TR-501 MIT Media Laboratory, June 17, 1999
Vision and Modeling Group Technical Report #501
Submitted to *Interacting with Computers*

This Computer Responds to User Frustration
Theory, Design, Results, and Implications

Jonathan Klein
IS Robotics, Inc.
Twin City Office Center
22 McGrath Highway, Suite 6
Somerville, MA 02143
jonathan@isr.com

Youngme Moon
Harvard Business School
Morgan Hall 197
Boston MA 02163
ymoon@hbs.edu

Rosalind W. Picard
MIT Media Laboratory
Room E15-383, 20 Ames Street,
Cambridge, MA
02139-4307
picard@media.mit.edu

**** NOTE: This is a submission draft, not for publication ****

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's text-only interaction used 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 thoughts 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-supporting 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. Implications of this solution are discussed, including possibilities for misuse or abuse, as well as potentially strong benefits when the approach is used correctly.

Keywords

User emotion, affective computing, social interface, frustration, human-centered design, empathetic interface

INTRODUCTION

Despite the best efforts of designers, computer interactions often make users 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 attention (Kitayama and Niedenthal 1994), memory retention (Kahneman 1973), learning (Lewis and Williams1989), thinking creatively (Isen, Daubman, and Nowicki 1997), 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 have a penchant for treating computers as if they were other people, as has recently been demonstrated (Reeves and Nass 1996), such negative experiences can alter perceptions of trust, cooperation and good faith on the part of the user. Worse, such experiences can 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. In addition, 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 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 (Myers 1989).

In the current research, the possibility that computers can be designed to respond meaningfully to common user emotions such as frustration is investigated. This paper presents a theory of how computers can address negative emotions such as frustration, together with the design and implementation of a software interaction agent, built by the first author, which implements this theory. The agent was evaluated using a study with 70 subjects and two control conditions, and was demonstrated to actively and significantly influence the behavior of frustrated users, apparently by helping relieve their frustration. Specifically, the system described herein provides active support for the user's own ability to manage and regulate his or her emotional states, particularly frustration, to the user's personal advantage, and with a number of side benefits. (Active supports, as well as passive ones, are defined below.) The study presents evidence that computer systems can do a great deal to support users when they experience frustration. In particular, 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. We describe a number of risks and benefits for this new genre of HCI.

BACKGROUND

Literature on emotion theory (e.g. Zajonc 1998, Fridja 1986, Izard 1990) identifies a number of possible functions for emotions: as barometers to internal state, external predicament and progress in meeting needs and goals; to help motivate people to action to meet needs and goals (Tomkins 1962); and to keep people from harm (Zajonc, 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's 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. In some cases, the latter may be more problematic 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 and Felsman & Vaillant 1987).

Humans are able to manage their own emotional states with varying degrees of success depending on the situation, their temperament (Kagan), 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 of cooperative human relationships, increased personal stress and severe depression (Gross and Muñoz), 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 deserves special mention because of recent attempts to incorporate it into the computing environment: *venting*, giving people the opportunity to express their feelings. Opinions are split over the benefits of unconstrained venting (or *catharsis,* in the psychology literature) when angry (Goleman, pp. 64-65). Some argue that venting has a calming effect, while others suggest that it can further intensify the emotion being expressed.

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 addressing or discussing the emotions themselves. These include media, activities, food and other substances. Interactions with people sometimes fall into this category. For example, playing a team sport with others may help reduce negative feelings, without any discussion or acknowledgment of such feelings. 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) may be described as providing sincere, non-judgmental feedback to an emotionally upset individual, with a focus on providing feedback of the emotional content itself. Active listening has its roots in Social Psychology (e.g. Raskin, N. J., and Rogers, C., Ch. 5 in Corsini and Wedding 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 theory (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).

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 empathy is “as I understand it, this is what you are going through. This is how I would feel *in your shoes.*”, sympathy’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 (Gordon).

Emotion regulation using computers

Current computer systems offer many ways of passively helping humans manage emotions: People use multimedia PCs routinely to change moods by playing music, displaying humorous 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 aren’t automatically offered by the system. In addition, these are all 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 programs that aim to bring about long-term positive change for the user’s psychological health. A discussion on the differences between these systems and active support for emotion regulation ensues below.

Computer and robotic pets such as Furby (Kirstner 1998) and Tamagocchi seem to foster some form of emotional communication, arguably offering passive support for emotional regulation. ActiMates’ Barney (Strommen 1998), Arthur and DW communicate affect to their young users, primarily in the form of affection and apparent attention to the child. However, no capability for actively addressing the user’s affective state has been designed into these systems.

Researchers have also begun to create interactive, computational models of affect for emotional communication with the user (Elliot 1990), to build “social skills” in robots (Scassellati 1998, Breazeal 1998), and to develop robots that can begin to “empathize” 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 the now-defunct "miCROsoFT sucKS! aND theSE peoPLE agREE!" (http://www.wamzle.com/blink/msft/msftagree.html) Such sites offer users the ability to publicly complain about their frustrating experiences and products. Often, other users with similar experiences will 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, 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, e.g., to make a convincing case in writing against a product.

Active Machine Support of User Frustration: An Illustration

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 been disabled in later versions of the software, and a sober-toned note to that effect has replaced the original response.) An undergraduate student at MIT named "Bob" was working in MATLAB 3.5 on a graphing problem set when he became frustrated to the point of despair over problems with MATLAB 3.5's syntax. In a fit of pique, unaware of the software's aforementioned response capability, Bob unwittingly typed in the word "f---" at MatLab's 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 this strong, negative state suddenly dissolved. Bob became intrigued by the phenomenon, played around with the feature for a while, contacted friends to share the revelation of this hidden treasure, and then settled back down to work.

There are many possible hypotheses for Bob's sudden turnaround, but some part of the truth is that this simple hack provided a response that was direct, immediate, and relevant 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 utilized a few components of active listening: Immediate, timely feedback, acknowledgment in relevant language (albeit loosely interpreted here), as well as conveyance of an empathetic message like "your negative emotional state is understood and accepted, and here's a humorous 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 arisen as a sense of connection to the *programmer*, and not the program. Yet the intriguing possibility remains that components of active support may have played a role in this dramatic turn-around of a frustrated user's emotions.

THEORY

A recent line of research has suggested that people display a natural propensity for interacting with machines as if they were other people (Reeves and Nass). 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). Indeed, social responses ranging from reciprocity (Fogg and Nass, 1997b) to politeness (Nass, Moon and Carney forthcoming) 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, we hypothesize that:

- 1) A computer system can help relieve negative emotions in a user by means of active support of emotion regulation, hitherto assumed to be a uniquely social function;
- 2) A computer system employing active emotion support can make a frustrated user feel better, relatively quickly, even when the computer is the source of the problem in the first place;
- 3) Such a system can be built using existing technology, without requiring strong AI.

