12th IEEE/RSJ Int. Conf. of Intelligent Robots and Systems, Vilamoura, Portugal, Oct. 2012

A Design Methodology for Expressing Emotion on Robot Faces

Mohammad Shayganfar, Charles Rich and Candace L. Sidner*

Abstract— **We have developed a systematic methodology for designing emotional facial expressions for humanoid robots, especially those with limited degrees of freedom. The methodology is firmly grounded in the psychological literature on human static and dynamic emotional facial expressions. We demonstrate the methodology by applying it to a recent humanoid robot and evaluate the results, confirming that the observed confusion matrix agrees qualitatively with the predictions of the methodology. We also investigate how robot facial emotion recognition compares for dynamic versus static expressions.**

I. INTRODUCTION

Facial expressions and, more specifically, emotional facial expressions play a key role in human interaction. They convey a person's motivational states and attitudes, which can help to improve collaboration. We therefore believe that the ability of robots to produce emotional facial expressions is similarly important for human-robot interaction. Some robot face designers, e.g., hansonrobotics.wordpress.com, have approached this problem by trying to build androids, i.e., robots whose faces duplicate human facial musculature and movements as closely as possible, which leads to very expensive mechanisms. However, most modern robots—even humanoid ones—have faces that are *much* simpler (fewer degrees of freedom) than human faces, often to the point of being schematic or cartoonish. Futhermore, because these robots vary widely in design, the approaches typically used to express emotions on their faces [1], [12], [18] have been ad hoc and highly robot-specific. In contrast, our work provides a systematic robot-independent design methodology that is firmly grounded in the psychological literature.

There is significant overlap, obviously, between the problems of designing emotional expressions for virtual agents, e.g., [13], and designing emotional expressions for robots. However, as a practical matter, since robotic faces have to be mechanically operated, their complexity lags far behind current graphical rendering techniques. Thus current work on emotional expression for virtual agents is not very useful for robots other than androids.

Most psychological research on emotional facial expression falls into two major camps. First, there are discrete (evolutionist) emotion theories, such as Ekman's [5], that focus on *static* facial patterns corresponding to each of the so-called basic emotions, such as happiness, sadness, anger, fear, surprise and disgust. There are also componential emotion models, such as Scherer's [9], [14], [21], that instead try to connect facial patterns to the results of an ongoing cognitive appraisal process, leading to a *dynamic* view of emotional expression. We will discuss these approaches in more detail in Section II.

The heart of our methodology, described in more detail in Section III, is to systematically explore, using the concept of confusion matrices, the mapping from standardized human facial action units to the degrees of freedom offered by a given robot. We demonstrate this methodology by applying it to our current robot, Melvin (see Fig. 1), constructed for us by Francois Michaud at U. Sherbrooke.

Fig. 1. Melvin

In Section IV, we describe a controlled study in which we show that our methodology qualitiatively predicts where there will be confusion in human recognition of Melvin's emotional facial expressions. Secondly, we investigate how robot facial emotion recognition compares for dynamic versus static expressions. We conclude in Section V with discussions of the implications and applications of this work.

II. BACKGROUND

The discrete (evolutionist) perspective on emotional facial expression makes five central claims: Emotional facial expressions (1) occur universally in emotionally arousing situations; (2) are linked with subjective experience; (3) are part of a coherent package of emotional responses; (4) are judged universally and discretely; and (5) have important social functions. This school of thought dates to Darwin and is recently associated most strongly with Ekman [4], [10].

For simplicity and to aid comparison with most other work in this area, we adopt Ekman's six basic emotions (happiness, sadness, fear, anger, surprise and disgust) as our starting point. However, it should be acknowledged that this notion of basic emotions is not accepted by all pyschologists [16].

A. Facial Action Coding System

The Facial Action Coding System (FACS) [7], and its later specialization EMFACS, were developed by Ekman and his colleagues as an objective method for coding facial expressions in terms of component movements (actions). In FACS, any observed facial expression is decomposed into standardized *action units* (AU's). An action unit may correspond to the activity of one or several muscles in the face [6]. Each action unit also has an associated intensity coded from A (low) to E (high).

Ekman's initial goal was to distinguish all possible visually distinguishable facial expressions without including behavioral interpretation. He then used this coding system to empirically determine which action units were universally used to express which emotions. This coding system is

*The authors are with the Department of Computer Science, Worcester Polytechnic Institute, Worcester, MA 01609, USA (mshayganfar|rich|sidner)@wpi.edu

thus an excellent basis for a systematic design methodology. FACS has been used by others [24] for designing robotic faces, but never in a systematic methodology.

