Correspondence

I Show You How I Like You—Can You Read it in My Face?

Lola Cañamero and Jakob Fredslund

***Abstract*—We report work on a LEGO robot that displays different emotional expressions in response to physical stimulation, for the purpose of social interaction with humans. This is a first step toward our longer-term goal of exploring believable emotional exchanges to achieve plausible interaction with a simple robot. Drawing inspiration from theories of human basic emotions, we have implemented several prototypical expressions in the robot's caricaturized face and conducted experiments to assess the recognizability of these expressions.**

***Index Terms*—Cooperative systems, psychology, robot programming, robot tactile systems, user centered design, user interfaces.**

I. 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] (Fig. 1): a 70-cm-tall "humanoid" LEGO robot that displays different facial (emotional) expressions in response to tactile stimulation. Our goal is to investigate emotion expression and activation in the context of social (human–robot) interaction, and therefore, we did not want this to be influenced by the robot performing a particular task. Feelix's only task is, thus, emotional expression for the purpose of social interaction with humans.

The motivation underlying the design of Feelix is twofold. First, we aimed at achieving a plausible way of interacting with a very simple robot. 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, often used in other expressive and social robots. Second, we wanted to endow our robot with a rich enough set of believable and easily understandable emotional responses (facial displays), while having to deal with the simplicity constraints imposed by Feelix's architecture and more generally by LEGO robots. Our purpose was thus 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 for this on the tendency humans have to anthropomorphize in their interactions with objects and technology presenting human-like features, as extensively argued by Reeves and Nass [20]. In this paper, we will focus on emotional expressions.

Previous work with Elektra [13], which is also a "humanoid" LEGO robot, showed that people, and in particular children, found it very natural to interpret the happy and sad expressions of Elektra's smiley-like face. However, more expressions are needed for richer interactions. Since we wanted emotional responses to be clearly recognizable, we decided to implement the subset known in the literature as "basic emotions," as they are hypothesized to have clear, "universal" associated expressions. The term "basic emotions" is still highly controversial among students of human emotions (see [7] and [17] for accounts of this controversy), as researchers do not agree neither on the particular emotions that can be considered as basic (classifications range from two to nine), nor in what sense they are so. A good characterization is provided in [7]. Some researchers have characterized emotions in terms of continuous dimensions rather than as discrete categories. From this perspective, emotions are described as points located on a small set of continuous scales; the most commonly used (names may vary) being *valence* (the fact of a stimulus being perceived as positive or negative) and *arousal* (the general level of activity). The category and dimension views of emotion are not incompatible, however, and in our model we use a combination of them. In fact, basic emotions can be easily placed in an emotional space defined by these dimensions (see, for instance, [18] and [20]). Current componential theories [21], [23], tend to reconcile both approaches and acknowledge the importance of both dimensions and categories, with dimensions being regarded as primary and categories as more elaborated information extracted in a second step of cognitive processing.

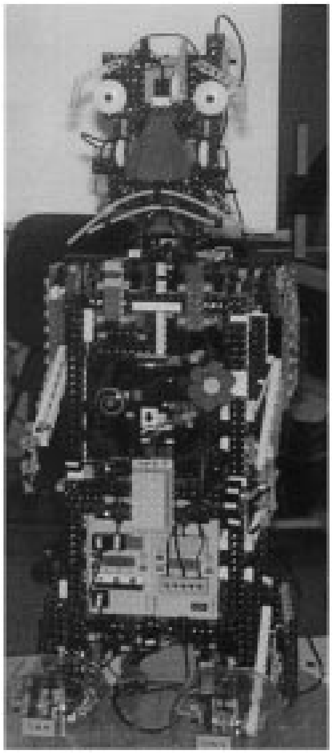

Fig. 1. Feelix.

In this study, we opted for the categorical approach, as we were aiming at a small set of easily recognizable emotions. The particular subset of basic emotions that we have adopted is (with the exception of disgust) the one proposed by Ekman in [6]—anger, disgust, fear, happiness, sadness, and surprise—as the main criterion used to define emotions as basic is their having distinctive prototypical facial expressions.

