Computers that Recognise and Respond to User Emotion: Theoretical and Practical Implications

Rosalind W. Picard
Room E15-020g
MIT Media Lab
20 Ames Street, Cambridge, MA
02139-4307
picard@media.mit.edu

Jonathan Klein
iRobot Corporation
Twin City Office Center
22 McGrath Highway, Suite 6
Somerville, MA 02143
jonathan@irobot.com

ABSTRACT
Prototypes of interactive computer systems have been built that can begin to detect and label aspects of human emotional expression, and that respond to users experiencing frustration and other negative emotions with emotionally supportive interactions, demonstrating components of human skills such as active listening, empathy, and sympathy. These working systems support the prediction that a computer can begin to undo some of the negative feelings it causes by helping a user manage his or her emotional state. This paper clarifies the philosophy of this new approach to human-computer interaction: deliberately recognising and responding to an individual user's emotions in ways that help users meet their needs. We define user needs in a broader perspective than has been hitherto discussed in the HCI community, to include emotional and social needs, and examine technology's emerging capability to address and support such needs. We raise and discuss potential concerns and objections regarding this technology, and describe several opportunities for future work.

KEYWORDS
User emotion, affective computing, social interface, frustration, human-centred design, empathetic interface, emotional needs

1 INTRODUCTION: EMOTION AND USERS

What is it that humans need out of life, to develop, live, and flourish, individually and as part of a larger society? Moreover, as these needs are identified, yet are known to go unmet due to a paucity of human support, what role can and should computational technology play in helping to address and even satisfy those needs? We began our work several years ago by asking, then endeavouring to answer, these enormous questions. The work presented in this issue of Interacting with Computers (Scheirer et al, 2001, Klein et al, 2001) represents some of the first findings resulting from this inquiry.

Clearly, these questions bear some reckoning. As for the first question, Maslow (1987) famously postulated the existence of a hierarchy of human needs, from the basic (i.e., survival, security) to the sublime (e.g., self-esteem and what he termed "self-actualisation"). Much has been done, before and since, to explore this problem space. For instance, scientists have learned of strong developmental needs such as for basic stimulation: kittens that are blindfolded early in life will lose their ability to see (Hubel and Wiesel, 1962). Human babies who are cleaned, fed and otherwise well cared for, but not held and touched affectionately for long periods, will often simply wither and die—the cause of death called "failure to thrive" for want of more understanding, as was noted in orphanages early in the last century (Spitz 1945). Those lacking in basic social-emotional skills tend not to fare as well as those with well-honed skills (Goleman 1995). Together, these findings form at once a broad extension to a theory of human needs, and a foundation for our inquiry.

As for the second question – a role for computers – the work has barely begun. For their part, educational technologists have sought to harness the power of computers in the classroom, whether by enabling new kinds of learning experiences (e.g. Resnick 1998), mastery of materials by means of simulation (Schank and Jonah, 1991), or by recognising that learners have needs that are different from other kinds of users (Soloway et al, 1996). This work is important, and indeed the educational domain is only one in which the role of computers has been creatively explored for attaining human goals and thereby meeting many kinds of human needs. For us, however, the challenge has been to "pop up a level", to begin to see the user as a complete being, and at the same time as someone with many more needs that require satisfaction than simply the accomplishment of productivity and efficiency goals.

Despite interest in a recognition of the user as a complete being (e.g. Muller et al, 1997), much of the HCI community has been limited for years to measuring and responding to aspects of human behaviour such as efficiency, memory capacity, and delay times, all of which essentially reflect a view of the user as part of a productivity equation. On the positive side, such Taylorism helped the field gain its initial legitimacy. But, to view a human in this way is to see a mere fraction of the real situation. Indeed, humans are affective beings, motivated to action by a complex system of emotions, drives, needs, and environmental conditioning in addition to cognitive factors (e.g., Myers 1989). HCI researchers have increasingly begun to acknowledge that how the user feels is relevant, as evinced by their use of evaluations of interfaces with user satisfaction ratings, to gather a little affective feedback. Good designers have usually considered the emotional impact of their design, at least at some level. However, often the aim has simply been one of marketing—if users are satisfied, or feel good about using the product, or like the way it looks and feels, then they may be more likely to buy it. These steps, while important, still fall short of honouring the full value and role of human emotion.

Scientific findings suggest an increasingly large number of important functions of emotion. Emotion contributes not only to irrational behaviour, but also evidently plays a significant role in producing rational behaviour and rational decision making (e.g., Damasio, 1994). Emotions impact health; for example, rates of respiratory infection and rates of clinical colds have been shown to be increased with increases in psychological stress (e.g., Cohen et al, 1991). Emotions impact morbidity; for example, tendency toward hostility was found to be a stronger predictor of early death than factors such as high blood pressure, high cholesterol, and smoking, in a study tracking physicians from medical school to age 50 (Barefoot et al. 1983; Williams 1987; also chap 11 of Goleman 1995 lists many studies examining the influences of stress, anger, and depression on health and morbidity.). Even a small change in emotional state can significantly impact creativity, problem solving, and willingness to lend a hand to others (e.g., Isen et al. 1987; Isen 2000). As people spend greater amounts of time interacting with computers, the interaction has increased potential to influence their well being beyond that of simply helping them get tasks accomplished in a satisfactory way.

The work we have presented in this issue comprises pieces of a broader effort to begin to directly address the human emotional component in human-computer interaction (HCI). One aspect of this emotional component is the consideration of human needs beyond efficiency and productivity: What do humans need to learn in life about emotions, both their own and others', to help them lead rich fulfilling lives? Further, what kinds of emotional needs do humans tend to have on a day-to-day basis that, if unmet, can significantly degrade quality of life (and, if met, can improve the quality of life)? And, to focus our goals and work, what can we, as designers of computational technology, do to help?

In our work, we have had to be clear about what we are trying to do in a largely unexplored area of inquiry. In this paper, we present a high-level look at the work we have undertaken, and what it may mean for users and societies in which the work develops. We will examine the high-level process we went through to arrive at our research goals, describe work that we see as related, and begin to discuss the ramifications of this new work. In particular, we detail some of the potential benefits that led us into this line of inquiry in the first place. Along with this new effort, however, there arise a number of potential theoretical and practical concerns that merit further discussion. This paper aims to present these concerns and some of the key questions they raise, in order to initiate dialogue in the HCI community and beyond.

In the remainder of this section we define two categories of emotional needs and provide examples of computing that supports these needs—both directly and indirectly. The next section discusses the theory of imitating human-human interaction, together with our philosophy about moving beyond imitating humans. The rest of the paper delves into specific implications of the proposed technology, discussing potential theoretical and practical concerns.

1.1 Exploring Emotional Needs

The first step in the process was to be concrete about what constituted the idea of a human emotional need, and then to examine what role computers and other media might play in supporting such needs. Our first task was to examine the literature from psychology and other social sciences to see if others had examined the problem before. What we found was a large body of evidence for emotional needs that tended to fit into two large categories—but that, to our knowledge, has yet to be assembled into a cohesive taxonomy. Thus, we set about defining these two categories, which are described below.

1.1.1 Emotional Skill Needs

The first of the two categories we will label as *emotional skill needs.* Much has been written in recent years about emotional skill needs, albeit under the potentially misleading label *Emotional Intelligence* (Salovey and Mayer, 1990; Goleman 1995). Emotional skill needs may be thought of as a set of basic skills for understanding and handling emotions in oneself and others[1]. Most humans learn these skills to some degree as a matter of rote; for example, children generally develop empathy for others in distress by age two (Kagan 1994). However, variations in development of these skills can spell the difference between a life of misery and one of success and personal fulfilment. Poorly developed emotional skills can severely limit one's ability for forming and maintaining fulfilling relationships. Well-honed skills, by contrast, can lead to success, both at home and at work. Examples of emotional skill needs include (Salovey and Mayer, 1990):

— Emotional self-awareness: An ability to accurately appraise and appropriately express what one is feeling;
— Managing emotions: Handling and regulating feelings so that they are suitable to the culture and context (Gross and Muñoz, 1995); helping others when their emotions are heading out of control;
— Self-motivation: Harnessing emotions in service of a goal, for example, setting up a reward of delayed gratification to help motivate completion of a hard task;
— Affect perception: Accurately appraising what another is feeling from observing their non-verbal and verbal expressions and reasoning about their situation;
— Empathy: Appreciating what others are feeling, and communicating this understanding accurately to them.

1.1.2 Experiential Emotional Needs

[1] Some might argue with our calling these "needs." For example, most people have a need for some degree of social approval, and the skills involved in regulating emotion in social contexts may be viewed not as needs, but as means that help satisfy a need for social approval. Nevertheless, because emotional skills are foundational to the execution of a large variety of important human interactions, and because many of the latter could arguably not be successfully accomplished without the former, we maintain this category.

