

Media Inequality in Conversation: How People Behave Differently When Interacting with Computers and People

Nicole Shechtman¹
Center for Technology in Learning
SRI International
333 Ravenswood Ave.
Menlo Park, CA 94025
nicole.shechtman@sri.com

Leonard M. Horowitz
Department of Psychology
Stanford University
Stanford, CA 94305
len@psych.stanford.edu

ABSTRACT

How is interacting with computer programs different from interacting with people? One answer in the literature is that these two types of interactions are similar. The present study challenges this perspective with a laboratory experiment grounded in the principles of Interpersonal Theory, a psychological approach to interpersonal dynamics. Participants had a text-based, structured conversation with a computer that gave scripted conversational responses. The main manipulation was whether participants were told that they were interacting with a computer program or a person in the room next door. Discourse analyses revealed a key difference in participants' behavior – when participants believed they were talking to a person, they showed many more of the kinds of behaviors associated with establishing the interpersonal nature of a relationship. This finding has important implications for the design of technologies intended to take on social roles or characteristics.

Keywords

Social interfaces, SRCT (social reactions to communication technology), Media Equation, personality, human-human interaction

INTRODUCTION

More and more, designers are building technologies intended to behave in ways that are social or take on roles that previously could only be performed by human beings. In order to create social technologies that are sound, effective, and appropriate, designers must have a basic understanding of human-human interaction and how this is similar and different from human-computer interaction. A strongly represented perspective in the literature [e.g., 20, 14] is that human-computer and human-human interaction are similar. The present study challenges this notion by demonstrating crucial differences in how people behave when conversing with computer programs and (what they

believe to be) people. These findings have important implications for user experience and behavior, as well as how designers should conceptualize and approach creating social interfaces.

The Media Equation [20] sums up the “social reactions to communication technology” perspective (SRCT; [13]): “Media equals real life. In short, we have found that individuals’ interactions with computers, television, and new media are fundamentally social and natural, just like interactions in real life.” In other words, these authors claim that people react socially to computers as they react to people. The underlying mechanism that they propose is that people respond “mindlessly” to social cues, no matter whether they come from other people or media behaving like people [14]. Their method for supporting this has been to take a robust finding from social psychology, replace a human actor with a computer actor, rerun the experiment, and show that the results are similar. Some examples of social psychological constructs they report targeting with success are politeness [17], in-group membership [15, 16], self-disclosure of personal information [12], and enjoyment of humor [13].

Our approach to this issue differs both theoretically and methodologically. Theoretically, our perspective is informed by a psychological framework called Interpersonal Theory (see discussion below), a study of the behavioral, cognitive, emotional, and motivational processes that occur between people during interpersonal interaction. We draw on some basic principles shown to underlie the mechanisms of human-human interaction and explore how these may or may not manifest in human-computer interaction. We believe that this approach allows for a richer and deeper understanding of the processes in question. Methodologically, while most previous SRCT studies have relied on self-report and nonconversational behaviors, our findings are grounded in discourse analyses of conversations.

What follows is a discussion of three key principles of Interpersonal Theory, a description of the experiment and findings, implications of the findings for design, and a discussion of questions for future research.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CHI 2003, April 5-10, 2003, Ft. Lauderdale, Florida, USA.
Copyright 2003 ACM 1-58113-630-7/03/0004...\$5.00

¹This study was done at Stanford University as part of this author's doctoral dissertation.

Interpersonal Theory

Interpersonal Theory [e.g., 3, 7, 8, 11, 21] is a psychological approach to interpersonal dynamics. The fundamental unit of analysis is the “interaction unit.” If Person A and Person B are in conversation, an interaction unit is one action taken by A and B’s subsequent reaction. Interpersonal Theory has produced thousands of quantitative and qualitative empirical studies, as well as perspectives on personality development, psychotherapy, and psychopathology. We focus on three basic principles: (1) behavior in conversation is driven by different types of goals; (2) there are two broad categories of “relationship goals”; and (3) individuals differ in the degree to which different relationship goals are important to them. Let us now discuss each principle in turn.

