Do Affect-Sensitive Machines Influence User Behavior?

Laurel D. Riek[1] and **Shazia Afzal** and **Peter Robinson**

Abstract. Machines are becoming more socially aware as the fields of affective computing and ambient intelligence advance. In the future such machines will start to become more commonplace in domestic and work environments. How will these machines affect people's behavior? Previous work shows that people both have a tendency to treat machines like humans, as well as to abuse them. We have designed an experiment to understand people's attitudes concerning affect-sensitive machines and their expressivity toward them.

1 INTRODUCTION

In the field of human-computer interaction, a paradigm shift has begun from "human factors to human actors" [4]. Researchers are now considering people's emotional experience to be a new dimension of usability as well as a measure of system success. In fact, the entire field of affective computing exists in part to help address some of the failings of traditional interactive systems that typically neglect affective state changes in users. The hope is that eventually machines will be sensitive to the affect of people interacting with them and able to adapt their behavior accordingly [19].

Affect-sensitive machines (ASM) becoming more prevalent in society raises a number of interesting questions. How are such machines going to change how people view and use technology? How transparent should the workings and reasoning of such systems be towards users? How might the behavior of people change when they are interacting with ASMs?

From previous work in human-computer interaction and human-robot interaction, it is clear people have pre-conceived opinions of and expectations toward the machines they interact with, and these beliefs are likely to influence their behavior. For example, Nass's Computers As Social Actors (CASA) paradigm [13, 14] suggests that cues of humanness are sufficient to encourage individuals to mindlessly apply social rules and expectations while interacting with media. Walters [18] showed that when people are interacting with robots they prefer them to be at the same "comfortable distance" exactly as they would another human, regardless of the robot's physical appearance. Kirby et al. [10] showed that people are far more likely to spend time interacting with an expressive robot as opposed to a neutral one. This effect was shown to be true regardless if the robot's affect was positive or negative. Interestingly, none of the aforementioned systems were sensitive to user affect, and yet people still interacted with the machines in ways similar to how they interact with other humans.

A few researchers have looked at people's attitudes toward ASMs. Axelrod and Hone [3] simulated real-time interaction with an ASM using a Wizard of Oz technique and found that users who were aware of the affect-sensitivity of the system portayed significantly more positive displays of affect than those who were unaware. Brave et al. [6] showed that embodied conversational agents that acted empathetically were viewed as more trustworthy, likeable, caring, and supportive than agents that were not empathetic. Riek and Robinson [16] found that people experienced more satisfaction when interacting with an intentionally empathetic robot compared to one that was mind-blind.

Video analysis of data obtained in a potential application setting, computer-based learning, reveals that emotional behaviour depends not only on individual differences and the task at hand, but is also influenced by people's attitudes. We ran a study with eight participants (six female, two male) and videotaped them doing two tasks: an interactive map-based geography tutorial and a card-matching game. See Figure 1 for exemplary facial displays users made during the experiment. Seven of the eight participants indicated in post-experimental interviews that they would probably interact differently if they knew the computer could respond to their affective state. For example, it's possible their gestures and facial expressions would be different. We hypothesize that this 'difference' might be an exaggeration of behavior that happens when one tries consciously communicate an emotion to other humans, such as pleasure [1].

Other research has revealed that when people interact with intelligent agents, they can be very abusive in their behavior [5]. We also saw a similar display of abusive behavior in our aforementioned study when one subject "gave the finger" to the computer while playing the card-matching game [1]. These abusive displays may be because the social consequences of behavior that apply toward human-human interaction are not necessarily applicable toward human-machine interaction.

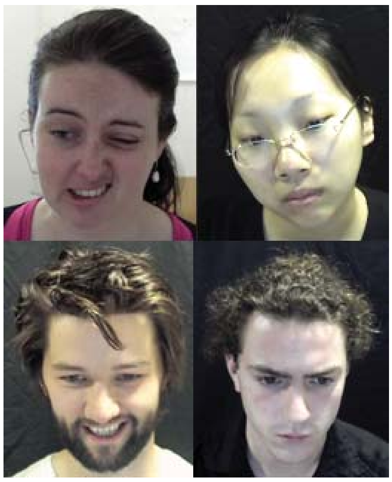

Figure 1. Both when completing a tutorial and playing a card game subjects unwittingly displayed a range of facial expressions.

[1] Computer Laboratory, University of Cambridge, United Kingdom, email: Laurel.Riek@cl.cam.ac.uk

People's tendency to either treat machines like humans or abuse them has lead us to question how we can measure people's attitudes toward ASMs. In particular, we are curious how their expressivity reveals these attitudes. Thus, we've designed an experiment that will allow us to explore these issues.