To test these hypotheses, a system was built that embodies an emotion-support strategy for helping a frustrated user alleviate his or her frustration. The system is based on theory gleaned from 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. We summarize this theory as a set of principles that a system would try to implement to actively support the user's own emotion regulation:

- Actively solicit information about the user's state, especially the user'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 user *has* emotions, and may be having an emotional reaction to stimuli.
- This solicitation must be timely; i.e. it must occur when the user is upset enough to have need of (and benefit from) the support that the system has to offer. Active listening is known to be ineffective in situations in which a user simply wants information, or wants to communicate information (Gordon 1970).
- Make sure the user is able to express what she is feeling. It does no use to solicit emotional feedback from a user who is feeling confused and then provide him with choices that read "I am: very sad | sad | neutral | happy | very happy about [product X]." Instead, provide:
  - Appropriate labels;
  - Recognition vs. recall of names for relevant emotional states (the user may or may not have the vocabulary at hand for describing her state);
  - A method for describing not only the state name, but also *how aroused* the user is;
  - Multiple selections for emotional expression (especially for discrete emotions—e.g. a user should be able to say that she is frustrated *and* angry about a product, not simply one or the other). If using a pre-made list of emotional state descriptions, try to include an optional text field where a user may enter her own description of her emotions.
- Provide feedback, especially on emotional content; show the user that his emotional state has been effectively communicated, ideally by paraphrasing what the user is trying to communicate (Nugent). The supportive agent must convey that the *idea* has been communicated, and not just a string of words or phrases that is parroted back.
- 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.
- Communicate a sense of empathy to the user. Since computers are not yet capable of genuine empathic understanding or feeling, this is a form of *artificial empathy*, but it is sincere in its designers' goal to try to make the user feel *empathized with.*
- Convey a sense of sympathy to the user as well. Along with empathy, sympathy 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 user can feel as if his emotions are understood, and heard.
- Convey to the user the sense that his emotional state is valid. Perhaps the computer might communicate that other users have felt this way in this situation. Note that this sentiment is different from either sympathy or empathy.

The principles above are hypothesized to be a cohesive set of effective methods for providing active emotion management support. Still, other promising support strategies may prove effective as well:

- When speaking to angry customers, for example, consumer affairs representatives are counseled to employ a variety of interesting strategies (Jenks illustrates ten of these). Some of these strategies are akin to active listening, but others include agreeing with the customer, and apologizing, albeit without accepting or placing blame. The model of the service relationship underlying the customer/consumer affairs interaction may share much in common with the current traditions of HCI, particularly in agent-based interaction (Lieberman and Maulsby 1996). Such strategies might be employed by an emotion-supportive agent, albeit with care and sensitivity to context, and ensuring that the computer's capabilities are not being misrepresented.
- MATLAB 3.5 (see above) arguably employed humor in the service of supporting the user's emotion management. Humor, when used appropriately, can be an effective means of lifting spirits. Morkes, et al (1998) found that humans responded socially to machines designed to express humor. Since a sense of humor may vary greatly from person to person, and appropriateness of humor can range widely from situation to situation, many sources in the literature shy away from using humor (e.g. Jenks, Gordon).

DESIGN OF THE AGENT

Following the theory proposed above, the interaction agent was designed with the following capabilities:

- It actively solicited information about the user's emotional state, focusing on emotions relating to the eliciting problem (see Table 2 in the Method section, below).
- It was designed to provide feedback within minutes following a frustrating situation (see the Method section, below, for a description of how this synchronization was achieved).
- Relevant labels were used to describe what users experienced (the study described below specifically prompted users to feel frustrated, and this label was employed).
- 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 in Method, below.)
- Paraphrased feedback of the user's emotional state was provided, and gauged to the level of emotion reported by the user (see Table 3).
- Means for repair were built in to the interaction to enable the user to clarify and/or correct the agent if the feedback was judged to be incorrect. The user was provided with up to three chances to clarify and/or adjust his reported level of emotion, with multiple options at every stage.
- Upon reaching a satisfactorily accurate statement of feedback from the system, the user was issued a statement of empathy relevant to the user's emotional state and reported intensity level.
- A statement of sympathy was also issued to the user, again relevant to his or her emotional state and reported intensity level.
- A sympathetic statement of apology was supplied to the user, as part of the sympathy statement.

METHOD

The principles described above were built into an interactive text-based agent, 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. All of the agent's responses were carefully scripted in an effort to reflect the theoretical conceptualization of active listening described above.

The effectiveness of this agent was tested in a 2x3 study of 70 subjects (see Table 1, below) who were led to believe that they were play-testing a new computer game. The experimental protocol and all collateral documents were approved by the XXXXX Committee On the Use of Humans as Experimental Subjects prior to the study, and all subjects were debriefed as to the true nature of the experiment after the study, and given the right to have their data withdrawn. (NOTE: *"XXXXX" refers to an academic institution, the name of which was withheld for this blind review process, and will appear in the final version.)*

Table 1: The six conditions in the 2 x 3 experiment

Questionnaire	NO DELAYS in Game 1	DELAYS in Game 1
IGNORE user state	N = 12 (6F, 6M)	N = 12 (6F, 6M)
Let user VENT	N = 11 (5F, 6M)	N = 12 (6F, 6M)
Affect-support	N = 11 (5F, 6M)	N = 12 (6F, 6M)

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 >= 5 females in each condition, to account for gender.

