ELSEVIER

Interacting with Computers 14 (2002) 119–140

Interacting with Computers

www.elsevier.com/locate/intcom

This computer responds to user frustration: Theory, design, and results

J. Klein[a,*], Y. Moon[b], R.W. Picard[c]

[a]*iRobot Corporation, Twin City Office Center, 22 McGrath Highway, Suite 6, Somerville, MA 02143, USA*
[b]*Harvard Business School, Morgan Hall 197, Boston, MA 02163, USA*
[c]*MIT Media Laboratory, Room E15-383, 20 Ames Street, Cambridge, MA 02139-4307, USA*

Abstract

Use of technology often has unpleasant side effects, which may include strong, negative emotional states that arise during interaction with computers. Frustration, confusion, anger, anxiety and similar emotional states can affect not only the interaction itself, but also productivity, learning, social relationships, and overall well-being. This paper suggests a new solution to this problem: designing human–computer interaction systems to actively support users in their ability to manage and recover from negative emotional states. An interactive affect–support agent was designed and built to test the proposed solution in a situation where users were feeling frustration. The agent, which used only text and buttons in a graphical user interface for its interaction, demonstrated components of active listening, empathy, and sympathy in an effort to support users in their ability to recover from frustration. The agent's effectiveness was evaluated against two control conditions, which were also text-based interactions: (1) users' emotions were ignored, and (2) users were able to report problems and 'vent' their feelings and concerns to the computer. Behavioral results showed that users chose to continue to interact with the system that had caused their frustration significantly longer after interacting with the affect–support agent, in comparison with the two controls. These results support the prediction that the computer can undo some of the negative feelings it causes by helping a user manage his or her emotional state.

Keywords: User emotion; Affective computing; Affect; Social interface; Frustration; Human-centred design; Empathetic interface

1. Introduction

Despite the best efforts of designers, computer interactions often leave users feeling frustrated. With elevated levels of adrenaline and other neurochemicals coursing through the body, a person feeling frustration not only has diminished abilities with respect to

* Corresponding author.
E-mail address: jonathan@irobot.com (J. Klein).

0953-5438/01/$ - see front matter
PII: S0953-5438(01)00053-4

attention (Kitayama and Niedenthal, 1994), memory retention (Kahneman, 1973), learning (Lewis and Williams, 1989), thinking creatively (Isen et al., 1987), and polite social interaction (Goleman, 1995), but a penchant for getting more frustrated in the immediate future as well. A frustrating interaction with a computer system can also leave a user feeling negatively disposed toward the system and its makers. If humans are predisposed to respond socially to computers, as has been demonstrated (Reeves and Nass, 1996), such negative experiences could alter perceptions of trust, cooperation and good faith on the part of the user. Worse, such experiences could injure what, to users, may be one of their most important working relationships.

What can and should a computer do when its user gets frustrated while using it? Frequently mentioned solutions include either (1) trying to determine and fix the problem that is causing the frustration, and/or (2) preemptively trying to prevent the problem from happening in the first place. Indeed, the latter has been the traditional, if tacit, approach of the Human–Computer Interaction (HCI) community for years (e.g. Norman, 1988).

However, while both approaches are important, neither has succeeded in eradicating the problem of user frustration, and there have been few attempts to search for alternatives to these approaches. More specifically, very little research has explored the possibility that a computer can be used to *address* the emotional state of its user. This is the case in spite of the fact that the HCI community consistently claims to put its primary emphasis on the user experience:

> Underlying all HCI research and design is the belief that the people using a computer system should come first. Their needs, capabilities and preferences for performing various activities should inform the ways in which systems are designed and implemented. People should *not* have to change radically to 'fit in with the system,' the system should be designed to match their requirements. From Human–Computer Interaction**,** by Preece et al. (1994), p. 15

Recently, a number of researchers have begun exploring these 'mismatches' in human–computer interaction by exploring how computer systems may be designed to support the fact that humans feel, as well as think and act (Picard, 1997). After all, humans are much more than information processors. Humans are affective beings, motivated to action by a complex system of emotions, drives, needs, and environmental conditioning (e.g. Myers, 1989).

In the current research, the possibility that computers can be designed to respond meaningfully to negative user emotions such as frustration is investigated. Using controlled experimentation, this research demonstrates that computers can relieve feelings of frustration by providing 'active support' for the user's own ability to manage and regulate his or her emotions. (Active supports, as well as passive ones, are defined below.) Indeed, the research suggests that a computer that arouses a high frustration level in its user can be designed such that the user emerges from the interaction feeling *positively*—rather than negatively—disposed toward the system. A discussion of a number of risks and benefits for this new genre of HCI may be found in Picard and Klein (2001), a companion piece in this issue.

2. Background

Literature on emotion theory (e.g. Zajonc, 1998; Fridja, 1986; Izard, 1993) identifies a number of possible functions for emotions: as barometers to internal state, external predicament and progress in meeting needs and goals; as motivators nudging people to action to meet needs and goals (Tomkins, 1962); and as internal sirens to keep people from harm (Zajonc, 1998; LeDoux, 1994; Kagan, 1994).

Lawson (1965), after Rosenzweig, defines frustration as an emotional state resulting from 'the occurrence of an obstacle that prevent[s] the satisfaction of a need.' *Need* in this case can be interpreted to mean either a need or a goal (Amsel, 1992). Frustration is often associated with anger (Oatley and Duncan, 1994), and may often precede it. What is more, people experiencing frustration may find themselves feeling doubly exasperated: They not only have to grapple with the *source* of their frustration, but they also have to deal with the emotional reaction itself. Often, the latter can be as problematic, or even more so, than the former.

Fortunately, humans possess skills and strategies for *emotion regulation* that can mediate frustration levels to varying degrees (Gross and Muñoz, 1995). Emotion regulation is at once an aptitude and a skill for modulating and managing one's emotional state. Also described as emotion self-management, it has been identified as a primary component of *emotional intelligence*, a set of essential emotional skills, the development of which have been argued to correlate more than IQ with a person's success in life (Goleman, 1995, citing Vaillant, 1977; Felsman and Vaillant, 1987).

Humans are able to manage their own emotional states with varying degrees of success depending on the situation, their temperament (Kagan, 1994), and their degree of emotion-management skill, among other things. Failure to manage one's emotions can have profound effects including decreased productivity, an inability to pay attention and learn, injury to cooperative human relationships, increased personal stress and severe depression (Gross and Muñoz, 1995), and even addiction (Cooper et al., 1995) and other health problems (see Chapter 11 of Goleman for an overview).

People use a variety of methods to help manage their emotions, such as interacting with media and/or other people, engaging in sports or work, meditating or praying, using positive thinking, and consuming foods and other substances such as alcohol, tobacco, and other drugs. One strategy, known as *venting*, or giving people the opportunity to express their feelings in an unstructured manner, deserves special mention because of recent attempts by many online companies to solicit customer feedback via computer: Opinions are split over the benefits of unconstrained venting (or *catharsis*, in the psychology literature) when angry (Goleman, 1995, pp. 64–65). Some argue that venting has a calming effect, while others suggest that it can further intensify the emotion being expressed.