2 EXPERIMENTAL DESIGN

Investigating these questions requires an experimental design that measures not only the frequency and occurence of emotional displays but also how people's attitudes affect their willingness to engage in emotional communication. Specifically, we are interested in the question of whether people make more facial expressions toward a machine that they believe to be sensitive to their affect.

Thus, we propose a within-subjects experiment that involves subjects playing a puzzle game. Subjects will be told that they are helping us design an intelligent game that adapts as they play. They will be told we are testing two types of automatic adaptation - one that is sensitive to their affect (AS condition) and the other that is based on game performance (GP condition). In reality, both modes of play will be identical, but we will be deceiving subjects to believe they're different. (See Section 2.1).

Our primary hypotheses are as follows:

(H1) People make more non-neutral emotional expressions in the AS condition vs. the GP condition

(H2) People make more facial expressions toward the beginning of the experiment vs. the end of the experiment

These hypotheses are motivated by several ideas. With regards to (H1), we think people may have a tendency to "game the system"; in other words, they may make exaggerated facial expressions in the AS mode in an attempt to affect the outcome of the game. (H2) is motivated by the idea that we expect people will habituate to the machines' perceived affect sensitivity, and make more facial expressions early on but then forget to as the game progresses. This result may largely depend on how effective we are at deceiving people that the game is in fact changing based on their facial expressions.

Additionally, we are also interested in whether:

(H3) People who are more expressive (as measured by the tests described in Section 2.2.2) will make more facial expressions

(H4) People who are more expressive will show a similar relative expressive pattern when interacting with a computer.

(H3) is motivated by the non-verbal behavior literature on how people vary in their emotional expressivity. (H4) is inspired by the work of Riggio and Riggio who showed that emotional expressiveness as a personal style is relatively consistent across situations [17].

2.1 Methodology

In our experiment we will be employing a social psychological method of emotion induction proposed by Harmon-Jones et al. [9]. This method involves using high-impact manipulation and deception to achieve a high level of psychological realism in a laboratory setting. The idea is to produce emotional responses by placing participants in psychologically involving situations.

Thus, we will tell subjects that we are evaluating two techniques for creating adaptive games. One technique is computer-vision based, and use the camera to monitor their emotional states as they play. The other is performance-based and uses complex mathematical techniques. They will know which mode they are in via an omni-present banner at the bottom of the game screen (see Figure 2).

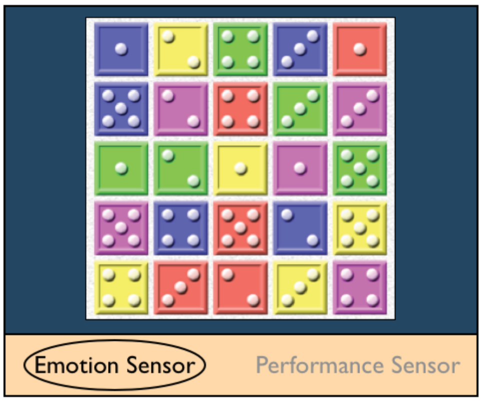

Figure 2. A screenshot from the game with a banner informing the user if the computer is currently monitoring the user's affect.

After the experiment, subjects will be appropriately debriefed and will be asked to sign an authorization form allowing the use of their video for research purposes.

2.2 Materials

2.2.1 *Pre-Experiment Questionnaire*

Before the experiment begins, we will ask subjects to complete a short questionnaire asking them demographic information, such as their age, gender, english fluency, and job title. To try to gauge user expertise, we will also ask them about the types of tasks they use a computer for (email, games, word processing, chat, etc), and the duration per day of such tasks.

We will also ask questions regarding people's cultural exposure to ASMs, such as particular films and books that may have an influence on their attitudes (i.e., the film *Wall-E* or the book *The Positronic Man*). We will also ask them directly about their attitudes toward hypothetical ASMs.

2.2.2 *Pre-Experiment Expressivity Tests*

Research in the field of non-verbal behavior indicates that there are differences in the manner and intensity by which people express their emotions. Self-report measures of nonverbal expressiveness assess such individual differences in the generation and/or expression of emotion. Such measures also assess a more general tendency in people to display affect spontaneously and across a wide range of situations [17].