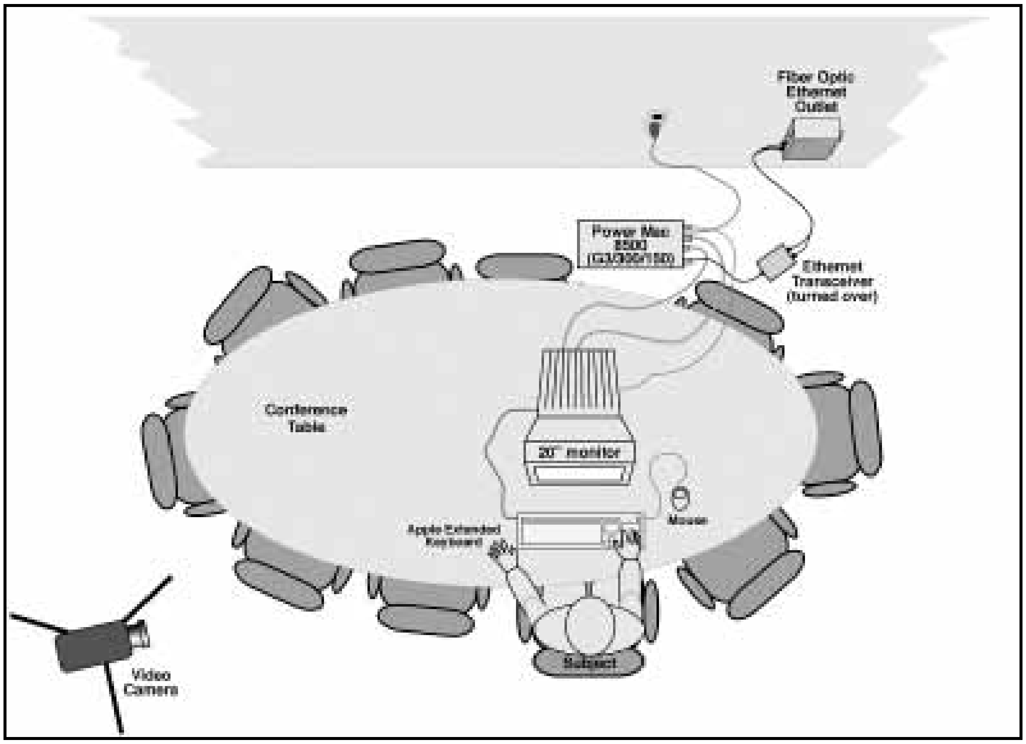

Figure 1: Schematic of the physical set-up for the experiment.

Procedure

Subjects were recruited via posted fliers 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, just 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 Figure 2).

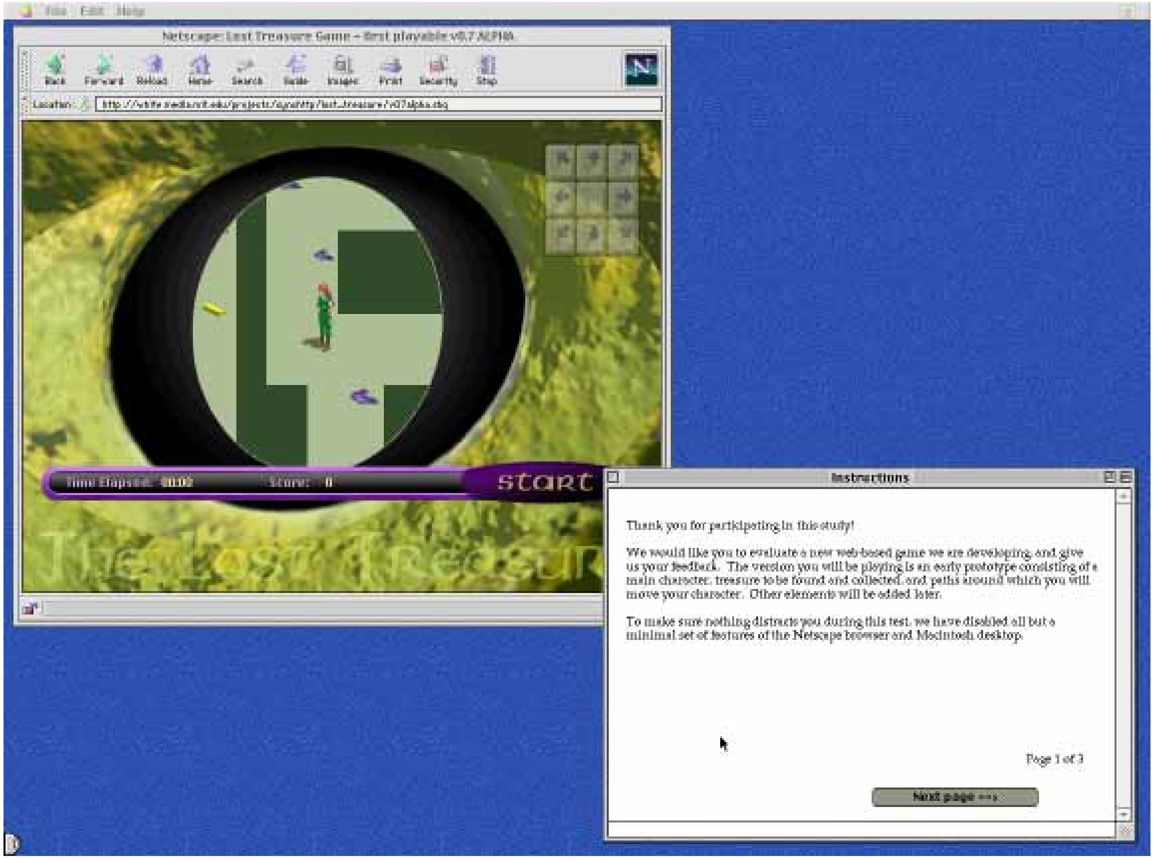

Figure 2: A screenshot of the entire screen that each subject sees for the first time upon entering the chamber.

Two level-of-frustration conditions were established, originally in an effort to compare a frustrating condition with a non-frustrating condition. Subjects in the DELAY condition played the game for 5 minutes, during which they experienced 9 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 minutes, but experienced no such delays. At 5 minutes, 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 actually was a sequentially-revealed questionnaire—a string of dialogs—asking questions that did not involve the subject's emotions, or any means for reporting a problem like screen delays (see Table 2). Responses were in the form of radio buttons with pre-defined labels. The VENT interaction was similar in form to the IGNORE, but gave subjects open-ended means to report the relevant problem, as well as the user's emotional state.

The AFFECT-SUPPORT interaction began in the same way as the VENT version, but when the computer prompted the user for how frustrated s/he was feeling, the computer gave feedback based on the reported frustration level. Feedback was structured after the principles summarized 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 empathy 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 sympathy statement (e.g., "Sorry to hear things didn't go so well."). Each feedback sentence was customized to every possible response of the user (see Table 3).

Table 2: Contents of the three questionnaire-interactions, noting similarities and differences across conditions

Questionnaire for control condition (no emotional vent)

First off, can you describe your age?

- ◯ 10-15
- ◯ 15-19
- ◯ 20-24
- ◯ 24-29
- ◯ 30-34
- ◯ 35-39
- ◯ 40-44
- ◯ 45-49
- ◯ 50-59
- ◯ 60-69
- ◯ 70 +

Second, what is your sex?

- ◯ Female
- ◯ Male

Okay. Now how often have you played computer games or video games before?

- ◯ 0 (Never)
- ◯ 1 (Once or twice in my life)
- ◯ 2 (Fewer than ten times)
- ◯ 3 (Between 10 and 20 times)
- ◯ 4 (Used to play regularly, but haven't in ages)
- ◯ 5 (About once a month)
- ◯ 6 (Once a week)
- ◯ 7 (Several times a week)
- ◯ 8 (Once a day)
- ◯ 9 (At least once a day)
- ◯ 10 (Many hours each day)

And how often (if ever) have you played games over a computer network?

(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 MUD's? Have you played any of these?

(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)

Okay. Now in this game, which of the following best describes the main character's appearance?

- ◯ 1 (Character too small)
- ◯ 2 (Character too large)
- ◯ 3 (Character not visible enough)
- ◯ 4 (Pattern of character's outfit needs adjustment)
- ◯ 5 (Color of character's outfit needs adjustment)

Got it. And which of the following treasure items appeared most often?