The second broad class of emotional needs may be termed *experiential emotional needs.* Experiential needs tend to be social in nature, as they are usually met via the assistance or presence of others. Chronically unmet experiential emotional needs can have the effect of degrading the quality of life. Long-term loneliness, for example, may be the result of extended isolation, and may be seen as the result of long-unmet needs for companionship or connection with others. Examples of experiential emotional needs include needs:

— For attention. Strong and constant in children, this need is more modulated in adulthood, although it by no means ceases. (Bowlby 1969);
— To feel that one's emotional state is understood by others, particularly if the state is intense (e.g. Ickes 1997);
— To feel that one's emotional responses, are acceptable by others (e.g. Gordon 1970);
— To feel that one's emotional experience and responses are considered "normal" or appropriate for a given situation (e.g. Myers 1989);
— To feel connected to others (e.g. Bowlby 1969);
— For companionship (as a basis for meeting many of these other experiential needs);
— For security (Maslow 1971);

1.2 Computers and the Emotional Needs of Users

Computers are beginning to address limited portions of the two categories of needs listed above. The system described by Scheirer et al. in this issue (Scheirer et al, 2001) describes an approach toward building computers that can begin to recognise aspects of user frustration. As an example of "affect perception"[2] by a computer, this system is designed to support a variety of emotional needs a user may have. For example, such a system may potentially be used to support the user's experiential emotional needs as the first step in enabling the user to feel as if his or her strong affective state has been effectively communicated. It should be noted that although that system is vastly more restricted in its scope than is human affect perception, the field is also very early in its development. As in other areas of computer recognition—speech, face, gesture, etc.—such systems are likely to gradually become more accurate and more robust. Whether or not they can attain 100% accuracy is an open question, but often there are applications where less than perfect accuracy is useful.

Systems that try to perceive user affect may find many potentially useful applications, for example:

- **As an experiential emotional aid:** Klein et al (2001) in this issue describes a system that was demonstrated to be effective in providing experiential emotional support of its users. The system employed knowledge of the user's emotional state to provide feedback to the user, together with cues from techniques for skilled listening that were borrowed from social psychology. Although the Klein, et al system does not use more sophisticated sensing of the user's emotional state than simple self-report, employment of an automated sensing system as described in Scheirer, et al is another logical approach.
- **As a pre-emptive tool:** Such a system might simply detect situations where the computer might improve its response in subsequent situations, thereby averting future user frustration.

[2] When we refer to "affect *perception*" or to "*recognition* of affect" we do not intend to imply that machines are conscious or human-like in how they perform such tasks. Our usage is one of convenience; we lack a better short phrase to replace what we really mean by machine affect recognition: "a computer system employing techniques such as signal sensing and detection, pattern analysis, probabilistic inference, and dynamic reasoning, in order to extract and characterise relevant patterns of sensed data in a way that produces a result similar to what a human would have produced if he or she had tried to observe and characterise the inputs according to their affective qualities."

- **As an emotional skill-building "mirror."** By providing feedback that a given strong emotional state has been sensed, the system may help the user to increase self-awareness of emotion, so as to hone emotional skills, and become better at managing emotion. In this area, researchers at the MIT Media Lab have built two systems that directly aim to help users in their need to learn about emotions. (1) A stuffed Tigger for children, which acts as a mirror of certain emotional states expressed by the child (Kirsch 1999); and (2) An interactive tutor for autistic children, which presents them with videos of characters engaged in emotional situations and prompts them to try to recognise the emotion that the character is likely to feel (Blocher, 1999). These applications are relatively non-controversial in that there is no deliberate attempt to support or manipulate the user's emotional state, only to educate and entertain. These systems are working examples of computers that aim to help humans develop skills for meeting emotional needs.

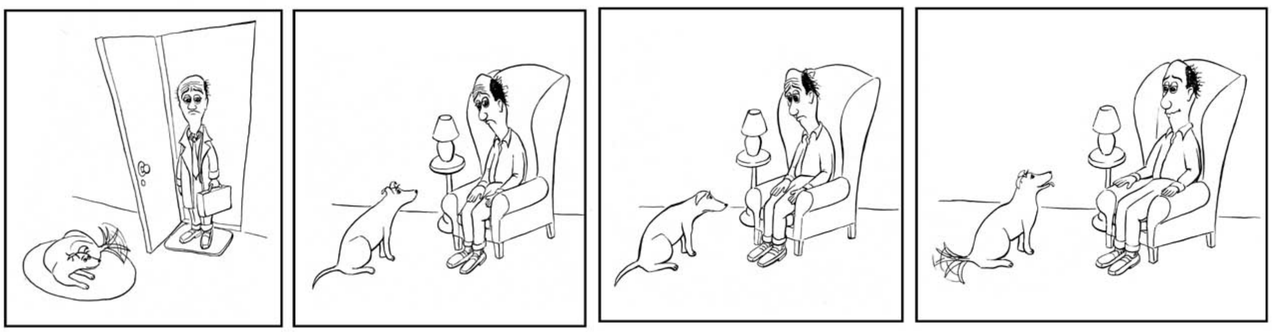

Figure 1: Example of an interaction with a dog that supports a human in meeting his experiential emotional needs.

1.2.1 Systems without sophisticated models

As described above, experiential emotional needs are most often social in nature, as other humans are usually involved in the satisfaction of these needs. Friends and family members are often part of "support structures" that are able to help meet such needs. Yet, while one may assume that experiential needs can only be met by other humans, we regard this perception as false. Many people routinely meet some of these needs via interaction with non-humans—for example, by interaction with pets, both real and computational. Such interactions provide evidence that some aspects of emotional needs can be met without humans, or much of what humans bring to these interactions, such as a sophisticated model or high-level understanding of what is going on.

Figure 1 illustrates a classic human-pet interaction. In the first panel, the dog happily greets his master coming home. His master is upset, and looks like he has had a bad day. The dog somehow recognises part of his master's affective state, perhaps that "something is not right" with his master. We know neither how nor what pets perceive in this situation. Perhaps the pet recognises some part or combination of the master's facial or gestural expressions, posture, vocal intonation, olfactory emissions, or behaviour. The dog reflects his recognition of his master's affect by putting his ears back and tail down. The master sees this expression of what one might call the dog's empathic response, and feels better. The dog sees (or somehow perceives) his master feeling better, then sits up happily, and wags his tail. The whole interaction, which does not appear to require great intelligence on the part of the dog, is very quick, taking much less time than this description of it. Nonetheless, the interaction has a powerful impact on the master's emotional state, enabling him to feel more connected, understood, and accepted by "others"—and in this way, helps him meet an aspect of his experiential emotional needs. Meanwhile, for all we know, the dog may be sitting there thinking, "I wonder when he'll feed me."

Sometimes we are asked how a computer might ever be able to recognise and respond to emotions given that humans do not even understand how this process works. Emotion theorists are still working toward a definition of emotion, and there is no available description for how it is communicated that is precise enough to implement in a machine. Thus, how can emotional skills be programmed if we do not know how to precisely describe their workings? Yet, without knowing *how* such interactions work, humans still navigate through them successfully much of the time.[3] A dog presumably has even less sophisticated modelling ability of emotions than does a human; nonetheless, the dog interaction, as illustrated above, successfully meets some human emotional needs (e.g. Cole and Gawlinski 1995, Beck and Meyers 1996, Zisselman et al, 1996).

At another extreme, there are examples where inanimate objects—a stuffed teddy bear, for example—may help meet certain human emotional needs. Where the computer will fall along this spectrum remains to be determined. However, we underscore that the computer does not have to attain the sophisticated abilities of a human in order for it to do better at meeting human emotional needs.

Although we do not support the goal of developing human-computer interactions to *replace* human-human interaction, we acknowledge that human-human interaction can leave many gaps in the ability of humans to meet the emotional needs of others, whether due to a dearth of others available to help, isolation, impaired socio-emotional skills on one side or another, or a multitude of other reasons. Such gaps present an opportunity where computers may be used to expand the space of possibilities for meeting these needs. For example, a key motivation behind the development of a computer emotion-support agent was this experience of Klein:

> My bride and I were en route to our long-awaited honeymoon aboard a major airline when I encountered a series of problems with my meal service. As I unwrapped the salad of my non-dairy vegetarian meal, I discovered two obviously dead insects on top. When the flight attendant finally came by, 10 minutes after I had called her, I started to explain the problem to her, pointing to the large, deceased flies on the salad. Rather than an acknowledgement of the problem or an apology on behalf of the food service, she whisked my food tray away, raised her index finger to her mouth, and said "Shhhhhhhh!" as if to tell me "Don't let anyone know!".
>
> The flies were bad but forgivable; the flight attendant's response was unacceptable. My wife and I were both upset by this experience, on our honeymoon no less, the occasion of which the flight attendant had already acknowledged. When I politely complained about the attendant's behaviour to the head flight attendant, her position was defensive: they could replace my dinner and that was all. Dumbfounded by their lack of concern for a customer's experience, I begrudgingly (and hungrily) accepted the replacement.
>
> The new meal, also labelled "non-dairy vegetarian" was eventually brought. This time the salad was insect-free. Partway through the salad, I began to have an allergic reaction, and noted that the salad dressing featured cheese (to which I am allergic) listed as an ingredient on the label, despite the non-dairy label stuck on top of

[3] This should not surprise us; after all, there are many examples where successful performance of some skill is not contingent upon knowing the mechanisms by which that skill operates. For example, people were able to reason correctly even before Aristotle extracted rules of correct reasoning. If for any operation to be intelligently executed, a prior theoretical operation had first to be performed and performed intelligently, it would be a logical impossibility for anyone to break into the circle. This argument has been articulated by the philosopher Ryle (1949) among others, who have written about how tacit knowledge is often evinced despite a lack of declarative knowledge.

