Designing for Affective Interactions

Carson Reynolds and Rosalind W. Picard

MIT Media Laboratory
20 Ames Street, Cambridge, MA 02139-4307
{carsonr,picard}@media.mit.edu

ABSTRACT

An affective human-computer interaction is one in which emotional information is communicated by the user in a natural and comfortable way, recognized by the computer, and used to help improve the interaction. Before the computer or its designer can adapt an interaction to better serve an individual user, feedback from that user must be associated with the actions of the machine: did a specific computer action please or displease the user? Did something in the interaction frustrate the user? One of the essential issues is sensing and recognizing the affective information communicated by the user in a way that is comfortable and reliable. This paper highlights several devices we have built that offer users means of communicating affective information to a computer.

1 INTRODUCTION

There is no one interface that will please all users, no matter how usable or well-designed. Instead, we are in an era where software is adapted to please a user through the effort of the individual user, who must hand-tweak dozens or hundreds of features to her liking. Software that self-adapts to the user is hard to design: the system must obtain a delicate balance between changing in the direction of what the user wants, while providing stability and predictability. A design that changes constantly will be irritating, but so too is one that doesn't change when it should. The hard part is how to know when to change, and how to know if the change has pleased the user or not. The right balance, we argue, is also user-dependent and situation-dependent.

How can a system attain the right balance in the interaction, to please the user? Adopting the theory that successful human-human interactions can be used to inform human-computer interactions (Reeves and Nass, 1996), we suggest addressing the problem of "when and how the computer should adapt" by considering when and how people decide to adapt. Consider, for example, when a person is explaining something to you and you start to look frustrated. In this case the person may carry on for a while, initially ignoring your expression; however, if you continue to look increasingly frustrated, they will eventually pause and (hopefully) politely try to get more information from you as to what is wrong. This is in contrast to most computers today, which don't notice that the user is frustrated, and don't have the skills to respond or to adapt.

Many existing so-called "smart" systems try to fix their behavior to better serve the user. However, in these systems the guiding principle is to automate tasks the user seems to be doing. The presumption is that efficiency is king, and that improving efficiency will please the user. But, when the system is wrong, say, presuming it knows how to fix the user's spelling or how to indent the user's entries, when in fact the user had it right the way he wanted it, then this is a problem where affective communication can help. We suggest that the guiding principle of pleasing the user is greater than the guiding principle of maximal efficiency. Thus, if a user expresses pleasure when the computer fixes his spelling, then the system should continue applying its spell-fixing behavior. If the user expresses annoyance, then the computer should consider asking the user if he would like that behavior turned off or adapted in some other way. Computers that presume to know a user's intentions, no matter how smart the computer is, will sometimes make mistakes that irritate and frustrate users. Thus, if a computer is designed to "fix its own behavior, " and does so in a way that causes the user to indicate increasing frustration, this should be a sign to the system that it needs to take more ameliorative action.

The importance of communicating affective information is often ignored in interaction design. Designers have built systems that try to help a user, but they have not built systems that see if offers of help irritate or please the user. In our minds, this lack of attention to the user's response is disrespectful and short-sighted. How is a system to adapt to a user if it can't even detect if its action is pleasing or displeasing the user? Adapting is not a one-size-fits-all problem. True respect for the user requires adapting to please him or her, and not merely to please some average user that doesn't exist. Even a dog, which presumably doesn't have any sophisticated intelligence about human-human interaction, can tell if it has pleased or displeased its owner. Such affect sensing is foundational to a learning system that interacts with a person– in order to learn, the system must be capable of receiving some valenced feedback – positive or negative – so that its positively-perceived actions can be reinforced and its negatively-perceived actions removed or adapted.

2 BACKGROUND

The idea of using sensors to get positive or negative feedback from an individual is not new. Twenty-five years ago, Sheridan discussed instrumentation of various public speaking locations (classrooms, council meetings, etc.) with switches so that the audience could respond and direct discussion by voting (Sheridan, 1975). However, one of the problems with using switches was that they required active effort on the part of the individual, which distracted them from the first task at hand. If a speaker stops and asks "do you like this?" then flipping a switch "yes" or "no" is natural. However, if your feelings change many times while listening to the speaker, you might prefer to give more continuous feedback, and not have to interrupt your train of thought to do so. Watching somebody's facial expressions subtly change from positive to negative is one way of sensing affective information that is passive – requiring no extra mental effort on the part of the person communicating the emotion.

A number of learning systems, especially robots, are starting to be given the ability to do passive sensing of displays of emotion. The robot Kismet is an example of a computer system that attends to qualities of human speech for indications of approval or disapproval – positive or negative affective information – with the goal of using this information to guide its interaction and help it learn from a person (Breazeal, 2000). By speaking in an approving way at the robot's action, a human "care-giver" serves to reinforce its action. The idea is that affect communication is important for enabling the robot to learn continuously while interacting with a human.