- ◯ 1 (Purple amethyst)
- ◯ 2 (Black onyx)
- ◯ 3 (Gold bar)
- ◯ 4 (Red ruby)
- ◯ 5 (Blue diamond)

Hmmm. What kind of treasure was easiest to spot?

(Treasure list again)

And what kind of treasure was hardest to spot?

(Treasure list again)

As this game develops, which of the following would you most like to see?

- ◯ 10 (More non-violent game play)
- ◯ 9 (Interaction/cooperation with other characters/players)
- ◯ 8 (Social interaction with other characters/players)
- ◯ 7 (A more fanciful or interesting main character)
- ◯ 6 (Puzzles to solve)
- ◯ 5 (The character wearing armor and wielding weapons)
- ◯ 4 (The ability to die and come back to life)
- ◯ 8 (A male character instead of a female one)
- ◯ 9 (Monsters to kill)
- ◯ 10 (Monsters pursuing your character)
- ◯ 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)

Questionnaire for "venting" condition

First off, can you describe your age?

- ◯ 10-15
- ◯ 15-19
- ◯ 20-24
- ◯ 24-29
- ◯ 30-34
- ◯ 35-39
- ◯ 40-44
- ◯ 45-49
- ◯ 50-59
- ◯ 60-69
- ◯ 70 +

Second, what is your sex?

- ◯ Female
- ◯ Male

Okay. Now how often have you played computer games or video games before?

- ◯ 0 (Never)
- ◯ 1 (Once or twice in my life)
- ◯ 2 (Fewer than ten times)
- ◯ 3 (Between 10 and 20 times)
- ◯ 4 (Used to play regularly, but haven't in ages)
- ◯ 5 (About once a month)
- ◯ 6 (Once a week)
- ◯ 7 (Several times a week)
- ◯ 8 (Once a day)
- ◯ 9 (At least once a day)
- ◯ 10 (Many hours each day)

And how often (if ever) have you played games over a computer network?

(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 MUD's? Have you played any of these?

(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)

Okay. Now how well do you think you did in this game?

- ◯ 10 (The best scorel)
- ◯ 0 (The worst score)

How was the speed and smoothness of the network while you played? Were there any delays?

- ◯ 10 (Awful; big delaysl)
- ◯ 0 (Smooth; no delays)

Hmm. If there were any delays, do you think they affected your game?

- ◯ 10 (Yes, greatlyl)
- ◯ 0 (No, not at all)

How frustrated do you think you got playing the game, all things considered?

- ◯ 10 (The most frustrated I have ever been in my life playing a gamel)
- ◯ 0 (Absolutely not frustrated at all)

And how much satisfaction did you experience as a result of playing?

(0-10 scale as above, using "satisfied")

Okay. Did playing the game make you feel anxious or tense? If so, how much?

(0-10 scale as above, using "anxious or tense")

Overall, how much fun did you have playing the game?

(0-10 scale, using "fun" [and "had" in place of "been"])

Did your experience playing the game make you feel angry? If so, how much?

(0-10 scale as above, using "satisfied")

Please describe your reaction to this experience.

(Subject can enter text in large field here)

Questionnaire portion (before CASPER agent) for CASPER condition

First off, can you describe your age?

- ◯ 10-15
- ◯ 15-19
- ◯ 20-24
- ◯ 24-29
- ◯ 30-34
- ◯ 35-39
- ◯ 40-44
- ◯ 45-49
- ◯ 50-59
- ◯ 60-69
- ◯ 70 +

Second, what is your sex?

- ◯ Female
- ◯ Male

Okay. Now how often have you played computer games or video games before?

- ◯ 0 (Never)
- ◯ 1 (Once or twice in my life)
- ◯ 2 (Fewer than ten times)
- ◯ 3 (Between 10 and 20 times)
- ◯ 4 (Used to play regularly, but haven't in ages)
- ◯ 5 (About once a month)
- ◯ 6 (Once a week)
- ◯ 7 (Several times a week)
- ◯ 8 (Once a day)
- ◯ 9 (At least once a day)
- ◯ 10 (Many hours each day)

And how often (if ever) have you played graphical adventure games on the Web, or graphical MUD's?"

(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 how well do you think you did in this game?

- ◯ 10 (The best scorel)
- ◯ 0 (The worst score)

How was the speed and smoothness of the network while you played? Were there any delays?

- ◯ 10 (Awful; big delaysl)
- ◯ 0 (Smooth; no delays)

Hmm. If there were any delays, do you think they affected your game?

- ◯ 10 (Yes, greatlyl)
- ◯ 0 (No, not at all)

How frustrated do you think you got playing the game, all things considered?

- ◯ 10 (The most frustrated I have ever been in my life playing a game)
- ◯ 9
- ◯ 8
- ◯ 7
- ◯ 6
- ◯ 5
- ◯ 4
- ◯ 3
- ◯ 2
- ◯ 1
- ◯ 0 (Absolutely not frustrated at all)

CASPER agent assumes control of questionnaire from here, adapting its content based on subject input.

(See Figure XX on next page for a flow diagram of CASPER agent interaction)

Table 3: A listing of all initial feedback sequences 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 felt a little frustrated you 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 you playing this game. Is that about right?" (Yes/No)

All subjects were then asked to play a non-delay version of the game for at least 3 minutes. At 3 minutes, the quit button became active, but subjects could play longer if they wished, up to 20 minutes. However, the game was designed to be boring to play, and subjects were given no incentive to play longer than 3 minutes. 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.

RESULTS

All analyses were based on either a full-factorial ANOVA or ANCOVA. When appropriate, planned orthogonal comparisons were conducted using Dunnett's $\underline{t}$ (see Winer, Brown, & Michels, 1991) which adjusts for the inflated significance levels associated with multiple comparisons. All results were based on two-tailed tests.

A manipulation check 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 (M = 4.8) rated their frustration level significantly higher after the first game than participants in the NO-DELAY condition (M = 3.56), $F(1, 64) = 4.54$, $p<.05$. 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 relief from this state immediately after experiencing the AFFECT-SUPPORT condition, 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 itself, and the networked computer system on which it was played. Subjects in the

DELAY/AFFECT-SUPPORT condition were therefore predicted to play longer, and perform better (i.e. score more points) than subjects in either the DELAY/IGNORE or DELAY/VENT conditions.

The results supported this prediction. Specifically, a full-factorial ANOVA revealed two main effects. First, there was a main effect for feedback-type, $F(2, 64) = 8.00$, $p<.01$. 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 ($t(45) = 2.63$, $p<.01$), or the VENT condition ($t(44) = 3.97$, $p<.01$). However, there was <u>no</u> significant difference in the game-playing behavior between participants in the IGNORE and VENT conditions. Subjects in the DELAY/AFFECT-SUPPORT condition also scored more points than any of the other conditions. However, upon analysis a close, linear correlation was found between amount of time users played and their score.