> the meal. The flight attendant gave a curt apology, nothing more, adding insult to injury.
>
> On the ground at our destination, I politely complained first to the agent at the airline ticket counter, who offered no apology or acknowledgement of my experience; all he offered was to refund the price of my meal. I struck out similarly with other attempts at other counters. I recognised that customer service people may, for one reason or another, not be able to provide what a customer needs. Nonetheless, half-joking, I thought to myself, I can build a *computer* that handles peoples' feelings better than this.

Indeed, a system for addressing some of the experiential needs of a user having had a frustrating experience was built, tested, and demonstrated to be successful in a study of 70 subjects with two control conditions (Klein et al., 2001). We know of no other attempt at building a system that directly addresses the experiential needs of a user feeling frustrated. No efforts prior to this work appear to be aimed at directly addressing the emotional needs of a user, although many carefully designed systems have sometimes had this effect indirectly.

1.2.2 What this Approach is not: Computer Psychotherapy

Although the approaches we advocate may be therapeutic, and may employ similar tools to those used in psychotherapy, our goal is different. The stated goal of approaches to computer psychotherapy (Turkle, 1995) is the same stated goal of most psychotherapy: to somehow bring about lasting amelioration of sustained pathological disorders using techniques such as psychoanalysis. Our work is not intended as a treatment to heal long-term psychopathological problems; our goal is the meeting of emotional needs. A technique such as Active Listening, which was originally developed by Carl Rogers but came to be used by non-Rogerian therapists, as well as by parents, teachers, and other non-professionals, is useful for a variety of goals. Our use of such techniques is 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.

1.3 Related Work Indirectly Addressing Experiential Emotional Needs

One interesting example of a system that indirectly addressed certain experiential needs was that developed by Card et al (Card et al, 1974) for interrogating patients about gastrointestinal symptoms. The computer communicated to users via text only, using an informal, friendly style, which, incidentally, included the use of the pronoun "I" by the computer. (This human-like self-reference by the computer was avoided in our research). To test the system, seventy-two patients were given two interviews each, either by a pair of human specialists, or by a human specialist and a computer.

The results were compared for all six possible pairings—comparing in pairs the four interrogators: the three humans and the computer—to assess accuracy of the information gathered and the overall user experience. The accuracy of the information gathered by the machine was less than that attained with the human specialists. Low accuracy was probably due to the fact that (1) the machine limited patients to "yes," "no" or "don't understand" answers, despite many patients' desires to qualify their answers, and (2) the computer could not vary the interaction based on tone of voice or attitude of the patient—two factors that the human specialists said influenced their interrogations. Experientially, however, the system was considered very successful. The majority of users of the computer system had positive feelings about the experience, with about 94% of them answering "no" to the questions "Would you mind if computers like this were used a lot in hospitals?" and "Do you think other people would mind?" Users reported specific experiential needs being met, for example:

100% of the users answered “yes” to the question “Did the computer seem to be paying attention to your answers? Did you feel that there was someone actually talking to you?”

Card et al. reported that, with their computer system, users experienced a feeling of rapport and a sense of being attended to and connected to others. In this sense, it may be said that this system was generally able to help meet some aspects of these users’ experiential emotional needs, per the examples above. Ten percent of patients reported that they thought the computer system interrogation was preferable to talking to a doctor: it was polite, gave them more time to answer, and they felt less nervous and embarrassed. (Card et al. noted that patients might have perceived a barrier between the doctor and their own "working class" status, a barrier that was not present with the chatty computer system.) The computer was also perceived as less judgmental, a finding that appears to play a role also in the study of Lucas, et al (Lucas et al., 1977), where a different group of patients, referred for assessment of alcohol-related illnesses, admitted to a consumption of alcohol that was 42% higher on average when interacting with a computer than when interacting with psychiatrists. Overall, the computer interrogation systems were perceived by the users as being generally very effective, despite the fact that these results were attained with early technology preceding general consumer use of PC’s.

1.3.1 Attitudes of Typical Users vs. Experts

Studies with medical question-and-answer systems have led to other findings of interest in our work, such as the dramatically differing attitudes between experts who provide content for the system and the attitudes of users of the system. Consider a study (Dove et al., 1977) where sixty young women had their medical and social history taken by computer before having an interview with the doctor. The researchers claim that the computer demonstrated therapeutic benefits on the patients, and enabled the doctor and patient to subsequently communicate better about the purpose of the visit. However, in their study, Dove et al. reported finding a noticeable difference between the attitudes of the patients who tested the system and the attitudes of the doctors, when they tested the system.

The system of Dove et al., unlike the one of Card et al., did not pretend to be a self-contained computer personality talking to the patient, but instead posed the questions as a conduit or assistant, as if the questions were coming from the doctor. The patients, after being introduced to the system by the doctor and interacting with it for 90 minutes on average, perceived the questions as friendly and the interaction as freeing of inhibitions. The patients believed that the answers they gave would be seen by the doctor, and this conception is believed to have paved the way for smoother discussion in the subsequent face-to-face meeting with the doctor. This phenomenon was observed despite the fact that in many cases (because of experimental error) the doctor did not see the answers, a fact that was unknown to the patient.

Several doctors and medical staff were asked to interact with the system in a way similar to the patients, so that they could see what the patients would be exposed to. In contrast with the patients, the doctors in this study expressed disdain for the computer communication, perceiving it as impersonal. Indeed, most of the doctors, medical students and paramedical staff who answered the computer’s questions were critical of the content and style of the program, and did not believe it could be of help, some of them showing initial hostility toward the system.

The reasons for the medical experts’ feelings are unclear. It is possible that some of them felt that their jobs were threatened by the system, but in general, one would expect that they realised the system was designed to compliment their effort, saving them precious time. Several members of the medical staff appeared to be concerned about the privacy of their answers, despite being told that no record was being kept. Of course, the staff’s use of the system was for testing, and thus was not as authentic as that of the patients, who used the system in its intended context for genuine medical reasons. Thus, the doctors had a different purpose and different expectations than the patients in using the system. Although the authors of the study did not report personality or self-esteem characteristics for the different groups of subjects, it is also possible that the “personality” of the system was more irritating to the medical staff than it was to the patients. Since the authors had noted in related work that the patient groups perceived class distinctions between themselves and the doctors, it is not unreasonable to speculate that a different style of interaction might be favoured by the experts as a whole, than by the patients whom they were serving. It has been shown that matching the interface’s style (of dominant or submissively worded text) to that of a dominant or submissive personality can lead to greater satisfaction with an interaction (Nass et al. 1995). Although the reasons for the strong difference in feelings about the system remain unknown, it is significant to note a contrast in the attitude of the “experts”, even those who administer the system and reap some of its benefits, vs. the primary users of the system.

Just as theatrical performers try to create experiences onstage that may or may not reflect the performers’ personal sensibilities, experts in computer-human interaction design should try to identify and set aside personal attitudes or prejudices they may have about emotional needs when trying to assess and meet the needs of the users they serve. Experts may “see through” a user-interface element that genuinely assists users in meeting needs, or may otherwise worry about its sincerity and its deceptive properties. These are valid worries, which we will return to below. Nevertheless, we wish to emphasise that what is perceived as sincerity in an interface is not that clear-cut. Attitudes about these factors can vary with differences in personal needs and in personal perspective. Just because experts do not like something does not mean that users won’t like it and willingly benefit from interacting with it.

2 HUMAN-HUMAN INTERACTION AS A GUIDE

Rubinstein and Hersh (1984) suggest that when language is used by a person to interact with a machine, the rules for person-to-person communication apply. The Media Equation of Reeves and Nass (1996) and many articles from their research argue that human-machine interaction is inherently natural and social, so that the rules of human-human interaction apply to human-machine interaction. Whether or not the rules in the human-machine condition are *exactly* the same as the human-human case has not been shown; however, dozens of studies have led to results that suggest the rules are similar; consequently, this theory has profound implications for how we think about designing affective systems. With the Media Equation theory as a launching point, we suggest that in evaluating a potential human -computer interaction, that designers first consider the closest human-human analogy to that interaction.

Some researchers have argued that users dislike when a machine interaction imitates human-like behaviour; consequently, it is interesting to revisit some of the early findings that might be seen as supporting such arguments, and apply the theory of the Media Equation. Consider for example the work of Resnik and Lammers (1985), which, among other things, concluded that subjects high in self-esteem generated more negative cognitive responses and made fewer errors when faced with human-like rather than machine-like feedback from a computer. Their study compared three categories of machine responses, so-called (a) “machine-like”, (b) “neutral,” and (c) “human-like.” Examples of these responses when a user made an error are, respectively, (a) “Alphanumerics illegal”, (b) “Use numbers only”, and (c) “I don’t understand these letters.” However, the theory of Nass and Reeves suggests that a user interacts with each of these responses as if it came from a person. Consequently, we might re-label the responses as *three styles of human-like* responses, perhaps: (a) “geek-like jargon,” (b) “command-like, direct and to the point,” and (c) “conversational English that could be interpretted as chatty and submissive”. We present these labels to suggest an alternate interpretation of the “machine-like” vs. “human-like” emphasis in the Resnik and Lammers study. With this new interpretation, their study supports a different conclusion: that the high-self-esteem business school subjects generated more negative responses toward the chatty submissive style, preferring instead the responses that were direct and to the point. Interestingly, the low-self-esteem subjects in the Resnik and Lammers study showed a reverse effect – generating significantly fewer negative responses to themore chatty and submissive interactions. Thus, we question when a study such as this is cited to argue that "human-like" is bad design. What is bad may be the mismatch in style and personality of the interaction, not the fact that it uses human-like expressions in its language.