Sensors that detect the expression of emotion can be grouped into two categories: those which passively collect data from the user and those that require the user's intent. WinWhatWhere™ is an example of a passive sensor: it records keystrokes and other data without the user's intent. Passive sensors are convenient, requiring minimal effort on the user's part, but also raise a host of legitimate concerns. When we canvassed non-technical staff around MIT as to whether we could place cameras and microphones or other sensors in their offices or on the surface of their skin to record how they naturally expressed frustration, they all indicated that they would prefer to intentionally express frustration to an interface on the computer, vs. having it sensed passively. Their discomfort with the latter in many cases was because of a feeling of invasion of privacy and loss of control. Thus, although people often use passive methods to sense the affect of another person, we decided to honor users' wishes and try to build means of sensing frustration that were active – requiring them to fill out a web form, adjust a software frustrometer, or click on a specific icon.

In the affective computing group at MIT, we have built a variety of passive and active systems that attempt to gather affective information from a user. Our earliest work focused on wearable computer systems, where there was no traditional keyboard or mouse, but there was contact with the surface of the skin. With these affordances in mind, we tried to develop methods that would passively detect patterns of change in physiology; for example, we built a system that attempted to sense when a user might be frustrated based on measuring changes in skin conductivity and blood volume pressure (Fernandez and Picard, 1998). We are also working on methods that would try to detect such feedback from the face and voice of the user. The rest of this paper focuses on more intentional or active means of sensing, highlighting a few of the prototypes we have built that require no cameras, microphones, or physiological sensing, but that work within the scope of more traditional computer keyboard-monitor-mouse interfaces.

3 DEVICES FOR COMMUNICATING AFFECTIVE FEEDBACK

We have prototyped a number of different designs intended to help people express frustration with their computing systems (Figure 1). The sensors are designed to detect user behaviors that might signal frustration, and to associate the user behavior with the current state of the computer.

Figure 1: Prototypes of tangible feedback devices: pressure-sensing mice and acceleration-sensing voodoo dolls.

Examples of our prototypes include a wireless "voodoo doll" that detects acceleration when picked up and thrown, many different pressure sensitive devices that detect intentional and unintentional muscle tension, and also a microphone which detects a raised voice. Numerous software interfaces ranging from a simple textbox, to thumbs-up/thumbs-down icons, to elaborate graphical displays were also developed.

3.1 Pressure Sensitive Mice

One possibility that emerged early in our inquiry was using force-sensitive resistors and conductive foam in combination with a pointing devices. When coupled with a low cost analog-to-digital conversion (ADC) circuit board, these sensors can be used as inexpensive grip pressure sensors.

Previously, Kirsch had examined using a mouse modeled after Manfred Clynes' "Sentograph" to detect emotional states (Kirsch, 1997). The Squeezemouse (as we call it) focuses less on directional pressure and more on detecting muscle tension and on providing a tangible interface for expressing dissatisfaction. To that end, emphasis was placed upon discovering comfortable locations for pressure sensors.

After building and testing designs in which the entire surface was made to be pressure sensitive, we settled on simpler designs in which only a clearly defined part of the mice was actively sensing. This arrangement had two distinct advantages: it had a fairly linear and predictable response, and it was less prone to being accidentally triggered.

Figure 2: Early Squeezemouse prototype in action.

Of course the design of the sensor is incomplete without a software user interface. We identified several tasks the interface might perform by conducting informal surveys of users. One thing we learned is that people like to know how data collected about when they are frustrated is used. It is not enough to provide users with tools so that they may easily submit usability bug reports; these tools must also give users information about what information is collected, and when. Furthermore, these interfaces need to allow users to edit content transmitted about the machine state so that any private information may be removed. Consequently, the interface for the squeeze mouse performs a lot of tasks:

Capture data from analog to digital conversion board.
Give user feedback from the frustration sensor
Allow the user to record and transmit usability incidents.
Allow the user to control exactly what sort of information is reported, without invading privacy.

We settled on an unobtrusive "bulb" that appears on the user's taskbar (a system-wide menu that is available at the bottom of the screen for Windows users). As the user squeezes the mouse harder, the bulb flashes from green to yellow to red (bottom of Figure 3). When a threshold is exceeded, the system records a screenshot, along with a textual dump of the open windows. This is recorded with a running average of the recent pressure on the mouse. When the user clicks on the "bulb" a menu appears allowing the user to edit and send recent frustration events back to usability professionals. A graph can also be displayed for visual feedback.