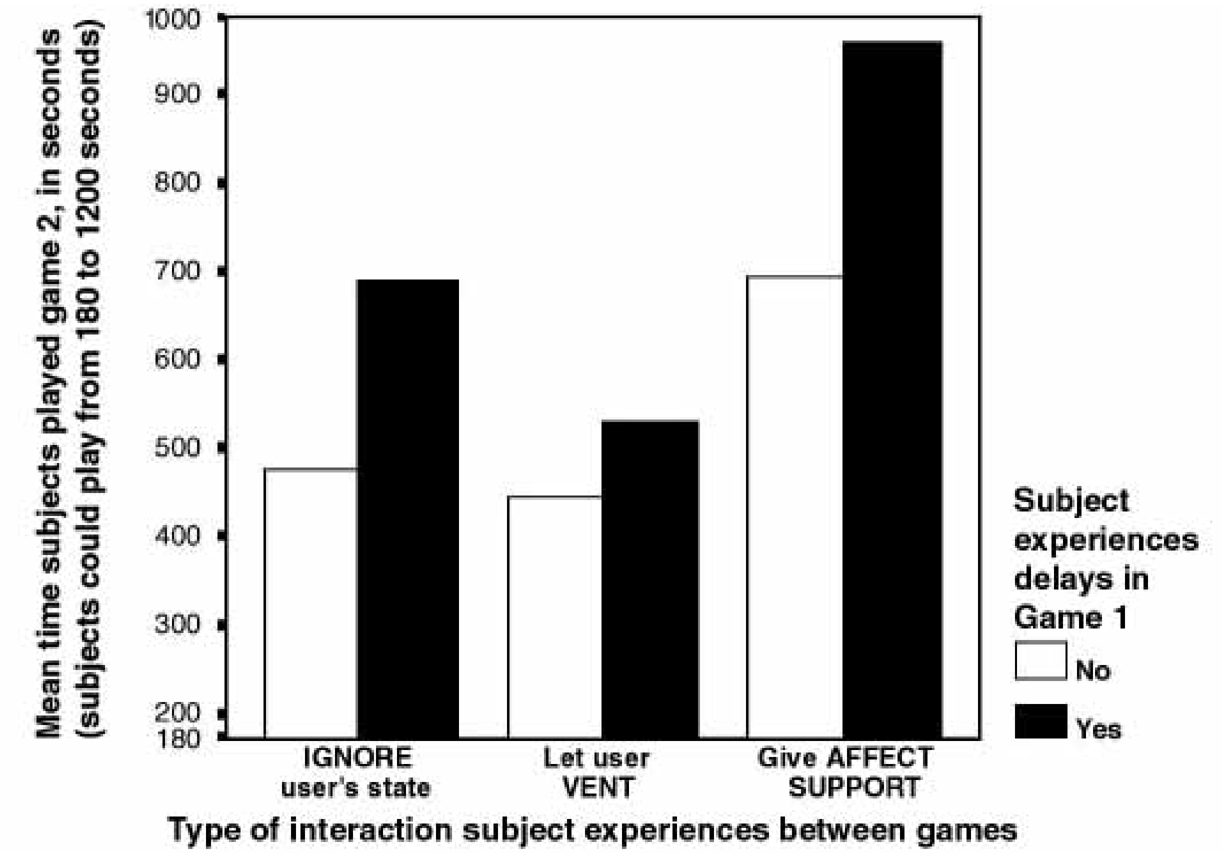

Figure 3: Mean times that subjects played game 2, by response type and DELAY/NO DELAY conditions. (Subjects could play for 180 >= T >= 1200 seconds—3 minutes to 20 minutes).

There was also a main effect for frustration, such that participants in the DELAY condition ($M = 735$) played Game 2 for a significantly longer time than participants in the NO-DELAY condition ($M = 545$), $F(1, 64) = 9.20$, $p<.001$. The interaction effect was not significant with respect to this variable. In other words, play times for all the DELAY conditions were consistently higher than those for the NO-DELAY conditions (see Figure 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 correlation was found for any of these.

DISCUSSION

The manipulation check was successful: subjects reported significantly higher levels of frustration in the DELAY condition than in the NO DELAY condition. However, the NO DELAY condition was not “non-frustrating” as we had planned, but still produced a low level of frustration. Consequently, the DELAY and NO DELAY conditions corresponded to high and low frustration conditions, respectively. Under statistical analysis, the behavioral measures of those in the DELAY condition revealed no meaningful difference between the amount of time that subjects in the IGNORE and VENT conditions played Game 2. However, those who experienced the AFFECT-SUPPORT agent displayed a significant improvement in the time they spent playing Game 2. Indeed, subjects in the DELAY/AFFECT-SUPPORT condition played much longer—over 5 and a half minutes longer, on average, than the other two conditions.

Primary Behavioral Results

Several explanations have been considered for the main finding that the AFFECT-SUPPORT condition led to longer play the second time around. 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 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 experienced relatively more excitement, better action, 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.

We also considered alternate explanations, such as that subjects might have played longer because the AFFECT-SUPPORT aggravated them, and Game 2 then became an example of passive support. Careful analysis of the data suggests that this explanation (and many others; see **** 1999) is not supported in general. Subjects were asked on the Exit Questionnaire to report their level of frustration, as well as other emotions, at key intervals during the study. These reports tended to follow the following pattern: Subjects generally reported feeling little or no frustration upon entering the room, then reported much higher frustration levels after Game 1, then slightly lower frustration levels (or the same) after the interactive questionnaire, then lower or the same levels after Game 2. On average, subjects in the DELAY/AFFECT-SUPPORT condition reported feeling either less frustrated by the interactive questionnaire than they did at the end of Game 1, or the same level of frustration for both. These reports were similar to those in all of the other conditions, with no statistically significant differences found between conditions. Scrutinizing the data, we found 2 of the 12 subjects in the DELAY/AFFECT-SUPPORT condition who reported feeling more frustrated after the questionnaire than after Game 1, but the other 10 followed the pattern described above.

Notably, the two subjects who deviated from the typical pattern each subsequently played Game 2 for the maximum allowable time (17 minutes). There are a number of possible reasons why these subjects may have played longer, the most obvious being well-known problems with self-report data, as discussed below. However, another possible reason may be that these subjects represent a type of personality that was untested for in this study—a type that is resistant, for one reason or another, to either the HCI implementation of the social cues, the cues themselves, or some other factor. Overall, however, no significant correlation was found between self-report of frustration level after answering the interactive questionnaire, and time spent playing Game 2. Finally, subjects were asked at the end of the Exit Questionnaire whether the interaction between games made subjects feel better or worse. Although there was wide, uncorrelated variation between 0 (no difference) and +5 (much better), no subjects reported feeling worse from interacting with it.

