GETTING PERSONAL WITH COMPUTERS: HOW TO DESIGN PERSONALITIES FOR AGENTS

D. CHRISTOPHER DRYER
IBM Research, San Jose, California, USA

Recent research indicates that people respond socially to computers and perceive them as having personalities. Software agents are artifacts that particularly embody those qualities most likely to elicit social responses: fulfilling a social role, using language, and exhibiting contingent behavior. People's disposition to respond socially can be so strong that they may perceive software agents as having a personality, even when none was intended. The following is a discussion about intentionally designing personalities for social agents. To design personalities, it is necessary to consider the nature of personality and its role in interactions between people and artifacts. In addition, a case study of designing a social software agent is presented. The conclusions from this experience are summarized as guidelines for future agent developers.

Personality is a fundamental linchpin of social relationships. In the context of human interaction, people automatically and unintentionally organize the behavior of their partners into simplifying traits (Uleman et al., 1996), and people tend to agree about which partners are best described by particular traits (Moskowitz, 1988). Beyond categorization, personality shapes the very nature of social relationships, even impacting how satisfying an interaction is for the participants (Dryer & Horowitz, 1997).

It is not surprising therefore that researchers and engineers have started thinking about personality as essential to the design of social agents for user interfaces. Indeed, people may perceive automatically a personality in any social agent, regardless of whether the agent's designer intended a personality. This suggestion presents some new opportunities for human-machine interaction designers. Rather than have unintended personalities shaping

Part of this article appeared as "Ghosts in the machine: Personalities for socially adroit software agents," by D. Christopher Dryer, pp. 31–36, in K. Dautenhahn (ed.), *Socially Intelligent Agents*, AAAI Press Technical Report FS-97-02, American Association for Artificial Intelligence, Menlo Park, California, 1997.

This work was in part the product of the CSLI interface laboratory at Stanford University, the New Technology team at the Walt Disney Company, and the WarpGuide team at IBM, including Ron Barber, Roger Didio, Steve Idhe, Bob Kamper, Bob Kelley, Marion Lindsey, Craig Moser, Paul Rogers, Ted Selker, John Tyler, Les Wilson, and Julie Wright. In addition, Wendy Ark, Renee Fadiman, Ronaldo Mendoza, and three anonymous reviewers made important contributions to an earlier draft. IBM and OS/2 Warp are trademarks of the IBM Corporation.

Address correspondence to D. Christopher Dryer, International Business Machines Corporation, Almaden Research Center, 650 Harry Road, San Jose, CA 95120, USA. E-mail: dryer@almaden.ibm.com

Applied Artificial Intelligence, 13:273–295, 1999

0883-9514/99 $12.00 + .00

people's interactions with machines, designers can intentionally create personalities that support satisfying social interactions. This task depends on understanding the nature of personality, the role of personality in people's interactions with artificial agents, and how agent personalities can be created.

THEORETICAL BACKGROUND: THE NATURE OF PERSONALITY

A personality is the collection of individual differences, dispositions, and temperaments that are observed to have some consistency across situations and time. Beyond this basic definition of personality, different researchers have thought about the phenomenon in various ways. For example, personality can be seen as only the measurable consistencies in behavior (e.g., Mischel, 1968). Personality also can be seen as the perceived consistencies in people, independent of whether those consistencies are actual (e.g., Shweder, 1975). Alternatively, personality can be seen as an interaction between the expectations of an observer and the behavior of the observed (e.g., Cantor & Mischel, 1979). Where intentionally created personalities are concerned, however, the behavior of the observed has been designed with the expectations of an observer in mind. The personality of a user interface agent does not matter outside of the perception of an observer. From this perspective, the actual and perceived personalities can be assumed to be the same.

People unconsciously use personality as a tool to organize otherwise overwhelming information about social partners. Successful social behavior negotiates the goals, beliefs, and emotions of multiple social partners. These negotiations occur in most domains of human life, including collaborative work, recreation, and family life. People use their perceptions of their partner's personality to make predictions about what that person will do and how that is different from what other people will do. Without personality and the ability to use personality to organize these predictions, social behavior would break down.

The variety of human personalities is infinite. Mapping the space of all of the possible personalities seems like an impossible task. Indeed, there are thousands of nonsynonymous trait adjectives, such as shy, adventurous, obsessive, domineering, feminine, and anxious.

The task, however, is simpler than it seems. While there are thousands of different facets to personalities, only a few factors really matter at a more abstract level. The five important factors (for review, see McCrae and John, 1992) can be described as (1) agreeable (cooperative to competitive), (2) extroverted (outgoing to withdrawn), (3) neurotic (anxious to calm), (4) conscientious (organized to lax), and (5) open (curious to closed minded). Certainly, there are other things to know about a partner, but nearly all of them covary with one of these five dimensions or some combination of them.

These factors result from empirical investigations into which sets of individual differences (among samples of all possible individual differences) covary. All of the theoretically derived personality schemes that have been tested, such as the Myers-Briggs (Croom et al., 1989), can be shown to be simply rotations of some or all of the Big Five dimensions (Furnham, 1996).

Of the five factors, the most important for social interactions are those that concern individual differences in social behavior, namely, agreeableness and extroversion or their common rotations, "friendliness" (friendly to cold) and "dominance" (authoritative to submissive) (Wiggins, 1979). These two factors, proposed by Aristotle in *Rhetoric*, have been reinvented and renamed by various theorists (e.g., Leary, 1957; Benjamin, 1974; Strong & Hills, 1986).

The factors and some typical labels are shown in Figure 1: the two-dimensional interpersonal space. This two-dimensional model has been proposed as a "nomological net" uniting various interpersonal research programs (Gurtman, 1991; Wiggins & Broughton, 1985). For example, researchers have found the two-dimensional structure to be fundamental to interpersonal traits (Conte & Plutchik, 1981; Wiggins, 1979), needs (Wiggins & Broughton, 1985), problems (Horowitz, et al., 1988), relations (Wish et al., 1976; Dryer et al., 1994), emotions (Knutson, 1996), and goals (Dryer & Horowitz, 1997).

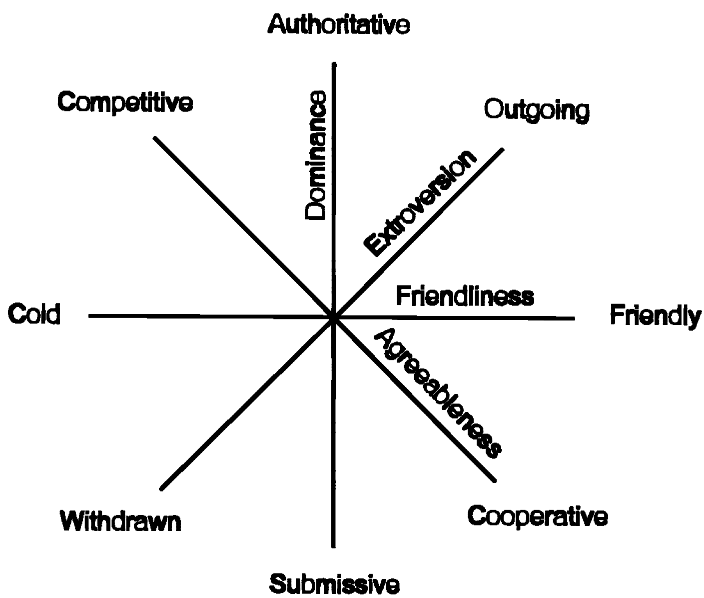

FIGURE 1. The two-dimensional interpersonal space.