Manuscript received January 4, 2000; revised June 1, 2001. This work was supported by the LEGO Lab, Department of Computer Science, University of Aarhus.

L. Cañamero is with the Department of Computer Science, University of Hertfordshire, Hertfordshire, U.K. (e-mail: l.canamero@herts.ac.uk).

J. Fredslund is with the LEGO Lab, Department of Computer Science, University of Aarhus, Aarhus, Denmark (e-mail: chili@daimi.au.dk).

Publisher Item Identifier S 1083-4427(01)07730-X.

[1] FEELIX: FEEL, Interact, eXpress.

II. EXPRESSIVE ARTEFACTS

The use of emotions and their facial expressions in the context of human–robot interaction is receiving increasing attention. Here we give an overview of selected work representative of different models of facial expression.

The Affective Tigger [14] is an expressive toy developed by Kirsch as a tool for the social and emotional awareness education of small (aged two to five) children. Affective Tigger's expressive facial features are a mouth—open or closed—and ears—pointing upwards or downwards. As in the case of Elektra, the predecessor of Feelix, the face of Tigger has only 2 degrees of freedom (DoF) that allow it to express two emotions—happiness and sadness—plus a neutral face. Facial and vocal expressions reflect the emotional state of the toy as a response to the child's physical manipulation. Emotion recognition from the face alone is reported to be rather poor, probably due to the rigid shape of the mouth. Although this two-emotion model can be too simple for some applications, it is, however, appropriate to teach small children about emotion valence and valenced reactions to their "nice" or "mean" behaviors.

At the opposite end of the complexity scale lies Kismet, which was developed by Breazeal [1] as a testbed for learning social interactions in situations involving an infant (the robot) and her caretaker (a human). To design Kismet's facial expressions, Breazeal has drawn inspiration from componential approaches, in particular [21] and [23]. This robot is a head with active stereo vision and configurable expressive features—controllable eyebrows, ears, eyeballs, eyelids, a mouth with two lips, and a neck that can pan and tilt—with 18 DoF. All these features, together with an expressive vocalization system, allow Kismet to display a wide variety of emotional expressions that can be mapped onto a three–dimensional (3-D) space with dimensions arousal, valence, and stance. Various experiments to assess how well humans can recognize Kismet's expressions are reported in [1], with good recognition results. However, Breazeal does not consider these results as conclusive, given that the number of subjects tested was not sufficient.

The next two expressive robots are closer to ours in the features and number of DoF used to express emotion. Minerva [24], which was developed by Thrun, is an interactive tourguide robot that displays four basic expressions—neutral, happy, sad, and angry—using a caricaturized face and simple speech. As in our case, the robot's face is a caricature with two expressive features and 4 DoF—one to control each eyebrow and two to control the mouth. Emotional states arise as a consequence of travel-related interaction, and their expressions aim at affecting this interaction toward achieving the robot's goals. 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.

Sparky is a social robot developed by Scheeff and colleagues with the aim of exploring new ideas in human–computer interface and interactive robotics [22]. It uses facial expression, gesture, motion, and sound to be social with humans in the immediate vicinity, focusing on emotional expression for the purposes of social interaction with people. As in the case of Feelix, the robot's only task is thus emotional expression in the context of social interaction. Unlike the other robots presented here, Sparky is not autonomous but teleoperated, in order to create richer social interactions than what the current state-of-the-art in autonomous robots allows, as well as analyze human responses during interaction. Sparky's face has 4 DoF to control three expressive features—eyebrows (1 DoF), eyelids (1 DoF), and lips (1 DoF each). In addition, the robot makes extensive use of its 6 DoF body for expressive purposes. Spontaneous interaction with Sparky has been studied in laboratory and public settings with subjects aged 6 to 81.

We finally mention a surprising experiment conducted by Elliott to test a computer's ability to express emotions by having humans recognize them [11]. 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. Subjects 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.

