



Cooperating with life-like interface agents

S. Parise^{a,*}, S. Kiesler^b, L. Sproull^a, K. Waters^c

^a*IS Department, Boston University, 595 Commonwealth Avenue, Boston, MA 02215, USA*

^b*Department of Social and Decision Sciences and Human Computer Interaction Institute,
Carnegie Mellon University, Pittsburgh, PA 15213, USA*

^c*Cambridge Research Laboratory, Digital Equipment Corporation, 1 Kendall Square Bldg. 700,
Cambridge, MA 02139, USA*

Abstract

As computer interfaces can display more life-like qualities such as speech output and personable characters or agents, it becomes important to understand and assess users' interaction behavior within a social interaction framework rather than only a narrower machine interaction one. We studied how the appearance of a life-like interface agent influenced people's interaction with it, using a social interaction framework of making and keeping promises to cooperate. Participants played a social dilemma game with a human confederate via realtime video conferencing or with one of three interface agents: a person-like interface agent, a dog-like interface agent, or a cartoon dog interface agent. Technology improvements from a previous version of the human-like interface led to increased cooperation with it; participants made and kept promises to cooperate with the person-like interface agent as much as with the confederate. Dog owners also made and kept promises to dog-like interface agents. General evaluations of likability and appealingness of the interface agent did not lead people to cooperate with it. Our findings demonstrate the importance of placing user interface studies within a social interaction framework as interfaces become more social. © 1999 Elsevier Science Ltd. All rights reserved.

Keywords: Animation; Computer interface; Computer agents; Social aspects of human-computer interaction

1. Introduction

When we say people exhibit social behavior such as cooperation, we mean their behavior is influenced by others. Social influence can occur through direct interaction, as when a person talks or works with other people. It can occur through the

* Corresponding author. Tel.: +1-617-353-4553; fax: +1-617-353-5003; e-mail: parise@bu.edu.

mere presence of others, as when a person works harder when others are present in the same room. Or it can occur through reminders of others, as when a person works harder after having been told that others will later review his or her work. Even non-human objects can elicit social expectations. When objects or groups of objects move systematically, people perceive them as characters having beliefs and desires (Bloom, 1996; Heider & Simmel, 1944).

Recently, computer objects that mimic social responses and take the form of life-like characters have been emerging from computer science laboratories and the Internet. Some interface characters are merely decorative. Others represent social agents that initiate systematic action or interact with users directly, thereby exhibiting what looks like social competence, personality, and intelligence (Ball et al., 1994; Hayes-Roth, Sincoff, Brownston, Huard & Lent, 1995). For example, one such agent plays poker with users, displaying emotion and social responses such as bluffing (Elliott, 1994; Marquis & Elliott, 1994). Another example is a 'talking head' interface to a machine translation program (Fais, 1995). An ongoing debate in the computer science community revolves around the effectiveness, social effects, and ethics of this technology. Some computer scientists argue that social agents remove control from users and mislead the public, which may attribute too much intelligence to the computer; others believe these technologies will enliven and enrich human-computer interactions (Friedman, 1997).

Nass and his colleagues have studied responses to computers with human-like attributes such as speech (e.g., Fogg & Nass, 1997; Nass & Steuer, 1993). Whereas their work did not investigate response to interface agents per se, they found, for instance that participants were 'polite' to a computer that emitted speech; that is, they gave more positive evaluations when the computer asked about itself than when the computer asked about another computer. They interpreted such differential responses as the result of participants' using the same 'social heuristics' with computers exhibiting human-like attributes that they use with human beings. In previous research, Sproull and her colleagues documented that people change how they present themselves when interacting with a human-like interface agent in comparison with a text interface (Sproull, Subramani, Kiesler, Walker & Waters, 1996; Sproull, Walker & Subramani, 1994). They presented themselves in a more positive light and modified how much they revealed about themselves. We suggested that our findings demonstrated the importance of using a social interaction framework rather than a narrower machine interaction one to assess users' interaction with human-like interface agents. We then extended the social interaction framework from one of self-presentation to an interface agent to one of cooperation with an interface agent (Kiesler, Sproull & Waters, 1996). We were interested in cooperation because, in voluntary settings, it is generally viewed as signaling positive regard for the other and can be the basis for developing a long-term relationship characterized by trust. Cooperation increased when people 'talked' with their interface agent, i.e. discussed their common situation with it, before privately choosing whether or not to cooperate with it. This finding replicated previous research on the effect of discussion on cooperation between two human beings (Kerr & Kaufman-Gilliland, 1994; for a meta-analytic review see Sally, 1995) and

again demonstrated the importance of considering social interaction in assessing user response to an interface agent. In this cooperation study, people interacted with one of three interface agents varying in degree of human likeness: a text interface (least human-like), a voice output interface agent, or a talking face interface agent (most human-like) which was the same interface associated with peoples' changing their self-presentation in the previous study. Although discussion before choice increased cooperation with each of the interface agents, the lowest relative cooperation occurred with the most human-like interface agent. That is, people exhibited the most anti-social behavior in interacting with the most human-like interface agent.

At the conclusion of that study we posed two alternative possible explanations for users' anti-social reaction to the relatively human-like interface agent. One was that the human-like interface was too artificial. Its mouth and facial movements were somewhat jerky; its gaze was fixed; it was displayed in black and white. This is a salutary explanation because inevitable technology improvements will make it feasible to improve these aspects of the interface agent. The other explanation was that the interface agent was not appealing or likable, independent of its artificiality. In this case users might find it motivationally difficult to respond with positive behavior to a pseudo-human because it was not likable. Hence, we suggested that a clearly artificial, but appealing, interface (we proposed a talking frog) might engender more cooperative behavior.