B. Static Facial Expression

The leftmost two columns of TABLE I show Ekman's coding of the six basic emotions. The action units shown in the second column are the *required* action units for each emotion. There are also optional action units for each emotion that are listed in [7] and will be referenced in Section III-D. The light horizontal lines in some boxes indicate alternative codings for that emotion. The difference between bold and non-bold action units has to do with the relationship between this column and the third column of the table, which will be explained in Section III-A.

Notice that this approach is based on the analysis of static photos or schematic drawings of facial expressions at their apex of their component movements, e.g., when the eyebrows are the most raised and the mouth is the most open, etc.

C. Dynamic Facial Expression

Other psychologists [3], [8], [23] point out that it is rare for emotional judgments to take place on the basis of a face caught in a single snapshot of time. They found that in natural dynamic situations, people start responding to emotional faces *before* the component motions have reached their apex. There are also human brain studies [15] showing more brain regions are active while a person is observing dynamic as compared to static facial expressions. One of the hypotheses we evaluated in our controlled study (Section IV) is that, at least for some emotions, recognition rates will be higher for dynamic facial expressions than for static.

There have been many different approaches to analyzing dynamism in emotional facial expressions. Scherer [16], [17] has focused on the temporal order and continuity of discrete action units. Others [11] have tried to represent emotional facial expressions in a continuous multi-dimensional space. Because it is more compatible with Ekman's framework, we are essentially following Scherer's approach in this work.

Thus we are concerned with the coordinated temporal trajectories of each action unit composing a basic emotion from its neutral position to its apex position. This temporal information has two aspects. First there is the order in which the action units occur. For example, for some emotions, such as happiness, the lower face (mouth and related areas) moves before the upper face (eyes and related areas). For other emotions, such as anger, the upper face moves first [2], [19], [20]. There are also durational differences, e.g., surprise is typically held for a short time, while sadness and fear are held longer. Space does not allow us to present the complete timing details for all of the basic emotions we tested. Fig. 4 shows an example of the temporal dynamics of the transition from a neutral to a sad face on Melvin.

III. Design Methodology

A design methodology is a sequence of steps that a designer can follow to solve a problem. Each step specifies a question to ask, an artifact to build, a rule of thumb to apply, etc. Methodologies are important because they are a way of passing knowledge from experts to novices in a design field and are more efficient than ad hoc, trial and error techniques.

Our methodology for designing emotional facial expressions for a given robot has four basic steps, which we will demonstrate below by applying them to Melvin (depicted in Fig. 1). The first step involves independently mapping action units to the degrees of freedom (DOF) on the robot's face. The second step involves applying this mapping to the action units for each basic emotion to obtain an initial set of servo motor commands for each emotion (which will almost certainly contain conflicts). Step three involves analyzing the conflicts and making tradeoffs with the help of a confusion matrix. Step four involves using optional action units and other ideas (such as cartoons) to mitigate remaining problems in the confusion matrix. In our evaluation of the methodology (Section IV) we will show that the final confusion matrix does a good job of predicting the confusions that humans will have in distinguishing the robot's facial expressions.

A. Mapping Action Units to DOF

This is the first step of the methodology. The inputs to this step are the degrees of freedom for the given robot, e.g., Fig. 2 for Melvin, and the second column (titled "Ekman's EMFACS") of TABLE I, which shows all possible action units used in the six basic emotions. The output of this step is the assignment of one or more DOF (including a +/- direction for each DOF) to each action unit, if possible. This output for Melvin is illustrated in TABLE II. Notice that, for simplicity, we are ignoring the AU intensity information here and just considering two values for each DOF, i.e., the corresponding servo motor is fully turned in one or the other direction.

The basic idea of this step is to think about how to most closely approximate each human action unit with the facilities available on the given robot. Studying a standard reference, such as [7], with pictures of the action units (for humans) is helpful for this process. In this step, you should think about each action unit *independently*, without worrying about how they will combine into emotions.

Some choices in this mapping process will be obvious, others will be less ideal. The goal here is to make the best approximation you can (there will be a chance to change your mind in the next step). For some robots, there will be no reasonable choice for some AU's.