2.1. Active, social approaches to emotion regulation

We define two varieties of support for emotion regulation: passive supports and active ones. Passive supports are those used by people to manipulate moods, without necessarily addressing or discussing the emotions themselves. These include media, activities, food

and other substances. Interactions with people can fall into either category. For example, playing a team sport with others may help reduce negative feelings, without any discussion or acknowledgement of such feelings; this would be an example of passive support. In contrast, *active* support occurs when people discuss or otherwise address their emotions (and the emotion elicitors) directly, as a means of managing them. Talking to a parent or caring friend about what is upsetting, and how that makes the person feel, is an example of active support.

Active listening (e.g. Myers, 1989; Nugent and Halvorson, 1995) is perhaps the best known example of providing active support. It may be described as providing sincere, non-judgemental feedback to an emotionally distressed individual, with a focus on providing feedback of the emotional content itself. Active listening has its roots in social psychology (e.g. Raskin and Rogers, 1995), although psychotherapists from many other schools have adopted it as well. One need not be a psychotherapist to practice this skill; indeed, it is used by laypeople in such diverse areas as parenting (e.g. Gordon, 1970), education, crisis counseling, consumer affairs (e.g. Jenks, 1993) and other corporate venues. The key ingredient of active listening is believed to be paraphrased, emotional-content feedback, which involves letting the emotionally upset person know that his or her emotional state has been effectively communicated (Nugent and Halvorson, 1995).

Active listening, when practiced effectively, demonstrates elements of both *empathy* and *sympathy*. These terms are often used interchangeably, but the meanings are subtly different. Whereas the message of sympathy is 'as I understand it, this is what you are going through. This is how I would feel *in your shoes*,' empathy's message is more like 'this is how *I* feel *about* what you're going through. Here is *my emotional response* to *your* predicament' (Ickes, 1997). The tone of effective active listening feedback reflects these attributes and thus conveys that the person's emotional experience is understandable, acceptable and indeed, accepted by the listener. When active listening is practiced with these attributes, it is known to relieve strong, negative emotional states quickly and efficiently (e.g. Gordon, 1970).

2.2. Emotion regulation using computers

Current computer systems offer many ways of passively helping humans manage their emotions: People use multimedia PCs routinely to change moods by playing music, displaying humourous comic strips and animated movies, and playing games (Brennan and Shaver, 1995; Catanzaro and Greenwood, 1994; Goleman, 1995). Networked PC's also enable people to converse electronically, potentially providing active emotional support. All of these examples support users in managing their emotional states, yet they are often time-consuming, and people have to *seek* these interactions—they are not automatically offered by the system. In addition, these are generally forms of passive support; most computers today offer virtually no active support for personal emotional regulation.

Systems have been built that have been able to communicate with users in ways that involve their emotions in one form or another. *Eliza* (Weizenbaum, 1966) employed simple techniques for emotional interaction, some of which were borrowed from Rogerian active listening. Weizenbaum's stated purpose for building the system was solely to

explore natural language processing, yet Eliza is often appreciated for the illusion of intelligence and attention it tenders. Turkle (1995) discusses the related phenomenon of computer-based psychotherapy programmes that aim to bring about long-term positive change for the user's psychological health. Computer and robotic pets such as Furby (Kirsner, 1998) and Tamagocchi seem to foster some form of emotional communication, arguably offering passive support for emotional regulation. Other toys such as ActiMates' *Barney* (Strommen, 1998) and Hasbro's *My Real Baby* doll communicate affect to their young users, primarily in the form of affection and apparent attention to the child. While Eliza unintentionally approaches the capability for actively addressing the user's affective state, arguably none of these systems either do so intentionally or offer active support for emotion regulation at all.

Researchers have also begun to create interactive, computational models of affect for emotional communication with the user (Elliot, 1992; Tosa and Nakatsu, 1996), to build 'social skills' in robots (Scassellati, 1998; Breazeal, 1998), and to develop robots that can begin to 'empathise' with humans and make them feel 'more comfortable as they read emotional changes expressed in our faces' (Strauss, 1997). Such developments show potential as platforms for active emotion regulation support; however, none currently offer such support.

Several Internet sites and newsgroups have been created where users angry or frustrated with products and services can go and complain in public, e.g. alt.unixhaters, alt.fan.billgates, and such websites as the National Society for Microsoft Haters at http://www.tradepage.co.za/nsmh/. Such sites may offer users the ability to publicly complain about their frustrating experiences and products. Other users with similar experiences may chime in and reinforce the original user's complaint. Although these sites enable users to address their frustrations specifically by engaging in cathartic, public venting, as will be shown later these systems may not offer the most effective strategy for emotional support. Further, these interactions must be initiated by users and often require considerable time and effort, for example to make a convincing case in writing against a product. Since none of these systems offers users direct, structured engagement *about* their emotional state designed for the times when users experience strong, negative emotions, none may be considered to provide active support for emotion regulation.

2.3. Active machine support for user frustration?

In MATLAB 3.5, an older version of a high-level mathematics software package produced by MathWorks, Inc., a user could type in the popular expletive 'f***' at the software's command line, and the application would respond with a relevant quip, randomly chosen from a list of quips, such as 'Your place or mine?' This functionality has since been disabled in later versions of the software, and a sober-toned note to that effect has replaced the witty response. However, the story is germane to our discussion of methods of alleviating frustration because of the effect it has had on some users. A colleague recently related the experience of a student at MIT named 'Bob' who, while working in MATLAB 3.5 on a problem set, became frustrated to the point of despair over problems with MATLAB's syntax. In a fit of pique, unaware of the software's aforementioned response capability, Bob unwittingly typed in the word 'f***' at the command line.

When the machine responded with 'Your place or mine?', Bob's mood changed instantly. From utter frustration came surprise and delight. Here was a witty, unexpected 'hack', a response to a word considered taboo in formal elements of this culture, a response programmed long ago by one or more of the software's developers. It did absolutely nothing to address the problem that elicited Bob's frustrated state, yet the strong, negative state that Bob was experiencing suddenly dissolved. Bob became intrigued by the phenomenon, played around with the feature for a while, contacting friends to share the revelation of this hidden treasure, and then settled back down to work.

Hypotheses for the sudden turnaround in Bob's disposition abound: The experience was unexpected, humourous, ribald, playful, irreverent. The interaction's 'designers' probably gave little thought to planning this response, yet the fact remains that in the midst of experiencing a strong, negatively-valenced emotional reaction, Bob made a bid to communicate this state to the software and got a reaction. The machine's response was dumb—it was one of several simple strings pre-programmed into the system to write to the screen in response to a user typing in a certain string. However, for Bob it was also a direct, immediate, relevant response to his frustrated emotional state—at or near a peak of frustration—from 'another' to whom Bob expressed his feeling state (albeit crudely). Arguably, this simple hack produced some of the same components of feedback as are found in active listening: Immediate feedback of emotional information as well, perhaps, as conveyance of a message like 'your negative emotional state is understood and accepted, and here's a humourous remark that says 'it's okay'. One may argue that this event represents an example of purely passive emotional management support (i.e. 'it was just a diversion'). The possibility also exists that some of Bob's amusement may have come from a sense of connection to the *programmer*, and not the programme. Yet the intriguing possibility remains that components of active support may have been at least in part responsible for Bob's sudden change of mood, as well as the coincident interest in playing with the system—and his soon settling back down to work with MATLAB again.