Because personality helps people organize the behavior of others, it also guides people's responses to their social partners. People differ in the kinds of personalities they want their partners to have, and the match between two people's personalities impacts how enjoyable and productive the interaction is (Dryer & Horowitz, 1997; Horowitz et al., 1997). In general, people like partners that have "positive" personalities. Each of the five personality dimensions has a positive (i.e., socially desirable) pole and a negative pole. Personalities that are described at the positive end of the pole are better liked, and vice versa; all else being equal, people would rather interact with cooperative, outgoing, calm, organized, and curious partners than competitive, withdrawn, anxious, lax, and close-minded partners.

A personality, however, can be too positive. Partners with personalities that are extremely positive on all five dimensions may be less attractive than partners who have a negative personality attribute along at least one dimension. People find it easier to identify with a character who has some weakness or flaw that makes them seem human; "perfect" characters are unnatural and seem less lifelike (Thomas & Johnson, 1981). Similarly, researchers (Aronson et al., 1966) have found that partners with a foible are liked better than those without.

Personalities that are similar to a person's own are also well-liked (Byrne, 1971). For example, friendly people find friendly partners more attractive than cold people do, and cold people find cold partners more attractive than friendly people do.

There are exceptions, however, to the maxim that "birds of a feather flock together;" sometimes, "opposites attract." When similarity is paired with undesirable characteristics (such as unattractiveness or evidence of mental disturbance), people actually prefer dissimilar partners over similar partners (Novak & Lerner, 1968; Taylor & Mettee, 1971). Likewise, personalities are liked if they make up for something positive that you do not have. People prefer partners who want to take on dissimilar rather than similar roles in a relationship (Krokoff et al., 1988). Also, when a task requires one person to take charge, a dominant person will prefer to work with a less dominant partner, and a submissive person will prefer to work with a more dominant partner (Dryer & Horowitz, 1997). In special situations like these, people prefer partners who are complementary to themselves rather than similar to themselves.

EMPIRICAL INVESTIGATIONS: THE ROLE OF PERSONALITY IN PEOPLE'S INTERACTIONS WITH ARTIFACTS

Through a number of research programs, psychologists have characterized the ways in which people fundamentally perceive their social partners along the dimensions of friendliness and dominance. One interesting research question is whether people perceive the personalities of artificial agents in the same way.

This hypothesis can be considered within the framework of social interface theory (for an overview, see Reeves & Nass, 1996). Social interface theory is built on the results of various studies demonstrating that people respond socially to machines. These studies suggest that people's responses to technologically medicated social partners can be the same as their responses to natural partners. People are inclined to treat everything as social and natural. Therefore people automatically and subconsciously leverage whenever possible what they know about their natural and social experiences to help them with their technological experiences.

Software agents in particular are likely to encourage social responses because of three important features. (1) They can use full-sentence text, in addition to the typical user interface forms of communication, like menus, controls, and icons. Full-sentence text is more natural, especially for people with a more verbal than nonverbal cognitive style (Horn & Cattell, 1966). (2) Software agents can embody task knowledge as well as use artificial intelligence to reason about when and how to engage a person in interaction. This gives them a compelling kind of contingent behavior. (3) Software agents can autonomously perform actions on a person's behalf. These features- natural language (text), contingent behavior (intelligence), and social role (autonomous assistance)- are the three factors that predispose people to respond socially (Reeves & Nass 1996).

Various tests of social interface theory have provided evidence that people apply apparently irrelevant knowledge about social interactions to their interactions with machines. For example, people are polite to agents (Nass et al., 1997b), use social rules to explain their behavior (Nass et al., 1994), apply social stereotypes to them (Nass et al., 1997a), treat them as teammates (Nass et al., 1996a), and look for social motives when agents flatter or criticize (Fogg & Nass, 1997). Given the importance of personality in people's social behavior, we wondered whether perceiving a personality is another social response that people have when interacting with an agent.

This question was first addressed with a laboratory study concerning social responses to computers (Dryer et al., 1993a). In this study, 44 computer-literate participants interacted with personal computers through a traditional-style graphical user interface. A computer first tutored participants on a series of topics (social customs, mass media, and computers). The participants were tested with a number of questions on the topics, ostensibly to determine how successful the tutoring was.

The participants then had a final interaction with a computer. In this interaction, a computer reviewed each of the test questions. For each question, the computer first indicated whether the participant had answered the question correctly and offered a critique of how successful the tutoring had been for this question. In all cases the critique concerned the performance of the tutor and not the performance of the participant. The experimental manipulations occurred during this final interaction.

The first manipulation concerned whether the evaluating and tutoring computers were the same or different. For one-half of the participants, the evaluating computer was the same as the tutoring computer; that is, the tutoring and the evaluating interactions both occurred with one and the same machine. For the other half, the evaluating computer was different from the tutoring computer. In this case, there were two computers sitting side by side that were identical in kind and appearance; participants were tutored on one machine, and then, for the evaluation interaction, they moved over to the other machine.

The second manipulation concerned whether the evaluation was typically positive (i.e., the tutoring was judged to be "successful") or negative (i.e., the tutoring was judged to be "unsuccessful"). For all the participants in the study, the review of the test questions was rigged so that everybody answered the same number of questions "correctly," and thus the evaluation of the tutoring was arbitrary. For one-half of the participants, the evaluation was positive, and for the other half the evaluation was negative. After this interaction, participants' perceptions of the dominance and friendliness of the evaluating computer were assessed with two ten-point Likert items on a pencil-and-paper questionnaire; the two items were "authoritative" and "friendly," with the ends of the scale anchored by the labels "describes very poorly" and "describes very well."

If this interaction involved people rather than machines, the manipulations could be described as self-evaluation versus other-evalution and praise versus derogation, terms that connote certain dispositions and social motivations. Rationally, machines clearly do not have motivations that involve a "self" or dispositions toward social relationships. The goal of this study, however, was to investigate whether people perceive personalities in machines. If people were to see the machines as having personalities, the conditions theoretically could be called self-praise, self-derogation, other-praise, and other-derogation. In all cases, the only actual differences among the conditions were whether the evaluating computer was the same as or different from the tutoring computer and whether the evaluation was positive or negative; all other aspects of their behavior were constant across the four conditions.

Researchers have used various methods (including behavioral measures) to determine how people use the interpersonal dimensions to organize behaviors like evaluation, resulting in catalogs of the theoretical locations in the interpersonal space of a number of behaviors (Kiesler, 1983; Strong & Hills, 1986; Horowitz et al., 1991). Specifically, other-derogation is authoritative and cold, other-praise is friendly, self-derogation is submissive, and self-praise is cold. To determine whether participants assessed the personalities of the machine as they would a person, we plotted their perceptions of the computer in the interpersonal space defined by authoritativeness (y) and friendliness (x). To construct this plot, condition means were calculated from the two questionnaire items (assessing authoritativeness and friendliness), and mean substraction was used to set the origin at zero. As illustrated in Figure 2, the mean personality ratings were as follows: other-derogation (− 0.6, 1.3), other-praise (1.1, 0.0), self-derogation (0.7, − 1.6), and self-praise (− 1.1, 0.3). The product moment correlation between the empirical and the theoretical placements of these traits in the two-dimensional interpersonal space was $r = 0.94$, two tailed $p < 0.001$ (for an explanation of the statistical analysis, see Hildebrand, 1986). In other words, we found good evidence that people use the personality traits of dominance and friendliness to organize the behavior of machines and people in the same way, even though such personality constructs may be rationally inapplicable to computers.