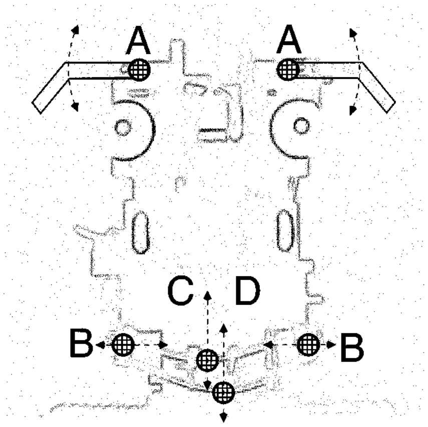

Fig. 2. Four DoF of Feelix's face.

III. FEELIX THE ROBOT

We have built Feelix using the LEGO Mindstorms™ robotic construction kit. Feelix is controlled by two LEGO Mindstorms RCX™ computers—one controls the facial (emotional) expressions, the other emotion activation. An RCX has a Hitachi H8/300 CPU, 32 K RAM, three input ports, and three output ports. It can have limited infrared communication with other RCXs or with a PC.

Feelix's face (Fig. 2) has 4 DoF and makes different emotional expressions by means of two eyebrows (1 DoF) and two lips (3 DoF). The eyebrows are two slightly bent LEGO parts resembling the shape of human eyebrows. They are attached at their long end to a shaft (labeled *A* in Fig. 2), around which they rotate symmetrically. They are controlled using an angle sensor and one motor. Lips are flexible rubber tubes that can independently curve both ways (DoF *C* and *D* in Fig. 2). Each lip is controlled by an angle sensor and a motor. The mouth can be made narrow or wide by symmetrically moving its corners inwards or outwards (DoF *B* in Fig. 2) using another motor. Each angle sensor is connected to an input port of the RCX. To control the 4 DoF of Feelix's face by means of only three output ports, we have arranged the four motors in two pairs: eyebrows/mouth width and upper lip/lower lip. A fifth motor switches control between these two pairs. With more motors, it would have been possible to build a more expressive face, either by increasing the number of DoF of the existing elements or by adding other expressive elements. In that case, however, the face would have been significantly bigger, much heavier, requiring a bigger body, and its increased complexity would have required an additional RCX to control it, which would have had a negative impact on its performance in case of noisy communication between RCXs. By limiting the DoF to four, the face can distinctively display the five basic emotions we chose while being controlled by one RCX.

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 spe-

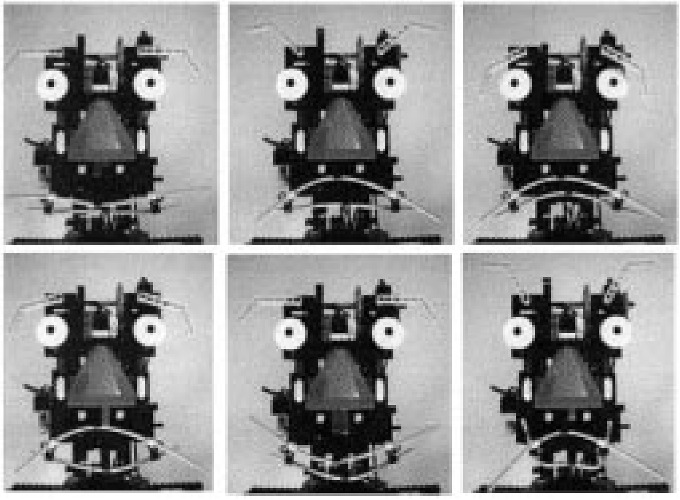

Fig. 3. Emotional expressions displayed by Feelix. From left to right and top to bottom: neutral, anger, sadness, fear, happiness, and surprise.

cial feet for Feelix using touch-friendly (smooth, large, and rounded) LEGO parts. Underneath each foot is a binary touch sensor.

IV. MODELING EMOTIONS IN FEELIX

The number and types of emotions and expressive features to be integrated in the robot depends on the kind of interaction and user we are aiming at, as illustrated by the examples in Section II. An interesting idea put forward in the design of Sparky [22] is that the progression from a nonrealistic to a realistic representation of a living thing is nonlinear, reaching an "uncanny valley" when similarity becomes almost, but not quite perfect. A caricaturized representation of a face can be more acceptable and believable to humans than a realistic one, which can present distracting elements that render emotion recognition more difficult and where subtle imperfections can be much more disturbing. This idea underlies the design of Feelix.