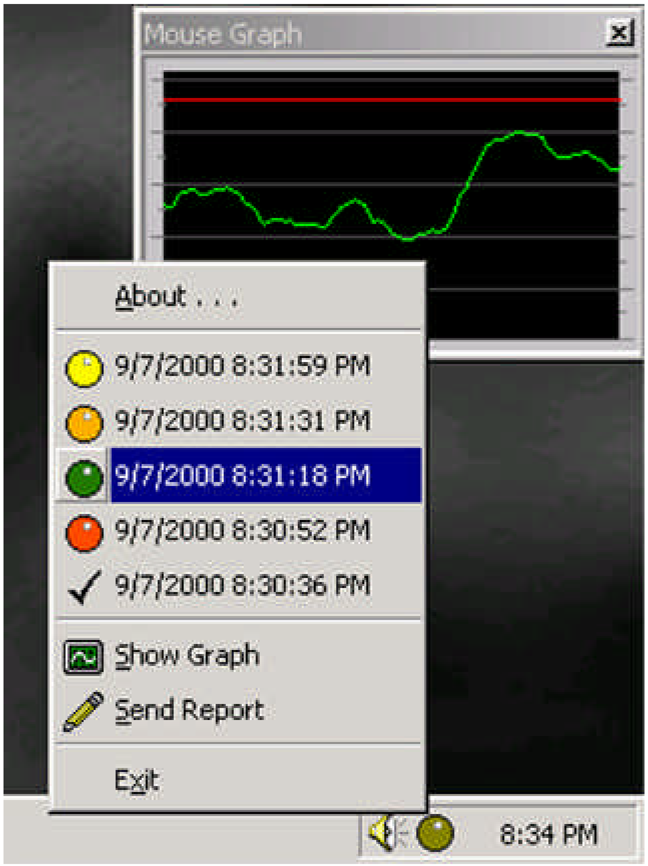

Figure 3: interface for squeezemouse shows level of pressure (graph) and record of events where user triggered feedback, color-coded by severity (red=most severe; green=mildest frustration).

3.2 Frustration Feedback Widgets

Some less elaborate, but quite useful "sensors" we have tested are user interface widgets designed for quickly and easily communicating frustration and usability issues. These have the distinct advantage of not requiring any elaborate hardware, such as special mice and boards. While the widgets don't allow for the same sort of physical expression of frustration, they do allow a greater degree of user control. They do not facilitate passive, unintentional detection.

Figure 4: thumbs-up and thumbs-down icons

One of the first prototypes we tried involved the use of thumbs-up and thumbs-down icons so that the user can register pleasure or displeasure with a particular system (Figure 4). But this early prototype did not provide a mechanism to communicate the severity of the usability incident, only whether a favorable or unfavorable event had occurred. Judith Ramey of the University of Washington had once mentioned that in usability tests, a cardboard "frustrometer" was used to help users express themselves [Ramey, personal communication]. Borrowing from this idea, we developed a software version. This interface allowed for a severity scale to be communicated (Figure 5).

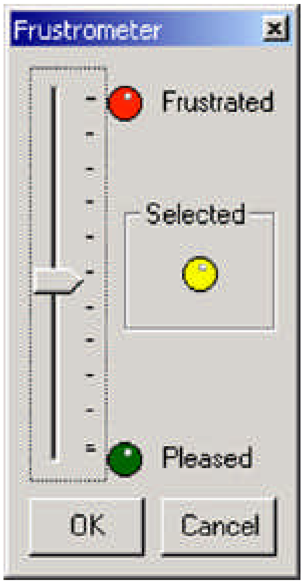

Figure 5: The Frustrometer

4 EVALUATION

As we iterated these interfaces it became increasingly clear that in order to be successful they would have to be considerably more comfortable than the standard mechanism that people use to communicate frustration about their computers: a web form. Consequently, we conducted a pilot study experiments to obtain subjective evaluations of the different sensors and interfaces we've designed. The methodology and preliminary results are sketched out below.

4.1 Methodology

To assess the utility of these sorts of sensors, we designed an experiment to compare two different frustration sensor designs and a more traditional customer-feedback form. After a bit of discussion we agreed that a web feedback form, like those currently in use on many websites represents one commonly used feedback mechanism. Consequently we designed a simple form for the control group to use.

Nine participants solicited from around MIT were asked to fill out a six-page registration sequence to a popular job search site. Additionally, they were asked to send feedback about what they liked and disliked on the website using the Squeezemouse, Frustrometer, or the control interface: a standard web based text form.

What the participants did not know until after the study was that the web pages were designed to be especially frustrating. This was achieved by violating known usability heuristics. Studies have shown that users respond poorly to varied, slow response times (Butler, 1983). Consequently, some pages were made to load especially slowly. To further exacerbate problems, certain long forms were designed so that no matter what sort of information was entered, the form would report errors that needed to be corrected, and forced the user to start filling in the page from scratch (Neilsen, 1999).

