

# Face to InterFace: Facial Affect in (Hu)Man and Machine

Diane J. Schiano, Sheryl M. Ehrlich, Krisnawan Rahardja & Kyle Sheridan

Interval Research Corporation

1801 Page Mill Road, Palo Alto, CA 94304

+ 1 650 842 6099

{schiano, ehrlich, rahardja, sheridan}@interval.com

## ABSTRACT

Facial expression of emotion (or "facial affect") is rapidly becoming an area of intense interest in the computer science and interaction design communities. Ironically, this interest comes at a time when the classic findings on perception of human facial affect are being challenged in the psychological research literature, largely on methodological grounds. This paper presents two studies on perception of facial affect. Experiment 1 provides new data on the recognition of human facial expressions, using experimental methods and analyses designed to systematically address the criticisms and help resolve this controversy. Experiment 2 is a user study on affect in a prototype robot face; the results are compared to the human data of Experiment 1. Together they provide a demonstration of how basic and more applied research can mutually contribute to this rapidly developing field.

## Keywords

Affective computing, facial affect, facial expression of emotion, affect, emotion, face, nonverbal communication.

## INTRODUCTION

Emotion (or "affect") is central to human experience, and facial expressions are our primary means of communicating emotion. Face-to-face communication is inherently natural and social for human-human interactions, and substantial evidence suggests this may also be true for human-computer interactions. That is, people appear to regard computers as social agents with whom "face-to-interface" interaction may be most easy and efficacious [10, 12]. In addition, human (or human-like) faces have been found to provide natural and compelling computer interfaces [e.g., 7, 9, 18]. These findings, together with advances in display and recognition technologies, have produced a surge of interest in facial affect by researchers in human-computer interaction (HCI) and artificial intelligence (AI) alike. Indeed, the burgeoning new field of "affective computing" focuses on computational modeling of human perception and display of emotion, and on

the design of affect-based computer interfaces [9, 13].

The question of how to best characterize perception of facial expressions has clearly become an important concern for many researchers in affective computing. Ironically, this growing applied interest is coming at a time when the established wisdom on human facial affect is being strongly challenged in the basic research literature. In particular, recent methodological criticisms have thrown suspicion on a large body of long-accepted data.

The classic psychological research on facial expression of emotion was performed by psychologist Paul Ekman and colleagues, beginning in the 1960s [see 5 for a review]. A substantial body of evidence has been gathered in over three decades, identifying a small number of so-called "basic" emotions: anger, disgust, fear, happiness, sadness and surprise (contempt was tentatively added only recently). In Ekman's theory, the basic emotions are considered to be the building blocks of more complex feeling states [5]. Ekman's data showed that each of these emotions was recognized cross-culturally with substantial consensus among study participants. Ekman and Friesen [4] developed the "facial action coding system" (FACS), a method for quantifying facial movement in terms of component muscle actions. The FACS is a highly complex coding system which requires extensive training to use appropriately. Recently automated [1], the FACS remains the single most comprehensive and commonly accepted method for measuring emotion from the visual observation of faces.

In the past few years, psychologist James Russell and colleagues [14, 15] have strongly challenged the classic data, largely on methodological grounds. Russell argues that emotion in general (and facial expression of emotion in particular) can be best characterized in terms of a multi-dimensional affect space, rather than discrete emotion categories (such as Ekman's "fear" or "happiness"). More specifically, Russell claims that two dimensions--"pleasure" (or "valence") and "arousal"--are sufficient to characterize facial affect space [13, 15]. He calls for new research on perception of facial affect using improved methods and multi-dimensional analyses [15]. The results of such a research program could have profound implications for the psychological understanding of facial affect and direct

implications for the field of affective computing. For example, a two-dimensional (2D) characterization of emotion related to Russell's model is assumed in many studies on affective computing, especially those employing physiological measures of arousal [see 11]. On the other hand, many AI models of facial affect recognition rely upon the FACS and Ekman's classic data as the means by which to assess performance [see 9].

This paper presents two studies on perception of facial affect. Experiment 1 provides new data on the recognition of human facial expressions, using experimental methods designed to systematically address Russell's criticisms and help resolve this controversy. Experiment 2 is a user study on affect in a prototype robot face; the results are compared to the human data of Experiment 1. Taken together, they provide a demonstration of how basic and more applied research can mutually contribute to this rapidly developing field.