Thus, we suggest that it is important to consider the closest human-human analogy to the human-computer interaction being designed, and ask how that interaction might make a person feel. This consideration is especially urged when designing with the goal of meeting emotional needs. The Klein et al. system that supports the user in trying to reduce frustration took into careful consideration the question, "How would I want a person that I had not interacted with before to support me in reducing my frustration?" If the basic idea didn't pass the person-supporting-you test, we would not expect it to pass the computer-supporting-you test.

2.1 Subtleties Can Ruin the Interaction

The subtleties of implementing the rules of person-to-person communication are extremely important; if ignored, they can destroy the success of an interaction. It would be a mistake to implement just a few human-like features and consider a system to be improved simply because it is "more human-like." A few human-like features poorly implemented can be much worse than no human-like features at all. This phenomenon is already well known in computer graphics, where interaction with a supposedly realistic humanoid character still leaves the human viewer with a more eerie and disturbing impression than a corresponding interaction with an intentionally non-realistic character. Some HCI efforts have also noted problems with incautiously adding some human-like features (Walker et al., 1994; Kiesler et al., 1996). However, this point is clearly not obvious to many HCI researchers, who have spent costly time and money taking these ideas to products. Here are examples that tend to annoy many people (although it is rare that *everybody* reacts the same way to a given system.):

— A car that talks to you and tells you what to do but doesn't listen to you. (Analogy: a person in your car that tells you what to do but never listens to you.)
— An animated paperclip winking at you every time you click on it to go away. (Analogy: a person who insists on winking at you every time you ask them to leave your office.)
— A computer that plays a triumphant fanfare every time you begrudgingly re-boot it. (Analogy: a colleague who whistles a triumphant ditty whenever you tell him or her they've messed up your work and now you have to take time to start over.)

These are just a few examples where appropriate attention has not been given to how people would *feel* during an analogous human-human interaction. It is no wonder so many attempts to build human-like features into machines seem cute at first, but quickly become annoying and irritating. There are many subtleties that designers do not address, for example, variations in an automated response based on whether it is the second or fifth time a problem occurs, versus the first time. Different styles of language and flow should be used if the goal is to appear professional and somewhat detached, as in many financial transactions, or to appear friendly and interested, as in many personal service transactions. When the subtle "personality" of the interaction doesn't feel right, even if it's only in the choice of words or font style, and even if there's no human face or voice or other human-like features, then the customer is likely to feel uncomfortable and look for alternatives.

Making human-human analogies in the spirit of the Media Equation theory allows us to predict an answer to several important questions, such as "will the emotion-support system we built continue to work over time?" Even before long-term studies are conducted with computerised emotion support, we can expect to find certain effects. For example, the first time a system causes you frustration and you go to customer service for help and get successful support (probably including active listening, empathy and sympathy) then you are likely to feel less frustrated. This may hold for another visit or two, as long as the text of the interaction is appropriately varied as one would expect with polite human-human interaction. However, if on the third or fourth time, they give you the same feeling-support, but your problem is still not fixed, then you are likely to change your opinion dramatically. We can expect the same repeated use problems with the system of Klein et al. as we can with the most similar human-human interaction: the technique works once or twice; however, after a while, there must be a different kind of response, on a level of helping with the problem, or all the efforts to date will seem disingenuous and possibly even manipulative.

Thinking about the equivalent “human-human” interaction is a critical part of designing a successful system. If you wouldn’t like an office colleague interrupting you in the midst of a time-crunch crisis just to point out that you’re stressed, then you’re probably not going to tolerate it if an office computer does that. If a customer service person tries to assuage your feelings without first listening to you and convincing you that they’ve understood you, you’re probably going to feel unheard and manipulated. Similar feelings can be expected if a computer tries to soothe you without learning anything of your pain. At the same time, the theory suggests that turning to a person, or to a machine, outside the source of the original problem, may also be beneficial, therefore opening up many potential applications beyond the scenario considered in the Klein et al. experiment.

Does a computer have to “feel pain” before it can really help you? On occasion, we have heard people with whom we’ve discussed this research say things like, “If it (the computer) could feel hurt like I do, then I would accept its sympathy.” In addition to some serious unfeasibility problems with giving machines feelings in the sense that we have them, we think that this approach is not necessary. Like a man who tries to empathise with a woman in labour, or like a therapist who hasn’t been sexually abused trying to help someone who has, there is plenty of evidence to suggest that an interaction can help meet emotional needs without the helper having to have the personal experience of the one getting help. Certainly when one can feel what another has felt it can be very powerful, but this is not necessary in the human-human interaction, and we would not expect it to be in the human-computer interaction.

2.2 Beyond Imitation of Humans

Does the principle that the “rules of human-human interaction apply” mean that building machines that are as much like people as possible is our primary goal? No. We are not advocating building humanoid machines; our purpose is to address human emotional needs in the interaction. Technology that aims to imitate humans has a long track record of being bettered (at least in some sense) by technology that is merely *inspired by* trying to imitate humans (Mumford, 1963; Shneiderman, 1998). For example, computer vision algorithms do not enable machines to “see” the way people see, but they can sometimes be designed to “see” things (like product defects) more reliably than people can spot these things. The same theory holds for computer audition algorithms, and many other applications of computing and so-called artificial intelligence.

The difference in our philosophy is a subtle but important one: the human way of addressing the problem may or may not provide the ultimate design solution, depending upon the desired goal. Just because humans are the best example we know when it comes to emotional interaction doesn’t mean that we have to duplicate their emotional abilities in machines, which may not even be possible. One cannot yet say whether such a solution is necessary, or even doable. The history of technology is full of cases where the best example known at the time was later bettered by a creative leap, freed of the constraints inherent in imitating humans or other living systems. In transportation, for example, wheels can get you most places much faster than legs, and airplane wings that do not flap have prevailed over those that did. Affective interaction is a new domain where little work has been done with machines; one simply cannot say what will work best, and we should not limit our thinking by presuming that imitation of humans is the only way to proceed.

3 THEORETICAL AND PRACTICAL IMPLICATIONS

Computers are gradually getting better at perceiving information related to user affect, although it is a very hard research problem and there is a long way to go. As designers are able to give machines better affect recognition abilities, this raises many potential implications, both positive and negative. In the future, we might expect computers to recognise vocal inflection and attitudinal characteristics of individual users, which, for example, might allow the computer to tailor its questions more like a human medical information-gatherer would do, an advantage noted above. Ultimately, recognising affect should greatly facilitate the ability of computers to heed the rules of human-human communication.

Despite many potential benefits, affect-recognising computers also pose a number of potential concerns, many of which have already been described (Picard, 1997, Chapter 4). In particular, the issue of privacy is usually raised when talking about machines that can sense whether we are acting frustrated. It is the belief of the authors that affect should only be sensed by computers when the user is aware that the computer has this ability, when he or she has control over the sensing, and when he or she chooses to allow this channel of communication because of a perceived benefit. In our research, we practice a policy of privacy and steer our applications toward those that benefit the user and are under his or her control (Affective Computing Group, 1999). That said, the rest of the privacy issues are ones that are not particular to affective computing, but rather to *any* personal information. Thus, although such issues are very important, we will set them aside in the rest of this paper's discussion, and focus instead on implications more specific to new technologies that recognise and respond to user affect.

In the Klein et al. (2001) paper, the authors described several practical implications for designers, such as how the system could be built with present day technology, at low cost, and with largely scripted interactions, and how merely designing sites that let customers vent strong feelings can possibly be just as bad as ignoring their emotions. This paper extends that discussion by considering the generality of interactions that provide emotion support and the implications of these for HCI. We also wish to address problems inherent in thinking anthropomorphically instead of affectively, as well as issues related to deception, human displacement, and mood manipulation, the value of certain negative emotions, and potential concerns about artificial empathy and sympathy.

The following sections address issues arising from and moving beyond the work of Klein et al (2001) and Scheirer et al (2001). How might computers meet emotional needs, perhaps even better than people do? How do we avoid manipulation and maintain integrity? Must support be "genuine?" Could such systems undermine real emotions? What about trying to make computers be friends? How might future research proceed, given these concerns?

3.1 Beyond Support of Frustrated Game Players

The finding that a computer can perform an active role in emotion support has potential implications for just about every kind of interactive software system imaginable. This strategy is not limited to the game-playing context in which it was tested, nor is it limited to the case where the computer that offers support was also the one where the problem arose. Consider another situation -- difficulty with a printer – together with two strategies that the computer might use if it sensed frustration from the user:

USER: Argh!! I can't get this thing to print! Oh God. I'm not going to make it.

COMPUTER: Having problems?

USER: Yes! I'm late for a deadline, and I can't get this stupid file to print. Grrr...

COMPUTER: Geez, that sounds really frustrating.

USER: You got that right. If I don't get this file to print, I may not get my paper into the conference on time.

COMPUTER: Well, I'm just a dumb computer, but I'm sorry you're in such a spot. It sounds really difficult.

USER: Thanks. Gee, this printer's obviously hosed. Hmm, maybe I can print this thing to another printer that's not being used...

Here's another approach to the same situation:

USER: Argh!! I can't get this thing to print! Oh God. I'm not going to make it.

COMPUTER: Having problems?

USER: Yes! I'm late for a deadline, and I can't get this stupid file to print. Grrr...

COMPUTER: Hmm. You sound frustrated. You know, Shelly and Bob both seemed to experience frustration while using this printer yesterday.