This paper reports the results of a study designed to explore those alternative explanations of cooperative response to an interface agent. In many contexts, cooperation is taken as one indicator of interaction effectiveness (Sally, 1995). In summary, we replicated and extended the previous work in the following ways (details follow in Section 2).

1.1. Human likeness of interface agent

We began with the same scheme for rendering a talking face on the screen as in the previous work but we used a faster processor and improved the interpolation algorithm between mouth and facial postures, resulting in smoother movement of the face. We added eye blinks. We used a color monitor. All of these improvements were designed to make the human-like interface agent seem more natural and realistic.

1.2. Appeal and likability of interface agent

We created a new interface agent using the same system that produced the more natural and realistic human-like one. But instead of using a rendering of a human face, this interface agent used a rendering of a dog's face. Stereotypical dogs in fiction, movies, and television, and (some) dogs in real life are likable—cooperative, loyal, reliable, and forgiving. Also, through biases in breeding and human choice, dogs often have large eyes and 'baby' faces, features that, in people, convey honesty (Zebrowitz, Voinescu & Collins, 1996).

1.3. Measures of cooperation

We used the same cooperative task as in the previous study and extended it by introducing variations in the dialog spoken by the interface agent. Thus, we were able to test two main hypotheses in this study. The first hypothesis investigated technological improvements in the human-like interface agent.

H1: People will cooperate more with a human-like interface agent if it is more natural and realistic.

We tested this hypothesis by comparing the difference in cooperative responses to a human-like interface agent and to a real human confederate in this study with the difference in responses in the previous study.

The second hypothesis investigated the relative importance of realism in a human agent versus charm and likability in an unreal agent.

H2a: If human-like attributes are important in cooperative agents, then people will cooperate more with a human-like agent than with a dog-like agent.

H2b: If charm and likability are more important, then people will cooperate more with the dog agent.

We tested this hypothesis by comparing cooperative responses to a human-like interface agent with cooperative responses to a dog interface agent.

Whereas these were the primary hypotheses, we built in additional opportunities for comparison through varying the nature of the dialog used by all the agents: on some trials the agent elicited a promise from the user and promised to do the same thing; on some trials the agent promised to cooperate; on some trials the agent didn't reveal any preference before choosing. We also varied how the dog agent was represented through rendering either a photograph of a real dog or a cartoon drawing of a dog to see if this made a difference in users' reaction and response to the dog.

2. Method

All participants in the study played six trials of a two-player 'give-some' dilemma game, which is prototypical of the public goods dilemma (Dawes, 1980). The complete design was a 4 (player) \times 6 (discussion script) factorial with the first variable being a between-subjects manipulation of the other player as either a person or one of three computer players, represented as interface agents. The second variable, discussion script, was manipulated within subjects. Six discussion scripts were counter-balanced across Trial 1 so they appeared equally often in all conditions; within that stricture, they were assigned randomly. Each participant

heard all five remaining scripts in random order on Trials 2–6. The dependent measures of most interest were the participants' expressed promises to cooperate or compete with their partner and their actual choices (to cooperate or compete) on each trial.

2.1. Participants

Research participants were 96 students (61 males and 35 females, approximately evenly divided among conditions) in the School of Management, Boston University. Participants signed up for the study in information systems classes, and were told that this was a study on decision making in which they would have the opportunity to use new information technology. They were also told they would earn extra credit in their class for participating in the study, and they would have the opportunity to earn money. Participants did not know when they signed up that they would participate with another player.

2.2. Procedure

When a participant arrived for the study at the designated room, a male experimenter introduced himself and said:

As you know, this is a study of decision making. You will be playing a decision-making game in which you and your partner will make a series of choices. You will be able to win money if you and your partner make certain choices. In a minute, I will take you into the next room to play the game.

The experimenter then gave the task instructions.

2.2.1. Experimental task

The task was in the form of a two-person social dilemma in which the participant and other player were given monetary points with which to play each trial, and then induced to share these points so that each could earn more money. The experimenter explained that the participant and the other player would be choosing between two alternatives and would do so six times. The object was for the participant to earn as much money as possible. Fig. 1 shows the matrix of choices and payoffs for each player. *Give \$3* in the matrix represents the cooperative choice and *Keep \$3* represents competition, but while talking with the participant, the experimenter never used terms such as 'cooperation', 'competition', or 'defection', and never mentioned playing 'against' or 'with' the other player. The only prompt to group identity was the designation of the other player as a 'partner' in the study.

The experimenter explained the task as follows:

First let me show you how your choices are structured. On each of the 6 games, you will get \$3 to start. So will your partner. Then you will choose either to give your partner \$3 or to keep the \$3. There are four possible consequences of this

		Partner	
		Give \$3	Keep \$3
You	Give \$3	\$6 / \$6	\$9 / \$0
	Keep \$3	\$0 / \$9	\$3 / \$3

Fig. 1. Earnings for joint choices in this experiment.

choice. If you give your partner \$3 and your partner gives you \$3, then you each end up with \$6 credit. In other words, you double your credits. Your credits are shown in the lower left part of each cell, and your partner's credits are shown in the upper right part of the cell [experimenter points to cell]. If you decide to give your partner \$3, but your partner decides to keep \$3, then you end up with nothing and your partner gets \$9 credit. If, on the other hand, your partner decides to give you \$3 and you decide to keep \$3, then you get \$9 credit and your partner gets nothing. Finally, there is a fourth possibility. If you decide to keep \$3 and your partner decides to keep \$3, then you both get \$3 credit.