In this study the machines' personalities were determined by their behavior alone. Frequently, however, designers have given personalities to agents by manipulating their representation rather than their behavior. Anthropomorphic agents are obvious examples of intentionally created personalities. Typically, anthropomorphic agents are represented with first-person communication styles and human-like interface characters.

We therefore were interested in whether people think about anthropomorphic agents fundamentally in terms of personality, just as people do with human partners and with text-based agents. To examine this, we conducted

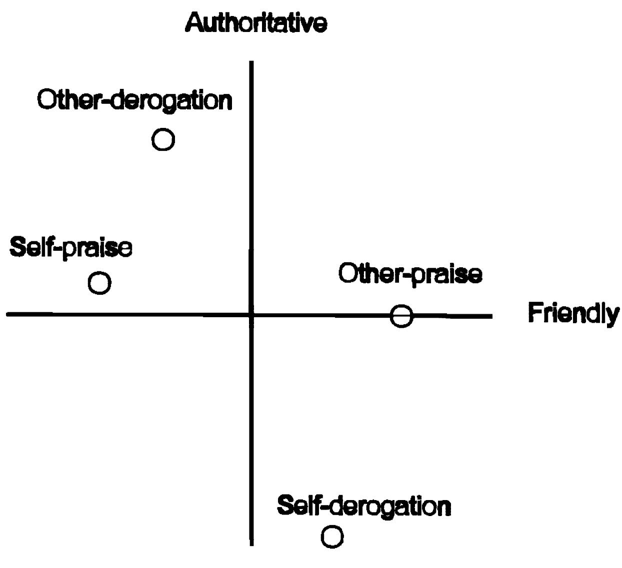

FIGURE 2. Personalities of evaluating computers.

a study (Dryer, 1995) that differed from the tutoring computer study in two important respects. First, we represented agents with animated characters rather than a text-based, traditional user interface. Second, we used a different method to assess participants' perceptions of the agents' personalities. In the earlier study, we assessed participants' perceptions along the established personality dimensions. It is possible that our use of these dimensions biased participants to distinguish the agents primarily in terms of personality. Without this method, would participants preferentially perceive agents differently perhaps relying on qualities other than personality?

To get at the qualities of the characters that were conspicuous to participants, we used a method called multidimensional scaling (MDS; for a description, see Young (1987)), which mathematically extracts whatever dimensional structure participants use to distinguish objects of judgment. For this study, 31 participants rated the similarities of 37 animated characters that were potential user interface agent representations. The mean similarity ratings were treated as ordinal numbers, and the distances as symmetric. We used a Euclidean multidimensional scaling method, deriving solutions for one through five dimensions. For each solution, we examined two measures of goodness of fit: Kruskal's stress (s) and the squared correlation of distances (RSQ) (Young, 1987). The goodness of fit values for the respective five-dimensional to one-dimensional solutions were $s = 0.594$ and $\mathrm{RSQ} = 0.766$, $s = 0.447$ and $\mathrm{RSQ} = 0.813$, $s = 0.568$ and $\mathrm{RSQ} = 0.692$, $s = 0.744$ and $\mathrm{RSQ} = 0.528$, and $s = 0.881$ and $\mathrm{RSQ} = 0.355$. The four-dimensional solution showed a relatively good stress value and accounted for a good proportion of the variance. In addition, we found that the four-dimensional solution was easily interpretable. Namely, the four-dimensional MDS solution revealed that participants perceived differences among characters along these dimensions: (1) cooperative versus competitive, (2) outgoing versus withdrawn, (3) sophisticated versus "cartoon-like," and (4) human versus animal.

Two of the dimensions clearly had nothing to do with personality. The dimension we called sophisticated versus cartoon-like represented the artistic style of the agents' visual representation. Some representations were more realistic, while others were designed with more of a cartoon style. Another dimension, which we called human versus animal, reflected the fact that the agents were a collection of fictional humans and anthropomorphized animals and objects. Therefore the artistic style of the representation as well as the kind of thing (human or not) that was represented were dimensions along which participants distinguished the agents.

The other two dimensions concerned the social nature of the agents. In fact, like the previous study, participants distinguished these agents within the two-dimensional interpersonal space, specifically along the dimensions of extroversion and agreeableness. When we inspected the placement of the characters within the interpersonal space, we observed that some distinctive visual elements were associated with clusters of similar characters. Friendlier characters (i.e., high on agreeableness and extroversion) tended to be represented by rounder shapes, bigger faces, and happier expressions than were colder characters. Similarly, bold colors, big bodies, and erect postures were more typical of dominant characters (i.e., high on extroversion and low on agreeableness) than they were of submissive characters.

Thus this study provided evidence that personality is fundamental to the way in which people perceive social agents represented by animated characters. In interpersonal interactions, perceiving personality is important because it impacts, among other things, how much people like interacting with each other. In the case of animated characters, a perceived personality could determine how well people will like a social agent. To examine this question, we conducted another study with the same set of 37 animated characters (Dryer, 1995). This study was designed to explore a number of factors that might impact a participant's liking for a character: the character's personality (friendliness and dominance), the strength of the character's personality (subtle versus extreme), the participant's own personality, the participant's level of experience with computers, and the participant's gender.

For this study, 40 participants indicated their liking for each of the 37 characters. In addition, we assessed the participants' own personalities, computer experiences, and genders, as well as their perceptions of the characters' personalities. (Personality measures were z-score standardized to a mean of zero and unit variance; for a description of z-scores, see Hildebrand (1986).) Using these measures, we divided the characters into groups. The most dominant half of the characters was assigned to a "dominant" group, and the least dominant to a "submissive" group. Similarly, the friendliest half was assigned to a "friendly" group, and the least friendly to a "cold" group.

Personality strength was defined as the sum of the absolute values of the dominance and friendliness measures, so that very cold or very friendly and very submissive or very dominant characters would score high on this measure and characters in between would score low. Then, for each of the four possible personality types (dominant friendly, dominant cold, submissive friendly, and submissive cold), the four characters with the highest personality strength score were assigned to a "extreme personality" group, and the four characters with the lowest personality strength score were assigned to a "subtle personality" group.

We examined the impact of these factors on the participants' liking of the characters with a repeated measures analysis of variance.[1] This analysis revealed that the character's personality was associated with how well liked the character was. Overall, dominant characters were better liked than submissive characters, $F(1, 24) = 11.15$, $p < 0.01$; friendly characters were better liked than cold characters, $F(1, 24) = 8.83$, $p < 0.01$; and characters with extreme personalities were better liked than those with subtle personalities, $F(1, 24) = 15.33$, $p < 0.001$. (for a description of repeated measures, analysis of variance, and the F statistic and its use in analysis of variance, see Hildebrand [1986]).

Consistent with responses to human personalities, characters with positive personality attributes were well liked. Independent of this effect, characters with strong personalities were also well liked. That is, not only were very dominant and very friendly characters relatively attractive, very submissive and very cold characters were relatively attractive as well. It may be better to be extreme on the socially undesirable end of a dimension than to be in the middle. Many intentionally created personalities are probably too subtle. Personalities are rarely so extreme that they appear unnatural and are disliked; in practice, it is difficult to create a personality that is too extreme (cf. Thomas & Johnston, 1981).