3. Theory

A recent line of research has suggested that people display a natural propensity for interacting with machines as if they were other people; this is the so-called *Media Equation* of Reeves and Nass (1996). In one study, researchers found that users respond to praise and criticism from computers the same way they respond to similar feedback from humans (Fogg and Nass, 1997a). In another series of studies, users responded to computer 'personalities' the same way they respond to human personalities (Moon and Nass, 1996; Morkes et al., 1998; Nass et al., 2001). Indeed, social responses ranging from reciprocity (Fogg and Nass, 1997b) to politeness (Nass et al., 1999) have been found in human–computer interaction.

To date, none of this research has explored the possibility that users will respond positively to a computer that exhibits emotion-support skills. However, based on this literature, the primary hypothesis of this study is as follows:

> When a computer system creates frustration on the part of the user, it can relieve that frustration quickly and effectively by means of active emotion support.

To test this hypothesis, a system was built that directly addressed users' frustration using active emotion support. The system incorporated a variety of active support strategies that have been demonstrated to be effective in a variety of social contexts, as documented in the literature in social psychology, developmental and clinical psychology, sociology, education, emotion theory, parenting theory, communications, research in empathy and crisis management, and consumer affairs studies. Specifically, research has suggested that the following strategies can significantly improve an individual's ability to regulate his or her emotions in a positive way:

1. *Actively solicit information about the individual's state*, especially the individual's emotional state and feelings about the eliciting problem (not just a cognitive assessment of the problem). This query serves in part to acknowledge that the individual *has* emotions, and may be having an emotional reaction to stimuli (Nugent and Halvorson, 1995).
2. *Solicit information in a timely fashion*, i.e. the information solicitation must occur when the individual is upset enough to have the need for (and benefit from) the support that the system has to offer. Active listening is known to be ineffective in situations in which an individual simply wants information or wants to communicate information (Gordon, 1970).
3. *Make sure the user is able to express what she is feeling*. It does no use to solicit emotional feedback from an individual who is feeling *confused* and then provide him with choices that read 'I am: very sad|sad|neutral|happy|very happy about [product X]'.
4. *Provide feedback, especially on emotional content*. Show the individual that his emotional state has been effectively communicated, ideally by paraphrasing what the individual is trying to communicate (Nugent and Halvorson, 1995). The supportive agent must convey that the *idea* of the individual's emotions has been communicated, and not just parrot back the words used by the individual.
5. *Allow for repair if the feedback is judged wrong*. This step is critical in active listening, to give the individual a chance to clarify her initial response and make her feel as if her state is accurately conveyed to the listener. After all, especially at times of high arousal, language can easily fail to convey precise meaning, yet at such times it is often all the more important to the aroused person that she convey her situation and state accurately (Gordon, 1970).
6. *Convey a sense of sympathy to the user*. Miriam-Webster defines sympathy as 'the act or capacity of entering into or sharing the feelings or interests of another.' In this case, the goal is to encourage the individual to feel sympathised with and that the user is not alone in feeling the way s/he feels (Gordon, 1970).
7. *Communicate a sense of empathy to the user as well*. Along with sympathy, empathy can encourage an upset person to feel a sense of understanding and emotional connection with the listener, arguably of critical importance in establishing an atmosphere in which the individual can feel as if his emotions are understood and heard (Ickes, 1997).
8. *Convey to the user the sense that his emotional state is valid*. In practice, this sentiment is often expressed more through a calm, sympathetic tone of voice when paraphrasing the emotional-content feedback than any specific kind of statement. Note that this sentiment is different from either sympathy or empathy but can sometimes be

accomplished through either of those means. Sometimes, simply naming what the person is probably feeling is a way of communicating the validity of that feeling (Gordon, 1970).

4. Design of the agent

In an attempt to implement all the strategies proposed above, the interaction agent was designed with the following capabilities:

1. Following Strategy #1, above, the agent was designed to actively solicit information about the user's emotional state, focusing on emotions relating to the eliciting problem (see also Table 2).
2. It was designed to provide feedback within minutes following a frustrating situation, following Strategy #2, above (see also the Section 5, which describes how this synchronization was achieved).
3. Following Strategy #3, relevant labels were used to describe what users experienced (subjects in the study described below reported feeling elevated levels of frustration, and this label was employed—see also Table 2).
4. Also per Strategy #3, users were queried about their frustration level on a scale of emotional intensity ranging from 'Not frustrated in the least' to 'The most frustration I have ever experienced in my life playing a game' (see Table 3).
5. Following Strategy #4, above, paraphrased feedback of the user's emotional state was provided, and gauged to the level of emotion reported by the user (Table 3), helping to convey a sense of validation for the user's emotional state (and thus addressing Strategy #8, above).
6. Means for repair were built in to the interaction to enable the user to clarify and/or correct the agent if the system's feedback was judged to be incorrect, per Strategy #5, above. The user was provided with up to three chances to clarify and/or adjust his reported level of emotion.
7. Upon reaching a satisfactorily accurate statement of feedback from the system, the user was issued a statement of sympathy relevant to the user's emotional state and reported intensity level. As an implementation of Strategy #6 (above) this statement is perhaps the most problematic one given by a computer, since an expression of sympathy really is an expression of feeling, and the computer is incapable of truly feeling anything the user might feel. Since computers are not yet capable of genuine sympathic understanding or feeling, this is a form of *artificial sympathy*, but it is sincere in the sense that its designers' goal is to try to enable the user to feel *sympathised with*.
8. Also per Strategy #6, a sympathetic statement of apology was supplied to the user, as part of the sympathy statement.
9. A statement of empathy was also issued to the user, again relevant to emotional state and reported intensity level, addressing Strategy #7, above. This statement also serves to address Strategy #8, as do the sympathy responses described above.

Table 1
The six conditions in the 2 × 3 experiment

Questionnaire type	No delays in Game 1	Delays in Game 1
Ignore the user's state	$N = 12$ (6F, 6M)	$N = 12$ (6F, 6M)
Let the user Vent	$N = 11$ (5F, 6M)	$N = 12$ (6F, 6M)
Provide AFFECT–SUPPORT	$N = 11$ (5F, 6M)	$N = 12$ (6F, 6M)

5. Method

The first author built the principles described above into an interactive text-based agent (embedded in a graphical user interface, or GUI), which will hereafter be referred to as the 'AFFECT–SUPPORT' agent. This agent was designed to interact with users who were experiencing various amounts of frustration. The range of the agent's responses was carefully scripted in an effort to reflect the theoretical conceptualization of active listening described above.

The effectiveness of this agent was tested in a 2 × 3 study of 71 subjects (see Table 1) who were led to believe that they were play-testing a new computer game. Elements of