After the users completed the registration sequence, they were interviewed and asked to fill out a questionnaire for their condition. After being interviewed and filling out a brief questionnaire, users were debriefed and told of the deception carried out. It was emphasized that the deception was necessary, since it is very difficult to elicit emotional states like frustration if subjects know you are trying to frustrate them.

4.2 Preliminary Results

All participants were asked on the questionnaire about the usability and responsiveness of the registration sequence. The questionnaire presented a seven point scale from (Very Easy) to (Very Hard). For the purposes of this paper, we've chosen to label (Very Easy) as 1 and (Very Hard) as 7. The questions and mean responses (in brackets) are shown below:

How hard was the job registration web form to use?
(Very Easy) ● ● (3.08) ● ● ● ● (Very Hard)

How responsive was the job registration website?

(Very Fast) ● ● ● ● (4.63) ● ● (Very Slow)

Since we were actually interested in the performance of our various frustration feedback sensors, the remainder of the questionnaire dealt more specifically with the sensors. For instance, the participants were asked about how difficult it was to send feedback. The participants were asked about the feedback device for their condition specifically:

- Did you like using the [Web Form, Squeezemouse, Frustrometer] feedback device?
- Did sending feedback interfere with filling out the form?
- How interested are you in using the [Web Form, Squeezemouse, Frustrometer] again?

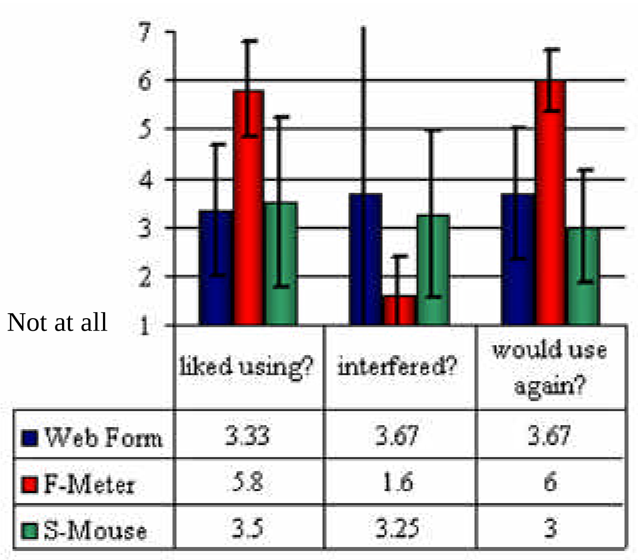

	liked using?	interfered?	would use again?
■ Web Form	3.33	3.67	3.67
■ F-Meter	5.8	1.6	6
■ S-Mouse	3.5	3.25	3

Figure 6: questionnaire mean results

Each condition had similar questions. For instance, the users of the web feedback form were asked "Did you like using the feedback page?" instead of "Did you like using the Squeezemouse feedback device?" The responses for each condition to these questions are summarized below in Figure 6.

After performing a single-tailed t-test on this data we found that participants were significantly more likely to say that they would use the Frustrometer again and that they liked using it. So why did the Frustrometer perform better than the Squeezemouse or the web form? We theorize that it is because the Frustrometer was the most accessible, most straightforward interface, and consequently seemed to require less cognitive overhead when reporting feedback. In short: it was less frustrating and distracting than the other options.

Overall users were enthusiastic about being able to send in feedback. One participant noted "I think being able to send feedback while in the middle of a process is cool and sort of prevented me from really losing my temper." Another noted, "The feedback option gave me a sense of power, in the sense that I could complain or compliment about features I dislike or like." Most users seemed to respond positively to the convenience of an accessible and easy to use feedback mechanism: "I liked it being set up such that as soon as I realized there was a problem, I could gripe."

REFERENCES

Butler, T. W. (1983) Computer Response Time and User Performance, in Proceedings of CHI '83.

Breazeal, C. (2000) Sociable Machines: Expressive Social Exchange Between Humans And Robots. PhD Thesis, Massachusetts Institute of Technology.

Fernandez, R. and Picard, R. W. (1998), "Signal Processing for Recognition of Human Frustration," *Proc. IEEE ICASSP '98,* Seattle, WA.

Kirsch, D. (1997) The Sentic Mouse. <http://www.media.mit.edu/affect/AC_research/projects/sentic_mouse.html>.

Neilsen, J. (1999) Top Ten New Mistakes of Web Design. <http://www.useit.com/alertbox/990530.html>

Reeves, B. & Nass, C. (1996). The media equation : how people treat computers, television, and new media like real people and places. Cambridge University Press, New York.

Sheridan, T. (1975) Community Dialog Technology. Proceedings of the IEEE 63, 3, 463-475.

WinWhatWhere Investigator. <http://www.winwhatwhere.com/index.htm>.