The other effects we examined concerned the participants' characteristics. The participants' level of computer experience and gender were not associated with the participants' liking of the characters; for experience $F(1, 24) = 2.89$ and for gender $F(1, 24) = 0.54$, both p values not significant. Similarly, the participants' personalities did not impact how much they liked the characters overall; for dominance $F(1, 24) = 0.20$, and for friendliness $F(1, 24) = 2.03$, both p values not significant. However, participants did tend to like those characters that had personalities that were similar to their own, at least with respect to friendliness. Relative to participants with cold personalities, participants with friendly personalities tended to like friendly characters and, relative to participants with friendly personalities, participants with cold personalities tended to like cold characters; for the effect on liking the interaction between the participant's friendliness and the character's friendliness, $F(1, 24) = 11.50$, $p < 0.01$. (The interaction between the participant's dominance and the character's dominance was not significant; $F(1, 24) = 0.90$.) Thus anthropomorphic agents can be represented graphically in ways that impact the personality that people perceive the agent having. Moreover, this personality can impact how well people like the agent.

Giving an agent an anthropomorphic representation can have other consequences as well. One common hypothesis is that anthropomorphic agents increase people's expectations of what a system can do. For example, people may expect that a system with a believable, lifelike character will be capable of moving like a human being or understanding natural language. In other words, if people's expectations for a system's personality are met or exceeded, then they subsequently may set higher expectations for other human-like proficiencies.

In the case of animated characters, it may be technically challenging to meet people's expectations for human movement. People are very good at detecting even the slightest deviations from normal human movement. As an example, researchers (Reeves & Nass, 1996) have studied the consequences of creating a mismatch between a picture of a face speaking and the sound of the words being spoken. In this case, the mismatch was subtle enough not to be noticed consciously; the audio was presented one-sixth of a second before the video. Even when the mismatch was unnoticeable, participants had significantly more negative reactions to the asynchronous faces than they did to the naturally moving faces. Participants did not fault the technology; rather, they judged the character to be awkward and unpleasant. This result indicates that people will like an agent less when its representation fails to meet their expectations for natural movement.

The influence of people's expectations for an agent's intelligence also may be important. In one study (King & Ohya, 1996), participants viewed 20 possible agent representations, ranging from simple geometric shapes to very high resolution, animated three-dimensional models of human forms that included random eye blinking movements. Participants were asked to indicate whether each representation was an agent, object, or event, and to rate the "intelligence or potential intelligence" of the thing represented. In general, the participants rated the more human-like representations as being more like agents and more intelligent than other representations. Other researchers (Kiesler & Sproull, 1997) had participants interact with human partners, software agents represented by a face, and software agents represented by a dog. They found that human faces elicited social responses of trust and cooperation that were comparable to the social responses elicited by human partners and stronger than the social responses elicited by the dog-like interface characters.

If people perceive an agent to be intelligent, are they disappointed if it betrays their expectations? In a study (Brennan & Ohaeri, 1994) of a simulated airline reservations agent, embodied in a text-based user interface without a character, participants received one of three message styles: telegraphic, fluent, or anthropomorphic. Participants expended the most effort on their language responses to the anthropomorphic agent, treating it as if it were intelligent. The participants, however, did not have expectations that the anthropomorphic agent would behave more intelligently than the other agents would. People may see anthropomorphic agents as intelligent without requiring them to behave any more competently than nonanthropomorphic agents. A number of researchers (Maes & Kozierok, 1993; Koda, 1996; Wexelblat, 1998) who have examined participants' expectations or frustrations have found very few differences in participants' responses to anthropomorphic agents versus nonanthropomorphic agents. The only difference that is commonly found is that participants find anthropomorphic agents more enjoyable than nonanthropomorphic agents.

Anthropomorphic agents may have some other interesting side effects as well. For example, researchers studying pedagogical agents have found that the presence of a lifelike character can have a positive impact on students' perceptions of their learning experience (Lester et al., 1997). There may be costs to an engaging personality, however; specifically, attractive characters might be taking student's attention away from learning. Other researchers have investigated whether anthropomorphic agents are distracting.

Walker et al. (1994), for example, compared a system with an animated and expressive face to a system with no face. They reported that the face helped to engage people and that it also required some of their attention. Moreover, when the face had a stern expression, people were most productive and had the most negative reaction to the system. Similarly, other researchers have reported that a realistically animated face is both entertaining and distracts people from their tasks (Takeuchi & Nagao, 1993; Takeuchi & Naito, 1994). In general, critics have argued that systems designed with anthropomorphic elements are unnecessarily inefficient (e.g., see the arguments presented by Don et al. (1992), Lanier (1995), and Shneiderman (1995)).

Anthropomorphic agents may be a powerful means of expressing a personality, and they may come with certain liabilities. When and how anthropomorphic elements are best used remain open questions. Are animated characters enjoyable when the goal is entertainment, and are they distracting when the goal is efficient productivity? Are interface characters distracting when they are represented by animation or complex human faces, and are they natural when they are kept simple? It is possible that interface characters are likable inasmuch as they provide a clear personality for a social agent. They are not, however, the only means of expressing personality. Useful personalities can be created for agents without making the agents anthropomorphic.

Our next study concerned how nonanthropomorphic agents might be given useful personalities. The interface character study provided evidence that the personality of an agent can determine how well the agent is liked. Moreover, participants preferred agents with personalities that were similar to their own. To examine this issue further, we studied the match between participants' personalities and agents' personalities. In this study (Nass et al., 1995a), 48 participants worked with computers on a joint problem-solving task. Participants were preselected to have a personality that was either dominant or submissive. The computers were endowed with properties associated with either dominance or submissiveness.

Specifically, we used the phrasing of text, the computer's confidence, the order of turn taking, and the computer's name to operationalize dominance or submissiveness. The computers communicated through full-sentence text displayed on the monitor. In the dominant-computer condition the text consisted primarily of assertions and commands (e.g., "You should definitely rate the flashlight higher. It is your only reliable night signaling device.''). In the submissive-computer condition the text was made up of questions and suggestions (e.g., "Perhaps the flashlight should be rated higher? It may be your only reliable night signaling device.'') During the interaction, both the computer and the human partner gave explicit ratings of their confidence with their proposed solution. In the dominant-computer condition the confidence ratings were high (average ratings on a 10-point scale of 8.0 with a standard deviation of 0.8), and for the submissive-computer condition they were low (average ratings of 3.0 with a standard deviation of 0.8). The interaction was designed so that for each round either the computer or the human participant would take a turn first. In the dominant-computer condition the computer always went first, and in the submissive-computer condition the computer always went last. Finally, the computers were labeled with names; in the dominant-computer condition the computer's name was Max, and in the submissive-computer condition it was Linus. These names were identified as optimally dominant and submissive in a pretest sample.

In order to isolate the impact of dominance, the computers were not designed to differ in friendliness. Indeed, our assessments of the participant's perceptions of the computer's friendliness revealed no significant differences between the submissive-computer and the dominant-computer conditions.