### EXPERIMENT 1: HUMAN FACIAL AFFECT

Russell attacks the methods used in the classic studies on facial affect on various grounds. He points out that a great deal of the data from a large number of studies was generated using a single corpus of (fairly unnatural) stimuli. Russell also indicates certain experimental design flaws in the previous research (e.g., failure to properly randomize stimuli, biasing procedures for practice trials, small numbers of trials) and makes an argument against the common reliance on "within-subject" designs. Russell's primary criticism, however, concerns response format. The vast majority of previous studies employed the forced-choice response format, in which participants were presented with a list of emotion labels and were asked to pick the one that best matches the expression on the stimulus image. Russell's critique of the forced-choice format is that: 1) it tends to be implemented in an all-or-none fashion, which is insensitive to perceived differences in intensity of emotion, and 2) if participants were free to pick multiple responses or to provide their own emotion terms, the results might well look very different.

Our laboratory has recently begun to focus its efforts on generating a credible new body of data on the perception of human facial affect, using improved experimental methods and a variety of stimulus and viewing conditions. An additional goal is to empirically assess the Ekman-Russell controversy: Will the methodological improvements suggested by Russell make the data on perception of human facial affect differ fundamentally from Ekman's classic findings or not? In two previous studies [see 3], we asked independent groups of participants to judge images of the six basic emotions, using "multiple-choice" (i.e., more than one emotion label could be chosen from the response list) and "open-ended" (i.e., participants freely provided the emotion label) response formats, as suggested by Russell. In Experiment 1 of this paper, we complete the response-

format assessment by performing a final study using the forced-choice format. All three studies were identical except for the response format manipulation, and they incorporated various methodological improvements over previous research. Several experimental design flaws were eliminated and appropriate techniques for stimulus randomization and presentation were used. We constructed a new (more naturalistic) stimulus set rather than re-using the standard corpus. A rating scale was provided with each forced-choice response so that degrees of perceived emotional intensity could be indicated. Finally, the response format studies comprise a "between-subjects" manipulation, which permits direct comparison of results to the classic data, and additional independent analyses.


### Method

#### *Participants*

Eighteen Stanford University students between the ages of 18 and 35 participated in this study. All were naïve to the purposes of the study.

#### *Stimuli*

Four drama students (2 female, 2 male) produced the facial expression stimuli. To promote naturalness, the actors were briefly familiarized with the facial expressions of interest, and then simply instructed to imagine a time when each emotion was felt strongly [see 3, 5]. Each actor provided a total of 14 different front-view exemplars of the six basic emotions (anger, disgust, fear, happiness, sadness, surprise). Figure 1 provides an exemplar of each emotion (in alphabetical order), and also shows each of the actors' faces. Each participant viewed all the stimuli created by two (randomly chosen) actors. The high-resolution digital images were shown on a 13" Panasonic color TV monitor connected to a PowerMac computer.



**Figure 1: Sample stimulus images for each emotion in Expt. 1: Anger, disgust, fear, happiness, sadness and surprise.**

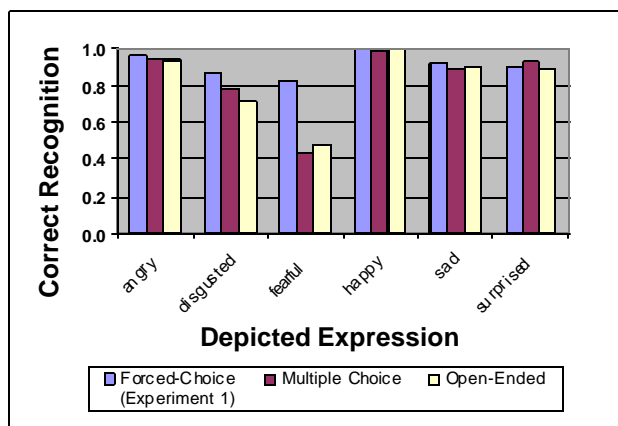
#### *Procedure*

Participants viewed stimulus images depicting emotional expressions and responded with an emotion label for each image, using the forced-choice (plus intensity ratings) response format. The images were viewed in random order, at a distance of about 30 inches. In each trial, the participant was shown the stimulus image on the TV monitor and simultaneously viewed a response screen on the computer

monitor. The response format was implemented in the following manner: An alphabetized list of labels for the basic emotions was displayed on the computer monitor. Participants chose the one label that best corresponded to the depicted emotion in each image, and then rated the degree to which that emotion was present in the image on a scale of 0 (not at all) to 6 (extremely high). The rating scale appeared as radio buttons adjacent to the selected emotion label. Ten initial (non-feedback) practice trials used randomly selected images from actors not viewed during the test trials. The experimental protocol was implemented in HyperCard. Participants proceeded at their own pace; the entire procedure lasted under one hour.