Compared to a human face, Melvin's facial features are extremely sparse and their motions have many constraints. His face has no skin, so AU's involving furrows and other subtle movements are simply not possible. Melvin has four servo motors (DOF 1-4) designed to flex his rubber mouth parts into smiles, frowns, etc. Using these DOF, in the appropriate directions, are thus a good choice for AU 12 (Lip

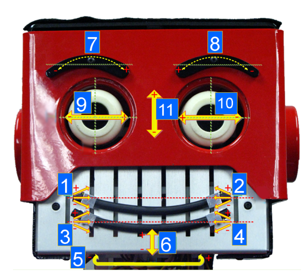

Fig. 2. DOF on Melvin's face

Emotion	Ekman's FACS	Melvin's Translation	Improved Melvin's Translation	Melvin
Happiness	**AU12 (Lip Corner Puller)**	**1- , 2- , 3- , 4-**	**1- , 2- , 3- , 4- ,** *7- , 8- ,* *6+,* *11+*	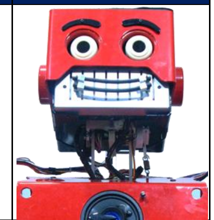
	AU6 (Cheek Raiser), **AU12 (Lip Corner Puller)**	**1- , 2- , 3- , 4-**		
	AU7 (Lid Tightener), **AU12 (Lip Corner Puller)**	**1- , 2- , 3- , 4-**		
Sadness	**AU1 (Inner Brow Raiser)**	**7- , 8-**	**7- , 8- ,** *1+ , 2+ , 3+ , 4+ ,* *5+ , 6- ,* *9+ , 10-*	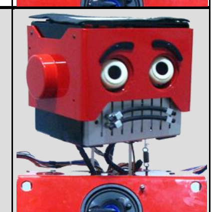
	AU1 (Inner Brow Raiser), AU4 (Brow Lowerer)	**7- , 8-**		
Fear	**AU1 (Inner Brow Raiser)**, AU2 (Outer Brow Raiser), AU4 (Brow Lowerer)	**7- , 8-**	**7- , 8- ,** *1- , 2+ , 3- , 4+ ,* *6+,* *11+*	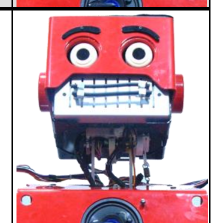
	AU20 (Lip Stretcher)			
Disgust	AU9 (Nose Wrinkler)		**1+ ,** *7+ ,* *5- , 6+ ,* *9- , 10+*	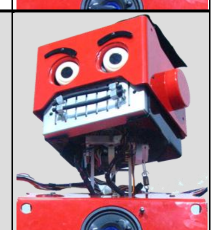
	AU10 (Upper Lip Raiser)	**1+**		
Anger	**AU4 (Brow Lowerer)**, AU5 (Upper Lid Raiser)	**7+ , 8+**	**7+ , 8+ ,** *6- ,* *4-*	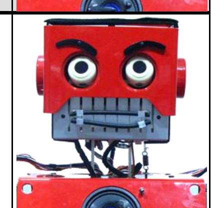
Surprise	**AU1 (Inner Brow Raiser)**, AU2 (Outer Brow Raiser), AU5 [low] (Upper Lid Raiser)	**7- , 8-**	**7- , 8- ,** **1+ , 2+ , 3- , 4- ,** *6+ ,* *11+*	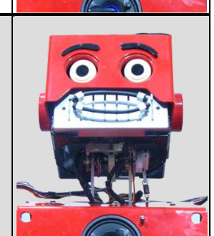
	AU1 (Inner Brow Raiser), AU2 (Outer Brow Raiser), **AU26 (Jaw Drop)**	**7- , 8- ,** **1+ , 2+ , 3- , 4-**		
	AU1 (Inner Brow Raiser), AU2 (Outer Brow Raiser), AU5 [low] (Upper Lid Raiser), **AU26 (Jaw Drop)**	**7- , 8- ,** **1+ , 2+ , 3- , 4-**		
			Neutral Face:	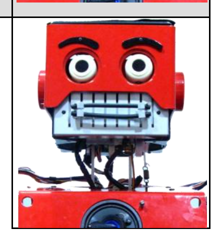

TABLE I

SUMMARY OF DESIGN METHODOLOGY

	AU	1	2	3	4	5	6	7	8	9	10	11
Inner Brow Raiser	AU 1							-	-			
Brow Lowerer	AU 4							+	+			
Right Upper Lip Raiser	AU L10	+										
Lip Corner Puller	AU 12	-	-	-	-							
Jaw Drop	AU 26	+	+	-	-							

TABLE II

MAPPING ACTION UNITS TO MELVIN'S DOF

Corner Puller). On the other hand, notice that AU 9 (Nose Wrinkler) does not appear in TABLE II, because Melvin has nothing approximating a nose. Action units in column two of TABLE I that do appear in TABLE II are indicated in bold; the non-bold action units had no mapping for Melvin.