In terms of dominance, however, when the computer and the participant shared similar personality types, (1) the computer was perceived as more competent, $F(1, 44) = 11.60$, $p < 0.001$; (2) the computer was perceived as more likeable, $F(1, 44) = 4.28$, $p < 0.05$; (3) the interaction was more satisfying and more beneficial, $F(1, 44) = 5.18$, $p < 0.05$, and $F(1, 44) = 6.35$, $p < 0.01$, respectively; and (4) the participants were more satisfied with their own performance, $F(1, 44) = 9.02$, $p < 0.005$. In other words, the "best" personality for the social agent depended on the personality of the human partner.

In a follow-up to this study, we were interested in how purely formal attributes of an artifact might interact with people's personalities. We therefore conducted a laboratory experiment (Detenber & Dryer, 1995) to assess the responses of participants with different personalities to various system-controlled visual displays. A machine presented 60 images to 125 participants. For some of the participants the display size was comparable to that of a home television set (roughly, a 10° horizontal visual angle). For other participants the display size was comparable to a movie theater screen (roughly, a 41° horizontal visual angle). Participants' emotional responses to the interaction were assessed along three dimensions: valence ("pleasant'' to "unpleasant''), arousal ("excited'' to "calm''), and dominance ("in control'' to "out of control''). The results indicated that participants with different personalities preferred interactions with different formal attributes. In particular, participants with highly neurotic personalities felt more unpleasant emotions, $F(2, 124) = 3.29$, $p < 0.05$, and more out-of-control emotions, $F(2, 124) = 4.01$, $p < 0.05$, while interacting with the small display than they did with the large display. Participants without this personality type showed the reverse pattern. It may be that people with highly neurotic personalities are anxious to get information about their environment. When people with this predisposition interact with a small display, they may feel frustrated that they do not have better access to the information that they seek (cf. Sparks & Spirek, 1988).

A machine's interaction behaviors need not differ in content; instead, the form alone of the human-machine interaction may have an impact on how people feel. In the study by Nass et al. (1995a) the form of the interaction determined whether the machine or the participant acted first. In the study by Detenber and Dryer (1995) the machine controlled how much information the participant received in the interaction. In both cases, this kind of control can be a social cue for a certain level of interpersonal dominance. The studies suggest that agent personalities can be built out of a certain visual appearance, the phrasing of text, or even the technological form that the interface with the agent takes.

Other researchers have explored how agent personalities might be best expressed. Because people prefer strong personalities, agent personalities work best when they are identifiable and consistent. Interesting personalities tend to be expressed consistently in three ways: (1) cohesiveness, or the tendency for all the expressions of personality (such as behavior, language style, and appearance) to express the same personality; (2) temporal stability, or the tendency for the same personality to be expressed at different times; and (3) cross-situation generality, or the tendency for the same personality to be expressed in different situations (cf. Moskowitz, 1988).

However, personalities also can be interesting when they change; what is important is that the change is clear and meaningful. When a dominant partner follows your lead rather than challenging your control, it has meaning because it communicates something, such as a respect for your authority. Personality "inconsistencies" work when they turn out to be clearly "consistent" after all (cf. Moon & Nass, 1996). Moreover, personality consistency does not require that a partner do the same thing all the time; that would be unnatural. Instead, interesting personalities get expressed in myriad, unique, and even surprising ways.

Finally, perceived personalities are abstractions constructed from a collection of specific behaviors. The five-factors model captures the structure of our perceptions of personality but not how personality is naturally expressed. For example, saying that someone is "open" is nearly meaningless; the experience of personality occurs at a more basic level. An open person might be someone who is always curious to learn new things, fond of wine tasting, and frequently changing jobs. In fact, our liking of social partners is associated with the degree to which we see them as having many, complex facets to their personality (Sande et al., 1988). Part of the artistry of designing a personality involves translating the abstract description of personality into a specific expression.

A CASE STUDY: HOW AGENT PERSONALITIES CAN BE CREATED

> Direct manipulation has its place, and in many respects is part of the joys of life: sports, food, sex, and for some, driving. But wouldn't you really prefer to run your home and office life with a gaggle of well-trained butlers [...] and, on some occasions, cooks, gardeners, and chauffeurs when there were too many guests, weeds, or cars on the road? (Negroponte, 1990)

Software agents have begun to have widespread application. In this new territory, researchers, designers, and developers are finding new ways to help people get the most from their machines by delegating some of the work to agents. Ongoing work in the area of software agents concerns the application of artificial intelligence (AI) technologies to the problems of human-machine interactions (Miller et al., 1991). As future software agents become more widespread, social interface theory suggests that the need for designing and managing personalities will become critical.

If people respond to an agent socially, they will perceive that agent as having a personality. The perceived personalities, however, are likely to be there by chance rather than by design. Ten different error messages may have been written by ten different authors, each with a different style. The result is not the absence of personality; instead, the agent seems to have a confusing or vague personality at best, and a frustratingly schizophrenic personality at worst. People will see a personality in the machine regardless of whether one was intended. These observations suggest some new and exciting opportunities for human-machine interaction designers. Specifically, how can software agents be designed so that their personalities support satisfying social interactions? This work has been started in various laboratories (e.g., Bates, 1994; Hayes-Roth et al., 1995; Maes, 1995; Nass et al., 1995b; Oren et al., 1990) with agents intentionally created to have specific emotions, states, roles, and personalities.

A recent effort to design a personality involved a software agent named "WarpGuide." WarpGuide is a task mentor in IBM's operating system, OS/2 Warp 4 (Dryer, 1997; Selker, 1994). To create a personality for WarpGuide, our interdisciplinary team used the theory and research from the fields of intelligent agents and social psychology to inform our design decisions.

WarpGuide uses software-agent technology to guide people through certain system tasks, helping to prevent them from making errors. WarpGuide communicates in full-sentence text and responds to voice commands (like the other components of OS/2 Warp 4). Because of WarpGuide's social role, use of language, and perceived intelligence, we considered people's possible social responses. Specifically, we addressed the questions (1) would WarpGuide's function be designed best as the behavior of a single social partner or many, (2) what personality would be best for WarpGuide, and (3) how was that personality best expressed?

We first considered the argument for multiple social personalities. People perceive whatever is present as the social partner (Reeves & Nass, 1996). A user interface, however, is not necessarily like a single physical object or single social partner. Rather, a user interface is like a proscenium to a stage populated by people, places, and things (Laurel, 1986). The stage may contain a single actor or an entire troupe. It is easy to see how icons in a graphical user interface might be like distinct physical objects. In the same way, autonomous interface personalities might be like distinct social partners.

As an example, researchers have demonstrated that social entities and voices are matched one to one. Two computers with one voice were perceived as a single social partner; one computer with two voices was perceived as two social entities (Nass & Steuer, 1993; Dryer et al., 1993b). In other words, voice is one clue that people use to integrate and distinguish collections of behaviors into discrete social entities.

In our case, we considered whether WarpGuide's behavior was best used to distinguish it as a social partner, separate from other entities. In general, computer system behaviors tend to be purely reactive responses to user actions. In contrast, WarpGuide works to help people accomplish tasks, behaving at times autonomously and proactively. Moreover, people's perception of WarpGuide as a discrete social entity might be advantageous. Separate social personalities may be used in a "good cop/bad cop" routine. That is, one partner could be made the scapegoat for all of the emotionally negative aspects of an interaction, leaving interactions with another partner emotionally positive (cf. Reeves & Nass, 1996).