## Results and Discussion

Figure 2 presents the "correct recognition" scores (i.e., the proportion of trials on which the participant responded with the expression portrayed by the actor) for each expression in Experiment 1 (forced-choice response format). These results are shown in context with comparable data from the alternative response format studies (in which the highest rated response was used). Correct recognition scores are the standard form of data presentation in the classic studies. As the Figure shows, correct recognition in Experiment 1 was highest for happiness (mean = 99%) and lowest for fear (mean = 82%), with the remaining emotions falling in between. Performance was uniformly high. These findings closely replicate Ekman's classic results, both in terms of relative pattern and absolute levels of performance. The similarity of results is especially impressive considering the differences not only in methods but also in the stimuli used in this research. The classic stimuli were created in a painstaking fashion by highly trained actors moving specific muscle groups; ours were made in a much more natural way.



**Figure 2: Recognition scores for each emotional expression by response format. Forced-Choice format was used in Expt. 1.**


The rating scale data serve as a manipulation check in this study, ensuring that the recognized emotion was in fact seen as present in the stimulus image to at least a moderate degree. If Russell's critique of the forced-choice method is correct and the recognition scores are inflated due to constrained response options, extremely low ratings might be expected for at least some expressions. However, mean ratings for the emotional expressions was 3.93 (out of 6) overall; the mean ratings did not fall below the moderate level for any of the expressions. This suggests that, overall, the participants did see the depicted emotion in the images to at least a moderate degree.

A brief comparison across the datasets for the forced-choice (Experiment 1), multiple-choice and open-ended response format studies is illuminating. (In the multiple-choice study, participants could respond with--and rate--more than one emotion label on the response list, or could press "other" and type in a preferred term. In the open-ended study, participants were given an open text window and simply asked to type in their own responses [see 3]. Contrary to Russell's predictions, the results for the three response formats show a strikingly similar pattern. Correct recognition was generally quite high, highest for happiness and lowest for fear. The one exception to this rule is found in the scores for fear using the multiple-choice and open-ended formats. When the alternative response formats were used, fear was often "misrecognized" as surprise (mean = 25%) or sadness (mean = 17%). This confusion pattern is consistent with similarities in the FACS codes for these emotions. Why fear alone should show such a performance decrement with response format is not quite clear. Fear may be the least compelling emotion under posed conditions, and because it is one of the more ambiguous expressions in terms of FACS code overlap, when observers are encouraged to give multiple responses, they may tend to do so more for fear. Still, even the "misrecognition" of fear was highly systematic, not simply showing greater variability. Taken together, the results of these studies generally support Ekman's classic findings. (The results of an extensive series of additional analyses of the data from these studies are available upon request.)

To perform multi-dimensional scaling (MDS) analyses on the data from Experiment 1, a confusion matrix was generated from the number of times each emotion was mistakenly recognized as any other emotion. This was used to create the similarity space for the analysis. In addition, we asked a trained FACS coder to independently create a FACS confusability index for the basic emotions, based on the degree of overlap of FACS codes between all pairs of emotional expressions (number of overlapping codes over total number of codes for the two emotions, x 100%). The specific expressions were those of the "Directed Facial Action" (DFA) task [8]. While this confusability index is a fairly crude measure (e.g., it does not permit differential

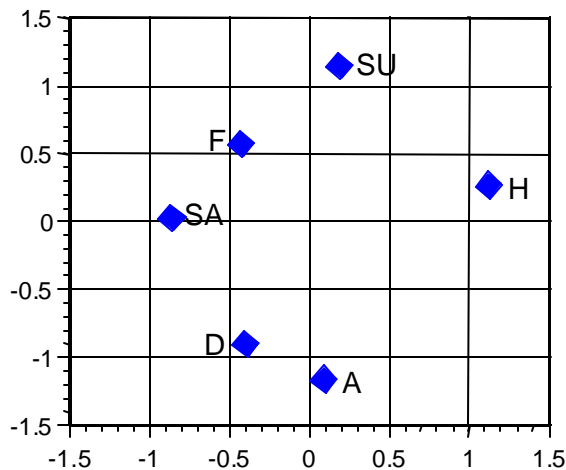
weighting of codes), such quantification permits the specification of a FACS prediction pattern that can serve as a benchmark for comparison with our data.