A. Facial Expression of Emotions

Each of Feelix's emotional states has an associated distinctive prototypical facial expression. To define the "primitives" for each emotion (positions of lips and eyebrows) we have adopted the features concerning eyebrows and lips usually found in the literature, which can be described in terms of action units (AUs) using the facial action coding system [9]. It must be noted, however, that the technical constraints imposed by the robot's design (see Section III) do not permit the exact reproduction of the AUs involved in all of the expressions (e.g., inner brows cannot be raised in Feelix); in those cases, we had to adopt the best possible approximation that Feelix's face allowed. Feelix's face is thus closer to a caricature than to a realistic model of a human face. Concerning an observer's perception of emotional expressions, we have adopted the hypothesis (proposed, for example, in [5] and [16]) that the upper and lower parts of the face function as the building blocks at the basis of emotion perception. Prototypical facial expressions displayed by Feelix[2] (Fig. 3) are as follows.

- *Anger*: raised[3] eyebrows, moderately open wide mouth with upper lip curved downwards, and straight lower lip;
- *Sadness*: maximally lowered eyebrows, closed mouth curved downwards;
- *Fear*: lowered outer brows, moderately open wide mouth;
- *Happiness*: straight eyebrows, mouth curved upwards;
- *Surprise*: highly raised eyebrows, very open narrow mouth.

Although it is possible to combine two different expressions in Feelix's face, we have, for now, adopted a winner-takes–all strategy[4] based on the level of emotion activation to select and display the emotional state of the robot.

B. Emotion Activation

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 by Izard under the categories of neurochemical, sensorimotor, motivational, and cognitive [15]. Some of these elicitors are emotion-specific, but emotions also show a certain degree of generality [25], e.g., of object and time. This accounts for the fact that a person can experience the same emotion under different circumstances and with different objects, but if emotions show this generality, what accounts for the activation of different affects? Activation theories that only take into account the arousal and valence properties of emotions are not able to fully account for their differential activation. Unlike previous work that used specific stimuli to elicit emotions [2], we have adopted here the generic model postulated by Tomkins [25], which proposes three variants of a single principle.

1) A sudden increase in the level of stimulation can activate both positive (e.g., interest) and negative (e.g., startle, fear) emotions.
2) A sustained high level of stimulation (overstimulation) activates negative emotions such as distress or anger.
3) A sudden stimulation decrease following a high stimulation level only activates positive emotions such as joy.

We have complemented Tomkins' model with two more principles drawn from a homeostatic regulation approach, to cover two cases that the original model did not account for.

1) A low stimulation level sustained over time produces negative emotions such as sadness (understimulation).
2) A moderate stimulation level produces positive emotions such as happiness (well-being).

This model proved particularly appropriate for Feelix for two main reasons. First, the coarse sensory capabilities of the robot considerably limit the number of different types of stimuli it can recognize; the use of particular stimuli or of an appraisal-based model (see, for instance, [4]) to elicit different emotions was thus unrealistic. Second, the generality of the model makes it particularly suited for implementation using different sensory modalities, allowing us to assess the significance and adequacy of each modality in human–robot interactions.

Feelix's emotions are activated by tactile stimulation on the feet. They are assigned different intensities calculated on the grounds of stimulation patterns designed on the above principles. To distinguish between different kinds of stimuli using only binary touch sensors, we use *duration* and *frequency* of presses. For duration, we defined three types of stimuli:

1) short (less than 0.4 s);
2) long (up to 5 s);
3) or very long (over 5 s).

Frequency (the rate at which presses follow each other) is calculated on the basis of a minimal time unit that we call *chunk*, defined here to last 2 s. When a chunk ends, information about stimuli—their number and type—is analyzed and the different emotions are assigned intensity

---

[2] A video can be seen at http://www.daimi.au.dk/~chili/feelix/feelix_home.htm.