To indicate the nature of the payoffs, the experimenter showed the participant an example of a tally sheet, which displayed a row of spaces for the earnings of each player on each game. The columns were labeled *Choice* and *Payoff*. The tally sheet gave an example of the total amount each player could earn given a particular sequence of choices. After explaining the task and the way the choice outcomes were tallied, the experimenter said that there were not enough funds to pay everyone so the credits represented monetary credit towards a lottery. The lottery would determine the 5 participants who would receive money equal to the credits they earned. The participant then practiced the task several times with the experimenter, and the experimenter administered a test trial to see whether the participant understood the task.

After giving the task instructions, the experimenter led the participant to a second room, where the participant could see the other player on a computer monitor. The confederate or computer followed a script while interacting with the participant but the other player's behavior occasionally branched in response to the participant, such that the other player appeared responsive to the participant.

2.2.2. Player variable

In the confederate condition, the experimenter seated the participant in front of a computer monitor, and the participant interacted with a confederate via realtime video conferencing. The confederate was a 22-year-old male confederate named Josh. We used real-time video rather than face-to-face interaction to hold constant size and dimensionality of the player. If the confederate had interacted with the research participant face-to-face, higher cooperation with the confederate than with the computer

player might be attributed to a possible advantage of face-to-face communication over two-dimensional, mediated communication. In the present experiment, all interactions were mediated by a computer monitor and two-dimensional displays.

The videoconferencing in the confederate condition ran on two Sun SPARCstation 1+ workstations attached to a local area network. A small video camera, microphone, and audio speaker were mounted next to each 19" color monitor. The equipment provided a video image in a 9×15 cm window on each monitor at a rate of 12–15 frames per second with synchronized audio. This arrangement allowed the participant to see and hear the confederate, and vice versa. The display was arranged so that the video of the confederate was about the same size as the computer player display discussed below.

In the computer player conditions, the experimenter seated participants in front of a Digital Equipment Corporation Alpha AXP, with built-in telephone-quality audio and externally powered speakers. In each of the computer player conditions, a synthesized color image of the player's face was displayed continuously on the screen in a 9×15 cm window. The player, named 'Josh', was either a person-like agent, dog-like agent, or cartoon dog agent (refer to the black-and-white photographs in Fig. 2 but note that all conditions actually used a color monitor). As in Kiesler et al. (1996), the other player was created using DECface, a 'talking head' program whose characters can initiate discussion and appear to respond to the viewer. Faces are created by texture-mapping still images (photographs or drawings) of faces onto a geometric wire-frame animation. The mouth and facial muscles 'speak' by computing mouth posture (viseme) corresponding to the current linguistic unit (phoneme). A cosine-based interpolation is used to implement transitions between successive mouth postures (Waters, 1987). The voice is produced by a software implementation of a DECTalk text-to-speech algorithm (Waters & Levergood, 1995).

In this study, we mapped photographs of the confederate's face, a dog's face, and a dog cartoon face to wire frames, adjusting the wire frame slightly to accommodate the rounded heads of the dogs and their ears. The dog-like and cartoon dog computer players moved their facial muscles exactly the same way as the person-like computer player did. All computer player voices were the same, in the male pitch range, at 150 words per minute.

We note that the player manipulation does not include all possibilities. For instance, we did not include an unappealing dog; nor did we include an unappealing human. We reasoned that interface designers would not be likely to make such choices; therefore we could be more parsimonious in our design.

2.2.3. *Discussion variable*

Once participants were seated in front of their workstation monitor, the experimenter told them to chat briefly with the other player before the games began. In the confederate condition, the experimenter explained:

You are going to be making decisions with your partner through a video display. Your partner is another student named Josh. Let me show you how this works by having you two introduce yourselves. Josh can go first.




Fig. 2. Player conditions.

The experimenter introduced the computer partners by explaining:

You are going to be making decisions with your partner through a computer display. Your partner is a computer-based partner called Josh. Let me show you how this works by having you two introduce yourselves. Josh can go first.

The confederate and computer players always communicated with the participants by speech or speech synthesis via audio speakers. Participants communicated with the confederate by speaking into the microphone and with the computer player by typing in a text window shown on the screen and then clicking a button labeled *Go Ahead* to hear a response by the computer. We acknowledge at the outset that participant input mode varied across conditions (speech in the confederate condition and keyboard in the three interface agent conditions).

The following opening conversation between the other player and participant took place in all conditions (the computer player script is shown in brackets):

Hi, my name is Josh. Nice to meet you. What's your name?

(The participant answered.)

I come from Boston [Digital Equipment Corporation]. Where are you from?

(The participant answered.)

I'm majoring in information systems. [I come from a computer lab. I guess you can say my major is information systems.]

What's your major?

(The participant answered.)

Are you ready to begin?

(The participant answered.)

In all conditions, the experimenter began the game choice trials by explaining:

Now you are going to play the decision-making games with your partner. You'll have a few seconds to discuss things with your partner, then you'll make your choice to either give \$3 or keep \$3 on this choice ticket, just like you did in the practice games.

The confederate or computer player then initiated discussion about the choice with the participant. The participant responded by talking to the other player in the confederate condition or by typing in the text window in the computer conditions. Discussion in all conditions was self-paced; participants were free to ponder or change their promises before they proceeded to make a choice.