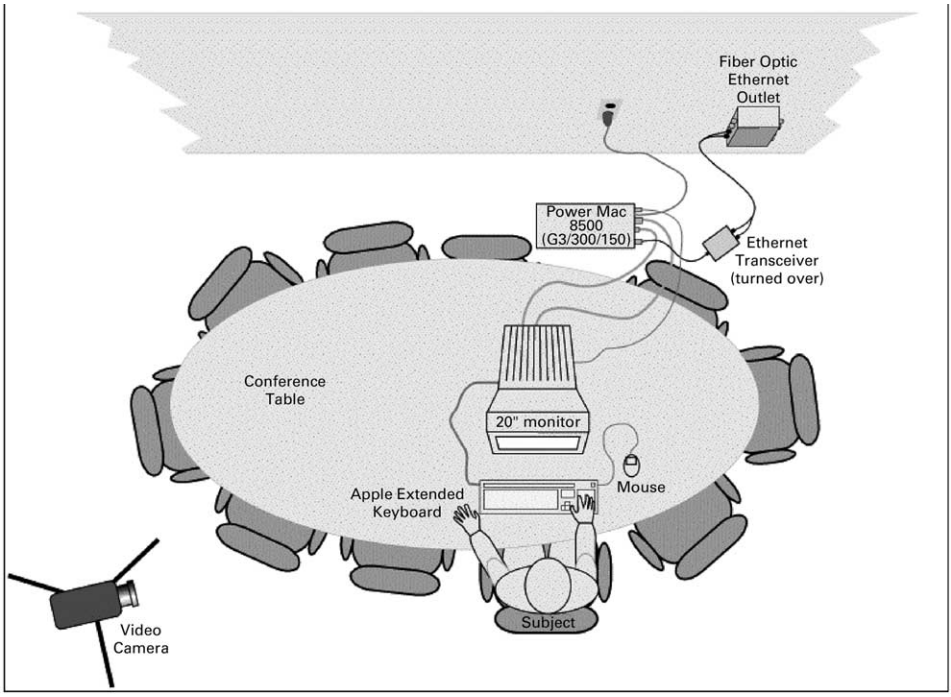

Fig. 1. Schematic of the physical set-up for the experimental chamber.

the protocol and cover story were adapted from the study created by Scheirer et al. (2001). After the experiment, all subjects were debriefed as to the true nature of the experiment and given the right to have their data withdrawn. One of the 71 subjects chose to have their data withdrawn, and that subject's data was not included in this study (Fig. 1).

5.1. Participants

Seventy human subjects were recruited, aged 14–44 (median 20–24). $N = 11$ or 12 for each of the six conditions, which included $n = 6$ males and $n \geqslant 5$ females in each condition, to account for gender in the study.

5.2. Procedure

Subjects were recruited via fliers posted around a college campus that offered $10 and a chance to win $100 for play-testing a new kind of game. Subjects were emailed a personality inventory questionnaire (Mehrabian, 1995) 'as part of the play-testing procedure', which were all filled out and returned at least one day in advance. Subjects came to the testing site individually, were read a briefing script, and were seated at an Apple Power Mac 8500/G3 computer with a 20″ display, keyboard and mouse. Two windows, barely overlapping, were displayed on the computer desktop: a Netscape browser with a simple, novel, graphical adventure game interface, and a text window displaying directions for the game (see Fig. 2).

Two level-of-frustration conditions were established. Subjects in the DELAY condition played the game for 5 min, during which they experienced nine seemingly-random 'web delays', in which the character froze on screen, but the on-screen timer continued to advance. Those in the NO DELAY condition also played for 5 min, but experienced no such delays. At 5 min, the game stopped automatically for all subjects. Each subject was then exposed to one of three different interactive sessions, under the guise of filling out a questionnaire about the game. The interactions corresponded to the three different conditions (IGNORE, VENT or AFFECT–SUPPORT).

The IGNORE interaction was actually a sequentially-revealed questionnaire—a string of dialogs—that contained questions that did not involve the subject's emotions, or any means of reporting a problem like screen delays (see Table 2). Responses were in the form of radio buttons with pre-defined labels. The VENT questionnaire was similar in form to the IGNORE questionnaire, except that it gave subjects open-ended means to report the relevant problem, as well as the subject's emotional state.

The AFFECT–SUPPORT interaction began in the same way as the VENT questionnaire, but when the computer prompted the subject for how frustrated s/he was feeling, the computer gave feedback based on the reported frustration level. Feedback was structured after the principles summarised above: offer emotional-content feedback and means for repair and correction, (e.g. 'It sounds like you felt fairly frustrated playing this game. Is that about right?'), as well as one statement of sympathy gauged to the reported frustration level (e.g. 'It sounds like you didn't have the best experience, then. That's not much fun') and one empathy statement (e.g. 'Sorry to hear things didn't go so well'). Each feedback sentence was customised to every possible response of the subject (see Table 3).

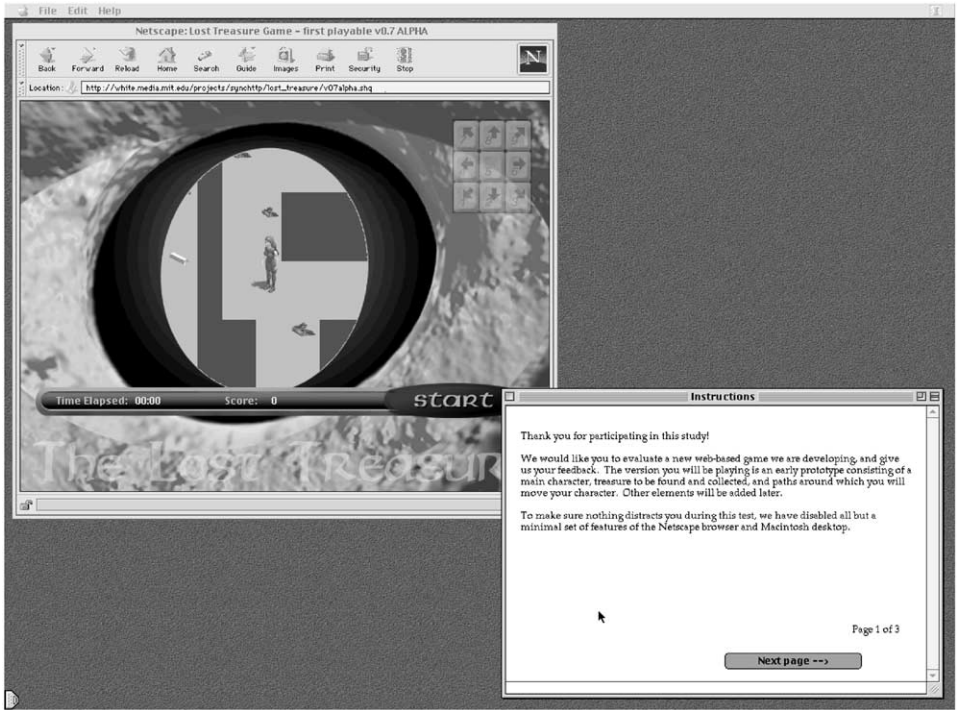

Fig. 2. A screenshot of the entire screen that each subject sees for the first time upon entering the chamber. Note that all elements in the display were tightly controlled for the experiment, as everything seen (including the grayed-out menu bar at the top) was actually specially created using Macromedia Director 6.0 for the Macintosh.

All subjects were then asked to play a non-delay version of the game for at least 3 min. At 3 min, a 'Quit' button (grayed out and disabled at the start of the second game) became active, but subjects could play longer if they wished, up to 20 min. However, the game was designed to be boring to play (there was no variation in reward schedule, no surprises, no new puzzles introduced, and no additional challenges or skill requirements beyond what was initially offered in Game 1), and subjects were given no incentive to play longer than 3 min. Upon finishing the second game, subjects were handed a paper questionnaire asking how they felt during various points of the whole experience. All self-report data was logged and recorded, and behavior was measured by automatically recording the time each subject played Game 2, as well as the score.[1]

[1] Physiological measures, such as those used in Scheirer et al. (2001), were not considered for this experiment, since they were beyond the scope of the driving questions of this research. The data from such measures might have been useful in related work being conducted by neighboring researchers, for example for correlating studies of reported frustration levels. It was decided, however, that adding to the complexity of this experiment by introducing such potentially curious measures was not desirable—and of course these measures were not supported by the cover story told to subjects prior to their participation.