Russell's model of affect space is a "circumplex" about two axes, identified as pleasure and arousal. A schematization of this model (using Russell's terms) is shown in Figure 3.



**Figure 3: Russell's circumplex model of affect space.**

The results of the MDS analysis showed that a two-dimensional (2D) solution accounted for 85% of the variance of the Experiment 1 dataset (and 76% of the FACS dataset). Figure 4 shows the 2D solution for the Experiment 1 dataset (the FACS pattern is highly similar). At first glance, this 2D space looks similar to Russell's model. The datapoints show a roughly circular arrangement. Since rigid rotations of the dataset are permissible in MDS, rotating our results for optimal fit with Russell's schematization does find a fair degree of overlap.




**Figure 4: 2D MDS solution for human (Expt. 1) data**

However, the ordering of emotions on the circle does not quite match; in particular, the relative ordering of anger and sadness--and the question of where to put disgust--is

problematic. The relative ordering of the datapoints is the primary result of the 2D MDS solution, and determines the interpretation of the orthogonal dimensions. Our space could perhaps be interpreted as showing an axis corresponding to pleasure, but identifying an arousal dimension is less plausible.

A 3D MDS solution accounted for 96% of the variance of the Experiment 1 data and 90% of the FACS dataset, a substantial improvement over the 2D approach. Figure 5 presents the 3D solutions for both datasets. A strikingly similar pattern is shown, which Russell's 2D model cannot match. Similar MDS analyses on the multiple-choice and open-ended format data show generally similar results. The fact that the FACS codes are based on physical facial features and movements suggests that the facial affect space's dimensions may correspond more to physical or image parameters than to feeling states (or at least to salient physical parameters associated with feeling states). However, the FACS codes are extremely complex, and at this point it is difficult to speculate on exactly what the dimensions may refer to. Further research is needed. One approach would be to explore facial affect under simplified conditions, which may shed some light on the most salient cues for emotional expression. This is the approach taken in the study described below.



**Figure 5: 3D MDS solution for human (Expt 1) & FACS data**

**EXPERIMENT 2: ROBOT FACIAL AFFECT**

Simplified or schematic faces are used to express emotion in many forms of interface, from icons to virtual pets. In a pilot study on icon design, we found that manipulations of an extremely simple facial feature set--which were not intended to express emotion--were almost universally interpreted in

emotional terms. Recent work elsewhere shows that cartoon-like facial icons are sufficient to serve as an affective interface for interactions with avatars and autonomous computer agents [e.g., 7]. The present study was performed in the context of user-testing an early design prototype of a mechanical robot face. While made of metal, the face was intended to be cartoon-like in nature. Created by an independent group of researchers at Interval Research [Tow & Scheeff, personal communication; see 17], it was designed largely through intuition inspired by cartoon animation principles [6] and some reading of Darwin [2] on emotions. This initial prototype was constructed largely as a "proof of concept" to demonstrate that a very simple robot face could effectively express affect. A similar face would be incorporated into a later, more complete prototype, capable of whole-body movement and expression.

The initial prototype consisted of a box-like face containing eyes with moveable lids, tilting eyebrows, and an upper and lower lip which could be independently raised or lowered from the center. Figure 6 shows the face displaying various emotional expressions. Compared to the highly detailed human faces of Figure 1, the robot facial features are extremely sparse and their motion is highly constrained. The face has no skin, so the telltale folds, lines and wrinkles specifying many FACS codes are simply not available. And the motion of the features (especially the eyebrows and lips) is only schematically related to human facial muscle movements. Experiment 2 was conducted in collaboration with the robot design team. They primarily wanted some assurance that users felt satisfied with the robot's ability to express a range of basic emotions. Another goal was to obtain feature settings for various emotional displays, to be stored as templates so that (at least in principle) the prototype could be set up to quickly display appropriate emotional responses as needed. We were especially interested in comparing the robot "affect space" with our human data, after verifying that the displays were indeed correctly recognized by an independent group of observers.

## METHOD


### Participants

Eighteen participants between the ages of 18 and 35 were involved in this study. Nine Interval Research Corporation employees participated in the first condition of the study, and 9 Stanford University students in the second condition. All were naïve to the purposes of the study.