USER: Really. (Leans out the door). Hey is Shelly or Bob still around?

BOB: Yeah, I'm here, what's up?

USER: I'm having trouble getting output from Printer 4.

BOB: Yeah, I think the page sensor is flaky. Send it again as max colour output. That fixed it for us. We'll get it repaired on Monday.

USER: Great! That worked. Thanks! (Turns to monitor) And thank *you.*

The two scenarios above have several things in common. They are examples of strategies for human-computer interactions during times when the user is experiencing emotional duress: The user may be feeling some combination of frustration, anger, anxiety, and/or other emotions, some or all of which may be directed at the computer or the printer. In both cases, however, the computer's responses actively acknowledge and address the user's emotional state. Neither strategy uses advanced artificial intelligence (AI) in the service of its goals, and only the second tries to determine or address the *cause* of the user's emotional state, and that only in the sense of noting that the printer was playing a part in the user's goal coincident with the timing of the frustration. Both strategies aim to support the user in managing his or her feelings, using principles known in a number of communities, from social psychology to crisis management, from parenting theory to consumer affairs.

The strategies above may not prove to be the most effective ones for the situation, and some elements of these approaches are obviously questionable: The computer referring to itself as "I", for example, or the computer offering that it's "sorry" for the user. Such approaches may be found to mislead the user as to the computer's capabilities, or present the computer as "caring" when no such feeling state exists in the machine. For that matter, the ability of the computer to accurately and reliably discern frustration is still a research topic. The computer's report of the emotional experiences of others (as in the second scenario) raises privacy issues regarding with whom you would want certain kinds of information shared. Some of these approaches may also simply prove ineffective in helping the user to manage her feelings, and recover from strong, negative emotional states. Further, we expect different personalities to prefer different kinds of responses from the machine (Nass et al., 1995). The strategies illustrated above may be a breath of fresh air for one user, while radically annoying for another.

Still, the very idea that both systems are able to address and respond to the user's emotional state represents an important departure in HCI, both in research and in practice. The prevailing attitude in HCI has been to make a one-size-fits-all solution that frustrates the fewest people, or to believe that "as soon as machines are intelligent" the problem will go away. But if machines were as intelligent as people the problem would still not go away, because even intelligent people frustrate and annoy each other from time to time. A key ingredient that people bring to an interaction is an ability to discern and address the other person's emotions, especially negative ones, once they arise. These abilities reflect a respect for emotional needs, which have been almost completely ignored in the design of most computer-human interactions.

3.1.1 Soothing, Salving, Modelling: For the Best?

As discussed in Klein et al. (2001), frustration in humans has many unpleasant side-effects: Increased ability to become more negatively aroused, increased likelihood of getting angry, decreased ability to pay attention, think (and problem-solve) creatively, and interact harmoniously with others. Notably, though, frustration is often regarded, along with the other emotions judged as negative, to be painful or, at the very least, removing pleasure from one's experience. The agency demonstrated in the Klein et al. project had the effect that people who interacted with it during times of emotional duress showed signs of significant improvement in their emotional state. This kind of effect has direct implications, then, for actively addressing debilitating effects of frustration: it demonstrates the possibility of a computer helping humans to better manage difficult events, thereby rendering day-to-day existence less stressful[4] and, perhaps, more productive and pleasurable. Similar effects may potentially be shown for:

- Improving harmony in interactions with work mates and other "cooperative parties" (other humans, as well as the computer systems with which the user interacts);
- Increasing one's ability to think creatively and generate solutions to problems with greater ease (and improving, in the process, ones sense of autonomy and control—Isen et al.1987);
- Decreasing the likelihood that subsequent stimuli, particularly those that may be mildly annoying, will be perceived as frustrating, which otherwise may provoke an even stronger negative reaction by virtue of a cascade effect. This effect should, over the long term, act to preserve—or even possibly improve—one's sense of self-control.

3.2 Online Social-Emotional Skills

To some extent, human-computer interactions are social, and may therefore help meet or aggravate certain social human needs. One problem of online communication, which might be a contributing factor to currently identified issues of increasing isolation (Myers 1993) and Internet-based malaise (Edupage 1998; Kraut et al., 1998), is the time spent communicating with limited affective bandwidth, which characterises most forms of online talk. In fact, a number of autistics[5] have described how they

[4] We use "stressful" in the colloquial negative sense here. In general, some amount of stress is necessary to function (too little stress and the body falls asleep). Optimal functioning is generally considered to occur in the presence of neither too little nor too much stress.

[5] There is a lot of variety among autistics, and what holds for one autistic person may not hold for another. In general, autism includes a difficulty recognising, understanding, and predicting emotions and the emotional significance of situations.

not only really like communicating online, but that it levels the playing field for them. In a sense, everyone is autistic online: today's systems limit your ability to see facial expressions, hear tone of voice, and sense those non-verbal gestures and behaviours that might otherwise help you disambiguate a hastily-sent non-angry message from a genuinely angry one.

To the extent that people spend more time communicating with each other via technology *without* sufficient affect channels, they may actually be reducing some of their emotional skills – a kind of "use it or lose it" opportunity cost. It would not be surprising to see an actual decline in emotional skills over long-term computer use, although we know of no studies that examine the growth or attrition of affective skills for computer users, controlling for personality and other potentially significant variables. On the other hand, some people's ability to communicate affect through text may show great improvement with increased email use and practice, especially if they are given clear feedback on the perceived tone of their email. If future technology is to help facilitate the meeting of certain human experiential emotional needs by improving a sense of social connection to others, then researchers need to develop new ways for computers to emulate the affective bandwidth that would occur in person-to-person communication.

Implications of our ideas go beyond facilitating more affective bandwidth online. Evidence from a number of quarters (Myers 1989, Goleman 1995), suggests that, in U.S. culture at least, precious few humans seem to possess solid, effective, non-judgmental active-listening skills. Indeed, positive, constructive communication skills are very much appreciated when encountered—well-developed skills of which are something of a rarity. One possible reason for this lack of such skills in the U.S. may be a kind of positive feedback loop: Few people are available to practice such skills in day-to-day life; consequently, there are few people to serve as role models for this behaviour. Hence, fewer people pick up these skills through ordinary social interaction. Computers that appropriately support users in handling their emotions might not only provide direct benefits to those in need, but might also serve as positive models for subsequent human-human interactions

3.3 Emotion Support vs. Emotion Manipulation

Providing support for a person to regulate her own emotions, as was the intent in the Klein et al system, may be seen as a tame version of 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 in the Klein et al. work: As described in (Klein, 1999), the stated goal of the system is to support the *user's own emotion management.* Ideally, it should be up to the user to employ the system in this manner and in this role, and to receive any benefit at all from it.

However, such a deliberate intention may not always be the case—as was demonstrated in the Klein et al. study used to evaluate the emotion-support system. Whether the process at the heart of this interaction is voluntary is of critical importance, 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 experimental evaluation run by Klein et al. involves just such a deception, and one that was apparently not uncovered by most, if not all, subjects.

The potential for manipulation without the user's consent can be affected by many variables. It should be noted that, as a first-of-a-genre device, the system's capabilities were not well known, and subsequently might become better known. Further, humans tend to develop ways of resisting efforts that are perceived as manipulative. On the other hand, many people gratefully accept music, coffee, humour, and other sources that manipulate their mood, often without conscious awareness of the manipulative factor. Manipulation is thus not inherently bad – it can be used to improve ones wellbeing. In the Klein et al. study, subjects demonstrated strong behavioural 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. Nonetheless, self-benefit is not the only factor – most people want to believe that they are in control of something that is aimed at changing them, even if the change is supposed to be for the better.

Because manipulation may be achieved without a subject's awareness, it is possible that this approach to emotion regulation could be used in an involuntary manner—and examples in which such uses are unscrupulous are not hard to imagine. At the same time, simple social graces such as saying "please" and "thank you" can also manipulate the feelings of those around you without their awareness, but these forms seem harmless; they are usually motivated by good will. Ultimately, though, we believe that even such manipulation will have its limits. If a person causes you trouble time and again, even if he is polite and apologetic at every turn, you may forgive a few more iniquities than if he were rude or brusque, but soon enough your patience will run out.

3.3.1 Heartstrings and Pursestrings: For the Worst?

Visitors from one of the world's largest and best-known computer companies described for us how they were not surprised at the strong findings we obtained with the subjects who used our emotion-support computer system. They had conducted a large study of customers of their product, comparing customers who had bought their system and had no problems with it, to customers who had bought their system, had problems, and received great support. Which group was more likely to buy their system again? The answer was clear: those who had problems and received great support were significantly more likely to keep buying their brand.

This finding led to a discussion of an ethical dilemma: do you deliberately design the product to cause people problems, and then craft great service, so that you can engender greater loyalty? In short: frustrate the customer, then make them feel better, and they'll come back to buy more. The visitors did not disclose to us how they dealt with this question, but the answer is apparent on a much larger scale if you look at general practice: the common behaviour of most high-tech companies these days is not to produce products that won't cause problems for the customer, but to produce products that get out there *fast*. Most products are put out for consumption long before they are ready, with an "80% is good enough" and "just get it out there before our competition beats us" attitude. The bad news is an increase in customer frustration, with loss of productivity, increased stress, and correspondingly increased health costs, but these numbers do not show up on the price-performance curves for new products. The good news for companies is that good support for customers' needs can be a remedy. However, lest companies mistakenly think the solution is simple, consider this: it is known that treating employees well leads to higher productivity and morale, but if they know you are only treating them nicely to manipulate them into better service, then the treatment can have the opposite effect. In short, both the intentions and the behaviour of the company matter: if buyers thought that emotion support was being provided so that the company could make more money at the expense of their experience with the product, then the strategy would be likely to backfire, causing the company to lose not merely sales, but also integrity.