We will ask subjects to complete three short, self-report measures of their dispositional (nonverbal) expressivity: the Berkeley Expressivity Questionnaire (BEQ) [8], the Emotional Expressivity Test (EES) [11], and the Affective Communication Test (ACT) [7]. We've selected these tests on the basis of their short administration time, easy availability, reliability, and internal consistency, as well as how they conceptualize emotion. The scores from these tests will allow us the ability to compare subjects with one another, as well as to interpret our results.

2.2.3 *Post-Experiment Interview*

After the experiment we will get subjective reports from our participants via semi-structured interviews in order to assess how conscious they were of the experimental manipulation and whether they changed their emotional behavior across the two conditions. Specifically, we will ask them whether they noticed the system adapting content based on their emotional state and whether/how they changed their facial expressions or emotional displays during the two conditions. We will also ask them whether they perceived any change in the way they interact with humans vs. machines and how likely ASMs might change their behavior.

2.3 The Game

We will be using a logic game called Boxit in our experiment, as shown in Figure 2. This game was created by Frank Hollwitz in 2005. The goal of the game is to remove tokens as many tokens as possible from the board by replacing one token with another. Tokens can be replaced if they are in the same row or column, are of the same color (red, blue, green, or yellow), or are of the same number (1 - 5). The game ends when no moves are left.

We've selected this game for several reasons. First, we wanted a game that was open source so we could easily modify it to automatically control for play duration, difficulty level, and play mode (AS vs. GP). We also wanted to be able to easily add a banner to the screen indicating the mode of play. Second, we wanted a game that did not rely on reflexes or speed because success at such games requires a significant amount of practice time which would make the experiment much longer. Third, we wanted a game that was vague in terms of its level of difficulty, so that we could be more successful at manipulating subjects to believe the affect-sensitive machine was altering game play.

2.4 Measures

To measure (H1) and (H2) we will simply count the number of non-neutral facial expressions subjects made during the experiment. To measure (H3) and (H4) we will correlate this count with scores obtained from the three expressivity tests. To get a quantitative estimate of the overall expressivity during the experiment we will further annotate the videos using six global dimensions of expressivity drawn from speech annotation [12, 2]:

1. Overall activation: amount of activity - {Static/Passive, Neutral, Animated, Engaged}
2. Spatial extent: amplitude of movements - {Contracted, Normal, Expanded}
3. Temporal extent: duration of movements - {Slow/Sustained, Normal, Quick/Fast}
4. Fluidity: continuity and smoothness of movement - {Smooth, Normal, Jerky}
5. Power: strength and dynamics of movements - {Weak/Relaxed, Normal, Strong/Tense}
6. Repetitivity: repetition of same expression/gesture several times - {Low, Normal, High}

2.5 Procedure

After being briefed and completing the pre-experiment questionnaire and expressivity tests, subjects will be seated at the computer and given the opportunity to learn how to play the game. All subjects will partake in a training session lasting 5 minutes in duration.

Following the training tasks, subjects will take a short break (e.g., viewing a nature video for a few minutes). They will then be assigned to either the AS or GP condition (counter-balanced across subjects). The will play in the first mode for 5 minutes, then have a short break, then play in the second mode for the same duration.

Following the experiment subjects will be given a post-experimental interview and appropriately debriefed.

2.5.1 *Subjects*

We will first run several small pilot studies with subjects from across the University. If we see an effect in our data, we will continue with a larger sample. Subjects will be recruited via email lists, bulletin board postings, etc.

3 DISCUSSION

We described details of an experiment we will run in the coming months regarding how people alter their behavior when faced with a machine that is seemingly sensitive to their affect. We anticipate very interesting data to come from this experiment and look forward to reviewing it. It is our hope that by employing a combination of quantitative and qualitative measures we will glean an understanding of some underlying attitudes people hold about affective machines.

If we find that people do act significantly differently when faced with an ASM, a number of interesting issues are raised. First, it means that people researching affective computing and ambient intelligence need to consider the problem that users may try to "game the system" during interaction. Therefore, it is increasingly important to carefully consider one's assumptions when designing affective-aware systems.

Second, such a result would also help to inform debate in the affective computing community regarding the use of naturalistic vs. non-naturalistic data. It would seem both sets of data may prove useful from an emotion-recognition perspective because it is likely for users to engage in both types of behavior when interacting with a system. And, further, that said modes of interaction will change depending on how people habituate to interacting with such systems and how that alters their expressivity.

From an ethical perspective, we believe it is important that the existence and workings of ASMs are made as transparent as possible to users. This stance is in line with one of the fundamental principles of Human-Centered Design - users should always know what a machine is and what it's doing [15]. In other words, users should always know what a machine's behavior and role will be during interaction. This is particularly important for ASMs, as people typically don't expect their behavior to be monitored.