### Materials

The robot face consisted of a 12 cm x 14 cm mechanical metal (primarily aluminum) face with independently moveable eyelids, eyebrows, upper and lower lips (see Figure 6). The eyelids were small metal sheets that could move up or down. The eyebrows were metal bars, placed with a pivot point towards each side of the face, to allow

rotations between horizontal and vertical positions. Each lip consisted of a spring fixed at both ends and with a tie in the center that could be pulled up or down (stretching the spring on both sides). Each feature was controlled by a computerized motor with 255 possible positions. For all features except the eyelids, the "neutral" position was in the center of the range of motion. The neutral position for the eyelids was fully open (or up).



**Figure 6: "Average" robot face display for each emotion in Expt 2: Anger, disgust, fear, happiness, sadness, surprise.**

### Procedure

*Condition 1: Feature Setting Adjustments.* In each of two blocks of trials, instructions on the computer monitor asked the participant to set the features of the robot face to express each of the 6 basic emotions (angry, disgusted, fearful, happy, sad, or surprised) at each of 3 degrees of intensity (slightly, moderately, or very), twice, in random order. Participants adjusted feature settings by pressing keys on a computer keyboard; "up" and "down" keys were labeled on the keyboard for each of the 4 features. Upon completion of each expression, participants rated their overall satisfaction with the expression on a scale of 1 (least) to 5 (most). Each trial began with the features in the neutral position, except for the eyelids, which were closed. Each testing session began with 10 randomly chosen (non-feedback) practice trials. Participants proceeded at their own pace, and the entire procedure took less than one hour. The robot face was attached to a Toshiba laptop PC, which implemented the experimental protocol.

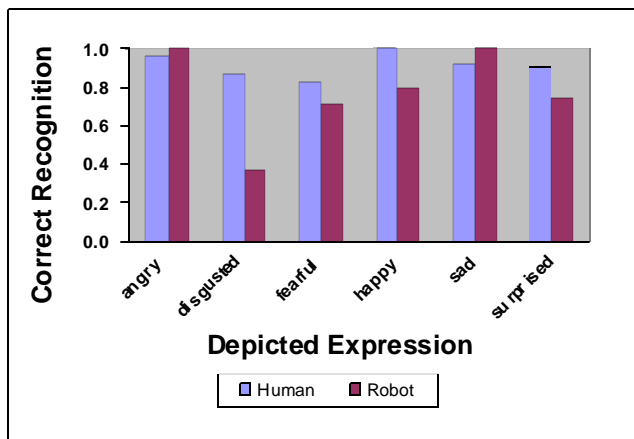
*Condition 2: Recognition Validation.* Participants viewed the robot facial expressions obtained from Condition 1. The expressions were given by the mean feature settings for each emotion (angry, disgusted, fearful, happy, sad, or surprised) at each of the three degrees of intensity (slightly, moderately, or very), plus the "average" setting. The average setting for each emotion is depicted in Figure 6. In all, 4 exemplars of each emotion were shown, 3 times, in random order. For each trial, participants used the forced-choice (plus ratings) response method of Experiment 1, choosing the one term from an alphabetized list of basic

emotion labels that best described the expression of the robot face. They then rated the degree to which the emotion was present in the robot face on a scale of 0 (not at all) to 6 (extremely high). In this study, the robot facial expressions were controlled by a Toshiba laptop PC while the rest of the experimental protocol was implemented in HyperCard on a PowerMac computer. Participants proceeded at their own pace; the entire procedure took about 30 minutes.

## RESULTS AND DISCUSSION

The results for Condition 1 are given in terms of numerical setting values for each feature, which are difficult to summarize succinctly except pictorially. Figure 6 illustrates the "average" display for each emotion, derived from the mean feature settings for each participant for each degree of emotion. The figure does suggest that the interface was capable of expressing various emotions, although perhaps not all equally well. While participants' satisfaction ratings were fairly high overall (3.7 out of 5, for all degrees of intensity), satisfaction with disgust, in particular, was fairly low (2.9). The FACS codes characterize disgust by the drawing up of the nasal-labial muscles, producing striking patterns of wrinkles and folds around the mouth and nose. However, the robot face has neither nose nor skin. That disgust was found to be particularly difficult to express was not especially surprising.