We therefore decided that WarpGuide ideally should be distinct from the "system homunculus." The idea was that, psychologically, people might have a troubled relationship with their primary partner, the system. As a partner, the system sometimes is hard to get along with. People's relationship with WarpGuide, however, could be different; WarpGuide would be a distinct social partner, whose only job was helping with tasks. It would be like a voice over the shoulder, guiding people through their interactions with the system. This way, WarpGuide would be "blameless" for whatever else the system did.

Our next decision concerned whether people would be interacting with (1) the system and a single personality called "WarpGuide" or (2) the system and multiple personalities called "WarpGuides." Here, we had to consider two competing arguments. On the one hand, separate social personalities can have the advantage of being "specialists." By virtue of having a narrow range of expertise, people both perceive specialists as being better than "generalists" for that domain and tend to have lower expectations for their behavior outside of that domain (regardless of whether these assumptions are true) (Nass et al., 1996b). On the other hand, people generally like to keep things simple. Managing multiple social relationships can get complicated, turning a potentially simple one-on-one interaction into a group relationship.

WarpGuide would lead people through many different tasks. Certainly, technologists talk about one agent for this task and another for that. Was WarpGuide best thought of as a collection of entities, each being an expert at a single task?

Before we answered that question, another question occurred to us about WarpGuide's behavior. WarpGuide would behave in different ways depending on the task. We had both wizard technology (SmartGuide) and guide technology (COACH) available to us.[2] Because each technology has its advantages for different tasks (Attkinson et al., 1995; Dryer, 1997; Wilson, 1995), we wanted WarpGuide to have both behaviors. Did this difference in behavior mean that we had to create WarpGuide as two separate social personalities?

We decided that what people really wanted was a single social partner that they could turn to for help. WarpGuide would be like that expert down the hall, someone you turn to whenever you have a problem with your computer. Thus, to simplify the interaction, people would have a single, social relationship with one task mentor. WarpGuide would be a skilled mentor, an expert with knowledge that extended across a set of tasks but no further. Moreover, people would not care what underlying technology (guide or wizard) WarpGuide was using for guidance; they would just want the best guidance for the particular task. So WarpGuide was conceived as a single personality with two roles: guide and wizard.

To complete our work, we needed to integrate WarpGuide's roles and behaviors into a single social partner. For OS/2 Warp 4, we did not have the option of giving WarpGuide a single voice, at least not one that could be heard. Instead, we needed to find other ways to create a single, consistent personality. We started with the current research and considered what kind of personality would be best for WarpGuide. In general, people prefer friendly personalities. Similarly, we decided WarpGuide needed to be authoritative but not so dominant that it was perceived as intrusive. As a task assistant, people might prefer either a learning partner (who would be similar to users with respect to computer experience) or an intelligent helper (who would be complementary to users). We could have made WarpGuide an inexperienced agent who would seem to learn along with people. Instead, we choose to embody in WarpGuide the expertise that people would seek. Therefore we defined the personality as friendly, authoritative yet unobtrusive, and intelligent. We then explored ways to create this personality.

To communicate a personality, an entity's behavior needs to be consistent. We therefore endeavored to design WarpGuide's form of interaction to be consistent. As one example, WarpGuide offers a consistent starting point for many different tasks. It presents a collection of "guidance objects" organized in a single location. When a guidance object is opened, WarpGuide brings up the appropriate task interface along with the accompanying assistance. WarpGuide also provides a way to navigate through tasks from beginning to end, even across multiple task interfaces that normally may be scattered throughout the user interface. WarpGuide works the same way across both roles and across a range of tasks, unifying these roles and behaviors into a single personality.

Visual appearance is another clue to personality. WarpGuide is represented graphically in the user interface by a set of objects, windows, icons, and annotations. As shown in Figure 3, all these elements are unified by common visual themes: in particular, a characteristic cue-card form, an annotated dialog, certain colors, and iconic "eyes" (suggesting its social role). The colors and forms are designed to be "warm" yet authoritative, representing friendliness and dominance.

We used "friendlier" rounder shapes, such as the rounded rather than angular corners of the cue card, and we went after a happy, "friendly" look in the design of the iconic eyes in the title bar control, attempting to use their form to represent the appropriate personality. All of these elements are

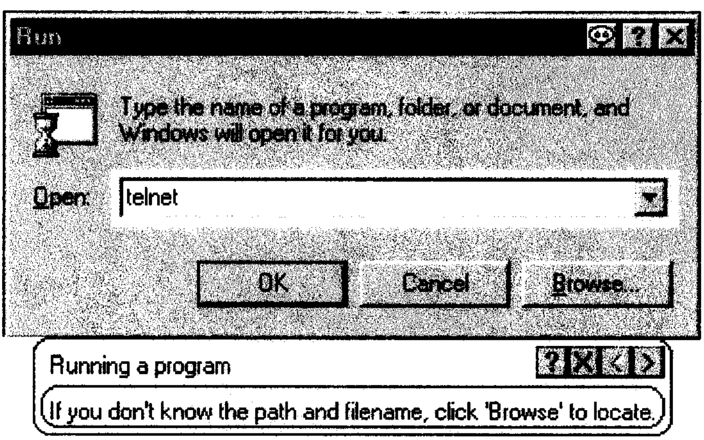

FIGURE 3. WarpGuide appearance.

unified to some extent by a bold color. Because WarpGuide's form is a set of visual elements that are variations on a single theme, that form helps to express a single personality.

The next step in building WarpGuide's personality was to consider its message content. Research has shown that text alone is a powerful cue to personality (Nass et al., 1995a). We therefore fashioned a style of speech around strict guidelines. We developed templates for sentence structure and wrote each text article to fit the templates. In addition, multiple team members reviewed each article, to ensure that (1) the content was intelligent (2) the tone and phrasing were friendly and authoritative, and (3) the structure was consistent, concise, and unobtrusive. As an example, an article for finding a file dialog is "Type as much of the name as you know, and use an asterisk (*) for the rest." A version of the sort, "You can type only as much as you know, and use an asterisk (*) for the rest" makes the software agent seem less authoritative. A version of the sort, "If you don't know the whole name, type what you know and use a asterisk for what you don't know" makes the software agent seem less friendly. Of all the design decisions, the phrasing of the text was probably the most important to establish WarpGuide's personality.

Finally, although we used behavior, images, and text phrasing deliberately to construct a personality, we also chose to avoid anthropomorphism and the possible overuse of social metaphors. In some contexts, people may prefer explicit social characterizations, such as the use of interface characters. Our goal, however, was to make the most of social responses that people would have, not artificially create a social relationship. We therefore attempted to strike a balance between a lifelessly bland representation and one that would be annoyingly affected. In particular, we avoided a full face or figure and did not use first-person pronouns in the text. WarpGuide was not meant to be "cute," nor was it meant to portray itself as "human." Instead, our intent was that it would behave socially in ways that make sense given its context.

GENERALIZATIONS: TIPS FOR DESIGNING PERSONALITIES

As designers become more aware of the impact that social responses have in human-machine interactions, they will face some of these same issues. Our experience can be summarized with the following guidelines:

1. Language use, fulfilling a social role (like helping) and perceived intelligence are likely to encourage social responses and the perception of a personality.