Principle #1: Three Tracks of Conversational Goals

Like all human behavior, behavior in conversation is driven by goals. In the complex process of human-human conversation, individuals generally have multiple goals, some conscious, some unconscious, some public, some private, some shared, and some unshared. Many authors [e.g., 4, 6] suggest that each of the different goals that people can have in a conversation fall into three main categories. First, there are *task goals*. These are the goals relevant to the task people have come together to accomplish, the purpose of the activity they are jointly involved in, or a plan they are jointly evolving. A second type of goal is *communication goals*. These are the goals aimed at making sure that the communication itself goes smoothly and everyone understands each other. A third type of goal is *relationship goals*. These are the goals that drive people to set and maintain the tone of the conversation or relationship, how much the interaction may be friendly, polite, hostile, reciprocal, conflicted, professional, intimate, formal, informal, and so on. Each of these types of goals – task, communication, and relationship – contributes to determining what will happen over the course of a conversation.

To simplify this discussion, we refer to these types of goals in terms of a metaphor of information on parallel tracks on a tape (see Figure 1). On the task track is all the behavior pertaining to the task at hand, on the communication track is all the behavior pertaining to clear communication, and on the relationship track is all the behavior pertaining to the nature of the relationship.




Figure 1. Three tracks of conversational goals

Principle #2: Two Broad Categories of Relationship Goals

Interpersonal theorists have found in more than fifty studies [e.g., 1, 10, 11] that two dimensions are prominent among

behaviors on the relationship track. The first dimension, generally labeled “communion,” describes behaviors oriented toward connecting with another (e.g., smiling, listening, sharing) or becoming disconnected from another (e.g., ignoring, turning away from, avoiding). The second dimension, generally labeled “agency,” describes behaviors oriented toward influence. At one end, the behaviors may be about exerting influence (e.g., dominating, bossing around, asserting an opinion), while at the other end, the behaviors may be about yielding to influence (e.g., agreeing, giving in, asking for help). These two axes correspond to two fundamental motives that people have with each other in all types of relationships – establishing degree of connectedness and exercising influence [2, 7]. The space formed by these dimensions is illustrated with a two-dimensional figure (see Figure 2).

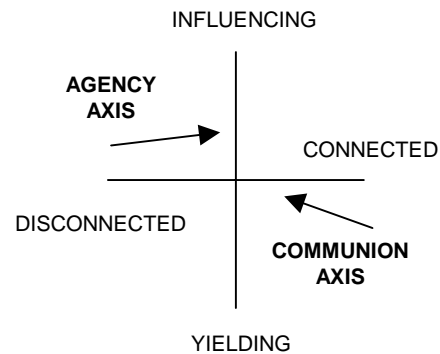


Figure 2. The relationship space

Principle #3: Individual Differences in Emphasis of Relationship Goals

While many factors affect what happens on the relationship track in a given conversation, one important determinant is individuals’ preferences. Specifically, a highly assertive individual may be particularly invested in being influential and may prefer not to yield. Less assertive individuals, in contrast, may be satisfied with allowing someone else to take the lead. As [5, 19] show, these individual differences can significantly affect what occurs on the relationship track in a given interaction.

The Relationship Track in Human-Computer Interaction

Let us now consider what happens on these tracks during human-computer interaction. We assume that the task and communication tracks should be full of activity. People typically use computers as tools to *do* things, and clear communication between people and their tools is essential. The specifics of these tracks, however, are left to future research.

What might happen on the relationship track? The SRCT perspective, arguing that people react to computers as they react to people, would predict that this track would be filled with the same activity in human-computer interaction as it is in human-human interaction. However, as the SRCT perspective has never specifically investigated conversational processes, the question is still open.

The present study focuses on the relationship track. It investigates “communal” and “agentic” behavior during human-computer and (apparently) human-human interaction.

PRESENT STUDY

Overview and Design

This experiment compared how people behave when they believe they are interacting with a computer program versus a person. The paradigm juxtaposed those in previous studies of human-computer interaction [13, 18] and human-human interaction [5, 19]. Participants had a text-based discussion with “a partner” they believed to be either a computer program or a person. The discussion was highly structured with scripted responses crafted to appear conversational.