The results for Condition 2 are shown in Figure 7, in terms of mean correct recognition scores for each of the emotions (averaged over degree of intensity). To aid comparison, the human data from Experiment 1 are also provided in the figure. The scores for the robot are generally somewhat lower than those for human faces (especially for disgust), but this is not very surprising. First, the schematic nature of the robot face (as described above) should have made it more difficult to express emotion than human faces.

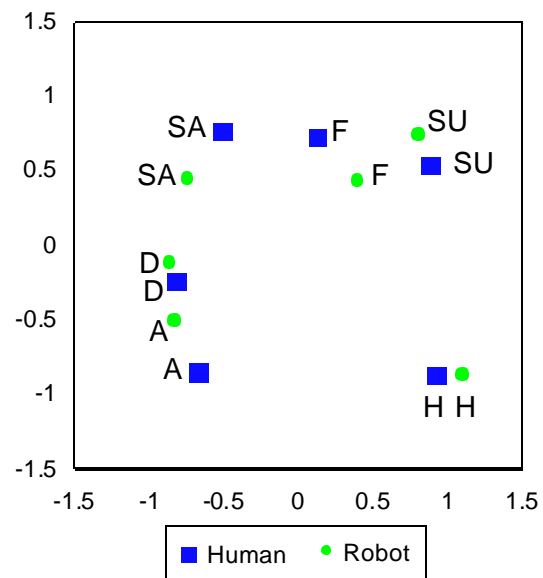


**Figure 7: Recognition scores for each emotional expression for human (Expt. 1) and robot (Expt. 2) stimuli.**

Secondly, these scores were averaged over stimuli intended to depict emotion at varying intensity levels, while the human actors presumably at least intended to create stimuli that showed each emotion to a high degree of intensity. Third, due to time constraints, our sample size was small and so the dataset is fairly variable. That the human and robot results are nonetheless so close is noteworthy.

As in Experiment 1, intensity ratings served as a manipulation check in this experiment. The mean ratings over all emotions were moderately high (mean = 4.34 out of 6). Further analyses (available upon request), found that the ratings did generally vary with the intensity of the depicted emotion. And, as expected, recognition of the emotions tended to increase with rated intensity.

MDS analyses were performed on the robot feature setting data. The mean direction and amount of movement of each of the 4 facial features for each emotion (taken from the average of all expressions) were used to generate 4D vectors; the distance between each vector gave the (dis-)similarity matrix for the MDS analysis of the robot dataset. Ninety-seven percent of the variance of the robot data was accounted for by a 2D MDS solution. Figure 8 presents the 2D solutions for the robot and human (Experiment 1) datasets. The close similarity of the patterns is immediately obvious, with ordering of emotions identical for the two datasets.



**Figure 8: 2D MDS solution for human (Expt 1) & robot data**

Ninety-nine percent of the variance of the robot dataset was accounted for by a 3D MDS solution. Figure 9 shows the 3D solutions for the human and robot data. The similarity of the patterns is remarkable when considering the disparity of the stimuli and the fact that the robot data were plotted directly from the feature setting parameters. As in the case of the human data, the robot results map easily onto the

FACS index pattern (see Figure 5) but not so easily onto Russell's model (see Figure 3). Moreover, the similarity of findings across the human, FACS and robot datasets further supports the notion that the dimensions of facial affect space might correspond most closely to physical or image parameters; indeed, to very simple ones [see also 19]. Our initial speculation is that the primary axis may correspond to concavity/convexity of the lips, the second to the upward/downward tilt of the eyebrows. The third dimension is less clear, but may be related to the set of the mouth, perhaps its degree of openness (note that many of our disgust stimuli had open mouths). Further research is clearly needed, but these results do suggest implications for many applications in which the complexity of the face is constrained or compressed. We are currently looking at human facial affect under a variety of compressed image conditions, to see if a similar affect space is found.




Figure 9: 3D MDS solution for human (Expt 1) & robot data

### GENERAL DISCUSSION

Experiment 1 provides new baseline data on human facial affect recognition, using improved experimental methods and somewhat more naturalistic stimuli than those of the classic studies. The pattern of results for the forced-choice response format closely replicate Ekman's classic findings, and (except for fear), this was generally true for the alternative response formats as well. Thus, on the whole, Russell's criticisms are not borne out by the data. Our MDS analyses suggest that 3 dimensions are substantially better than 2 in specifying facial affect space. However, even the 2D solution does not match Russell's model. Indeed, our data match the FACS-based solutions much more closely.

We find this intriguing, suggesting that the dimensions of facial affect may be based more on physical or image parameters than on feeling states (such as pleasure and arousal) per se.