2. People perceive the personalities of artificial agents along the same dimensions (especially friendliness and dominance) that they perceive human personalities.
3. Agent personalities can be communicated through behaviors (such as praising others), linguistic phrasing, appearance, and interaction form.
4. Agents that are cooperative, outgoing, calm, organized, or curious are liked better than those that are competitive, withdrawn, anxious, lax, or close-minded.
5. Personalities with a foible are liked.
6. Strongly expressed personalities are better liked than subtle ones.
7. In general, an agent with a personality that is similar to that of a user is liked better than an agent with a personality dissimilar to a user.
8. Sometimes a complementary personality is better than a similar one.
9. Consistent or meaningfully changing personalities are liked.
10. People prefer personalities that are expressed specifically rather than vaguely.
11. Multiple personalities can exist in a user interface. If you are likely to have a social response, consider designing one or more discrete personalities to manage that relationship.
12. If people are likely to be frustrated, consider designing a personality that they can blame and another that can help them.
13. People generally like simplicity; integrate as much behavior as possible into a single personality. Technological distinctions (like guide versus wizard) may not be psychologically relevant.
14. Create a social partner in a way that makes sense for the context. The goal is not to duplicate mindlessly the natural and social world in cyberspace; animation, faces, and the pronoun "I" may not be appropriate.

END NOTES

1. The five characters that were not assigned to either of the two personality strength groups were dropped from the analysis. Consequently, the repeated measures analysis of variance had 2 (character dominance) × 2 (character friendliness) × 2 (character personality strength) × 4 (trial) within subject factor levels and 2 (participant dominance) × 2 (participant friendliness) × 2 (participant experience) × 2 (participant gender) between subject factor levels.
2. Wizards are a common type of user interface agent. Wizards offer task assistance to people by breaking a task into a (typically) linear series of steps and presenting the steps to a person one at a time. A well-designed wizard replaces a multipurpose user interface with a task-specific interface that guides a person along an efficient path to success, autonomously completing those steps of the task that do not require the person's attention. Guides are another kind of user interface agent. Typically, guides provide task assistance by presenting some kind of context-sensitive direction (Oren et al., 1990; Selker, 1994). In particular, a guide could annotate a graphical user interface in a manner that communicates how to perform the next step in a task and draw a person's attention to where the next step

will occur. The specific guide technology used in this research (COACH) monitors task interactions, develops a model of a person's experience with certain task steps, and then uses this model to determine when and how task annotations are presented.

REFERENCES

Aronson, E., B. Willerman, and J. Floyd. 1966. The effect of a pratfall on increasing interpersonal attractiveness. *Psychonomic Science* 4:227–228.

Attkinson, B., S. Brady, D. Gilbert, D. Levine, P. O'Connor, D. Osisek, S. Spagna and L. Wilson. 1995. IBM intelligent agents. In *UNICOM Seminar Proceedings*.

Bates, J. 1994. The role of emotion in believable characters. *Communications of the ACM* 37:7.

Benjamin, L. S. 1974. Structural analysis of behavior. *Psychological Review* 81:392–425.

Brennan, S. E., and J. O. Ohaeri. 1994. Effect of message style on users' attributions toward agents. In *CHI' 94 Conference: Companion Human Factors in Computing Systems*, 281–282. Boston, Mass.: ACM Press.

Byrne, D. 1971. *The attraction paradigm*. New York: Academic.

Cantor, N., and W. Mischel. 1979. Prototypes in personal perception. *Advances in Experimental Social Psychology* 12:3–52.

Conte, H. R., and R. Plutchik. 1981. A circumplex model for interpersonal personality traits. *Journal of Personality and Social Psychology* 40:701–711.

Croom, W. C., J. M. Wallace, and J. M. Schuerger. 1989. Jungian types from Cattellian variables. *Multivariate Experimental Clinical Research* 9(1):35–40.

Detenber, B., and D. C. Dryer. 1995. Eye of the beholder: Individual differences in responses to mediated images. Presented at the annual meeting of the International Communication Association, Sante Fe, N. Mex.

Don, A., S. Brennan, B. Laurel, and B. Shneiderman. 1992. Anthropomorphism: From Eliza to Terminator 2. In *Proceedings of CHI '92: Human Factors in Computing Systems*, 67–70. Monterey, Calif.: ACM Press.

Dryer, D. C. 1995. Personalities of user interface characters. Technical report for Walt Disney Pictures and Television, Burbank, Calif.

Dryer, D. C. 1997. Wizards, guides, and beyond: A rational and empirical method for selecting optimal intelligent user interface agents. In *Proceedings of the Intelligent User Interface Conference, Orlando, Fla.*, 265–268.

Dryer, D. C., and L. M. Horowitz. 1997. When do opposites attract? Interpersonal complementarity versus similarity. *Journal of Personality and Social Psychology* 72(3):592–603.

Dryer, D. C., T. Jacob, and V. Shoham-Saloman. 1994. The communication style typology of families. Presented at the Annual Meeting of the American Psychological Association, Los Angeles, Calif.

Dryer, D. C., C. Nass, J. S. Steuer, and L. Henriksen. 1993a. Interpersonal responses to computers programmed to function in social roles. Presented at the Annual Meeting of the Western Psychological Association, Phoenix, Ariz.

Dryer, D. C., J. Steuer, L. Henriksen, E. Tarber, and H. Reeder. 1993b. Computers are social actors II: Voice is the mirror of the soul. Presented at the Annual Meeting of the International Communication Association, Washington, D.C.

Fogg, B. J., and C. Nass. 1997. Silicon sycophants: The effects of computers that flatter. *International Journal of Human-Computer Studies* 46(5):551–561.

Furnham, A. 1996. The big five versus the big four: The relationship between the Myers-Briggs Type Indicator (MBTI) and NEO-PI five factor model of personality. *Personality and Individual Differences* 21:303–307.

Gurtman, M. B. 1991. Evaluating the interpersonalness of personality scales. *Personality and Social Psychology Bulletin* 17:670–677.

Hayes-Roth, B., L. Brownston, and E. Sincoff. 1995. Directed improvisation by computer characters. Knowledge Systems Laboratory, Stanford University, Technical Report KSL-95-04.

Hildebrand, D. K. 1986. *Statistical thinking for behavioral scientists*. Boston, Mass.: Duxbury Press.

Horn, J. L., and R. B. Cattell. 1966. Refinement and test of the theory of fluid and crystallized ability intelligences. *Journal of Educational Psychology* 57:253–270.

Horowitz, L. M., D. C. Dryer, and E. N. Krasnoperova. 1997. The circumplex structure of interpersonal problems. In *Circumplex models of personality and emotions*, eds. R. Plutchik and H. R. Conte. Washington, D.C.: American Psychological Association.

Horowitz, L. M., K. D. Locke, M. B. Morse, S. V. Waikar, D. C. Dryer, E. Tarnow, and J. Ghannam. 1991. Self-derogations and the interpersonal theory. *Journal of Personality and Social Psychology* 61:68–79.

Horowitz, L. M., S. E. Rosenberg, B. A. Baer, G. Ureno, and V. S. Villasenor. 1988. Inventory of interpersonal problems: Psychometric properties and clinical applications. *Journal of Consulting and Clinical Psychology* 61:549–560.

Kiesler, D. J. 1983. The 1982 interpersonal circle: A taxonomy for complementarity in human interactions. *Psychological Review* 90:185–214.

Kiesler, S., and L. Sproull. 1997. "Social" human-computer interaction. In *Human values and the design of computer technology*, ed. B. Friedman. Cambridge, Mass.: CSLI Publications.

King, W. J., and J. Ohya. 1996. The representation of agents: Anthropomorphism, agency, and intelligence. *In Proceedings of the CHI Conference, Vancouver, BC, Canada*.