[3] When we talk about raised or lowered eyebrows, it is in fact their external ends that are raised or lowered since the internal end of each eyebrow is attached to a shaft, and therefore, eyebrows can only rotate on that axis;

[4] 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.

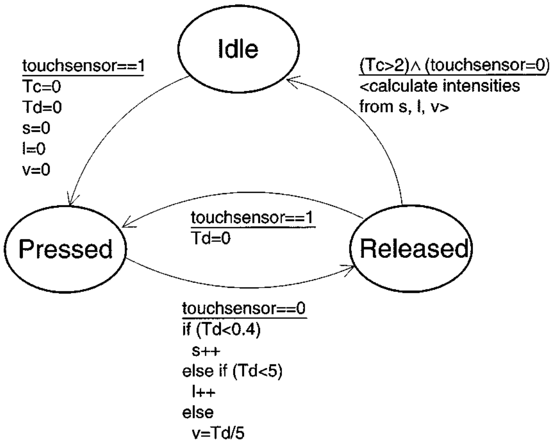

Fig. 4. Finite state machine used to implement emotion activation. T_c and T_d are timers measuring time chunks and duration of current press (in seconds), respectively. A third timer, not depicted here, measures 5-s periods. Variables s and l count the number of short and long presses, and v counts the number of 5-s periods a very long press has lasted. Based on the status of the binary touch sensor, the system is in one of three states—idle, pressed, or released. Associated with each transition is a condition. If the condition holds, the transition occurs and the associated code is executed.

levels according to the various stimulation patterns in our emotion activation model. We mapped the general stimulation patterns 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. Tomkins' model also incorporates a pattern for happiness in the sense of relief, i.e., the cessation of a too high, overwhelming stimulation level. This pattern has not been implemented in Feelix yet.
- Sadness is produced by a *sustained low level* of stimulation. In Feelix this corresponds to lack of (or very little) interaction for a long period.

This model of emotion activation is implemented by means of a timed finite state machine (FSM), as shown in Fig. 4, with three states:

1) idle;
2) pressed;
3) released.

Three timers are used to measure the chunks, the duration of a stimulus, and 5-s periods (the minimal length of a very long stimulus), respectively. During each chunk the FSM will circle between the pressed and released states, counting the number of short and long stimuli. The duration of a very long stimulus is calculated across chunks, by counting the number of 5-s periods that the stimulus lasted. The FSM returns to the idle state and awaits the next stimulus either at the end of a chunk (i.e., after 2 s) or when a stimulus longer than 2 s ends. Upon returning, the intensities of the different emotions are updated according to the number of short, long, and (when appropriate) very long stimuli.

The intensity levels of all emotions are thus recalculated every 2 s (every chunk) or, when appropriate, after a very long press ends. 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.[5] In that case, a message encoding Feelix's current emotional state and its intensity is sent to the RCX controlling the face, so that the emotional expression can be updated if necessary. It thus takes at least 2 s (one chunk) for Feelix to change its emotional state and expression, or for the intensity of the current emotional state to vary—increase or decrease. The amount of stimulation required to change Feelix's 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. When a new emotion becomes active, it temporarily inhibits all the other emotions by resetting their intensities to zero.

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 [7]. The intensity of the active emotion increases with appropriate stimulation depending on how long this emotion has been active. Intensity increases fast within an initial period that we have set to last five chunks (10 s) after the onset of the emotion, until it reaches a high level; the increase is then close to zero for a period of 20 chunks (40 s); the increase is negative thereafter until the intensity drops below the activation threshold. An emotion of which the intensity has just dropped below the threshold will therefore be easily reactivated if the same type of stimulation persists. This reflects the fact that emotions have a limited, short duration[6] (from a few seconds to few minutes in humans), although they can be repeatedly reactivated. 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.

V. EXPERIMENTS

We have investigated two aspects of emotions in Feelix: the recognizability of its facial expressions and the suitability of the interaction patterns. Emotion recognition tests are based on subjects' judgments of emotions expressed by faces, both in movement (the robot's face) and still (pictures of humans). Interactions assessed in informal observations are reported in [3].