This experiment had a 2x2x2 factorial design. The three factors were:

1. *Belief about Partner.* This was a cognitive manipulation of the participant’s belief about who or what the conversation “partner” was. Those in the “apparently-computer” condition were told their partner was a computer program, while those in the “apparently-human” condition were told their partner was another student in the room next door. In reality, participants in both conditions received identical scripted responses via their interface. Aside from this cognitive manipulation of framing, participants in the apparently-computer and apparently-human conditions were treated exactly the same.
2. *Participant Assertiveness.* This individual difference variable enabled a comparison of behavior of assertive and nonassertive individuals on the relationship track.
3. *Partner Assertiveness.* This manipulation of partner behavior enabled a comparison of reactions to assertive and nonassertive behavior.

The primary dependent measures were discourse analyses of participants’ conversational behavior.

Method

Participants

The participants were 130 (64 female, 66 male) undergraduate students receiving class credit or pay for their participation. Participants were preselected using a questionnaire on which students rated themselves on traits connoting assertiveness (e.g., aggressive, assertive, competitive, dominant). The 68 assertive participants scored 0.5 standard deviations above the mean (or higher), and the 64 nonassertive participants scored 0.5 standard deviations below the mean (or lower). There were no significant gender differences on assertiveness.

All 130 participants were randomly assigned to Belief about Partner and Partner Assertiveness conditions.

Participants in both the apparently-computer and apparently-human conditions were run in same-gender pairs

of strangers. While co-participants in the apparently-computer condition were told that they were working independently with their own computers, co-participants in the apparently-human condition were deceived into believing they were interacting with each other. Since co-participants never truly interacted with each other, pairing by level of assertiveness was irrelevant.

Procedure

As in previous studies [5, 18], the interaction was structured around the Desert Survival Problem (DSP; [9]). In the DSP, participants are asked to imagine a scenario in which a plane crash has stranded them in the desert and then to rank twelve salvaged items (e.g., a flashlight, a liter of water, a mirror) according to their importance for survival. Participants then exchange their rationales for these rankings and independently make a final ranking.

At the beginning of the session, co-participants, with minimal introduction to one another, were seated back-to-back to fill out consent forms. Next, the experimenter read aloud the DSP scenario and left the room so the participants could independently formulate initial rankings. The experimenter then returned and read aloud instructions for the discussion. These instructions introduced the Belief about Partner manipulation (see details below).

Co-participants were then taken to separate rooms and left alone for the actual discussion. This controlled the situation such that the simple presence of another human being in the room could not confound the results in either the apparently-computer or apparently-human condition.

Each room was equipped with a single PC running a Java applet interface (see Figure 3) that gave the illusion of a web-based conversation. The participant began by entering his or her rankings. The interface then displayed the participant’s rankings along with those of the “partner.” In reality, the partner rankings were a systematic transformation of the participant’s rankings.

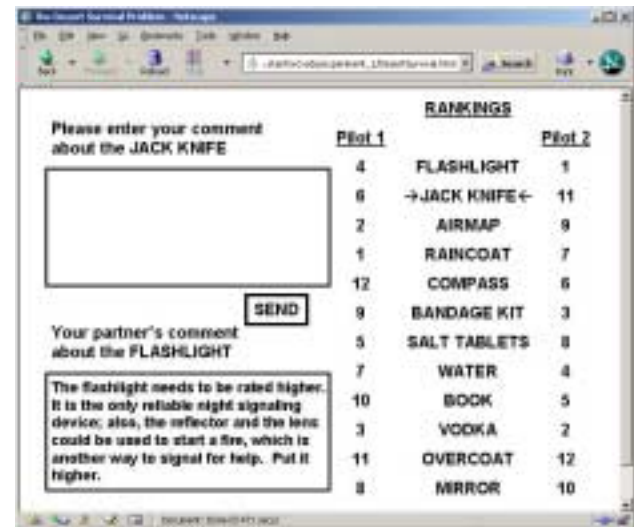


Figure 3. The DSP discussion interface

The interface structured the conversation such that there were twelve sequential exchanges, one for each item. Each exchange had the following sequence: (1) the participant typed thoughts about the item or item ranking in the top box and pressed “send” when finished, (2) there was a delay which helped to maintain the illusion that a partner was taking the time to read and respond (in the apparently-computer condition, participants reported in debriefing that they thought this delay was due to technical issues), (3) a response comment appeared in the lower box.