Table 2
Contents of the three questionnaire-interactions, noting similarities and differences across conditions. Note that the AFFECT-SUPPORT interaction (column 3) includes an abbreviated set of questions, as this condition was time-balanced to account for the agent interaction. Due to the branching nature of the agent interaction, a diagramme of the entire interaction is not described here, but is summarized in the text and can be seen in its entirety in Klein (1999).

Questionnaire for the IGNORE control condition (no vent or affect support)	Questionnaire for the VENT condition	Questionnaire portion (before agent) or AFFECT–SUPPORT condition
First off, can you describe your age?	First off, can you describe your age?	First off, can you describe your age?
○ 10–15	○ 10–15	○ 10–15
○ 15–19	○ 15–19	○ 15–19
○ 20–24	○ 20–24	○ 20–24
○ 25–29	○ 25–29	○ 25–29
○ 30–34	○ 30–34	○ 30–34
○ 35–39	○ 35–39	○ 35–39
○ 40–44	○ 40–44	○ 40–44
○ 45–49	○ 45–49	○ 45–49
○ 50–59	○ 50–59	○ 50–59
○ 60–69	○ 60–69	○ 60–69
○ 70 +	○ 70 +	○ 70 +
Second, what is your sex?	Second, what is your sex?	Second, what is your sex?
○ Female	○ Female	○ Female
○ Male	○ Male	○ Male
Okay. Now how often have you played computer games or video games before?	Okay. Now how often have you played computer games or video games before?	Okay. Now how often have you played computer games or video games before?
○ 0 (Never)	○ 0 (Never)	○ 0 (Never)
○ 1 (One or twice in my life)	○ 1 (One or twice in my life)	○ 1 (One or twice in my life)
○ 2 (Fewer than ten times)	○ 2 (Fewer than ten times)	○ 2 (Fewer than ten times)
○ 3 (Between 10 and 20 times)	○ 3 (Between 10 and 20 times)	○ 3 (Between 10 and 20 times)
○ 4 (Used to play regularly, but haven't in ages)	○ 4 (Used to play regularly, but haven't in ages)	○ 4 (Used to play regularly, but haven't in ages)
○ 5 (About one a month)	○ 5 (About one a month)	○ 5 (About one a month)
○ 6 (Once a week)	○ 6 (Once a week)	○ 6 (Once a week)
○ 7 (Several times a week)	○ 7 (Several times a week)	○ 7 (Several times a week)
○ 8 (one a day)	○ 8 (one a day)	○ 8 (one a day)
○ 9 (At least once a day)	○ 9 (At least once a day)	○ 9 (At least once a day)
○ 10 (Many hours each day)	○ 10 (Many hours each day)	○ 10 (Many hours each day)
And how often (if ever) have you played games over a computer network?	And how often (if ever) have you played games over a computer network?	
(*frequency scale again*)	(*frequency scale again*)	

Table 2 (*continued*)

Questionnaire for the IGNORE control condition (no vent or affect support)	Questionnaire for the VENT condition	Questionnaire portion (before agent) or AFFECT–SUPPORT condition
Have you played role-playing or adventure games before? On computer or not, it doesn't matter. If so, about how often? (*frequency scale again*)	Have you played role-playing or adventure games before? On computer or not, it doesn't matter. If so, about how often? (*frequency scale again*)	
Hmm. And what about graphical adventure games on the Web, or graphical MUDs? Have you played any of these? (*frequency scale again*)	Hmm. And what about graphical adventure games on the Web, or graphical MUDs? Have you played any of these? (*frequency scale again*)	And how often (if ever) have you played graphical adventure games on the Web, or graphical MUDs? (*frequency scale again*)
Interesting. Have you ever played computer or video games in which the main character was female? If so—again—how often? (*frequency scale again*)	Interesting. Have you ever played computer or video games in which the main character was female? If so—again—how often? (*frequency scale again*)	Hmm. Have you ever played computer or video games in which the main character was female? If so—again—how often? (*frequency scale again*)
Okay. Now in this game, which of the following best describes the main character's appearance?	Okay. Now how well do you think you did in this game?	Okay. Now how well do you think you did in this game?
○ 1 (Character too small)	○ 10 (The best score!)	○ 10 (The best score!)
○ 2 (Character too large)	●	●
○ 3 (Character not visible enough)	○ 0 (The worst score)	○ 0 (The worst score)
○ 4 (Pattern of character's outfit needs adjustment)	How was the speed and smoothness of the network while you played? Were there any delays?	How was the speed and smoothness of the network while you played? Were there any delays?
○ 5 (Color of character's outfit needs adjustment)	○ 10 (Awful; big delays!)	○ 10 (Awful; big delays!)
Got it. And which of the following treasure items appeared most often?	●	●
○ 1 (Purple amethyst)	○ 0 (Smooth; no delays)	○ 0 (Smooth; no delays)
○ 2 (Black onyx)	Hmm. If there were any delays, do you think they affected your game?	Hmm. If there were any delays, do you think they affected your game?
○ 3 (Gold bar)	○ 10 (Yes, greatly!)	○ 10 (Yes, greatly!)
○ 4 (Red ruby)	●	●
○ 5 (Blue diamond)	○ 0 (No, not at all)	○ 0 (No, not at all)
Hmmm. What kind of treasure was easiest to spot?	How frustrated do you think you got playing the game, all things considered?	How frustrated do you think you got playing the game, all things considered?

Table 2 (*continued*)

Questionnaire for the IGNORE control condition (no vent or affect support)	Questionnaire for the VENT condition	Questionnaire portion (before agent) or AFFECT–SUPPORT condition
(*treasure list again*)	○ 10 (The most frustrated I have ever been in my life playing a game!)	○ 10 (The most frustrated I have ever been in my life playing a game!)
And what kind of treasure was hardest to spot?	●	○ 9
(*treasure list again*)	○ 0 (Absolutely not frustrated at all)	○ 8
As this game develops, which of the following would you most like to see?	And how much satisfaction did you experience as a result of playing?	○ 7
○ 10 (More non-violent game play)	(*0–10 scale as above, using* '*satisfied*')	○ 6
○ 9 (Interaction/cooperation with other characters/players)	Okay. Did playing the game make you feel anxious or tense?	○ 5
○ 8 (Social interaction with other characters/players)	If so, how much? (*0–10 scale as above, using* '*anxious or tense*')	○ 4
○ 7 (A more fanciful or interesting main character)	Overall, how much fun did you have playing the game?	○ 3
○ 6 (Puzzles to solve)	(*0–10 scale, using* '*fun*' *{and* '*had*' *in place of* '*been*'*}*)	○ 2
○ 5 (The character wearing armor and wielding weapons)	Did you experience playing the game make you feel angry? If so, how much?	○ 1
○ 4 (The ability to die and come back to life)	(*0–10 scale as above, using* '*satisfied*')	○ 0 (Absolutely not frustrated at all)
○ 8 (A male character instead of a female one)	Please describe your reaction to this experience.	
○ 9 (Monsters to kill)	(*Subject can enter text in large field here*)	**AFFECT–SUPPORT agent assumed control from here, adapting its content based on subject input.**
○ 10 (Monsters pursuing your character)	(END OF VENT CONTROL QUESTIONNAIRE)	
○ 11 (Traps to avoid)		
Assuming your top choice were implemented, what would you most like to see after that?		
(*Development feature list again*)		
(END OF CONTROL QUESTIONNAIRE)		

6. Results

All analyses were based on either a full-factorial ANOVA or ANCOVA. When appropriate, planned orthogonal comparisons[2] were conducted using Dunnett's t (see Winer et al., 1991) which adjusts for the inflated significance levels associated with multiple comparisons. All results were based on two-tailed tests.