Experiment 2 was performed in the context of a user test; its primary aim was to inform the designers of the affective robot face. We succeeded both in demonstrating that the face was sufficient to communicate various emotional expressions, and in providing feature setting templates for specific emotions of varying intensities. The revised prototype of the affective robot incorporates a face very similar to the one we tested. The pattern of results for this experiment was strikingly similar to our human data, despite extreme schematization of the robot face. The similarity of the MDS solutions for robot, human and FACS-based data underscore the notion that physical or image-based parameters--perhaps very simple ones--could be used to interpret the dimensions of facial affect space. Some speculations on what those parameters may be were provided above, and we note that our research on this topic is continuing. In addition to exploring human facial affect under various compression conditions (using both still and moving images), we are also collaborating with another laboratory in training a neural-net AI model on our stimuli, to see what features it picks up [see 9]. Interestingly, 3-D models of affect have been suggested before, largely based on feeling states [e.g., 16; see 15], but no consensus in axis interpretation was achieved in that earlier research. This is an exciting time for new research on facial affect in both humans and machines. We hope this paper helps demonstrate the need for both basic and applied contributions to this rapidly developing field.

### ACKNOWLEDGEMENTS

We would like to thank Diane Beck, John Pinto, Mark Scheeff, and Rob Tow for their contributions to this project.

### REFERENCES

1. Bartlett, M.S., Hager, J.C., Ekman, P. & Sejnowski, T.J. *Psychophysiology*, 36 (1999), 253-263.
2. Darwin, C. *The Expression of Emotion in Man and Animals*. Univ. of Chicago Press: Chicago, IL, 1965 (originally 1872).
3. Ehrlich, S., Schiano, D.J., Sheridan, K., Beck, D. & Pinto, J. Facing the issues: Methods matter *Psychonomic Society* (Dallas, TX, May 1998).
4. Ekman, P. and Friesen, W.V. *The Facial Action Coding System*. Consulting Psychologists Press: Palo Alto, CA, 1978.
5. Ekman, P., Friesen, W.V. & Ellsworth, P. *Emotion in the Human Face*. Pergamon, New York, 1972.
6. Hamm, J. *Cartooning the Head and Figure*. Perigree, New York, 1982.

7. Kurlander, D., Skelly, T. & Salesin, D. "Comic chat," *Computer Graphics Proceedings*, (1996), ACM Press, 225-236.
8. Levenson, R. W., Ekman, P., Friesen, W. V. Voluntary facial action generates emotion-specific autonomic nervous system activity. *Psychophysiology*, 27 (1990), 363-384.
9. Lisetti, C.L. & Schiano, D.J. Automatic facial expression interpretation: Where Human-Computer Interaction, Artificial Intelligence and Cognitive Science intersect. *Pragmatics and Cognition* (Special Issue on Facial Information Processing, a Multidisciplinary Perspective) 1999, in press.
10. Nass, C.I., Steuer, J.S. & Tauber, E. "Computers are social actors", in *Proceedings of CHI'94* (Boston MA, April 1994), 72-78.
11. Picard, R. *Affective Computing*. MIT Press: Cambridge, MA, 1997.
12. Reeves, B. & Nass, C.I. *The Media Equation: How People Treat Computers, Television, and New Media Like Real People and Places*. Cambridge University Press: New York, 1996.
13. Russell, J.A. A circumplex model of affect. *Journal of Personality and Social Psychology*, 39(6) (1980), 1161-1178.
14. Russell, J.A. Is there universal recognition of emotion from facial expression? *Psychological Bulletin*, 95 (1994), 102-141.
15. Russell, J.A and Fernandez-Dols, J.M. *The Psychology of Facial Expression*. Cambridge University Press: New York, 1997.
16. Schlosberg, H. Three dimensions of emotion. *Psychological Review*, 69 (1954), 81-88.
17. Tow, R. "Affect-based Robot Communication Methods and Systems" - U.S. Patent No 5,832,189, November, 1998.
18. Walker, J.H., Sproull, L., Subramani, R. (1994). Using a human face in an interface, in *Companion of CHI'94 Conference on Human Factors in Computing Systems* (Boston MA, April 1994), 205.
19. Yamada, H., Matsuda, T., Watari, C. & Suenaga, T. Dimensions of visual information for categorizing facial expressions of emotion. *Japanese Psychological Research*, 35( 1993), 172-181.