by having humans recognize them (Elliot, 1997). The computer used both caricaturized facial expressions and voice inflection to convey different emotional states while saying sentences devoid of emotional content. As a control, he had an actor say the same sentences and express the same emotions. It turned out that humans performed substantially better when recognizing the emotions expressed by the computer (70% of success) than those expressed by the actor (50% of success). Elliott suggests that these results might be partly due to the use of caricaturized expressions.

What Features? What Interactions?

Are emotions better thought of in terms of continuous dimensions, or in terms of basic categories that can be combined to form more complex ones? This is an old—unresolved and, in our opinion, ill-posed, see (Cañamero, 2000) for a discussion—controversy in the human emotion literature. Interestingly, Kismet and Feelix implement these two "opposite" models, opening up the door to an investigation of this issue from a synthetic perspective—which model is better suited to build an expressive social robot? Our answer is that both present advantages and disadvantages, and the choice depends on the issue we intend to investigate with our robot.

In building Feelix, our purpose was to come up with a minimal set of features that would however allow "rich enough" emotional interaction—i.e. express and elicit at least the set of emotions which are considered basic and universal. We also wanted to test "how much emotion" can be expressed with a caricaturized face alone, and to explore the adequacy of different sensory modalities for emotional interactions. On the contrary, Kismet is a much more ambitious projet that proposes a testbed to investigate diverse infant-caretaker interactions, not only emotional ones. More expressive features, DoF, and sensory modalities are thus needed. Let us discuss some of our design choices in the light of the relevant "design constraints" proposed in (Breazeal & Forrest, 1999) for robots to achieve human-like interaction with humans.

Issue I: The robot should have a cute face to trigger the 'baby-scheme' and motivate people to interact with it. Although one can question the cuteness of Feelix, the robot does present some of the features that trigger the 'baby-scheme'[4], such as a big head, big round eyes, and short legs. However, none of these features is used to express or elicit emotions in Feelix. In fact, it was mostly Feelix' expressive behavior that elicited the baby-schema reaction.

Issue II: The robot's face needs several degrees of freedom to have a variety of different expressions, which must be understood by most people. The insufficient DoF of Elektra's face was one of our motivations to build Feelix. The question, however, is how many DoF are necessary to achieve a particular kind of interaction. Kismet's model doubtless allows to form a much wider range of expressions; however, not all of them are likely to convey an emotional meaning to the human. On the other hand, we think that Feelix' "prototypical" expressions associated to discrete emotional state (or to a combination of two of them) allow for easier emotion recognition (although of a more limited set) and association of a particular interaction with the emotion it elicits. This model also facilitates an incremental study of what features are really relevant to express or elicit different emotions. Indeed, our experiments showed that our features were insufficient to express fear, were body posture (e.g. the position of the neck) adds a lot of information.

Issue IV: The robot must convey intentionality to bootstrap meaningful social exchanges with the human. It is however questionable that "more complexity" conveys "more intentionality" and adds believability. As we observed with Feelix, very simple features can have humans put a lot on their side and anthropomorphize very easily.

Issue V: The robot needs regulatory responses so that it can avoid interactions that are either too intense or not intense enough. In Feelix' case, emotional expression itself acted as a regulatory mechanism influencing people's behavior—in particular sadness as a response to lack of interaction, and anger as a response to overstimulation.

Our conclusion is thus that the richness and complexity of Kismet's face is not necessary for the issue we wanted to investigate and the kind of interactions we aimed at. The converse question can be posed, though—could Kismet's emotional expression system be simpler and based on discrete emotion categories, and still achieve the rich possibilities of interaction it aims at? Although answering this question

---

[4] According to Irenäus Eibl-Eibesfeld, the baby-scheme is an "innate" response to treat as an infant every object presenting certain features present in children.

could be difficult, as it would require substantial reimplementation of the robot, a comparison of human performance on emotion recognition on both robot's faces would doubtless be a very interesting experiment to perform.

Conclusion

We have presented early work on Feelix, a humanoid-looking LEGO robot capable of displaying several emotional expressions in response to direct physical stimulation. Feelix implements two models of emotional interaction and expression inspired by psychological theories about emotions in humans. This makes Feelix not only very suitable for entertainment purposes, but also a proof-of-concept that these theories can be used with a synthetic approach that complements the analytic perspective for which they were conceived.