A manipulation check[3] confirmed that the DELAY condition produced significantly more frustration than the NO-DELAY condition. In the post-experiment questionnaire, participants responded to the question, 'How much frustration did you feel immediately after playing Game 1?'. A full-factorial ANOVA revealed that participants in the DELAY condition rated their frustration level significantly higher after the first game ($M = 4.8$) than did participants in the NO-DELAY condition ($M = 3.56$), $F(1, 64) = 4.54$, $P < 0.05$. There were no other significant effects with respect to this variable.

The key prediction in this experiment was that subjects who were experiencing high levels of frustration, resulting from playing Game 1 in the DELAY condition, would feel greater relief from this state immediately after experiencing the AFFECT–SUPPORT interaction, in comparison with the two control conditions. Based on this improvement, subjects were expected to feel more positive affect toward the task, as well as to the source of their frustration—the game and the networked computer system on which it was played. Subjects in the DELAY/AFFECT–SUPPORT condition were therefore predicted to play longer than subjects in either the DELAY/IGNORE or DELAY/VENT conditions.

The results supported this prediction. Not only was there a main effect with respect to the type of feedback subjects received, $F(2, 64) = 8.00$, $P < 0.01$, but planned orthogonal comparisons indicated that participants in the AFFECT–SUPPORT condition played Game 2 for a significantly longer time than participants in either the IGNORE condition ($t(45) = 2.63$, $P < 0.01$), or the VENT condition ($t(44) = 3.97$, $P < 0.01$). (There was *no* significant difference in the game-playing behavior between participants in the IGNORE and VENT conditions.) These differences also remained statistically significant upon separate analysis of the DELAY vs. NO DELAY conditions. That is, subjects in the DELAY group who experienced the AFFECT–SUPPORT agent played Game 2 *significantly longer* (about 5.5 min longer, on average) than those in either the DELAY/IGNORE or DELAY/VENT conditions.

Similarly, subjects in the NO DELAY condition who experienced AFFECT–SUPPORT played Game 2 significantly longer than did those in the NO DELAY/IGNORE and NO DELAY/VENT conditions. In short, the AFFECT–SUPPORT led to longer play

[2] A planned comparison is a standard way of testing the significance of a specific comparison of two or more variables, the relationship of which are predicted by an hypothesis. Such a comparison is said to be planned, rather than 'fishing around' in the data for significant results. The term orthogonal simply means that all such comparison tests are 'independent' from each other—i.e. the result of one comparison test does not depend on the results from another. Planned orthogonal comparisons are a well-accepted, robust means of conducting statistical analyzes.

[3] A manipulation check tests that the manipulation of the variables in the experiment had the desired effect—in this case, that the DELAY condition was significantly more frustrating for subjects than the NO-DELAY condition.

in Game 2 in both the low and high frustration conditions, as compared to the controls (VENT and IGNORE).

There was one other main effect. Participants in the DELAY conditions ($M = 735$) played Game 2 for a significantly longer time than participants in the NO-DELAY conditions ($M = 545$), $F(1, 64) = 9.20$, $P < 0.001$ (see Fig. 3).

Three-way ANOVAs were also performed to test for the significance of gender, emotional (trait) arousability, and prior game play experience on the main behavioral results. No significant effect was found for any of these.

7. Discussion

7.1. Primary behavioral results

Many people play games to relax, and subjects in the DELAY condition all played longer than did those in the NO-DELAY condition. However, this difference can be attributed to what is known as *the rebound effect:* subjects who experienced delays in the first game did *not* experience them in the second game. Instead of relative tedium in Game 2, the DELAY subjects all experienced relatively more excitement and better action in Game 2, and a blissful end to the problems in Game 1—a pleasant surprise, which led to their playing longer. The rebound effect makes it meaningless to compare times *across* the DELAY/NO-DELAY conditions. What is significant is that *within* the DELAY condition, those who experienced the AFFECT–SUPPORT agent played significantly longer than those who did not. The same result was significant within the low-frustration NO-DELAY condition: those who experienced the AFFECT–SUPPORT agent played significantly longer than those who did not.

Several alternative explanations have been considered for the primary finding that the AFFECT–SUPPORT condition led to longer play the second time around. For example, subjects might have played longer because the AFFECT–SUPPORT agent interaction aggravated them, and Game 2 then became an example of subjects using passive support to deal with aggravation caused by the interaction. However, careful analysis of the data does not support this explanation. Specifically, subjects were asked on the Exit Questionnaire to report their level of frustration, as well as other emotions, at key intervals during the study. Overall, no significant correlation was found between self-report of frustration level after answering the interactive questionnaire and time spent playing Game 2. In addition, it is important to remember that in the DELAY conditions, it was *the computer/network system itself* that was the apparent cause of the subject's frustration levels. That subjects would perform a task on a system that frustrated them and made them feel less patient and more agitated, *for longer than* those who were less agitated by the system, is a much harder argument to defend than that of the primary hypothesis: That subjects played longer because they felt better, not worse.

What about the possibility that the AFFECT–SUPPORT condition was more 'friendly' in nature? Could this account for the differences in user behavior across the various conditions? This is possible but unlikely, since friendly, 'interested,' conversational language was peppered throughout all the interactions (see Table 2).

Table 3
A listing of all initial feedback sequences (displayed in text on-screen) that the AFFECT–SUPPORT agent provides, based on user-reported level of frustration. The responses were designed to appear natural and lifelike.

Reported level of frustration	INITIAL emotion-content feedback response by the system
0 Absolutely not frustrated at all	'It sounds like you weren't frustrated in the least playing this game. Is that about right?' (Yes/No)
1	'It sounds like you felt ever so mildly frustrated playing this game. Is that about right?' (Yes/No)
2	'It sounds like you felt a little frustrated playing this game. Is that about right?' (Yes/No)
3	'It sounds like you felt fairly frustrated playing this game. Is that about right?' (Yes/No)
4	'It sounds like you felt somewhat frustrated playing this game. Is that about right?' (Yes/No)
5	'It sounds like you felt pretty frustrated playing this game. Is that about right?' (Yes/No)
6	'It sounds like you felt pretty darn frustrated playing this game. Is that about right?' (Yes/No)
7	'Hmmm. It sounds like you felt really frustrated playing this game. Is that about right?' (Yes/No)
8	'Wow! It sounds like you felt terribly frustrated playing this game. Is that about right?' (Yes/No)
9	'Wow! It sounds like you felt unbelievably frustrated playing this game. Is that about right?' (Yes/No)
10 The most frustrated I have ever felt in my life while playing a game	'Good heavens! It sounds like felt unbelievably frustrated and completely fed up playing this game. Is that about right?' (Yes/No)

The most robust explanation seems to be that a text-only interaction that carefully applies social-affective skills known to succeed in human–human interaction can also successfully be used by a computer. Such a human–computer interaction need not take more than a few minutes; nonetheless, it can provide relief of negative emotional states related to frustration, as manifest in subsequent user behavior toward the *object* of the negative emotion.