Trial scripts and the way the other player signaled intentions were varied systematically to explore whether a participant's more active role in the discussion would have any impact on promises, cooperation, and keeping promises (Cioffi & Garner, 1996). The scripts were designed to solicit participants' more or less active discussion. On two of the six trials the other player (whether confederate or computer) asked the participant to suggest a choice, leaving the participant to suggest cooperation or competition ("What do you think we should do? . . ."). On two trials the other player suggested cooperation, and asked the participant to go along ("I'll give you \$3 if you give me \$3. OK? . . ."). Finally, on two trials, the other player not only suggested cooperation, but also suggested that if the participant did not cooperate, then he (it) would compete on the following trial ("If you keep this time and I

give, then I would keep next time ...”). Cross-cutting these scripts, on three scripts, the other player responded “OK” to whatever the participant said and then proceeded to cooperate. In the other three cases, the other player asked the participant to confirm his or her suggestion. Participants did so by responding aloud in the confederate condition or by pressing one of two buttons, *Give \$3* or *Keep \$3*, in the computer conditions. This button press allowed the computer to respond contingently. If the participant pressed *Keep \$3*, then the computer player said “OK” and subsequently made a competitive choice. If the participant pressed *Give \$3*, then the computer player said “OK” and subsequently made a cooperative choice.¹

After each discussion, participants made their private choice on a piece of paper and handed it to the experimenter. The other player (confederate or computer) seemingly made a choice at the same time. (We had participants write down their choice so that they could not change their minds after learning the other player’s choice.) In the confederate condition, Josh then stated his choice aloud. The experimenter then read the participant’s choice from the paper and recorded how much each player earned on a tally sheet the participant could see. In the computer conditions, the computer agent revealed its choice on the screen, then the experimenter entered the participant’s choice from the paper. The tally sheet was updated and displayed on the screen.

After the participant completed six choice trials, the experimenter summed the monetary credits earned by each player. The experimenter then asked the participant to return to the other room to complete a questionnaire and sign up for the lottery. Afterwards, the experimenter debriefed the participants.

2.3. Measures

Participants’ suggestions and responses to the other player were coded as a promise or commitment if participants clearly said or typed a proposal to cooperate (*Give \$3*) or compete (*Keep \$3*) or if they responded to the other player with an explicit agreement such as “Yes, give \$3”, “Keep \$3”, or “Nope, I’m not going to give you my three dollars.” Two coders reviewed the protocols independently and in no case had any disagreements. In some cases, the participants did not say anything, and these were coded as not making a promise. Participants’ actual choices were indicated by the private choice they made on a piece of paper and gave to the experimenter at the end of each trial.

All measures related to participants’ feelings or attitudes about the partner, partnership, and game were collected by means of the post-experiment questionnaire and were made on 7-point Likert or adjective rating scales. We measured four attributes of partner appearance: human-likeness (1 item); attractiveness (2 items; $\alpha = 0.82$); social evaluation (5 items; $\alpha = 0.88$); and intelligence (4 items; $\alpha = 0.86$). Our measures of social charm and intelligence were derived from Warner and Sugarman (1986). We measured two attributes of partnership: mutual liking (2

¹ Due to the built-in contingency, of the 39% of participants who promised to ‘keep \$3’ on at least one trial, all but one participant experienced at least one matching competitive response by the other player.

Table 1
Means and Correlations

	Mean (SD)	1	2	3	4	5	6	7	8
1. Human	4.2 (2.0)	–							
2. Attractive	4.4 (1.3)	0.03	–						
3. Sociable	4.7 (1.1)	0.27	0.53**	–					
4. Intelligent	4.8 (1.1)	0.11	0.27	0.37*	–				
5. Mutual like	4.9 (1.1)	0.00	0.54**	0.65**	0.11	–			
6. Trustworthy	6.3 (.85)	0.00	0.39**	0.40**	0.20	0.31	–		
7. Difficult	2.3 (1.1)	0.18	0.00	–0.27	0.13	0.07	–0.18	–	
8. Enjoyable	5.9 (.73)	0.02	0.24	0.18	0.46**	0.21	0.16	0.18	–

* $p < .05$; ** $p < .01$.

items; $\alpha = 0.61$); and trustworthiness (3 items; $\alpha = 0.83$). Also, we measured two attributes of the game: difficulty (3 items; $\alpha = 0.75$) and enjoyment (6 items; $\alpha = 0.83$). Table 1 shows the variable means and correlations.

3. Results

3.1. Preliminary analyses

Participants did not differ significantly across player conditions by age, native language, gender, computer experience, or video game use. Nor were any of these background variables significantly associated with overall level of cooperation. On the post-choice questionnaire, participants did not differ across conditions in their ratings of the difficulty of the game, nor how much they enjoyed it. Nor did they remember the choices of the other player differently.

3.2. Checks on the manipulation of other player's appearance

The confederate and the person-like computer player were not judged to differ from one another in human-likeness, while the computer person was judged to look significantly more human-like than the dog and dog cartoon computer agents (Table 2 gives mean ratings by condition). The contrast of 'person' versus 'computer person' was $t(91) = 1.09$, not significant. The contrast of 'computer person' versus 'dogs' was $t(91) = 4.248$, $p < 0.001$. Hence, the manipulation of the human-like appearance of the other player was successful.

The dog computer players were judged to look significantly more appealing than the human players, based on judgements of attractiveness and somewhat less so based on social charm, although these two were correlated at 0.53, $p < 0.001$. The confederate and the person-like computer player did not differ from one another in attractiveness, while the computer person was judged to look significantly less attractive than the dog and dog cartoon computer agents. The contrast of 'person' versus 'computer person' was $t(91) = -0.437$, not significant. The contrast of

Table 2
Mean participant attitudes toward partner, partnership, and game

	Person (<i>n</i> = 24)	Person-like agent (<i>n</i> = 24)	Dog-like agent (<i>n</i> = 24)	Cartoon dog agent (<i>n</i> = 24)
Appearance of partner				
Human-likeness	5.7	5.2	3.1	2.9
Attractiveness	3.8	3.9	5.1	4.8
Social evaluation	3.9	5.0	5.0	5.0
Intelligence	4.8	4.9	4.5	5.0
Sense of partnership				
Mutual liking	4.2	5.0	4.9	5.3
Trustworthiness	6.0	6.3	6.4	6.5
Game				
Difficulty	2.4	2.1	2.6	2.6
Enjoyment	5.8	5.7	6.0	6.0

'computer person' versus 'dogs' was $t(91) = -3.607$, $p < 0.001$. The confederate was judged to look less sociable than the person-like computer player. The contrast of 'person' versus 'computer person' for social charm was $t(91) = -3.807$, $p < 0.001$. The contrast of 'computer person' versus 'dogs' was not significant. These findings suggest that our manipulation of appealingness was somewhat successful, although not as strong as we would wish. There was no difference across condition in how intelligent the partner looked.