Finally, notice that Melvin's neck and eye servos (DOF 5-6 and 9-11) are not used in TABLE II. These are available for use in the final step of the design methodology, in which optional action units and cartoon ideas are incorporated.

B. Six Basic Emotions

The second step of the methodology is algorithmic. Simply apply the mapping in TABLE II to all the bold action units in column two of TABLE I. In other words, wherever a bold action unit appears, substitute the corresponding DOF value(s). For example, for AU 10 in Disgust, substitute 1+. Ignore the non-bold action units.

Notice that this step may introduce conflicts whenever, after the mapping, the same DOF appears in a given emotion (alternative) with different direction settings. In other words, since there is only one servo, it cannot be in two positions at once. Such conflicts need to resolved in the next step of the methodology. In the case of Melvin, there were no conflicts.

C. Predicted Confusion Matrix

A confusion matrix (see examples in Tables III and IV) is a very useful tool for resolving DOF conflicts within a single emotion (see preceding section) and analyzing the tradeoffs between using different DOF for different emotions. The columns of a confusion matrix represent the emotions that the robot *intends* to express; the rows represent the emotions that the human *recognizes*. The contents of each cell in the matrix can be either a percentage recognition rate (each row and column sums to 100%), as in TABLE VI, or a qualitative value, as in TABLE III. The cells on the diagonal of the matrix represent correct recognition; all other cells represent confusion, e.g., if surprise is intended, but fear is recognized. Thus for an ideal system, all the diagonal cells are 100%.

The third step in the methodology is to populate a *predicted* confusion matrix based on the following two rules of thumb. First, the greater the proportion of required action units in a given emotion that can be mapped to DOF in the robot (i.e., are in bold), the better the recognition will be for that emotion. We call this the "coverage" rule. Second, the fewer DOF in a given emotion that overlap with other emotions, the better the recognition will be for that emotion (or said conversely, the more DOF overlapping between two emotions, the more likely the two emotions are to be confused). We call this the "overlap" rule.

Applying these two rules of thumb to Melvin results in TABLE III, showing three qualitative levels of predicted recognition: good, medium and poor. Also, the ?'s in this table indicate where we expect confusions to occur due to overlap. Notice that predicted confusion matrices are symmetric, because we have no basis on which to make different predictions, for example, between intending sadness and recognizing fear versus intending fear and recognizing sadness. Observed confusion matrices, such as Tables VI and VII, are not necessarily symmetric.

Happiness in TABLE III is predicted to have good recognition because the first alternative coding has perfect coverage, i.e., the only required action unit is mapped to DOF, and there are no expected confusions with other emotions due to overlap (no ?'s in the happiness column). Sadness and surprise have only medium predicted recognition because of the confusing overlaps on DOF 7 and 8. Anger has medium predicted recognition because it has only 1/2 coverage. Fear has even lower coverage and two confusing overlaps, so it is predicted to have poor recognition. Finally, disgust is a special case: it is predicted to have less than good (i.e., medium) recognition, even though one alternative is fully covered, because the DOF mapping in this case is very weak (servo 1 doesn't really curl up one side of the robot's mouth).

D. Optional AU's and Cartoon Ideas

The fourth and last step of the design methodology is the most open-ended and leaves the most room for creativity. The goal of this step is to improve the predicted confusion matrix. The basic strategy in this step is to reduce the degree of overlap between the confused pairs of emotions in the matrix (indicated by ?'s) by adding additional (different) DOF to each emotion. These additional DOF come from two sources: optional action units and cartoon ideas.

In addition to the required action units for each of the six basic emotions listed in TABLE I, Ekman [7] also reports optional action units for each emotion. These optional units include eye closings, blinks, winks, and other small

Intended

Recognized: Emotion	Happiness	Sadness	Anger	Fear	Surprise	Disgust
Happiness	Good recognition					
Sadness		Medium recognition		?	?	
Anger			Medium recognition			
Fear		?		Poor recognition	?	
Surprise		?		?	Medium recognition	
Disgust						Medium recognition

Good recognition Medium recognition Poor recognition

TABLE III

PREDICTED CONFUSION MATRIX

Intended

Recognized: Emotion	Happiness	Sadness	Anger	Fear	Surprise	Disgust
Happiness	Good recognition				?	
Sadness		Good recognition				
Anger			Medium recognition			
Fear				Poor recognition	?	
Surprise	?			?	Medium recognition	
Disgust						Medium recognition

Good recognition Medium recognition Poor recognition

TABLE IV

IMPROVED PREDICTED CONFUSION MATRIX

movements that are sometimes, but not always, associated with certain emotions. Unfortunately, Melvin does not have eyelids (many other similar robots do) or any of the other facial features used in the optional action units, so we were not able to draw on this source to help reduce confusion.