We have conducted experiments to assess how well humans can recognize emotional expressions in Feelix' face. Our results approach results reported in the literature on emotion recognition from pictures of human faces. They also show that the "core" basic emotions of anger, happiness, and sadness are most easily recognized, whereas fear was mostly interpreted as anxiety, sadness, or surprise. This result might be due to the need of additional expressive features. Rasmus B. Lunding, of the Danish Institute of Electro-acoustic Music in Aarhus, has composed sound patterns in order to enhance the different emotions conveyed (a video can be seen at http://www.daimi.au.dk/∼chili/feelix/feelix_home.htm), but we have not yet tested whether the addition of sound improves emotion recognition. Observation of spontaneous interactions with the robot showed that humans tend to mimic and empathize with Feelix in a quite natural way. This matches our initial expectations concerning "how much" humans anthropomorphize when interacting with objects with human-like features, i.e. how few of these (emotion-related, in this case) human-like features are needed to make the interaction believable.

In order to better assess the adequacy of tactile interactions and the generality of our emotion activation model, the model must be implemented using other sensory modalities. Sound is our next target, and special ears have already been developed to enable Feelix to respond to auditory stimulation (e.g., claps). Finally, to obtain a more sound analysis of Feelix' emotional expressions and a better comparison with studies of recognition of emotional expressions in human faces, we have started to analyze the robot's expressions with psychologists in De Bonis group at the University of Paris-XI.

Acknowledgments

We are grateful to Henrik Lund, Ole Caprani, Brian Mayoh, and the students at the LEGO-Lab for their insightful suggestions and their cheerful support. This research is funded by the LEGO-Lab, Dept. of Computer Science, University of Aarhus.

References

Breazeal, C. 1998. Early Experiments using Motivations to Regulate Human-Robot Interaction. In D. Cañamero, ed., *Emotional and Intelligent: The Tangled Knot of Cognition*. Papers from the 1998 AAAI Fall Symposium. AAAI Technical Report FS–98–03. Menlo Park, CA: AAAI Press.

Breazeal, C., Forrest, A. 1999. Schmoozing with Robots: Exploring the Boundary of the Original Wireless Network. In K. Cox, B. Gorayska, J. Marsh, eds., *Proc. 3rd. Intl. Cognitive Technology Conference*. Sillicon Valley, San Francisco, USA, August 11–14, 1999.

Cañamero, D. 1997. Modeling Motivations and Emotions as a Basis for Intelligent Behavior. In W. Lewis Johnson, ed., *Proceedings of the First International Conference on Autonomous Agents*, pp. 148–155. New York, NY: ACM Press.

Cañamero, D. 2000. Designing Emotions for Activity Selection in Autonomous Agents. In R. Trappl, ed., *Emotions in Humans and Artifacts*. Cambridge, MA: The MIT Press (in press).

Cañamero, D. and Fredslund, J. 2000. I Show You How I Like You: Human-Robot Interaction through Emotional Expression and Tactile Stimulation. DAIMI Technical Report 544, Dept. of Computer Science, University of Aarhus, Denmark, February 2000.

Cornelius, R.H. 1996. *The Science of Emotion*. Upper Saddle River, NJ: Simon & Schuster.

De Bonis, M., De Boek, P., Pérez-Díaz, F., Nahas, M. 1999. A Two-Process Theory of Facial Perception of Emotions. *C.R. Acad. Sci. Paris, Life Sciences* 322: 669–675.

Ekman, P. 1992. An Argument for Basic Emotions. *Cognition and Emotion* 6(3/4): 169–200.

Elliot, C. 1997. I Picked Up Catapia and Other Stories: A Multimodal Approach to Expressivity for 'Emotionally Intelligent' Agents. In W. Lewis Johnson, ed., *Proceedings of the First International Conference on Autonomous Agents*, pp. 451–457. New York, NY: ACM Press.

Fredslund, J. 1998. Elektra. Unpublished Technical Report. Department of Computer Science, University of Aarhus, Denmark. Available at http://www.daimi.au.dk/∼chili/elektra.html.

Geppert, U., Schmidt, D., Galinowski, I. 1997. Self-Evaluative Emotions Coding System (SEECS) Technical Manual. Paper 19/1997, Max Planck Institute for Psychological Research, Munich, Germany.

Izard, C.E. 1993. Four Systems for Emotion Activation: Cognitive and Non-Cognitive Processes. *Psychological Review* 100(1): 68–90.

Thrun, S. 1999. Spontaneous, Short-term Interaction with Mobile Robots in Public Places. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA).*

Tomkins, S.S. 1984. Affect Theory. In K.R. Scherer & P. Ekman, eds., *Approaches to Emotion*, pp. 163–195. Hillsdale, NJ: Lawrence Erlbaum Associates.