Finally, it will be interesting to see whether people adhere to social display rules when interacting with an apparently affect-sensitive machine. This will help us to understand if perceived emotional awareness in a machine engenders polite, or abusive, social behavior.

ACKNOWLEDGEMENTS

This work is supported by the QUALCOMM Research Studentship and the Gates Cambridge Trust.

REFERENCES

[1] S. Afzal, C. Morrison, and P. Robinson. Intentional affect: An alternative notion of affective interaction with a machine. In review, 2009.

[2] S. Afzal and P. Robinson, 'Dispositional expressivity and hci', in *Workshop on Emotion in HCI*. British HCI, (September 2008).

[3] L. Axelrod and K. Hone, 'Uncharted passions: User displays of positive affect with an adaptive affective system', in *Affective Computing and Intelligent Interaction*, pp. 890–897, (2005).

[4] L. Bannon, 'From human factors to human actors: the role of psychology and human-computer interaction studies in system design', in *Design at work: cooperative design of computer systems*, eds., J. Greenbaum and M. Kyng, 25–44, L. Erlbaum Associates Inc, Hillsdale, NJ, USA, (1992).

[5] S. Brahnam and A. De Angeli, 'Special issue on the abuse and misuse of social agents', *Interacting with Computers*, **20**(3), 287 – 291, (2008).

[6] S. Brave, C. Nass, and K. Hutchinson, 'Computers that care: investigating the effects of orientation of emotion exhibited by an embodied computer agent', *International Journal of Human-Computer Studies*, **62**(2), 161 – 178, (2005). Subtle expressivity for characters and robots.

[7] H. S. Friedman, L. M. Prince, R. E. Riggio, and M. R. DiMatteo, 'Understanding and assessing nonverbal expressiveness: The affective communication test', *Journal of Personality & Social Psychology*, **39**(2), 333–351, (1980).

[8] J.J. Gross and O.P. John, 'Facets of emotional expressivity: three self-report factors and their correlates', *Personality and Individual Differences*, **19**, 555–568, (October 1995).

[9] E. Harmon-Jones, D. M. Amodio, and L. R. Zinner, 'Social psychological methods of emotion elicitation', in *Handbook of Emotion Elicitation and Assessment*, eds., J. A. Coan and J. J. B. Allen, 91–105, Oxford University Press, New York, (2007).

[10] R. Kirby, J. Forlizzi, and R. Simmons, 'Interactions with a moody robot', in *ACM Conference on Human-Robot Interaction (HRi 06)*, pp. 186–193, (March 2006).

[11] A. Kring, D. A. Smith, and J. M. Neale, 'Individual differences in dispositional expressiveness: Development and validation of the emotional expressivity scale', *Journal of Personality & Social Psychology*, **66**, 934–949, (1994).

[12] J. Martin, S. Abrilian, L. Devillers, M. Lamolle, M. Mancini, and C. Pelachaud, 'Levels of representation in the annotation of emotion for the specification of expressivity in ecas', *Intelligent Virtual Agents*, 405–417, (2005).

[13] C. Nass and Y. Moon, 'Machines and mindlessness: Social responses to computers', *Journal of Social Issues*, **56**(1), 81–103, (2000).

[14] C. Nass, Y. Moon, and B. J. Fogg, 'Can computer personalities be human personalities?', *International Journal of Human-Computer Studies*, **43**, 223–239, (1995).

[15] D. A. Norman, *The Psychology of Everyday Things*, Basic Books, April 1988.

[16] L. D. Riek and P. Robinson, 'Real-time empathy: Facial mimicry on a robot', in *Workshop on Affective Interaction in Natural Environments (AFFINE) at the International ACM Conference on Multimodal Interfaces*. ACM, (2008).

[17] R. E. Riggio and H. R. Riggio, 'Self-report measures of emotional and nonverbal expressiveness', in *The sourcebook of nonverbal measures: Going beyond words*, ed., V. Manusov, 105–111, Erlbaum, Mahwaj, NJ:, (2005).

[18] M. L. Walters, *The Design Space for Robot Appearance and Behaviour for Social Robot Companions*, Ph.D. dissertation, University of Hertfordshire, March 2008.

[19] Z. Zeng, M. Pantic, G. Roisman, and T. Huang, 'A survey of affect recognition methods: Audio, visual, and spontaneous expressions', *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **31**(1), 39–58, (January 2009).