7.2. Low variance for VENT conditions

Somewhat surprising were the uniformly low mean times that subjects played Game 2 in the VENT condition, both in the DELAY and NO-DELAY conditions. It was noted earlier that there is conflict in the literature about the benefits of emotional venting. The current findings lend support to the argument that venting is an unreliable emotional regulation strategy.

7.3. Little statistically meaningful self-report data

Much self-report data was taken, especially in the Exit Questionnaire, to try to capture subjects' affective states at various points throughout the experiment. Yet, despite the strong behavioral results described above, very little of the self-report data reflected the

behavioral finding. Our findings are thus consistent with those of emotion theorists, who have long argued that self-report data tends to be unreliable, yielding answers that may be put forth for a variety of reasons (Zajonc, 1998).

7.4. Theoretical and practical implications

The current research presents the first prototype of an interactive software agent that seems capable of providing active emotional support to a frustrated user. This agent was of very simple construction, and used no technology more advanced than elements found in a standard graphical user interface—elements such as dialog boxes, radio buttons, and text. Interaction with the AFFECT–SUPPORT agent occurred solely through traditional means: monitor, keyboard, and mouse. The entire transaction amounted to roughly reading 6 or 8 short lines of 12-point text, pressing a few buttons, and typing an optional blurb of text, so that it took very little time or effort to effect. A bit of natural-language construction was used to make the text output seem more natural and life-like.

There are both practical and philosophical implications of this work. While a more thorough discussion of these issues may be found in the companion paper in this issue (Picard and Klein, 2001), some practical implications are summarised below.

1. Computers are capable of alleviating strong, negative emotions, even when they are the source of these negative emotions.
2. It is possible to build agents that have 'social' capabilities, without the explicit personification of those agents. The implications of this study support the argument that social agents need not be personified characters, or use advanced interaction techniques such as speech I/O in order to be effective.
3. Building sites that simply let people vent may be about as effective as doing nothing at all, in terms of ameliorating the bad feelings.

7.5. Implications for designers and industry

The results from this study suggest ways in which designers might think differently about how to address failures in the design of software:

- consider the user's emotional state as a continuous factor in the design process;
- appreciate that offering no support for the user's emotional state during or after a system failure is tantamount to ignoring emotion, which can lead to a loss of opportunity for continued interaction with the user;
- text and buttons may be an under-utilised and overlooked method for creating agency.

The AFFECT–SUPPORT agent was shown to be effective not only at relieving strong, negative affect, but in making subjects feel better about using the same computer system (and performing very much the same activity) that had frustrated users in the first place. These findings have practical implications for virtually every kind of system that interacts with people:

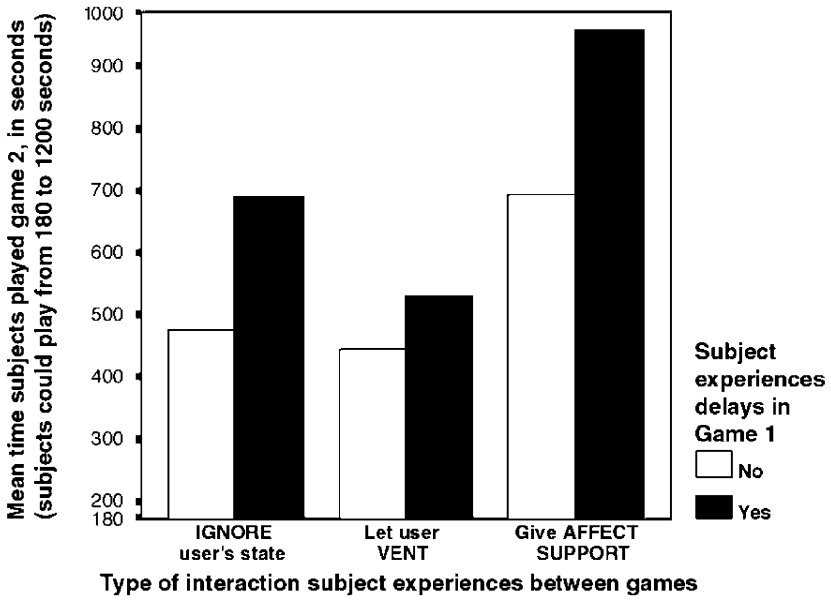

Fig. 3. Mean times that subjects played game 2, by response type and DELAY/NO DELAY conditions (subjects could play for $180 > T > 1200$ s–3–20 min)

- the method demonstrated here represents a cheap solution to addressing failure with an enormous amount of power; i.e. a big 'bang for the buck';
- the findings demonstrate that users who use a product that fails for some reason may actually wind up liking the product better with an AFFECT–SUPPORT agent onboard than without it;
- the idea of a positive outcome tied to the acknowledgement of a corporation's or product's failure (or even potential failure) may have important lessons for industry. The current assumption in corporate America seems to be that admitting failure, offering apologies, or taking any other kind of responsibility for poor performance is to be avoided at all costs—at least until a strong legal case is mounted against it. This work demonstrates that there may be benefit to customer satisfaction if such things as timely, sympathetic apologies are made, even by machines.

8. Future directions

One of the assumptions in this work is that the situation needs to be appropriate, and the time ripe, for the computer to actively broach the topic of frustration, or emotion in general. We tried to ensure this appropriateness in our experiment, by subjecting users to a situation that induced some level of frustration. One can also imagine an agent that the user might approach voluntarily when he or she was distressed, and simply wished to

express him or herself. However, if the agent were desired to be pro-active in responding to the user's emotion, it would need some ability to analyze the user's context and make predictions and/or queries to the user about his emotional state, or be able to recognise the user's emotional expressions. Although there has recently been an increasing amount of research toward this goal (e.g. see references in Picard, 1997, and in the companion paper by Scheirer et al., 2001), by and large it is an area for future work.

Clearly, work remains to be done to determine which components of the present emotion-support theory are most effective, which ones are less so, and which ones not tested might further improve the effectiveness of an AFFECT–SUPPORT-type agent. Is the paraphrasing feedback necessary and sufficient to yield effective results? Are both empathy and sympathy required to maintain effectiveness? Might humour or another technique be more effective in alleviating frustrating situations? Does it help, hinder, or maintain the agent's effectiveness to make more explicit the computer's inability to feel sympathy or other affective states? Does it affect the agent's effectiveness to make more explicit the idea that the computer itself is the social actor, and to try to remove all traces of (or for that matter, to try to boost) the connection between the agent (and other software of which it is a part) and its human designers? Does the idea that the information communicated by the user will eventually reach human eyes contribute to, hinder, or muddy the effect achieved by the AFFECT–SUPPORT agent? And how do all these vary with user personality, preference, situation, and other personal traits?

In the experiment, the AFFECT–SUPPORT condition was shown to be effective in relieving frustration, as well as the overall negative affect felt by subjects toward the source of their frustration—the computer. This relief occurred while subjects had no knowledge of the agent's existence, its goals or its method for achieving its goals. It remains to be seen if users will accept such an agent when they know about the agent beforehand, or if they use it more than once. These effects should be investigated in future tests of this agent and subsequent versions of it.

Acknowledgements

We would like to thank Clifford Nass and Sherry Turkle for their guidance and contributions toward this work. We should also like to extend thanks to Jocelyn Riseberg Schierer for her help and her experimental design, upon which some of this work is based, and to Raul Fernandez, Keith Oatley and Terry Winograd for their support and helpful comments.