3.4 Sapping Needed Wind from Sails

A problem that may be of great concern in this work is the possibility of diffusing emotions in one way that should have been diffused in another, or worse, a computer that tries to nip all negative emotions in the bud, a kind of computational Soma for people who aren't happy (Huxley, 1965). The problem is simply this: negative emotions are not necessarily bad emotions. Years ago, on a June day, Lee Iacocca, an icon of the American "can-do" spirit of change, spoke at MIT's commencement ceremony to thousands of graduates and their families as they gathered in the lovely outdoor courtyard to celebrate. For a long time after, people talked about what Iacocca said, how he banged his fist on the podium and charged the graduates, “You must get ANGRY! You must get ANGRY!” At first, the listeners were startled and a little annoyed – was he trying to ruin their emotional experience on this otherwise jubilant occasion? Eventually the message became clear: Iacocca was saying that anger is a great motivator of important change. Bad feelings do not necessarily cause bad things to happen: mild irritations that simmer and create distress or frustration, then boil into anger, can serve as impetus to find a better way. Therefore, diffusing frustration and anger prematurely may undermine a person's ability (if not her right) to perform a potentially unpleasant task, such as confronting the company that sold her a poorly-designed system and demanding her money back.

A system that tries to help diffuse strong negative user emotion performs a service for the manufacturer (lowering the number of complaints that mandate 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.

Such effect is not necessarily limited to the domain of commerce, either. It is not hard to imagine politicians, if not whole governments, falling prey to the seduction of employing masses of such devices with which to manipulate a disgruntled constituency. What better way for a president to try to reconcile an ailing image in the minds of angry voters, say, than to use an inexpensive tool that makes it seem as though he himself were performing personal acts of contrition? In the near future one might imagine, for example, an administration mass-mailing (or offering online) a cheap interactive system to voters that presents itself as a public-opinion poll questionnaire, but that can acknowledge and assuage the feelings of angry voters, thereby relieving some of their discontent.

Yet, it is also possible that this strategy would quickly run amuck. Let us take the corporate case as an example: users, as experts in social relationships, would soon see such a strategy as disingenuous, especially if it led to no supporting evidence of real change over time. Humans seem to have need of some semblance of authenticity in their interactions. If all a person ever receives in transaction with another is emotion assuaging (without the eliciting problem ever being addressed in some reasonable timeframe, for instance), the person would 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 reputation for trying to “pull a fast one” on users. Perhaps a bit slower than the company that offers no such consideration for the user’s feelings, but the reputation would evolve downward nonetheless.

The onus, then, remains on the manufacturer to maintain its integrity and public image. Still, it would seem that effective emotion-support devices might draw out toward infinity the period before the epiphany of disillusion, to the detriment of the consumer.

3.5 Artificial Caring

It may be argued that computers that offer signs of empathy, sympathy, or caring in any way, 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 behaviours without truly *knowing* what the problem is, and without truly experiencing empathy, sympathy, or really understanding the emotions that the person is experiencing.

As described earlier, the fact that this kind of interaction can still work is no surprise given a number of similarities between it and the daily interaction many humans undertake with pets, in which positive, ameliorative effects on the part of the human are measurable and significant. Humans have long benefited from pet interactions, including emotionally, even though it is likely that these animals have at best a very superficial understanding of the emotions of the human. Humans routinely use vague cues from non-human sources, and anthropomorphise them to meet emotional needs such as feeling empathised with, feeling accepted and understood—all despite the probability that the animal may have very different motivations and goals, and that there is probably no real empathy or understanding (as people conceive of it) present.

Even the appearance of empathy in another human is not necessarily as genuine or sincere as it may seem. Humans sometimes find themselves in supportive situations in which believability and credibility are seen as highly valued, yet the substance behind the interaction may be lacking. Consider, for example, a situation in which a psychotherapist or social worker is conducting a session with a client who is upset: crying, angry, or otherwise demonstrative of his feeling. The therapist is trying to do the hard work of sincerely listening to the client, while distracted by some thought or concern about the client's behaviour, progress, or even some notion unrelated to the client. In such a situation, in that instant the therapist may not actually be listening very much or very well to the client, yet with well-trained skills in providing the right kinds of feedback (such as active listening paraphrasing), is providing the client with cues to the effect that the therapist *is* listening and understanding. The therapist knows that one important goal of this interaction is for the therapist and client to make some kind of therapeutic progress. However, while the client is very upset, the therapist may be trying to get the client to calm down. The focus of the therapist at this instant is likely to be similar to the theory practised by the emotion-support agent: Enable the client to feel listened to, to feel that his emotional state is understood, acceptable, and accepted. As a result, *while it is being delivered*—while the therapist is distracted but offering solace anyway—such apparent empathy may differ from the therapist's genuine feelings. In this situation, the session may be successful if the immediate therapeutic effect was achieved—if the client came away from the interaction feeling heard, listened to, and understood. Certainly, interaction with an automated system that is greatly limited in its capabilities guarantees a qualitatively different experience than one involving a real, trained, human listener. Also, with the skilled human listener, the situation above is the exception and not the norm. Yet some of the benefits may be the same as in apparently authentic social interactions, and even the best-intentioned human listeners may sometimes employ similar illusions in their presentation.

The approach of the emotion-supportive computer is, at its heart, an idealised simulation of real empathy, real understanding, and real caring. It's effectiveness may lie in leveraging the very fact that when humans hear words of empathy or sympathy, they sound so honest and true that, in their rarity of utterance, perhaps intuitively they sound real and "heartfelt." It may be, in essence, a simulated transaction that, in its simplicity and boiled-down absence of fallible, ineffectual statements, is seductive—perhaps too much so to resist. This notion by itself is troubling. As is the case with simulation by definition (Starr 1994), the extent of the model tends to be unclear, and can lead to over-attribution of capabilities, as well as just plain wrong attributions.

Over-attribution is systemic in our society: arising not only from the use of language by machines, including non-emotional language, but also by the common depictions of machines in the media. Audiences love to see emotional machines and producers deliver, as evinced by the making of many films that feature the emotions of a machine, such as *2001: A Space Odyssey*, *Colossus: The Forbin Project*, *Star Wars*, *Bladerunner*, *Short Circuit*, *Star Trek*, and *The Bicentennial Man*. The inherent scientific problems and lack of evidence for how to actually give machines anything like human consciousness or genuine feeling experiences have largely been ignored by the media. The public has essentially come to expect that emotional capabilities are inevitable in computers, after all, this is what they have seen on TV and the movies for decades. Indeed, it is possible with current technology to make machines appear as if they have some rudimentary emotions, and to give them some expressions and internal states that might perform emotion-like functions. But lost in all of the hoopla about giving machines this subset of emotional abilities is the simple fact that it is only a very tiny subset of human emotional abilities. Scientists have not yet foreseen any means of bridging the deep chasm between what machines can do and the kind of experience, emotional or not, we as humans have continuously [Picard, 2000].

The argument that machines can and must be designed to display their capabilities and limitations to the user has long been a tenet of HCI (e.g. Shneiderman 1997). As computers improve at certain social-emotional tasks, even if only to the level of certain tasks of a dog, but much better than they are today, how will machines clearly communicate their capabilities without crossing over into illusions, false appearances, and unsatisfiable expectations? Pets don't make their capabilities and limitations clearly known, and consequently some pet owners attribute much greater emotional understanding to their animals than can be justified scientifically. What should designers do with users who find it beneficial to their goals to sustain an illusion? In addition, for a given style of interaction, how many users fall into this category?

Do the tools of computerised emotional support represent a contribution to the tools society has developed to hoodwink its citizens, or is it a meaningful, beneficial tool whose benefits to human health, productivity, and attainment of personal goals outweigh its possible maleficent use? Indeed, generating ways of envisioning this new era of HCI without reducing authenticity will be a challenge in the coming years.

3.6 Diminished Reality

As mentioned, one possible positive benefit of this technology is the routinization (albeit automated) of a currently rarely experienced, very beneficial conversation, and one that might help model positive, effective human-human interaction for its users. However, perhaps in successfully automating a portion of a genuine human response, the net effect may be to weaken the foundation of efficacy for both the real and the simulated conversation. Thus was the case with the famous cave paintings of Lascaux, which after they were duplicated (so that more people could enjoy them, without bringing harm to the originals), Baudrillard wrote, “The duplication is sufficient to render both artificial” (Baudrillard 1983). Thus, another possible implication of this work is that it may potentially render authentic, human-human interactions less effective. Whether humans become so jaded in interaction with emotionally savvy software that they become less receptive to real human empathy is an open question, but one of legitimate concern nonetheless.

3.7 Computers as Friends: The Next Generation?

Children have perennially grown up feeling emotionally attached to their play objects. Indeed, children have seemingly always used blankets, dolls and other toys to fulfil needs of one kind or another throughout their development. For example, toddlers bond with “transitional objects” to help them move smoothly from primal attachment with their primary caregiver to an independent identity (Winnicott 1971.) Young children subsequently use baby dolls to model and simulate relationships they perceive in the real world (e.g. Fischer and Biddell 1997). What difference will the new generation of increasingly “intelligent” interactive toys make to the generation of children that grows up with them?