The reason subjects chose to continue interacting with the game seems best described by a reduction in frustration, as strongly shown in the rebound effect, and in the conditions where they interacted with the AFFECT-SUPPORT agent. After all, the game was not exciting to play, and it had some bugs in its interaction that all subjects experienced, that made it fairly tedious. Perhaps most important, though, was the fact that for those in the DELAY conditions, it was *the computer/network system itself* that was the 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 central argument in this hypothesis: That subjects played longer because they felt better, not worse.

The time spent responding to the interactive questionnaires is interesting to look at, although it does not account for the behavioral differences. The average times, in minutes, between the end of Game 1 (with NO DELAY) and the start of Game 2 were 4.38 minutes, 4.78 minutes, and 5.11 minutes for IGNORE, VENT, and AFFECT-SUPPORT, respectively. In the DELAY condition, these times were 4.08, 6.08, and 4.64 minutes, respectively. This time interval includes time spent on the questionnaire as well as time spent reading directions for Game 2. Noteworthy is that subjects in the DELAY/VENT condition took a greater amount of time to express themselves than any other group—and then proceeded to display the least amount of patience, on average, with Game 2 (see Figure 3).

Did the possibility that the AFFECT-SUPPORT condition was more "friendly" in nature account for the differences in performance? Unlikely, since friendly, "interested", conversational language was controlled for: We carefully peppered it throughout all questionnaire-interactions (see Table 2).

Alternatively, we have been asked if perhaps subjects played longer because the AFFECT-SUPPORT agent was novel and fun to interact with. This explanation also does not hold, because it is not the agent that they chose to interact longer with; rather, it was Game 2, which remained the same across all conditions. It also may be possible that the agent was found to be so much more entertaining that it provided a kind of passive support instead of the intended active support; however, the fact that it was fully scripted and there was no opportunity for subjects to "play" with it weaken this explanation. Further, self-report data from the exit questionnaire do not support this hypothesis.

The best explanation to date is 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 can provide relief of negative emotional states related to frustration, as manifest in subsequent user *behavior* toward the *object* of the negative emotion.

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. Because there is conflict in the literature about the benefits of emotional venting (as described earlier), it was predicted that variance for the VENT conditions would be high, particularly in the DELAY condition—exposing the unreliability of venting as an emotional regulation strategy. This was not the case. For the subjects in this experiment, the VENT interaction proved slightly worse, on average, than doing nothing at all to support the user's emotion regulation. In this respect, our results support the arguments of those emotion researchers who suggest that pure venting may actually help users become *more* frustrated, perhaps by reinforcing their focus on the negative situation. Although on the whole the difference between offering support for venting and no help at all (the IGNORE condition) was not statistically meaningful, on average doing nothing seems to be a little better (although yielding higher variance in times played) than supporting venting alone. However, both seem to be no match compared with the results achieved by the AFFECT-SUPPORT approach.

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 showing, or indeed anything of statistical consequence beyond the strong measure verifying the manipulation check. Our findings are thus consistent with those of other 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).

For example, self-report data for anger experienced was low, despite evidence of a close, established relationship between anger and frustration (Oatley and Duncan). Possible reasons for this phenomenon might be a lack of an available, identifiable target for subjects' anger, or that lab-induced frustration doesn't reach the intensity of what a subject would call anger. Additionally, subjects may feel stigmatized admitting to anger, which for many in this culture is a sign of weakness.

THEORETICAL AND PRACTICAL IMPLICATIONS

The first prototype has been built of an interactive software agent that has been demonstrated to provide an emotionally supportive interaction, based on demonstrating components of skills such as active listening, empathy, and sympathy. This agent is of very simple construction, and uses no technology more advanced than elements found in a standard graphical user interface (GUI)—elements such as dialog boxes, radio buttons, and text. Interaction with the AFFECT-SUPPORT agent occurs 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. A bit of natural-language construction was used to make the text output of the repair screens (which few saw) seem more natural and life-like.

There are both practical and philosophical implications of this work. Some of these are summarized and discussed below.

- Alleviating strong, negative emotional states is something computers can actively do with present technology. This finding has, to our knowledge, never been demonstrated before.
- Building sites that let people vent, without any acknowledgment of the user's negative experience and resulting affective state, may be less effective than doing nothing at all, in terms of ameliorating the bad feelings.
- The AFFECT-SUPPORT agent described in this study is clearly a social agent, with demonstrated social interaction capability. Yet the approach used here involves no brightly animated characters, no introductions, no cheerful greetings, and no name, so that as a social agent, it appears to be an extremely crude one. 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.
- The work reported in this paper represents an extension of the Media Equation: Beyond indications such as that flattery, criticism and other kinds of social cues can affect human users and their perceptions about machines that "express" them, this work demonstrates that strong, negative emotional states can be actively modulated through the use of a computer system.

Implications for Designers and Industry

From the standpoint of the human-computer interaction designer, there are also many implications, some of which are as follows:

- 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 underutilized 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:

- 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 acknowledgment of a corporation's or product's 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. This work demonstrates that there may be benefit to customer satisfaction if such things as timely, sympathetic apologies are made, even by machines.
- There may be practical applications for this work in helping to automate customer service and complaint departments. The idea of talking to a machine in order to complain to a company about a product or service may sound hard to buy, yet there are a number of benefits to the use of such systems that many members of society might actually prefer—particularly people who are shy. Being able to complain about a product, feel heard and understood, yet not have to risk interaction with a service representative who may be confrontational, and may not be trained (or have the patience) to deal with frustrated customers, may actually be attractive to those who fear confrontation. Such an automated service would ideally be offered in tandem with a real, well-trained human service department, and not in place of one.

Sociological and Philosophical Implications

The AFFECT-SUPPORT agent represents the first of a new genre of human-computer interaction. Such devices carry with them a broad host of new implications, positive and negative alike. On a commercial level, issues include the ethical use of such devices, and the potential that manufacturers may feel a diminished incentive to develop high-quality products when there is a readily available and effective tool for making the customer feel better despite the quality of her experience. On a personal level, increased frustration increases the likelihood of getting angry, and decreases abilities to pay attention, think (and problem-solve) creatively, and interact harmoniously with others. By reducing frustration, one can potentially render day-to-day existence less stressful and, perhaps, more pleasurable. Clearly such longer-term effects have yet to be shown, yet the results of this work do point in this direction.