The design of experiments to test recognition of facial expressions of emotion is a rather problematic issue involving many methodological problems (see [26] for an overview). Many factors can bias the experiments, such as the fact of using still pictures of faces that only capture a snapshot of the expression rather than its development over time, and that can also reflect emotion blends rather than "pure" emotional expressions. The use of linguistic labels to force choices is another source of problems, although this method is usually preferred to free recognition without any guidance. In an attempt to avoid undesired biases as much as possible, we designed three different tests where we varied these parameters. Tests were performed in the order in which they are

---

[5] By setting this threshold higher or lower, we can make Feelix more "extroverted" or more "introverted."

[6] Affective states with a long duration are called *moods*. They are much less intense and have different elicitors and much weaker associated (behavioral, physiological, etc.) manifestations; they make more likely the onset of specific emotions. Emotional states with an intensity below the activation threshold can be seen as moods in Feelix.

TABLE I
RESULTS OF EMOTION RECOGNITION BY CHILDREN

Children (ages 9-10)	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 II
RESULTS OF EMOTION RECOGNITION BY ADULTS

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%

described. 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, and 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 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 expressing anger, sadness, happiness, fear, contempt, surprise, and disgust. For this, we used standard pictures common in a recognition test that we took from [19].

We have conducted experiments on 86 subjects—41 children (ages 9 to 10) and 45 adults (ages 15 to 57). Experiments were performed in four suites. Due to time constraints, only one group of 14 children could do the free test on human faces. The other two tests (on the robot's face) were performed by all subjects. Answers were considered to be correct when the subjects used the same descriptors we have employed or very close synonyms. Results are summarized in Table I (experiments with children) and Table II (experiments with adults). Average recognition of emotional expressions[7] was 58% for adults and 64% for children in the free test on Feelix's face, 55% for adults and 48% for children in the multiple-choice test on Feelix's 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's 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. Valence recognition was very high (close to 100%) in all cases except for the always controversial case of surprise, which was attributed a negative valence by about two thirds of the subjects, and a positive one by the rest. 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 be crucial factors to assess observed emotion [12].

The results we obtained are very congruent with those commonly reported in the emotion literature on recognition of facial expressions of basic emotions in crosscultural studies (see [8] for an overview), in particular in the free tests. For example, they show that the "core" basic emotions of anger, happiness, and sadness are most easily recognized, whereas fear was mostly interpreted as surprise, anxiety, or sadness. This matches findings from studies on human faces (see [8]) that fear and surprise are often confused, even indistinguishable to some preliterate cultures and that surprise is perceived differently from other emotions, not defining an exclusive category and devoid of a clear valence.

Our figures are, however, lower than those reported in the psychology literature, not only in the case of emotion recognition from Feelix's face, but also in the test with human pictures. This might be partly explained, in the case of the robot experiments, by the coarseness and limited number of expressive features, and in the case of the free tests, by our very strict criterion when grouping synonyms under emotion labels, e.g., we separated "anxiety" and "fear," often grouped in crosscultural studies.

VI. CONCLUSION

We have presented early work on Feelix: a humanoid-looking LEGO robot capable of displaying several emotional expressions in response to direct physical contact. 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 as proof–of–concept that these theories can be used within a synthetic approach that complements the analytic perspective for which they were conceived. We do not claim, however, that our work provides evidence regarding the scientific validity of these theories, as this is out of our scope.

We have conducted experiments to assess how well humans can recognize emotional expressions in Feelix's 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 latter result also confirms studies on emotion recognition from pictures of human faces, and we believe it might be due to structural similarities among those emotional expressions, i.e., shared AUs, and/or to the need of additional expressive features.

A question that naturally arises is whether Feelix's expressions would have been equally understandable using a componential (dimensional) perspective. While we do not intend to approach this question from the perspective of psychology studies of emotion recognition, it would be very interesting to investigate this issue in our robot from the perspective of human–robot interaction—in particular, the meaning attributed to different expressive units and their roles in the emotional expressions in which they appear.

Finally, in order to better assess the adequacy of the emotion activation model, a formal analysis of observed interactions and other tests must be performed. An implementation of the activation model using

---

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

other sensory modalities, e.g., sound, would also allow us to assess its generality.