Today's interactive Barneys, Tamagocchis, Furbies, and “Winnie the Pooh” dolls may be socialising a generation of children to not only having emotional relationships with artefacts, but believing, long after childhood wanes, that toys can really have feelings. If these toys could recognise even a few of the child's truly expressed feelings and reflect as much empathy as a dog, then the illusion may become as powerful as it is for many adults who swear their pet understands their feelings better than anyone.

One credible possibility is an entire generation of toys that are capable of this kind of emotional-content interaction with their young users—capable of soothing a crying child or of perhaps artificially preventing strong feelings of loneliness, sadness, frustration, and a host of other strong, negative emotions. Might such artefacts discourage their owners from fostering normal, healthy interactions with their parents and other children? There are certainly many adults who prefer interacting with their pets and computers to interacting with other people, and who are quite happy with this state of affairs. If such support for emotion regulation is provided too early by a non-human source, would this have a beneficial, educational effect, or might it possibly leave some children emotionally crippled, thwarting the development of the skills needed to interact successfully with other humans?

This question may be developed in the adult world as well: if such devices achieve popular success, and humans routinely use them to help manage their emotional states, what happens to the human's sense of his or her own self-control? One can imagine possible addictions similar to those for interactions with other inanimate objects: coffee, cigarettes, and chocolate: "I just need a quick break to be with my computer and then I'll feel better." As when interacting with a real person, the one who is in control becomes less clear – to what extent does control reside with you or with your confidante?[6] Where does it reside in a world in which humans may depend on emotional cyborg relationships for their emotional well being? Clearly this is not a problem yet; computers have a long way to go before they are accused of erring on the side of providing too much emotional well being. However, this is a foreseeable specific concern arising from new technologies that begin to assess and respond to expressed user feelings.

4 CONCLUSIONS AND FUTURE DIRECTIONS

The implications that arise from new affective devices and the approach they provide for problem solving are broad in scope, and range from the level of the individual to that of the culture. It is therefore imperative, scientifically as well as ethically, that this impact be explored as fully as possible—before such devices can responsibly and ethically be put into widespread use. At the level of the personal, issues include how humans may use (or abuse) such devices themselves; how such devices might change the nature of human-computer (and human-human) interactions; and how humans will define themselves in a world where such devices are regularly used. On a commercial level, issues include the ethical use of such devices and the incentive corporations may have to develop such products—as well as high-quality products in general, when incentive to release fine products is diminished by mounting pressure to release products early. Political issues run the gamut from public consensus on acceptable design and use of such devices, as well as the potential misuse and/or abuse of them, including the use of these devices to help maintain disciplined citizens and consumers. At the level of culture, might the advent of such devices be used to foster positive change on a society-wide basis, or might they be used as another means for manipulation and control, fostering the dismantling of a society that once held dear values of individuality, autonomy, and authenticity? And finally, on the global stage, how might widespread use of such devices help to enfranchise humans around the globe, and how might they be used in the steamrollering effect that Western culture seems to have on other, diverse cultures around the world?

This paper raises more questions than it answers, but we have deliberately kept the scope tethered to experimental results of current technology, which have shown that computers are beginning to be able to recognise and respond to certain limited expressions of user emotion. There are many

[6] Individual differences in perceived locus of control (internal vs. external) have been shown to have a significant impact on a user's emotional experience (anxiety level) with animated agents (Rickenberg and Reeves, 2000). These findings suggest that resolution of these issues will probably depend on trait differences, and that these issues are already relevant for designers of today's systems.

additional experiments and ideas that could be explored to examine the possible implications of this work. Some of the questions foremost on our minds include:

- How will user responses differ when they know the strategy of the system up front? The equivalent human-human interaction would suggest that the techniques would still work if appropriately practised, e.g., we know that counsellors trained in active listening, empathy, and sympathy still feel better when they go to another counsellor who uses these techniques with them. However, we also know that the attitude and expectations of a person can bias their experience, so this must be taken into consideration.
- How will users handle the inevitable failures of the technology? Computers fail in ways that are fundamentally different from human mistakes. Will such failures simply lead to momentary suspensions of any illusion the user has chosen to submit to, or something worse? And how should computers handle such failures – with a reminder to the user that "This is just a stupid machine" and "Please forgive this machine's incapability of really understanding you; if you want to proceed, it will also do so, in the best way it can." Or if this approach fosters a misleading impression then how can expectations be set straight?
- What effects do time and experience have on the effectiveness of such interactions? Is automated emotion-support a one-time-only phenomenon, akin to *fool me once, shame on me*, where later reflection by the recipient leads to a feeling that he or she has been hoodwinked, and should not let it happen again? Or will the human-human interaction prediction hold true, suggesting that humans will accept and/or benefit from the behaviour of such an agent on more than one occasion, if the interaction takes into account the same subtleties that a human would take into account over repeated interactions (recalling previous problems, showing steady progress toward a solution, while continuing to show elements of active listening, empathy, and sympathy.)
- When we queried users of a frustrating software package about what they directed their frustration toward – the computer, the software, its programmers, the company that made the software, the system admin, their boss, etc., the most common answer was the software (Norwood, 2000). How is the interaction affected when the computer makes clear that the dialogue originates "from the human software maker" vs. "from the computer or its software" vs. when the source of the dialogue is ambiguous?

4.1 Looking Ahead: A Future Scenario[7]

Company X has a notorious reputation, even among its loyal customers, for prematurely releasing products that have severe problems associated with them—then fixing the eliciting problem, but leaving their users feeling distinctly dissatisfied and angry.

Company X's operating system (OS), for example, was released and subsequently found to have critical errors in its design. At least three subsequent "service pack release" products were issued, fixing many of the bugs in the OS software. However, all these fixes did nothing to address users' frustration with the problem, or, for that matter, to acknowledge the inconvenience and cost in time to users. The authors have overheard more than one conversation about Company X and this product characterised by incredulity, sarcasm and derision—conversations that occurred many months after the "service pack release" was intended to fix all problems. Although we cannot state with certainty how the situation would differ if Company X had tried to address the consumer's experience, theory about human emotional needs suggests that Company X's lack of acknowledgement of and apology for the trouble they caused users contributed to the undermined consumer confidence people felt in the company and to the residual bad feelings many users have about Company X and their products.

[7] This story is true. Only the identifying information has been changed.

Imagine, if you will, a different scenario: Assume the same grievous error in prematurely releasing the software, which makes customers and users feel frustrated and manipulated, as if their time were not valued. However, this time imagine that the service pack release includes an emotion-savvy interaction that is simply text-based, and so computationally inexpensive that it would be capable of running effectively during the software installation process.

The interaction would engage users in a brief dialog about the product and its subsequent service pack release, and query users regarding their feelings about the experience. Depending on the user's response, the system would offer a statement offering a heartfelt apology appropriate to the level and type of affect expressed by the user, as well as some empathy and sympathy for the user's predicament, stated inconvenience, and loss of faith in the product. Then, this agent would encourage users to describe in their own words how they felt, and anything else they would like to add.

Imagine that, once the software was finished installing, that a dialog box appeared politely asking for the user's permission to automatically send the user's feedback via email to Company X headquarters, with specifics about who would see the information, how it would be used, and assurances that it was indeed valuable, confidential information that will contribute to better products in the future. The company then can keep track of its user base, maintain some idea of how its products are perceived in the marketplace (extremely valuable information), as well as possible leads for designers on new product innovations or bug fixes that should be made.

The user, who is free to engage or disengage with the interaction as he chooses, is left feeling *much* better—about the product specifically, the interaction in general, and about work overall. Indeed, the user may actually feel downright *warmly* toward a product and a company that would care so much to consider the user's feelings and that would make the user feel so heard, understood, and accepted. Perhaps best of all, the information sent to the company (depending of course on the company's ability to make sense of and use this information to make better products) helps maintain a sense of *authenticity* to the interaction.

This sense would help users accept this entire process—especially if the user is subsequently sent a brief, email note from a real person at Company X, thanking the user for the feedback, and giving some assurances that the information was heard (i.e. providing more paraphrasing feedback) and will represent a change in the ways that the company will do business in the future. Such a sense of authenticity would also greatly elevate the affect that the user feels toward the company, its products and services. Moreover, this element of authenticity would help make subsequent blunders on the part of the company more acceptable and forgivable in the future.

Of course, this scenario might be just as effective—if not more so—if the company were smaller, newer, and had oriented itself as a customer-centred design shop, complete with built-in mechanisms for accepting, dealing with, and responding to such feedback.

The research we've conducted scratches the surface of areas of human emotional needs that might begin to be addressed by affective technologies. For example, there is a growing interest in persuasive technologies (Fogg, 2000) and many of these could potentially benefit from greater affect sensitivities, especially since emotional distress is usually the number one reason why people who are undergoing behaviour modification (for drugs, smoking, dieting, and so forth) suffer relapses. Yet, here the criteria for success may be in direct conflict with the traditional HCI goals of minimising user frustration and maximising user satisfaction while helping them attain their goal (see below). Such cases further emphasise our point that the user should be in charge, directly involved in setting the system goals. Otherwise, the second time that system politely nags the user when she starts to pull out a cigarette, she's likely to toss the system out with the soon-to-be-empty cigarette package.