Artists, and especially cartoonists, know a lot about how to express emotions graphically. It is therefore worth studying an artist's reference, such as [22], for ideas that may be applicable to a cartoonish robot like Melvin.

Focusing on the confusions between sadness and fear or surprise in TABLE III, and the fact that Melvin's neck (5-6) and eye (9-11) DOF were not used in TABLE II we came up with the cartoonish idea of adding different neck and eye movements to Melvin's expressions of sadness, fear and surprise to decrease their degree of overlap with one another. We also decreased the degree of overlap between sadness and fear by adding some distinguishing mouth movements (DOF 1-4). The details can be seen numerically in the third column and pictorially in the rightmost column of TABLE I.

TABLE IV is Melvin's predicted confusion matrix after these improvements are incorporated. Since the two confusions have been mitigated, the predicted recognition of sadness has now been improved from medium to good. Also notice that, as a side effect of the changes made in this step of the design, the predicted confusion between sadness and surprise has "moved" to happiness and surprise. This is a good tradeoff because happiness has such strong coverage.

IV. Evaluation

We conducted a controlled study to evaluate whether our methodology qualitatively predicts which emotions will be most correctly recognized in static facial expressions and also to compare the recognition of dynamic versus static expressions. We first present the experimental procedure and results, and then discuss the results.

A. Experimental Procedure

The study was conducted in our laboratory space at WPI (see Fig. 3). Each participant was seated across a narrow table from Melvin. On the table at their right was an open laptop, on which the participants answered a questionnaire; at their left was a small keypad that they used to sequence to the next robot face.

There were total of 45 participants (average age 20 years) in the study, randomly and evenly assigned to two conditions (16 male and 4 female in the static; 16 male and 3 female in the dynamic). Each participant spent approximately 15 minutes in the study. Six participants' results were discarded because of a robot hardware failure during their sessions. All the participants were WPI students. The participants were told in advance that the study concerned "robot emotions," but were given no other details. There was no photography or video recording during the study, but an experimenter watched the participants via a ceiling-mounted camera from a hidden location nearby.

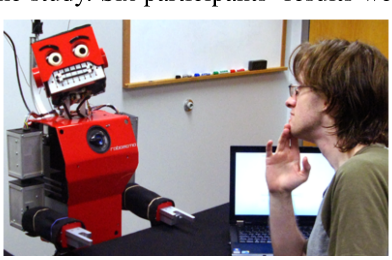

Fig. 3. Experimental setup

Each participant was presented with 18 faces, comprising three identical instances of each of the six basic emotions, in a random order with the constraint that the same intended emotion was never presented twice in a row. The presentation of each face was triggered by the participant pressing any key on the keypad. Melvin's face always starts out in the neutral face, shows an expression for some period of time (see below) and then returns to the neutral face. The participants were allowed to take as long as they wished to answer the questionnaire (see TABLE V) after the face was presented, but were not allowed have the face presented again.

1) Static and Dynamic Conditions: In the static condition, as soon as the participant presses the keypad, all of Melvin's servos are simultaneously commanded to move as quickly as possible to their final (apex) positions. Within 1/10 second all of the facial features (DOF 1-4, 7-11) are at apex; within 1/2 second for the neck DOF 5-6. The apex configuration is held for 1.5 seconds and then the face is returned as quickly as possible to the neutral face.

As discussed in Section II-C, the timing in the dynamic condition is much more complicated and different for each emotion. Fig. 4 shows an example of the temporal dynamics of Melvin's transition from the neutral to the sad face. The overall transition takes about 4 seconds. Also, the mouth (lower face) starts moving first, followed by the eyebrows (upper face). The neck starts moving somewhere in between. Melvin then holds the final pose for 2 seconds before returning as quickly as possible to the neutral face. The other emotions have comparable but different temporal dynamics, based as much as possible on the psychological literature and filling in missing details using our own intuitions.

2) Questionnaire: The results of the study described above validate our design methodology by showing that the (final improved) confusion matrix predicted by the methodology agrees qualitatively with the actual confusion matrix observed in the study (in the static condition). Results in the dynamic case are inconclusive.