Assuming that such an interaction constitutes a perception of a social exchange (for which there is some evidence in the literature), such interactions may have the effect of enabling humans to help meet their social needs—feeling a sense of social connection to others and combating loneliness—despite apparent large-scale cultural trends for the reverse: Increasing isolation (Myers 1993), Internet-based malaise (Edupage 1998), as well as increasing scarcity of human resources and human attention in a culture of divorce, two-income families, telecommuting, etc. Evidence suggests that, in this culture at least, precious few humans possess solid, effective, non-judgmental active-listening skills (Myers 1989, Goleman 1995. One possible reason for this lack of such skills in the larger culture may be a kind of positive feedback loop: Since so few people are available to practice such skills in day-to-day life, there are thus few people to serve as role-models for this behavior. An AFFECT-SUPPORT-like agent, however, is an easily-reproducible device that may, in some cases, act as a positive model for subsequent human-human interactions.

It may be argued that the kind of interaction of which the AFFECT-SUPPORT agent is capable may constitute a critical lack of authenticity, of believability and, therefore, credibility. This could be true for a variety of reasons, among which is that the machine is *simulating* human-like behaviors without truly *knowing* what the problem is, without experiencing empathy, sympathy, or really understanding the emotions that the person is experiencing. However, there is a kind of interaction with which humans partake every day, in which there a number of striking similarities with the AFFECT-SUPPORT approach, and in which positive, ameliorative effects on the part of the human are measurable and significant: Those that occur with pets, especially dogs and cats. Humans have long benefited from these interactions in many ways, including emotionally, even though it is clear that these animals may not understand the emotions of the human (Beck and Meyers1996).

Emotion Management Assistance vs. Emotion Manipulation

Providing active support for a person to regulate her own emotions may be seen as a tame version of far more nefarious concepts, such as involuntary emotion manipulation, brainwashing and mind control. While these concepts and their relationship to this research is a frank concern, there seems to be a major difference: The stated goal of this system is to support the *user's own emotion management.* It should, ideally, be incumbent upon the designers to support this goal, and for the user to use the system in this manner and in this role or not, and to receive any benefit at all from it or not.

However, such a deliberate intention may not always be the case—as was demonstrated in the design of this study. Whether or not the process at the heart of this interaction is voluntary is of critical importance to this work, since if it can be shown that a software agent can perform its work beyond the suspicions of the user, this indeed demonstrates a means for a kind of mind control. Troublingly, the experiment around which this research was based involves just such a deception, and one that was apparently not uncovered by most, if not all, subjects.

It should be noted that, as a first-of-a-genre device, this agent's capabilities might well subsequently become better known. Further, humans tend to develop ways of resisting efforts that are perceived as manipulative. Still, this research demonstrated strong behavioral effects that appear, from the self-report data, to have been almost undetected cognitively. Subjects appear to have felt much better, but were not aware of any specific manipulation. Therefore, it is possible that this approach to emotion regulation could be used in an involuntary manner—unscrupulous examples of which are not hard to imagine.

Another possible long-term pitfall of this work is that it may render the authentic, human-human interactions that it simulates less effective. Perhaps, in automating an artificially-derived human response, the net effect may be to wreck the foundation of efficacy for both the real and the simulated conversation. As Baudrillard wrote of a copy of the cave paintings of Lascaux, "the duplication is sufficient to render both artificial" (Baudrillard 1983). Whether humans become so jaded in interaction with AFFECT-SUPPORT-style software that they become less receptive to real human empathy is an open question, but one of great concern nonetheless.

Sapping Needed Wind From Sails

Another concern in this work is the possibility of diffusing emotions that need not and should not have been diffused. Since an important function of emotion is motivation, building frustration can often be seen as building fuel to overcome obstacles. Therefore, diffusing frustration and anger prematurely may act as a kind of "computational Prozac", undermining a person's ability (if not their right) to perform a potentially unpleasant task, such as confronting the company that sold her a poorly-designed system and demanding her money back.

In this situation, a system that tries to help diffuse strong user emotion performs a service for the manufacturer (lowering the number of complaints mandating better products), at the expense of the user. Indeed, it is not hard to imagine a system that convinces the user that it genuinely has the user's best interests "at heart", when the system is obviously undermining the user's goals by prematurely assuaging his emotional state.

Yet it is also possible that this strategy would quickly backfire. Users, as experts in social relationships, would soon see such a strategy as disingenuous; humans seem to have need of some semblance of authenticity in their interactions. If all a person ever receives in transaction with another is assuaging (without the eliciting problem ever being addressed in some reasonable timeframe, for instance), the person may become offended at this abuse of apparent trust. Similarly, a manufacturer that offers nothing but poor product after poor product, albeit with wonderful apologies after each, would eventually develop a bad reputation, albeit perhaps a bit more slowly than the company that offers no such consideration for the user's feelings.

What This Approach is Not: Computer Psychotherapy

Emphatically, this approach has little to do with computer psychotherapy. The stated goal of approaches to computer psychotherapy (outlined in Turkle) is the same stated goal of most psychotherapies: to somehow bring about healing from lasting pathological disorders via psychotherapeutic techniques. While the approach taken in this thesis is informed in part by techniques adapted from social psychology such as Active Listening (Ch. 5 of Corsini and Wedding), this approach is not intended as a treatment to heal long-term psychopathological problems. Rather, it is only intended to make it easier for the user to modulate his or her own emotional state in the very short term, in an environment that is otherwise unconducive to such support.

Further-Reaching Design Implications

This work has implications for user-centered design, in particular the design of systems that are able to interact with their users as social and affective beings, not just as information processors. Beyond helping users recover quickly and efficiently from strong, negative emotional states such as frustration, this work may also lead to side effects from its social-interaction nature, both positive and negative. Positive side effects may include a sense of rapport with computers, perhaps fostering an increased sense of cooperation and good feeling in the user, leading to increased productivity and job satisfaction. Negative side effects may include emotional and/or cognitive objections to perceived false (or imitation) caring on the part of the machine that uses empathy as an emotional management support strategy. Good interface design, however, includes communicating the system's capabilities and "intentions" clearly to the user (e.g. Preece), and this system is no exception. Proper design of an AFFECT-SUPPORT-type agent should go far toward mitigating negative responses, and no such objections were found in the self-report data collected in this study. However, objections (as well as other, unforeseen complications) may be found in subsequent tests of the AFFECT-SUPPORT approach, in particular when subjects are informed of the agent's existence, and goals, up front.

There are many more implications for this work than have been described here, both in interaction design and in the philosophy of how humans interact with computers. This work opens up new possibilities for human-computer communication, which in turn influences the way people form relationships with machines, and ultimately, the way that they think about these machines.

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 address 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 would only approach if he or she wished to talk about emotions. However, if the agent were desired to be pro-active about 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 recognize the user's emotional expressions. Although there has recently been some research toward this goal (e.g., Picard), 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 skills of an AFFECT-SUPPORT-type agent. Is the paraphrasing feedback necessary and sufficient to yield the same results? Are both empathy and sympathy required to maintain effectiveness? Might humor be more effective in frustrating situations, as was suggested by the MATLAB anecdote, or less so? Does it help, hinder, or maintain the agent's effectiveness to make more explicit the computer's inability to feel empathy 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 apart) 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 whether 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 the agent.