References

Amsel, A., 1992. Frustration Theory. Cambridge University Press, Canada.

Breazeal, C., 1998. Regulating human–robot interaction using 'emotions', 'drives' and facial expressions. Presented at Autonomous Agents 1998 workshop 'Agents in Interaction—Acquiring Competence through Imitation', Minneapolis/St Paul, May.

Brennan, K.A., Shaver, P.R., 1995. Dimensions of adult attachment, affect regulation, and romantic relationship functioning. Personality and Social Psychology Bulletin 21 (3), 267–283.

Catanzaro, S.J., Greenwood, G., 1994. Expectancies for negative mood regulation, coping and dysphoria among college students. Journal of Counseling Psychology 41 (1), 34–44.

Cooper, M.L., Frone, M.R., Russell, M., Mudar, P., 1995. Drinking to regulate positive and negative emotions: a motivational model of alcohol use. 69 69 (5), 990–1005.

Elliot, C., 1992. The affective reasoner: a process model of emotions in a multi-agent system. PhD Thesis, Northwestern University, May. The Institute for Learning Sciences, Technical Report #32.

Felsman, J.K., Vaillant, G.E., 1987. Resilient children as adults: a 40-year study. In: Anderson, E.J., Cohler, B.J. (Eds.). The Invulnerable Child. Guilford Press, New York.

Fogg, B.J., Nass, C.I., 1997. Silicon sycophants: the effects of computers that flatter. International Journal of Human-Computer Studies.

Fogg, B.J., Nass, C.I., 1997. How users reciprocate to computers: an experiment that demonstrates behavior change. Proceedings of ACM CHI 97 Conference on Human Factors in Computing Systems, vol. 2. , pp. 331–332.

Goleman, D., 1995. Emotional Intelligence. Bantam Books, New York.

Gordon, T., 1970. PET: Parent Effectiveness Training. Plume Books (New American Library), New York.

Gross, J., Muñoz, R.F., 1995. Emotion regulation and mental health. Clinical Psychology Science and Practice 2 (2), 151–164.

Ickes, W. (Ed.), 1997. Empathic Accuracy. Guilford Press, New York.

Isen, A.M., Daubman, K.A., Nowicki, G.P., 1987. Positive affect facilitates creative problem solving. Journal of Personality and Social Psychology 52 (6), 1122–1131.

Izard, C.E., 1993. Four systems for emotion activation: Cognitive and noncognitive processes. Psychological Review 100 (1), 68–90.

Jenks, M., 1993. Ten easy steps to defusing anger. Customer Relationship Management 12 (2), 36.

Kagan, J., 1994. The Nature of the Child: Tenth Anniversary Edition. Basic Books, New York.

Kahneman, D., 1973. Arousal and attention. Attention and Affort. Prentice Hall, Englewood Cliffs, NJ, pp. 28–49 Ch. 3.

Kirsner, S., 1998. Moody furballs and the developers who love them. Wired, September, pp. 138.

Kitayama, S., Niedenthal, P.M. (Eds.), 1994. Heart's Eye: Emotional Influences in Perception and Attention. Academic Press, New York.

Klein, J., 1999. Computer Response to User Frustration. MIT Master's Thesis, January.

LeDoux, J.E., 1994. Emotion, memory and the brain. Scientific American June, 50–57.

Lewis, V.E., Williams, R.N., 1989. Mood-congruent vs. mood-state-dependent learning: implications for a view of emotion. In: Kuiken, D. (Ed.), Mood and Memory: Theory, Research, and Applications, Volume 4 of Special Issue of the Journal of Social Behavior and Personality (2), 157–171.

Mehrabian, A., 1995. Theory and evidence bearing on a scale of trait arousability. Current Psychology 14 (1), 4–28.

Moon, Y., Nass, C.I., 1996. How 'real' are computer personalities? Psychological responses to personality types in human–computer interaction. Communication Research 23 (6), 651–674.

Morkes, J., Kernal, H., Nass, C., 1998. Effects of humor in computer-mediated communication and human–computer interaction. In: Proceedings of the Conference of Human Factors in Computer Systems (CHI98 Summary), CHI 98, Los Angeles.

Myers, D.G., 1989. Psychology. 2nd ed Worth Publishers, New York.

Nass, C.I., Moon, Y., Carney, P., 1999. Are people polite to computers? Responses to computer-based interviewing systems. Journal of Applied Social Psychology 29 (5), 1093–1110.

Nass, C.I., Fogg, B.J., Moon, Y., 2001. Can computers be teammates?. International Journal of Human-Computer Studies (in press).

Norman, D., 1988. The Psychology of Everyday Things (a.k.a. The Design of Everyday Things). Basic Books, New York.

Nugent, W.R., Halvorson, H., 1995. Testing the effects of active listening. Research on Social Work Practice 5 (2), 152–175.

Oatley, K., Duncan, E., 1994. The experience of emotions in everyday life. Cognition and Emotion 8 (4), 369–381.

Picard, R.W., 1997. Affective Computing. MIT Press, Cambridge, MA.

Picard, R. W., Klein, J., 2001. Computers that recognise and respond to user emotion: theoretical and practical implications. In this issue of Interacting with Computers.

Preece, J., Rogers, Y., Sharp, H., Benyon, D., Holland, S., Carey, T., 1994. Human–Computer Interaction. Addison-Wesley, New York.

Raskin, J., Rogers, C., 1995. Person-centered therapy. In: Corsini, R.J., Wedding, D. (Eds.). Current Psychotherapies. 5th ed. F. E. Peacock Publishers, Inc, Itasca, IL.

Reeves, B., Nass, C.I., 1996. The Media Equation: How People Treat Computers, Television, and New Media Like Real People and Places. Cambridge University Press.

Scassellati, B., 2001. Imitation and mechanisms of joint attention: a developmental structure for building social skills on a humanoid robot. To appear as part of a Springer-Verlag series.

Scheirer, J., Fernandez, R., Klein, J., Picard, R.W., 2001. Frustrating the user on purpose: a step toward building an affective computer. In this issue of Interacting with Computers.

Strauss, S., 1997. How to make a robot smile. MIT Technology Review 100 (7), 14.

Strommen, E., 1998. When the interface is a talking dinosaur: learning across media with ActiMates Barney. In: Proceedings of the Conference of Human Factors in Computer Systems, CHI 98, Los Angeles.

Tomkins, S.S., 1962. Affect, Imagery, Consciousness, volumes 1 & 2: The Positive Affects and The Negative Affects. Springer, New York.

Tosa, N., Nakatsu, R., 1996. Life-like communication agent—emotion-sensing character 'MIC' and feeling session character 'MUSE'. In: Proceedings of the 1996 International Conference on Multimedia Computing and Systems (Multimedia '96).

Turkle, S., 1995. Life on the Screen. Simon & Schuster, New York.

Vaillant, G.E., 1977. Adaptation to Life. Little, Brown, Boston.

Weizenbaum, J., 1966. A computer program for the study of natural language communication between man and machine. Communications of the Association of Computing Machinery 9, 36–45.

Winer, B.J., Brown, D.R., Michels, K.M., 1991. Statistical principles in experimental design. 3rd ed. McGraw-Hill, New York.

Zajonc, R.B., 1998. Emotions. In: Filbert, D.T., Fiske, S.T., Lindzey, G. (Eds.). The Handbook of Social Psychology. 4th ed. The Handbook of Social Psychology, vol. 1., pp. 591–632.