Knutson, B. 1996. Facial expressions of emotion influence: Interpersonal trait influences. *Journal of Nonverbal Behavior* 20:165–182.

Koda, T. 1996. Agents with faces: The effects of personification. In *Proceedings of Human-Computer Interaction '96, London*.

Krokoff, L. J., J. M. Gottman, and A. K. Roy. 1988. Blue-collar marital interaction and a companionate philosophy of marriage. *Journal of Personal and Social Relationships* 5:201–222.

Lanier, J. 1995. Agents of alienation. *Interactions* 2:3.

Laurel, B. 1986. Interfaces as mimesis. In *User centered design: New perspectives on human-computer interface*, eds. D. A. Norman and S. W. Draper. Hillsdale N.J.: Lawrence Erlbaum.

Leary, T. F. 1957. *Interpersonal diagnosis of personality*. New York: Ronald Press.

Lester, J. C., S. A. Converse, S. E. Kahler, S. T. Barlow, B. A. Stone, and R. S. Bhogal. 1997. The persona effect: Affective impact of animated pedagogical agents. In *Proceedings of CHI '97*, 359–366.

Maes, P. 1995. Artificial life meets entertainment: Lifelike autonomous agents. *Communications of the ACM* 38(11):108–114.

Maes, P., and R. Kozierok. 1993. Learning interface agents. In *Proceedings of AAAI'93*. Menlo Park, Calif.: International Joint Conference on Artificial Intelligence, Inc.

McCrae, R. R., and O. P. John. 1992. An introduction to the five-factor model and its applications. *Journal of Personality* 60:175–215.

Miller, J. R., J. W. Sullivan, and S. W. Tyler. 1991. Introduction. In *Intelligent user interfaces*, eds. J. W. Sullivan and S. W. Tyler. New York: Addison Wesley.

Mischel, W. 1968. *Personality and assessment*. New York: Wiley.

Moon, Y., and C. Nass. 1996. Adaptive agents and personality change: Complementarity versus similarity as forms of adaptation. In *Proceedings of the CHI Conference, Vancouver, BC, Canada*.

Moskowitz, D. S. 1988. Cross-situational generality in the laboratory: Dominance and friendliness. *Journal of Personality and Social Psychology* 54:829–839.

Nass, C., B. J. Fogg, and Y. Moon. 1996a. Can computers be teammates? *International Journal of Human-Computer Studies* 45(6):669–678.

Nass, C., Y. Moon B. J. Fogg, B. Reeves, and D. C. Dryer. 1995a. Can computer personalities be human personalities? *International Journal of Human-Computer Studies* 43:223–239.

Nass, C. I., Y. Moon, B. J. Fogg, B. Reeves, and C. Dryer. 1995b. Human-computer interaction as interpersonal interaction. In *Proceedings of the CHI Conference, Denver, Colo*.

Nass, C., Y. Moon, and N. Green. 1997a. Are machines gender neutral? Gender-stereotypic responses to computers with voices. *Journal of Applied Social Psychology* 27:864–876.

Nass, C. I., Y. Moon, J. Morkes, E. Kim, and B. J. Fogg. 1997b. Computers are social actors: A review of current research. In *Human values and the design of computer technology*, ed. B. Friedman. Cambridge, Mass.: CSLI Publications.

Nass, C., B. Reeves, and G. Leshner 1996b. Technology and roles: A tale of two TVs. *Journal of Communication* 46:121–128.

Nass, C., and J. S. Steuer. 1993. Computers, voices, and sources of messages: Computers are social actors. *Human Communication Research* 19:504–527.

Nass, C., J. S. Steuer, L. Henriksen and D. C. Dryer. 1994. Machines and social attributions: Performance assessments of computers subsequent to "self-" or "other-" evaluations. *International Journal of Human-Computer Studies* 40:543–559.

Negroponte, N. 1990. Hospital corners. In *The art of human-computer interface design*, ed. B. Laurel. Reading, Mass.: Addison Wesley.

Novak, D. W., and M. J. Lerner. 1968. Rejection as a function of perceived similarity. *Journal of Personality and Social Psychology* 9:147–152.

Oren, T., G. Salomon, K. Kreitman, and A. Don. 1990. Guides: Characterizing the interface. In *The art of human-computer interface design*, ed. B. Laurel. Reading, Mass.: Addison Wesley.

Reeves, B., and C. Nass. 1996. *The media equation*. Cambridge, Mass.: Cambridge University Press.

Sande, G. N., G. R. Goethals and C. E. Radloff. 1998. Perceiving one's own traits and others': The multifaceted self. *Journal of Personality and Social Psychology* 54:13–20.

Selker, T. 1994. Coach: A teaching agent that learns. *Communications of the ACM* 37(7):92–99.

Shneiderman, B. 1995. Looking for the bright side of user interface agents. *Interactions* 2:1.

Shweder, R. A. 1975. How relevant is an individual difference theory of personality? *Journal of Personality* 43:455–485.

Sparks, G. G., and M. M. Spirek. 1988. Individual differences in coping with stressful stimuli: An activation-arousal view. *Human Communication Research* 15:195–216.

Strong, S. R., and H. I. Hills. 1986. *Interpersonal communication rating scale*. Richmond Va: Virginia Commonwealth University.

Takeucki, A., and K. Nagao. 1993. Communicative facial displays as a new conversation modality. In *Proceedings of INTERCHI'93 Human Factors in Computing Systems*, 24–29 April, 187–193. Amsterdam, The Netherlands: ACM Press.

Takeuchi, A., and T. Naito. 1994. Situated facial displays: Towards social interaction. In *Proceedings of CHI'94: Human Factors in Computing Systems*, 24–28 April. Boston, Mass.: ACM Press.

Taylor, S. E., and D. Mettee. 1971. When similarity breeds contempt. *Journal of Personality and Social Psychology* 20:75–81.

Thomas, F., and O. Johnston. 1981. *Disney animation: The illusion of life*. New York: Abbeville Press.

Uleman, J. S., A. Hon, R. Roman, and G. B. Moskowitz. 1996. On-line evidence for spontaneous trait inferences at encoding. *Personality and Social Psychology Bulletin* 22:377–394.

Walker, J., L. Sproull, and R. Subramani. 1994. Using a human face in an interface. in *Proceedings of CHI'94: Human Factors in Computing Systems*, 24–28 April, 85–91. Boston, Mass.: ACM Press.

Wexelblat, A. 1998. Don't make that face: A report on anthropomorphizing an interface. In *Intelligent Environments*, 173–179. AAAI Press Technical Report SS-98-02.

Wiggins, J. S. 1979. A taxonomy of trait-descriptive terms: The interpersonal domain. *Journal of Personality and Social Psychology* 37:395–412.

Wiggins, J. S., and R. Broughton. 1985. The interpersonal circle: A structural model for the integration of personality research. *Perspectives in Personality* 1:1–47.

Wilson, L. 1995. Intelligent agents: A primer. *Personal Systems* (September/October):47–49.

Wish, M., M. Deutsch and S. J. Kaplan. 1976. Perceived dimensions of interpersonal domain. *Journal of Personality and Social Psychology* 33:409–420.

Young, F. W. 1987. *Multidimensional scaling: History, theory, and applications*, ed. R. M. Hamer. Hillside, N.J.: Lawrence Erlbaum.