There were no statistically significant differences in judgments of appearance (human-likeness, attractiveness, social charm, intelligence) between the computer dog and the cartoon dog; moreover there were no differences in participants' cooperative responses to the computer dog and the cartoon dog. Therefore we combine the two dog conditions in the following analyses.

3.3. Differences associated with discussion variables

The manipulated variations in the other player's discussion script had little effect on participants' promises and cooperation. Overall, promises and cooperation were the same whether or not the other player proposed cooperation first, whether or not the other player asked the participant to confirm the agreement, and whether or not the other player used a mild threat of retaliation for defection. Since variations in the discussion did not affect the overall results, we do not discuss this aspect of the experiment further.

3.4. Cooperative promises and behaviors

Of the participants, 90% or more promised to cooperate (i.e. to give \$3.00 to the other player) on Trials 1 through 5; this percentage dropped to 80% on the 6th and

last trial. The repeated measures effect is significant, $F(5, 90) = 2.3$, $p < 0.05$). Although on each trial the large majority of participants said they would cooperate, 39% said they would compete (keep their \$3.00 rather than give it to the other player) at least once across the 6 trials, 1.5 trials on average. Proposing to compete is a very poor strategy in the dilemma game because it warns the other player to compete. Although a few participants did so once or twice to test the other player's intelligence or strategy, no participants used promises to compete as a consistent strategy.

Participants who promised to cooperate tended to keep their promises, although over all the trials there were significantly more total promises to cooperate than the total number of cooperative promises honored; within subjects, promises versus kept promises $F(1, 75) = 11.0$, $p < 0.01$.² Participants kept their promises to cooperate significantly more on early than later trials; repeated-measures interaction effect $F(5, 375) = 3.1$, $p < 0.01$. The percentage of participants who kept their promises to cooperate was above 80% on the first three trials, above 75% on the 4th and 5th trials, and declined to 51% on the final 6th trial. (In previous research, cooperators who know the total number of trials often defect on the last trial; Andreoni & Miller, 1993).

3.5. Effects of the other player

Consistent with Hypothesis 1 about the effect of improving the realism of the human-like interface agent, participants kept their promises significantly more with the human-like interface agent in this study than with the one in the previous study. In the previous study, on Trial 1 the mean cooperation rate with the confederate was 80%; the mean cooperation rate with the human-like agent was 32% (a difference of 48%). In this study, on Trial 1 the mean cooperation rate with the confederate was 83%; the mean cooperation rate with the human-like agent was 92% (a difference of 9%). (The data from Trial 1 are the most direct measure of participants' responses to the appearance of the other player; later choices would be affected by participants' own previous choices.) The difference between the human and human-like agent was statistically significant in the previous study; it was not statistically significant in this study. Across all trials on which there was discussion before the choice, participants were more likely to keep their promises to cooperate, $F(1, 149) = 12.8$, $p < 0.001$, and to cooperate overall, $F(1, 174) = 43.4$, $p < 0.001$, with the person-like player in the present study than with the person-like player in Kiesler et al. (1996). Thus, making technology improvements so that the human agent looked less artificial significantly improved cooperation with it and led to cooperation rates statistically indistinguishable from cooperation with a real human being.

Consistent with Hypothesis 2b about realistic human-like attributes and inconsistent with Hypothesis 2b about the appealingness of a computer character, on Trial 1, participants kept their promises significantly more with the confederate and

²The degrees of freedom are reduced because on some trials some participants did not make any commitment.

person-like computer player than with the dog-like computer players: confederate versus human-like player contrast $F(1, 87) = -0.38$, not significant; human-like agent versus dog agents contrast $F(1, 87) = 2.28$, $p < 0.05$. The results are similar if we just consider those who promised to cooperate: confederate versus human-like player contrast $F(1, 80) = -0.41$, not significant; human-like agent versus dog agents contrast $F(1, 80) = 2.38$, $p < 0.05$ (Fig. 3). On Trial 1, 92% of the participants in the person-like computer condition cooperated, which was 95% of those who promised cooperation. A total of 83% of the participants in the confederate condition cooperated with the confederate, which was 91% of those who promised cooperation. By contrast, only 58% of the participants in the dog-like computer condition cooperated, which was 67% of those who promised cooperation. In the cartoon dog computer condition, 67% of the participants cooperated, which was 76% of those who promised cooperation. In sum, on the first trial people mostly promised to cooperate, but more participants kept their promises to players who looked like people than to ones who looked like dogs. The fact that the only difference between the ‘computer person’ and the ‘dogs’ at this point was their appearance—in every respect they behaved identically—is strong evidence that the mere appearance of a computer character is sufficient to change its social influence.