We think HCI in the 21st century should pay the utmost attention to the user as not just an information processor, but as a human being, having emotional needs as well as other goals such as productivity and efficiency. The arrival of computerised tools offering affect analysis and response, together with advances in understanding the natural ways in which people often interact socially with machines, lead us to suggest that the time is ripe for expanding the view of users to include their emotional needs. This expanded view extends to the human-computer interaction, and to the design of systems capable of interacting with the user in a way that honours and respects these needs. Indeed, important questions such as "how can we make the user's experience more productive" and "how can we get higher ratings of user satisfaction on our post-interaction surveys" might both be addressed by examining new questions such as, "what is the impact of each part of the interaction on user feelings?" and "is this system meeting or getting in the way of meeting human emotional needs?" We have suggested that scrutinising the human-computer interactions at the level of subtlety used to scrutinise human-human interactions can be beneficial—for example, helping to reveal changes in wording and in interaction style that better help satisfy experiential emotional needs, making the system seem smarter and more effective, without necessarily making the system seem "more human-like." Perhaps more importantly, we are beginning to make inroads into a new and exciting area of research into Human-Computer Interaction that, despite potential risks, holds the promise of supporting users in wholly new and beneficial ways.

ACKNOWLEDGMENTS

We would like to thank Sherry Turkle, Youngme Moon, Elisabeth Sylvan, Nancy Alvarado, and Ben Shneiderman for their discussions related to this work, as well as Jocelyn Scheirer for her role in the development of these ideas.

REFERENCES

Affective Computing Group (1999), Policy Page URL: http://www.media.mit.edu/affect/AC_policy.html.

Barefoot, J. C., Williams, R. B., Dahlstrom, W. G. (1983). Hostility, CHD incidence and total mortality: A 25-year follow-up study of 255 physicians. *Psychosom. Med*. 45: 59-63, 1983.

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.

Blocher, K. H. (1999). Affective Social Quest: Teaching Emotion Recognition with Interactive Media and Wireless Expressive Toys. MIT Master's Thesis, Cambridge, MA, June 1999.

Bowlby, J. (1969). *Attachment and Loss, Volume 1: Attachment. Second Edition.* Basic Books, New York.

Card, W. I., Nicholson, M., Crean, G. P., Watkinson, G., Evans, C.R., Wilson, J., and Russell, D. (1974). A Comparison of Doctor and Computer Interrogation of Patients, Int. J. of Bio-Medical Computing, 1974 (5), 175-187.

Cohen, S., Tyrell, D.A.J. and Smith, A.P. (1991). Psychological Stress and Susceptibility to the Common Cold. *The New England Journal of Medicine*. August 1991, pp. 606-612.

Cole, K. M., and Gawlinski, A. (1995). Animal-assisted therapy in the intensive-care unit—A staff nurse's dream comes true. *Nursing Clinics of North America* 30(3), September 1995, pp. 529.

Damasio, A.R. (1994). *Descartes' Error: Emotion, Reason, and the Human Brain*. Gosset/Putnam Press. New York.

Dove, G. A. W., Wigg, P., Clarke, J. H. C., Constantinidou, M., Royappa, B. A., Evans, C. R., Milne, J., Goss, C., Gordon, M., and de Wardener, H. E. (1977). The therapeutic effect of taking a patient's history by computer. *Journal of the Royal College of General Practitioners*, 27, 477-481.

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

Fischer, K. W. and Bidell, T. R. (1997). Dynamic Development of Psychological Structures in Action and Thought. In Lerner, R. (Ed. of volume), Volume 1: Theoretical Models of Human Development, in Damon, W. (Ed.), *Handbook of Child Psychology*, Wiley & Sons, New York.

Fogg, B.J. (2000). Persuasive Devices and Netsmart Devices. Chapter 11 in Bergman, E. (Ed.), *Information Appliances and Beyond.* Academic Press, San Diego, CA

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.

Hubel, D. H., and Wiesel, T. N. (1962). Receptive fields, binocular interaction, and functional architecture in the cat's visual cortex. *Journal of Physiology,* No. 160, 106-154.

Huxley, A. (1965). *Brave New World/ Brave New World Revisited*. Harper & Row. New York.

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.

Isen, A. M. (2000). Positive Affect and Decision Making, in Handbook of Emotions, Eds. M. Lewis and J. Haviland, Guilford Press, NY, 2000.Kagan, J. (1994). *The Nature of the Child: Tenth anniversary edition.* New York: Basic Books.

Kiesler, S. Sproull, L., and Waters, K. (1996). A Prisoner's Dilemma Experiment on Cooperation with People and Human-Like Computers. *J. of Personality and Social Psychology*, Vol. 70, No. 1, 47-65.

Kirsch, D. (1999). The Affective Tigger: A Study on the Construction of an Emotionally Reactive Toy. MIT Master's Thesis, June 1999.

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

Klein, J., Moon, Y., and Picard, R. W. (2001). This Computer Responds to User Frustration. *Interacting with Computers* (this issue).

Kraut, R. Kiesler, S. Mukhopadhyay, T., Scherlis, W., and Patterson, M., “Social Impact of the Internet: What Does it Mean?” *Communications of the ACM*, Dec 1998, Vol. 41, No. 12.

Lucas, R. W., Mullin, P. J., Luna, C. B. X., and McInroy, D. C. (1977). A comparison of psychiatrists and a computer as interrogators of patients with alcohol related illnesses. *British Journal of Psychiatry*, 131, 160-167.

Maslow, A. (1987). *Motivation and Personality, Third Edition.* Harper & Row, New York.

Muller, M. J., Wharton, C., McIver, W. J., and Laux, L. (1997). Toward an HCI Research and Practice Agenda Based on Human Needs and Social Responsibility. Proceedings of CHI ‘97 Human Factors in Computer Systems, Atlanta.

Mumford, L. (1963). *Technics and civilization.* Harcourt, Brace & World, New York.

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., Moon, Y., Fogg, B.J., Reeves, B., and Dryer, D. C., (1995) Can Computer Personalities Be Human Personalities? *International Journal of Human-Computer Studies*, 43, 223-239.

Norwood, M. R. (2000). *Affective Feedback Devices for Continuous Usability Assessment*, MIT Master of Engineering in Electrical Engineering and Computer Thesis.

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

Picard, R. W. (2000). What does it mean for a computer to “have” emotions?. Chapter to appear in *Emotions and Human Artifacts*, Eds., R. Trappl and P. Petta, MIT press, Cambridge, MA.

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.

Resnick, M. (1998). Technologies for Lifelong Kindergarten. *Educational Technology Research and Development,* Vol. 46, No. 4.

Resnik, P. V. and Lammers, H. B., (1985). "The Influence of Self-Esteem on Cognitive Responses Toward Machine-Like vs. Human-Like Computer Feedback," Journal of Social Psychology, 125,(6), 761-769.

Rickenberg, R. and Reeves, B. (2000).”The Effects of Animated Characters on Anxiety, Task Performance, and Evaluations of User Interfaces” CHI 2000, Apr 2000.

Rubenstein, R. and Hersh, H.M. (1984). *The Human Factor*. Digital Press. pp. 60-62.

Ryle, G. (1949). *The Concept of Mind*. Hutchinson & Co., London.

Salovey, P. and Mayer, J.D. (1990). Emotional Intelligence. *Imagination, Cognition and Personality*. Vol. 9, No. 3, pp. 185-211.

Schank, R.C. & Jona, M.Y. (1991). Empowering the student: New perspectives on the design of teaching systems. *The Journal of the Learning Sciences*, 1 (1), 7-35.

Shneiderman, B. (1997). Direct manipulation vs. agents: Paths to predictable, controllable, and comprehensible interfaces. In Bradshaw, J. M. (Ed.), *Software Agents.* AAAI/MIT Press, Cambridge, MA.

Shneiderman, B. (1998). *Designing the User Interface: Strategies for Effective Human Computer Interaction,* 3rd Edition. Addison Wesley Longman, Inc, Reading, MA.

Scheirer, J., Fernandez, R., Klein, J., and Picard, R. W. (2001). Frustrating the User on Purpose: A Step toward Building an Affective Computer. *Interacting with Computers,* (this issue).

Soloway, E., Jackson, S. L., Klein, J., Quintana, C., Reed, J., Spitulnik, J., Stratford, S. J., Studer, S., Jul, S., Eng, J., and Scala, N. (1996). Learning Theory in Practice: Case Studies of Learner-Centred Design. Proceedings of CHI '96 Human Factors in Computer Systems, Vancouver, 189 – 196.

Spitz, R. A. (1945). Hospitalism: An inquiry into the genesis of psychiatric conditions in early childhood. *Psychoanalytic Study of the Child.* Vol. 1, 53-74.

Starr, P. (1994). Seductions of Sim: Policy as a simulation game. *The American Prospect,* Spring 1994, 19-29.

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

Walker, J. H., Sproull, L. and Subramani, R. (1994). Using a Human Face in an Interface. CHI '94 Human Factors in Computing Systems, Boston MA, April 1994, 85-91.

Winnicott, D. W. (1971). Transitional Objects and Transitional Phenomena. Chapter 1 of *Playing and Reality.* Basic Books, New York.

Zisselman, M. H., Rovner, B. W., Shmuely, Y., and Ferrie, P. (1996). A pet therapy intervention with geriatric psychiatry inpatients. *American Journal of Occupational Therapy* 50 (1), January 1996, pp. 47-51.