3) Static Condition Results: TABLE VI is the confusion matrix observed in the static conditions. Notice that this matrix includes an extra row, labeled *Other*, that contains the percentage of participants in each case who selected the "Other emotion" choice in the questionnaire (see TABLE V).

Which of the following emotions _best_ describes the robot's
static condition: **current facial expression?**
dynamic condition: **facial movements that you just watched?**
Choose more than one emotion _only_ if you cannot distinguish between them.

☐ Anger ☐ Happiness ☐ Sadness
☐ Disgust ☐ Fear ☐ Surprise
☐ Other emotion (please specify in text box below)

TABLE V

Participant questionnaire after each face

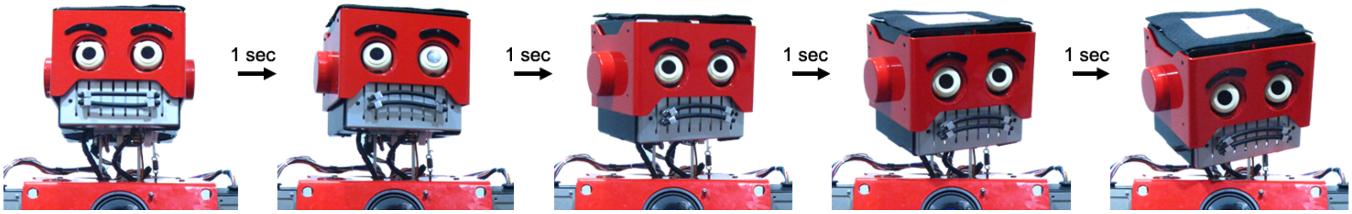

Fig. 4. Temporal dynamics of transition from neutral to sad face

The complete list of values entered into the text box associated with this choice is: no emotion, I do not know, confused, not sure, unsure, shock, worry, mild surprise, thoughtful, skeptisism, weirded out, disappointment, crazy, demented, intimidating, doubt, disbelief, disapproval.

In addition to the percentage recognition rates in the cells of this confusion matrix, we also show Cohen's Kappa coefficient for each intended emotion. Kappa is a standard way of evaluating "inter-coder reliability," which in this case, indicates how reliably participants correctly distinguished between the intended emotion and all the other choices. The greater the kappa value, the better the reliability, and a kappa above 0.75 is generally considered good reliability.

TABLE IV predicts that happiness and sadness will have good recognition, fear will have poor recognition, and the other three emotions (anger, surprise and disgust) will be in between. In TABLE VI we see that happiness and sadness both have kappa values in the good range, fear has the lowest kappa value, and the other three emotions are in between, just as predicted. This is the primary result of the evaluation confirming our design methodology. (Note that the kappa for disgust is much lower than the other two medium values, which is likely due to the lip curl problem discussed earlier.)

4) Dynamic Condition Results: TABLE VII is the confusion matrix observed in the dynamic condition. In this condition, the predictions of the methodology are only partly confirmed: happiness has a kappa value in the good range and fear has the lowest kappa value, as predicted, but all four of the other emotions are in the middle, with anger having an almost-good kappa that is higher than sadness.

Comparing the dynamic and static conditions, we see that four of the diagonal values in TABLE VII are lower than in TABLE VI, and two are higher. TABLE VIII looks at this data in another way, comparing the correct recognition percentages overall and for each emotion in the static versus dynamic conditions. The only one of these comparisons which is close to being significant (2-tailed t-test $p = .088$) is happiness, where the static case has better recognition.

Intended

Recognized: Emotion	Happiness	Sadness	Anger	Fear	Surprise	Disgust
Happiness	93	0	0	7	14	0
Sadness	0	75	0	3	0	2
Anger	2	3	68	0	1	17
Fear	0	7	0	23	5	17
Surprise	3	1	6	25	77	5
Disgust	0	12	18	17	1	47
Other	2	2	8	25	2	12
κ	0.84	0.8	0.66	0.2	0.64	0.38

TABLE VI

OBSERVED CONFUSION MATRIX (STATIC)

Intended

Recognized: Emotion	Happiness	Sadness	Anger	Fear	Surprise	Disgust
Happiness	85	0	0	5	6	0
Sadness	2	60	5	15	1	0
Anger	0	8	77	2	0	12
Fear	2	3	5	17	18	12
Surprise	6	0	2	20	74	3
Disgust	0	7	7	10	0	62
Other	5	22	4	31	1	11
κ	0.84	0.6	0.73	0.1	0.67	0.6

TABLE VII

OBSERVED CONFUSION MATRIX (DYNAMIC)

5) Successive Answers: TABLE IX shows the overall correct recognition percentages separately for the first, second and third instances that each participant saw the same intended face during their session (these three answers are averaged in the results presented above). Notice that the correctness gets *worse* on successive answers, with almost identical statistics in both conditions. Furthermore, a 2-tailed t-test indicated that the difference between the first and third answers is statistically significant ($p < .001$) in both conditions. The same effect shows up looking at the recognition statistics for each emotion individually.