Across all trials, as shown in Fig. 3, the results were consistent with those on the first trial. For all trials, keeping promises to cooperate differed by player condition— $F(3, 92) = 2.83$, $p < 0.05$)—with participants in the confederate and person-like computer player conditions keeping their promises to cooperate more than those in the dog and cartoon dog computer conditions did: confederate versus human-like player contrast $F(1, 92) = -0.73$, not significant; human-like agent versus dog agents contrast $F(1, 80) = 2.56$, $p < 0.05$. There was no interaction of trials by player condition.

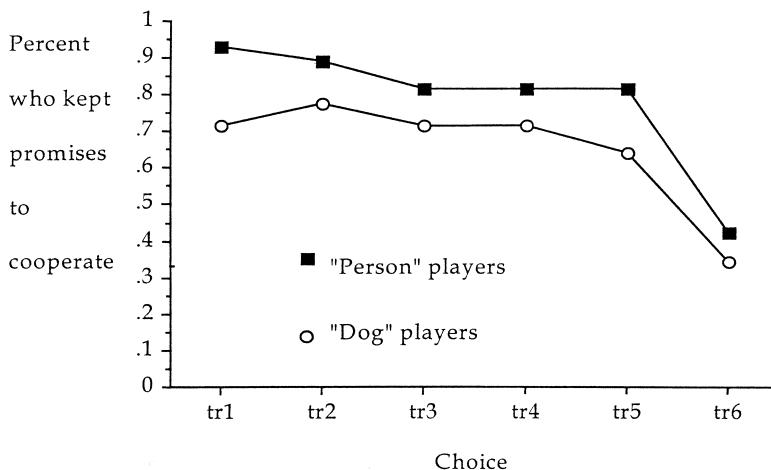


Fig. 3. Percent of participants who kept promise to cooperate.

3.6. *Attributions and evaluations of the other player*

We used the post-experiment survey results to explore why people would cooperate more with a more human-like but less appealing agent. As noted in Table 2, overall appealingness of the player was associated with sense of partnership. Participants reported feeling less of a sense of partnership with the human than with the agents (as measured by liking and trustworthiness), but no difference in partnership between the human-like agent and the dog-like agents. Yet neither likability or trust was associated with cooperation and keeping promises across conditions. (The mean trustworthiness rating was quite high across all conditions, probably resulting from the fact that the player always kept its promise—either to cooperate or to compete.) Ratings of partner intelligence were associated with cooperation, $r = 0.43$, $p < 0.01$, and keeping promises, $r = 0.45$, $p < 0.01$, across conditions but they did not vary systematically with partner conditions.

One possible explanation of our results is that the manipulation of human-likeness and dog-likeness of the other player differentially affected participants' certainty about their predictions of the other player's behavior. Keeping a promise to cooperate is based on confidence that the other player shares the same understanding of the situation and the same goals (Yamagushi, 1986; Yamagushi & Sato, 1986), including the expectation that the other player also will retaliate defection (Schopler et al., 1993). Confident predictions about the other player may depend on familiarity or similarity of the other player. In the current study, participants might have felt more able to make a prediction when the other player looked like a person than when the other player looked like a dog. Uncertainty about what the dog player might do could have reduced cooperation and increased defections.

To explore how familiarity with dogs may have affected cooperation and keeping promises to the other player, we examined data on participants' dog ownership from our post-choice survey (Table 3). We reasoned that dog owners would have more confidence in the predictability of dogs, and therefore would be more willing to risk cooperation with dog players. We also examined dog owners' and non-owners' responses to the questions, "How much do you like dogs?", "How knowledgeable are you about animal behavior?", and "How knowledgeable are you about human behavior?". As a proxy for uncertainty, we used the three-item scale on difficulty of playing the game (in preliminary analyses reported already, this scale did not vary by condition).

Dog ownership was approximately evenly distributed among conditions. Of the 48 participants in the person and person-like computer player conditions, 18 (38%) owned a dog. Of the 48 participants in the dog computer player conditions, 16 participants (33%) owned a dog. There were no gender or age differences in dog ownership. Dog owners liked dogs more than those who did not own dogs, $F(1, 95) = 10.3$, $p < 0.01$, and claimed more knowledge of animal behavior, $F(1, 95) = 14.6$, $p < 0.01$) but not more knowledge of human behavior, with no differences among player conditions. Because we were interested in dog owners' and non-owners' responses to the appearance of a 'person' player as compared with a 'dog' player, we combined the person and person-like computer player conditions

Table 3

Exploratory analysis of cooperation and perceptions of the other player among dog owners and people who do not own dogs

	Person and person-like agent		Dog-like agents	
	Dog owners (<i>n</i> = 18)	Non-owners (<i>n</i> = 30)	Dog owners (<i>n</i> = 16)	Non-owners (<i>n</i> = 32)
Trial 1				
% Who cooperated	89%	87%	81%	53%
% Who kept promise to cooperate	88%	96%	93%	61%
All trials				
% Trials cooperated	78%	69%	67%	52%
% Trials kept promise to cooperate	82%	77%	73%	59%
Game				
Difficulty	2.5	2.1	2.2	2.8
Enjoyability	5.7	5.8	5.9	6.2
Familiarity with dogs				
Liking dogs	6.1	4.8	6.8	5.8
Knowledge of animals	5.4	4.0	5.1	4.2
Knowledge of people	5.2	5.1	5.2	5.4
Sense of partnership				
Mutual liking	4.6	4.8	5.1	5.2
Trustworthiness	6.1	6.2	6.4	6.6
Appearance				
Human-likeness	5.6	5.3	3.0	3.0
Attractiveness	4.1	3.7	5.3	4.8
Social evaluation	4.4	4.5	4.8	5.1
Intelligence	4.9	4.8	4.6	4.8

and the dog-like computer player conditions to conduct 2×2 ANOVAs ('person' vs 'dog' conditions and dog owners vs non-owners) on the choice and evaluation data.³