CONCLUSIONS

This work establishes that computer interfaces can be designed to actively help users recover from strong, negative emotional states, especially those related to frustration—including, but not limited to, the frustration that can arise from using computer systems. Further, the work shows that such interfaces have the added benefit of making computer users feel better *about* the very computer systems and applications that were involved in eliciting frustration in the user in the first place.

The implications for this work span domains including human productivity, customer satisfaction, product design, emotion theory, human-computer interaction, the fields of human psychology, sociology, communications, business administration, and even philosophy—not to mention the quality of day-to-day life for the average computer user.

Along with important practical contributions, this work, and its demonstration of a new genre of human-computer interaction, carries some potentially powerful implications for the culture and society. Some of these implications are negative, perhaps even risks to the fabric of this and other cultures, and to the way humans think about themselves. Many of the implications, however, are positive---helping people better manage strong negative emotions, and reap the benefits of decreased stress and increased productivity and creative thinking.

Finally, efforts that teach us more about emotionally-assistive technologies and what makes them succeed ultimately help us to create a much more complete picture of the human user than is currently conceived in many disciplines, including human-computer interaction. Human emotional needs, and the fact that nearly every single human user has them, need no longer be ignored as a "human factor" in the design of modern technology.

ACKNOWLEDGMENTS

To be included in the final submision, after the blind review process.

REFERENCES

****. (1999). *NOTE: This reference is to the first author's thesis, and would give away the identity of the authors; it will be included in the final submission pending approval for publication.*

Amsel, A. (1992). *Frustration Theory.* Cambridge University Press, Canada.

Baudrillard, J. (1983). *Simulations*. Semiotext(e), New York.

Beck, A. M., and Meyers, N. M. (1996). Health enhancement and companion animal ownership. *Annual Review of Public Health 17*, 1996, pp. 247-257.

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 1998.

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

Cassidy, J. (1994). Emotion Regulation: Influences of Attachment Relationships. *Monographs of the Society for Research in Child Development,* Vol. 59(2-3), 1994, pp. 228-283.

Catanzaro, S. J., and Greenwood, G. (1994). Expectancies for Negative Mood Regulation, Coping and Dysphoria Among College Students. *Journal of Counseling Psychology,* Vol. 41(1), Jan. 1994, pp. 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. *Journal of Personality and Social Psychology*, Vol. 69(5), Nov. 1995, pp. 990-1005.

Corsini, R. J., and Wedding, D., Eds. (1995). *Current Psychotherapies: Fifth Edition.* F. E. Peacock Publishers, Inc., Itasca, IL.

Damasio, A. R. (1994). *Descartes' Error: Emotion, Reason and the Human Brain.* G. P. Putnam's Sons, New York.

Edupage (author unknown), (1998). The Lonely Net. From *Edupage* [an Internet-based education technology news service], August 30, 1998. http://webserv.educom.edu/edupage/98/edupage-0830.html

Elliot, C. (1992). *The Affective Reasoner: A process model of emotions in a multi-agent system.* Ph.D. thesis, Northwestern University, May 1992. The Institute for Learning Sciences, Technical Report #32.

Felsman, J. K., and Vaillant, G. E. (1987). Resilient Children as Adults: A 40-Year Study. In Anderson, E. J., and Cohler, B. J. (Eds.) *The Invulnerable Child.* Guilford Press, New York.

Fogg, B.J., and Nass, C.I. (1997a). Silicon Sycophants: The Effects of Computers that Flatter. *International Journal of Human-Computer Studies.*

Fogg, B.J., and Nass, C.I. (1997b). How Users Reciprocate to Computers: An Experiment that Demonstrates Behavior Change. *Proceedings of ACM CHI 97 Conference on Human Factors in Computing Systems* 1997, v.2, p.331-332.

Frijda , N. H. (1986). *The Emotions. Studies in Emotion and Social Interaction.* Cambridge University Press, Cambridge.

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., and Muñoz, R. F. (1995). Emotion Regulation and Mental Health. *Clinical Psychology Science and Practice*, Vol. 2(2), Summer 1995, pp. 151-164.

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

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

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

Jenks, M. (1993). Ten Easy Steps to Defusing Anger. *Customer Relationship Management,* Vol. 12, No. 2, (1993). p. 36.

Kagan, J. (1994). *The Nature of the Child: Tenth anniversary edition.* New York: Basic Books.

Kahneman, D. (1973). Arousal and Attention. Ch. 3 of *Attention and effort.* Prentice Hall, Englewood Cliffs, NJ. pp. 28-49.

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

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

LeDoux, J. E. (1994). Emotion, memory and the brain. *Scientific American*, June 1994, pp. 50-57.

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

Lieberman, H., and Maulsby, D. (1996). Instructible Agents: Software that just keeps getting better. *IBM Systems Journal,* Vol. 35, Nos. 3 and 4.

Mehrabian, A. (1995). Theory and evidence bearing on a scale of Trait Arousability. *Current Psychology,* 14(1) Spring 1995. pp. 4-28**.**

Moon, Y., and 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, Second Edition.* Worth Publishers, New York.

Myers, D. G. (1993). *The Pursuit of Happiness: Who Is Happy and Why.* Avon Books, New York.

Nass, C. I., Moon, Y., & Carney, P. (forthcoming). Are people polite to computers? Responses to computer-based interviewing systems. *Journal of Applied Social Psychology.*

Nass, C. I., Fogg, B.J., and Moon, Y. (in press). Can Computers Be Teammates? *International Journal of Human-Computer Studies.*

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

Nugent, W. R., and Halvorson, H. (1995) Testing the Effects of Active Listening. *Research on Social Work Practice* 5(2), 1995, pp. 152-175.

Oatley, K., and Duncan, E. (1994). The Experience of Emotions in Everyday Life. *Cognition and Emotion,* Vol. 8 (4), July 1994, pp. 369-381.

Picard, R. W. (1997). *Affective Computing.* MIT press, Cambridge, MA.

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

Reeves, B. and 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. (in press). 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.

Strauss, S. (1997). How to make a robot smile. *MIT Technology Review,* Vol. 100(7), Oct. 1997, p. 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.

Tannen, D. (1990). *You Just Don't Understand: Women and Men in Conversation.* Balantine Books, New York.

Tosa, N, and Nakatsu, R. (1996). Life-Like Communic*ation Agent—Emotion-Sensing Character "MIC" and Feeling Session Character "MUSE". Proceedings of the 1996 International Conference on Multimedia Computing and Systems* (Multimedia '96).

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

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 (1966), pp. 36-45.

Zajonc, R. B. (1998). Emotions. In Gilbert, D. T., Fiske, S. T., and Lindzey, G. (Eds.) The Handbook of Social Psychology, Vol. 1 (4th Edition). pp. 591-632.