ACKNOWLEDGMENT

The authors wish to thank H. Lund, O. Caprani, B. Mayoh, and the students at the LEGO Lab as well as two anonymous reviewers for their suggestions to improve an earlier version of this paper.

REFERENCES

[1] C. Breazeal, "Sociable machines: Expressive social exchange between humans and robots," Ph.D. dissertation, Dept. Elect. Eng. Comput. Sci. Mass. Inst. Technol., Cambridge, 2000.

[2] L. D. Cañamero, "Modeling motivations and emotions as a basis for intelligent behavior," in *Proc First Int. Conf. Autonom. Agents*, W. L. Johnson, Ed., New York, 1997, pp. 148–155.

[3] L. D. Cañamero and J. Fredslund, "How does it feel? Emotional interaction with a humanoid LEGO robot," in *Socially Intelligent Agents. Papers AAAI Fall Symp.*, Menlo Park, CA, 2000, Tech. Rep. FS-00-04, pp. 23–28.

[4] C. Castelfranchi, "Affective appraisal versus cognitive evaluation in social emotions and interactions," in *Affect in Interactions: Toward a New Generation of Interfaces*, A. Paiva, Ed. Heidelberg: Springer-Verlag, 2000, pp. 64–75.

[5] M. De Bonis *et al.*, "A two-process theory of facial perception of emotions," *C.R. Acad. Sci. Paris, Life Sci.*, vol. 322, pp. 669–675, 1999.

[6] P. Ekman, "An argument for basic emotions," *Cognit. Emotion*, vol. 6, no. 3/4, pp. 169–200, 1992.

[7] ——, "Basic emotions," in *Handbook of Cognition and Emotion*, T. Dalgleish and M. Power, Eds. New York: Wiley, 1999, pp. 45–60.

[8] ——, "Facial expressions," in *Handbook of Cognition and Emotion*, T. Dalgleish and M. Power, Eds. New York: Wiley, 1999, pp. 301–320.

[9] P. Ekman and W. V. Friesen, *Facial Action Coding System*. Palo Alto, CA: Consulting Psychol., 1976.

[10] P. Ekman, W. V. Friesen, and P. Ellsworth, "What emotion categories or dimensions can observers judge from facial behavior?," in *Emotion in the Human Face*, 2nd ed, P. Ekman, Ed. Cambridge, U.K.: Cambridge Univ. Press, 1982, pp. 39–55.

[11] C. Elliot, "I picked up catapia and other stories: A multimodal approach to expressivity for emotionally intelligent agents," in *Proc. First Int. Conf. Autonom. Agents*, W. L. Johnson, Ed., New York, 1997, pp. 451–457.

[12] J. M. Fernández-Dols and M. A. Ruiz-Belda, "Is the meaning perceived in facial expressions independent of its context?," in *The Psychology of Facial Expression*, J. A. Russell and J. M. Fernández-Dols, Eds. Cambridge, U.K.: Cambridge Univ. Press, 1999, pp. 275–294.

[13] J. Fredslund. (1998) Elektra. Unpublished Tech. Rep., Dept. Comput. Sci., Univ. Aarhus,, Aarhus, Denmark. [Online]. Available: http://www.daimi.au.dk/~chili/elektra.html.

[14] D. Kirsch, "The Affective Tigger: A study on the construction of an emotionally reactive toy," M.S. dissertation, Prog. Media Arts Sci., Mass. Inst. Technol., Cambridge, 1999.

[15] C. E. Izard, "Four systems for emotion activation: Cognitive and noncognitive processes," *Psychol. Rev.*, vol. 100, no. 1, pp. 68–90, 1993.

[16] D. W. Massaro, *Perceiving Talking Faces. From Speech Perception to a Behavioral Principle*. Cambridge, MA: MIT Press, 1998.

[17] A. Ortony, G. L. Clore, and A. Collins, *The Cognitive Structure of Emotions*. Cambridge, U.K.: Cambridge Univ. Press, 1988.