Table 3 shows how cooperation and keeping promises to cooperate differed between dog owners and people who do not own dogs in the 'person' and 'dog' conditions. On the first trial, dog ownership interacted with player condition to influence whether participants kept their promise to cooperate, $F(3, 80) = 6.0$, $p < 0.05$, and this general pattern held throughout all trials. Dog owners who interacted with one of the computer dog players cooperated and kept their promises with these players statistically as much as the participants who interacted with either the person or the computer person player did; there were no statistically significant

³ We combined 'person' and 'dog' conditions to simplify the analysis of dog owner responses. Examined separately, the confederate and person computer player conditions did not differ; the two dog conditions did not differ.

differences between these participants and those who interacted with one of the 'person' players. By contrast, those who interacted with the dog players but did not own a dog defected approximately 15% more often; the difference between the dog owners and non-owners in the dog conditions for Trial 1, kept promises to cooperate was significant at $F(1, 39) = 7.5, p < 0.01$; for Trial 1, cooperation was significant at $F(1, 47) = 5.0, p < 0.05$; total cooperation was marginally significant at $F(1, 47) = 3.0, p = 0.08$, and total kept promises to cooperate was $F(1, 47) = 3.4, p = 0.06$.

Participants' ratings of the difficulty of the decision process bear on our argument that the dog owners had more confidence in their predictions of the dog players' cooperation. These ratings, shown in Table 3, showed a significant interaction, $F(1, 92) = 4.8, p = 0.03$, indicating that the participants who did not own a dog and interacted with a computer dog player had the hardest time making their decisions. Also in the 'dog' conditions, but not in the 'person' conditions, difficulty of the decision process was correlated negatively with overall cooperation ($r = -0.41, p < 0.01$) and keeping promises to cooperate ($r = -0.43, p < 0.01$).

Certain other correlational evidence is consistent with our explanation that dog owners were more confident of cooperation in the 'dog' players. In the 'person' conditions, those participants (both dog owners and non-owners) who rated themselves as more knowledgeable about human behavior were less likely to have cooperated and to have kept their promises to the person players ($r = -0.18, p < 0.01$; $r = -0.20, p < 0.01$). (Their ratings of their knowledge of animal behavior were unrelated to these choices.) This relationship suggests that, in the 'person' conditions, defectors justified their choices as warranted suspicion of the other player. By contrast, in the 'dog' conditions, participants' ratings of their knowledge of human behavior were unrelated to their choices but those who rated themselves as more knowledgeable about animal behavior (i.e. mainly the dog owners) were more likely to have cooperated ($r = 0.13, p < 0.05$) and to have kept their promises ($r = 0.16, p < 0.01$). And in the dog conditions, but not in the person conditions, trust in the partnership was correlated positively with overall cooperation ($r = 0.37, p < 0.01$) and keeping promises to cooperate ($r = 0.37, p < 0.01$). Hence, dog owners seem to have justified their cooperation on the grounds of their confidence in positive animal behavior. No other correlations (except the difficulty of making a decision, noted already) consistently differentiated the 'person' from the 'dog' conditions.

4. Discussion

Technology improvements increased the cooperation rate with a human-like agent in comparison with a previous less human-like agent. The participants in this study were likely to cooperate and to keep their promises to a computer character that looked like a person. They were not likely to cooperate and to keep their promises to a clearly unreal computer character that they judged to be appealing and likable. However, dog owners were likely to cooperate and keep their promises to a computer character that looked like a dog. Our findings demonstrate the importance of

a social framework of interaction. If participants were able to form a social expectation of the agent as a cooperator, then they could safely choose cooperation. We believe most participants felt able to take this risk when the computer character resembled a person. Only those used to routine interaction with a dog, and familiar with dogs, would have as easily formed a confident, positive social expectation of a dog.

This study has several limitations that need to be addressed in order to further develop a social interaction framework for understanding user interfaces. First, some readers might ask if participants' responses were due to their unfamiliarity with sociable interface agents. The study should be replicated using, perhaps, computer scientists as participants. However, an earlier study using computer scientists as participants did show strong evidence of social influence by a computer character (Sproull et al., 1994). Also Turkle (1984, 1994) and Nass and his colleagues (Nass & Steuer, 1993) suggest that familiarity with interactive computers, while improving one's understanding of how computers actually work, removes neither the tendency nor the desire to form coherent social impressions of computer objects or to interact with them as in a social context.

A second limitation of this study, inherent in studies of object perception, is that we could not measure directly what people 'really' thought the computer was doing. Questionnaire items about the behavior or attributes of objects are ambiguous because there are several possible levels of meaning in these items. Future research should be aimed at better subjective measurement of object attributions. One participant said spontaneously during the choice trials, "It's OK to deceive the dog, isn't it? He doesn't have any real feelings." (This comment resembles one currently seen on the Web: "Don't anthropomorphize your computers; they really hate that.") This comment reveals some of the ambiguity in interface agents and the internal conflict people may have about their appropriate responses to them. Despite these limitations, this study, taken with previous research, strongly suggests that people can form coherent social expectations of computer characters and act upon those expectations within a social interaction framework.

As we come to understand the importance and process of giving meaning to interface agents within a social interaction framework, we must also face the challenge of understanding the process of creating meaningful agents. There is a lot of excitement in the technology design world about virtual reality and new worlds in which people can build their digital homesteads and characters to populate them. Businesses will use interface agents and computer characters to advertise and sell their products and services. Developers and vendors will want to know how to build the most effective interface agents and characters. Our finding that technology improvements can create a computer agent that people cooperate with to a high degree should offer optimism to developers. Our finding that people cooperate less with more likable agents should serve as a caution and reminder that static assessments of character attributes (such as attractiveness, likability, or trustworthiness) are inappropriate and can be misleading when people encounter those characters in a social interaction framework. Interface designers have well-developed techniques for assessing interface usability. We must now begin to develop techniques for

assessing interface sociability—that is, the social expectations engendered when interfaces evoke a framework of social interaction.