B. Discussion

The primary result of the study is a very strong one. Our design methodology qualitatively predicted the observed confusion matrix in the static condition. It is not surprising that the prediction of the confusion matrix in the dynamic condition was less successful, since there were many aspects

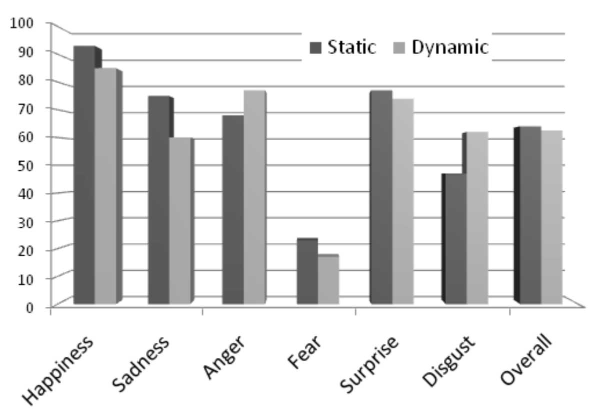

TABLE VIII

CORRECT RECOGNITION PERCENTAGES

Condition	1st Answer	2nd Answer	3rd Answer
Static	82	63	37
Dynamic	81	64	38

TABLE IX

OVERALL CORRECT RECOGNITION PERCENTAGES BY ANSWER of the dynamic expressions, such as the order of onset of action units and their durations, which were not explicitly considered in the design methodology.

We expected the overall recognition rates to be higher with dynamic expressions than with static expressions, because of the psychological literature and because dynamic expressions intuitively seem to have more "information." However, the data did not bear this out, for which we have no explanation other than to speculate that we did not get the dynamic temporal details right.

The decreasing correctness results on successive answers were a big surprise. We expected exactly the opposite, i.e., a typical learning curve in which the third answers would be the best. Our speculation on the cause of this effect is that, when participants see their first face, they are less aware of the possible overlaps between faces. As they see more faces, however, they become more confused by the expanding overlap possibilities, in a sort of "avalanche" effect. The implication of this effect would be that overlap between two faces hurts not only the recognition of the two faces involved, but overall recognition also.

V. Conclusion

We have developed and documented a step-by-step methodology that allows any roboticist to design static emotional facial expressions for his or her robot without having to spend a lot of time delving into the relevant pyschological literature. We demonstrated the application of this methodology to our own robot, Melvin, and showed that we could qualititively predict how people will recognize the resulting emotional expressions.

Fig. 5. Elvo

In future work, we plan to apply and evaluate our methodology with another robot in our laboratory, Elvo (see Fig. 5). Elvo has a similar number of DOF to Melvin, but a very different style of face. Instead of moveable fixed features, Elvo has simulated skin that is pushed and pulled from inside his head by motors. This will allow Elvo to perform a number of action units, such as wrinkling of the upper brow, that were not available to Melvin. Elvo even has a special DOF to curl up one side of his mouth, since this was one of Elvis Presley's hallmark gestures (which should be very useful for expressing disgust).

Finally, there is much room for future work extending our methodology to explicitly deal with the temporal details of dynamic expressions.

Acknowledgement: This work is supported in part by the National Science Foundation under awards IIS-0811942 and IIS-1012083. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation. We would like to thank Prof. Ryan Baker for his advice and Aaron Holroyd, Will Coon, Chris Conley and Bahador Nooraei for their advice, support and help.

References

[1] C. Bartneck and M. J. Lyons. Facial expression analysis, modeling and synthesis: Overcoming the limitations of artificial intelligence with the art of the soluble. In J. Vallverdu and D. Casacuberta, editors, *Handbook of Research on Synthetic Emotions and Sociable Robotics:New Applications in Affective Computing and Artificial Intelligence*, pages 33–53. IGI Global, 2009.

[2] J. N. Bassili. Emotion recognition: The role of facial movement and the relative importance of upper and lower areas of the face. *Personality and Social Psychology*, 37(11):2049–2058, 1979.