[18] R. Plutchik, "Emotions: A general psychoevolutionary theory," in *Approaches to Emotion*, K. Scherer and P. Ekman, Eds. Hillsboro, NJ: L. Erlbaum, 1984, pp. 197–219.

[19] M.R. Rosenzweig, A.L. Leiman, and S.M. Breedlove, *Biological Psychology. An Introduction to Behavioral, Cognitive, and Clinical Neuroscience, 2nd ed.*. Sunderland, MA: Sinauer, 1999.

[20] B. Reeves and C. Nass, *The Media Equation: How People Treat Computers, Television, and New Media Like Real People and Places*. Cambridge, U.K.: Cambridge Univ. Press, 1996.

[21] J. A. Russell, "Reading emotions from and into faces: Resurrecting a dimensional-contextual perspective," in *The Psychology of Facial Expression*, J. A. Russell and J. M. Fernández-Dols, Eds. Cambridge, U.K.: Cambridge Univ. Press, 1997, pp. 295–320.

[22] M. Scheeff *et al.*, "Experiences with Sparky, a social robot," in *Proc. Workshop Interactive Robot. Entertainment (WIRE)*, Pittsburgh, PA, Apr. 1–May 30 2000.

[23] C. A. Smith and H. S. Scott, "A componential approach to the meaning of facial expressions," in *The Psychology of Facial Expression*, J. A. Russell and J. M. Fernández-Dols, Eds: Cambridge Univ. Press, 1997, pp. 229–254.

[24] S. Thrun, "Spontaneous, short-term interaction with mobile robots in public places," in *Proc. IEEE Int. Conf. Robot. Automat.*, Detroit, MI, May 10–15.

[25] S. S. Tomkins, "Affect theory," in *Approaches to Emotion*, K. R. Scherer and P. Ekman, Eds. Hillsdale, NJ: L. Erlbaum, 1984, pp. 163–195.

[26] H. L. Wagner, "Methods for the study of facial behavior," in *The Psychology of Facial Expression*, J. A. Russell and J. M. Fernández-Dols, Eds. Cambridge, U.K.: Cambridge Univ. Press, 1997, pp. 31–54.

Diminishing Returns of Engineering Effort in Telerobotic Systems

Myra Wilson and Mark Neal

***Abstract*—Robotic systems range from teleoperated (where the human operator is in full control of all aspects of the robot's behavior) to fully autonomous (where no human intervention takes place). The word "telerobotic" describes robotic systems which, although guided by a human, have a degree of autonomous behavior.**

This paper examines the tradeoff between the increasing design and implementation effort necessary as the system moves through the continuum from teleoperated to autonomous and the amount of human intervention required. A case study of a human "shepherd" interacting with a robotic "sheepdog" which directs a robotic "sheep" is used.

***Index Terms*—Appropriate behavior, autonomy, engineering effort, teleoperation, telerobotic.**

I. INTRODUCTION

Telerobotics involves the remote interaction of an operator with a robot which has the capacity for a degree of autonomous action. The repertoire of appropriate autonomous behavior determines the amount of interaction required. There is a tradeoff between increasing design and implementation effort, and the degree of user interaction.

The work described in this paper examines this tradeoff using a case study inspired by a biological "telerobotic" task, that of a human shepherd herding a robotic sheep using a robotic dog. The implementation presented is well controlled and repeatable and allows a wide range of dog and shepherd interaction to be examined using an appropriate environment. Two sets of four robotic sheepdogs with different behavior repertoires ranging from purely teleoperated to those displaying a reasonable degree of autonomy are tested under a variety of experimental conditions.

A variety of sheep which exhibit behavior have been designed to elicit a significant degree of interaction between the dog and the shepherd. This interaction takes the form of a small set of overriding commands which are analogous to those used when interacting with a biological sheepdog. The interaction between the dog and the sheep also has a basis in the biological equivalent. The sheep moves away from the

---

Manuscript received December 18, 2000; revised April 23, 2001.

The authors are with the Department of Computer Science, University of Wales, Aberystwyth, U.K. (e-mail: mxw@aber.ac.uk).

Publisher Item Identifier S 1083-4427(01)07731-1.