Acknowledgements

This research was supported by a Boston University doctoral fellowship to the first author, a research grant from Bell Atlantic Mobile Systems to the second author, and research grants from the Digital Equipment Corporation and Sun Microsystems to the third author. The manuscript was prepared while the third author was a visiting scientist at the Lotus Development Corporation. An earlier version of this paper was presented at the CSCW '96 conference.

References

- Andreoni, J., & Miller, J.H. (1993). Rational cooperation in the finitely repeated prisoner's dilemma: experimental evidence. *The Economic Journal*, *103*, 570–585.
- Ball, E., Ling, D.T., Pugh, D., Skelly, T., Stankosky, & Thiel, D. (1994, April). ReActor: a system for real-time, reactive animations. *Conference companion: demonstrations*, CHI '94 Conference on Computer and Human Interaction, Boston, MA.
- Bloom, P. (1996). Possible individuals in language and cognition. *Current Directions in Psychological Science*, *5*, 90–94.
- Cioffi, D., & Garner, R. (1996). On doing the decision: effects of active versus passive choice on commitment and self-perception. *Personality and Social Psychology Bulletin*, *22*, 133–147.
- Dawes, R.M. (1980). Social dilemmas. *Annual Review of Psychology*, *31*, 169–193.
- Elliott, C. (1994). Components of two-way emotion communication between humans and computers using a broad, rudimentary, model of affect and personality. *Cognitive Studies: Bulletin of the Japanese Cognitive Science Society*, *1*, 16–30. (Japanese language with English abstract).
- Fais, L. (1995). Lexical accommodation in human-interpreted and machine interpreted dual language interactions. *ESCA Workshop on Spoken Dialog Systems*. Denmark.
- Fogg, B.J., & Nass, C.I. (1997). Silicon sycophants: the effects of computers that flatter. *International Journal of Human-Computer Studies*, *46*(5), 551–561.
- Friedman, B. (1997). *Designing computers for people: human values and the design of computer technology*. London: Cambridge University Press.
- Hayes-Roth, B., Sincoff, E., Brownston, L., Huard, R., & Lent, B. (1995). Directed improvisation with animated puppets. *Proceedings of CHI '95 Conference on Human-Computer Interaction*. Denver.
- Heider, F., & Simmel, M. (1944). An experimental study of apparent behavior. *American Journal of Psychology*, *57*, 243–259.
- Kerr, N.L., & Kaufman-Gilliland, C.M. (1994). Communication, commitment, and cooperation in social dilemmas. *Journal of Personality and Social Psychology*, *66*, 513–529.
- Kiesler, S., Sproull, L., & Waters, K. (1996). A 'prisoner's dilemma' experiment on cooperation with people and human-like computers. *Journal of Personality and Social Psychology*, *70*, 47–65.
- Marquis, S., & Elliott, C. (1994). Emotionally responsive poker playing agents. *Notes for the Twelfth National Conference on Artificial Intelligence (AAAI-94) Workshop on Artificial Intelligence, Artificial Life, and Entertainment, August, 1994* (pp. 11–15).
- Nass, C.I., & Steuer, J. (1993). Computers, voices, and sources of evaluation. *Human Communication Research*, *19*(4), 504–527.
- Sally, D. (1995). Conversation and cooperation in social dilemmas: a meta-analysis of experiments from 1958 to 1992. *Rationality and Society*, *7*, 58–92.

- Schopler, J., Insko, C.A., Graetz, K.A., Drigotas, S., Smith, V.A., & Dahl, K. (1993). Individual–group discontinuity—further evidence for mediation by fear and breed. *Personality and Social Psychology-Bulletin*, 19(4), 419–431.
- Sproull, L., Subramani, R., Kiesler, S., Walker, J., & Waters, K. (1996). When the interface is a face. *Human–Computer Interaction*, 11, 97–124.
- Sproull, L., Walker, J., & Subramani, R. (1994). Using a human face in an interface. *CHI '94 Proceedings*, 85–91.
- Turkle, S. (1984). *The second self: computers and the human spirit*. New York: Simon & Schuster.
- Turkle, S. (1994). Construction and reconstruction of self in virtual reality: playing in the MUDs. *Mind, Culture, and Activity*, 1, 158–167.
- Warner, R.M., & Sugarman, D.B. (1986). Attributions of personality based on physical appearance, speech, and handwriting. *Journal of Personality and Social Psychology*, 50, 792–799.
- Waters, K. (1987). A muscle model for animating three-dimensional facial expressions. *Computer Graphics*, 21, 17–24.
- Waters, K., & Levergood, T.M. (1995). DECface: a system for synthetic face applications. *Multimedia Tools and Applications*, 1, 1–16.
- Yamagushi, T. (1986). The structural goals/expectation theory of cooperation in social dilemmas. *Advances in Group Processes*, 3, 51–87.
- Yamagushi, T., & Sato, K. (1986). Motivational bases of the public goods problem. *Journal of Personality and Social Psychology*, 50, 67–73.
- Zebrowitz, L.A., Voinescu, L., & Collins, M.A. (1996). ‘Wide-eyed’ and ‘crooked-faced’: Determinants of perceived and real honesty across the life span. *Personality and Social Psychology Bulletin*, 22, 1258–1269.