[3] E. Bould and N. Morris. Role of motion signals in recognizing subtle facial expression of emotion. *British Journal of Psychology*, 99(2):167–189, 2008.

[4] P. Ekman. Facial expression and emotion. *American Psychologist*, 48(4):384–392, 1993.

[5] P. Ekman. *Emotions Revealed*. Times Books, 1st ed., 2003.

[6] P. Ekman and W. V. Friesen. Measuring facial movement. *Environmental Psychology and Nonverbal Behavior*, 1(1):56–75, 1976.

[7] P. Ekman, W. V. Friesen, and J. C. Hager. *Facial Action Coding System*. A Human Face, 2002.

[8] G. Horstmann and U. Ansorge. Visual search for facial expressions of emotions: a comparison of dynamic and static faces. *Emotion, Washington DC*, 9(1):29–38, 2009.

[9] S. Kaiser, T. Wehrle, and S. Schmidt. Emotional episodes, facial expressions and reported feelings in human-cmputer interaction. In *Proc. of the Xth Conf. of the Int. Society for Research on Emotions*, pages 82–86. ISRE Publications, 1998.

[10] D. Matsumoto and P. Ekman. Facial expression analysis. *Scholarpedia*, 3(5):4237, 2008.

[11] D. Matsumoto, D. Keltner, M. N. Shiota, M. O'Sullivan, and M. Frank. Facial expressions of emotion. In M. Lewis, J. M. Haviland-Jones, and L. F. Barrett, editors, *Handbook of emotions*, pages 211–234. Guilford Press, New York, London, third edition, 2008.

[12] M. Paleari, A. Grizard, and C. Lisetti. Adapting psychologically grounded facial emotional expressions to different anthropomorphic embodiment platforms. In *Proc. of the 20th Florida Artificial Intelligence Research Society Conf.*, FLAIRS '07, pages 565–570, Menlo Park, CA, USA, 2007. AAAI Press.

[13] C. Pelachaud. Modelling multimodal expression of emotion in a virtual agent. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 364(1535):3549–3557, 2009.

[14] G. M. Reevy. *Encyclopedia of Emotion*. Greenwood, 1st ed., 2010.

[15] E. D. Ross, C. I. Prodan, and M. M. Human facial expressions are organized functionally across the upper-lower facial axis. *The Neuroscientist*, 13(5):433–446, 2007.

[16] K. R. Scherer. On the nature and function of emotion: A component process approach. In K. R. Scherer and P. E. Ekman, editors, *Approaches to emotion*, pages 293–317. NJ: Erlbaum, Hillsdale, New Jersey, London, 1984.

[17] K. R. Scherer and H. Ellgring. Are facial expressions of emotion produced by categorical affect programs or dynamically driven appraisal? *Emotion*, 7(1):113–130, 2007.

[18] D. J. Schiano, S. M. Ehrlich, K. Rahardja, and K. Sheridan. Face to interface: Facial affect in (hu)man and machine. In *Proc. of the ACM CHI 2000 Conf. on Human Factors in Computing Systems*, CHI '00, pages 193–200, New York, NY, US, 2000. ACM Press.

[19] R. J. Sloan, B. Robinson, K. Scott-Brown, F. Moore, and M. Cook. Choreographing emotional facial expressions. *Computer Animation and Virtual Worlds*, 21(3-4):203–213, 2010.

[20] R. J. S. Sloan, M. Cook, and B. Robinson. Considerations for believable emotional facial expression animation. In *Proc. of the 2009 Second Int. Conf. in Visualisation*, VIZ '09, pages 61–66, Washington, DC, USA, 2009. IEEE Computer Society.

[21] E. Tanguy, P. J. Willis, and J. J. Bryson. A dynamic emotion representation model within a facial animation system. *Int. Journal of Humanoid Robotics*, 3(3):293–300, 2006.

[22] F. Thomas and J. O. *The illusion of life: Disney animation*. Disney Press, 1995.

[23] T. Wehrle, S. Kaiser, S. Schmidt, and K. R. Scherer. Studying the dynamics of emotional expression using snthesized facial muscle movements. *Journal of Personality and Social Psychology*, 78(1):105–119, 2000.

[24] T. Wu, N. Butko, P. Ruvulo, M. Bartlett, and J. Movellan. Learning to make facial expressions. In *Proc. 2009 IEEE 8th Int. Conf. on Development and Learning*, pages 1–6, Washington, DC, USA, 2009.