Response comments were chosen from a lookup table of scripted responses crafted to sound conversational (see examples below). This table was used previously by [19] and based closely on that used by [18]. A comment was chosen from the table on the basis of whether the partner ranking for the current item was higher or lower than the participant’s. The content of the comment addressed the discrepancy in rankings. This facilitated the illusion that the comment was responsive to the participant’s statement.

After the discussion, participants made a final ranking of the items and filled out a paper-and-pencil self-report questionnaire. Participants were debriefed and probed thoroughly for suspicions.

Manipulation of Belief about Partner

Belief about Partner was manipulated during the instructions for the discussion. The instructions began, “The goal of the next phase of this study is to use the computer to improve your rankings. In order to do this, I would like you to discuss your rankings with [PARTNER] via the network.” In the apparently-computer condition, [PARTNER] was filled in with, “a computer program.” In the apparently-human condition, [PARTNER] was filled in with, “each other.” Participants in both conditions were otherwise treated identically.

Manipulation of Partner Assertiveness

While the desert survival content of the comments was held constant across conditions, Partner Assertiveness was manipulated by controlling the phrasing of the scripted responses. In the assertive condition, the comments were crafted to seem commanding, leading, and dominating. In the nonassertive condition, the comments were crafted to seem polite and deferential. For example, suggesting the flashlight should be rated higher, the scripts in the two conditions read:

Assertive: The flashlight needs to be rated higher. It is the only reliable night signaling device; also, the reflector and the lens could be used to start a fire, which is another way to signal for help. Put it higher.

Nonassertive: Do you think the flashlight should maybe be rated higher? It may be a pretty reliable night signaling device. Also, maybe the reflector and lens could be used to start a fire, which could possibly be another way to signal for help.

Measures

The primary dependent measures were discourse analyses of participants’ conversational behavior (see next section for detail). Also, two self-report measures examined how assertive the partner seemed (Perceived Partner-Assertiveness) and how expert the partner seemed (Partner-Expertise).

Results and Discussion

Manipulation Checks and Expertise

Two manipulations were checked for their effectiveness. The first was whether participants believed their Belief about Partner instructions. Of 142 initial participants, 2 were excluded from the apparently-computer sample because they suspected they *were* talking to a person, and 10 were excluded from the apparently-human sample because they suspected they were *not* talking to a person. None of the 130 participants in the final sample expressed suspicion about the framing they had been given.

The second manipulation check was to verify that the assertive partner was perceived as more assertive than the nonassertive partner. On a 7-point scale, assertive and nonassertive partners received mean ratings of 6.6 (SD=1.64) and 4.0 (SD=1.77) respectively on Perceived Partner-Assertiveness. This difference was significant (F(1,122)=68.7, p<0.001) and similar in both the apparently-computer and apparently-human conditions.

Participants perceived apparently-computer and apparently-human partners as equally expert in the task. A 2x2x2 ANOVA of Partner-Expertise showed no significant difference across groups.

Analyses of Conversational Behavior

Broad Differences Across Belief about Partner Conditions

The first measure was general effort, how much work and time participants put into the conversations. Overall, participants used more words and spent more time writing comments to an apparently-human partner (see Table 1).

Table 1. Effort in conversation

	Apparently- Computer	Apparently- Human
Words per comment	27.9 (SD=14.5)	34.9 ^a (SD=10.7)
Seconds per comment	71.8 (SD=27.4)	89.0 ^b (SD=26.0)
^a (F(1,121)=9.6, p<0.01); ^b (F(1,121)=13.2, p<0.001)		

A next step was to determine whether there were general salient discourse differences that might be evident to third-party observers. Two female coders, blind to condition and uninformed about the differences in effort, read the transcripts from each conversation. Without receiving any a priori criteria, they were told to make a dichotomous guess for each participant about whether they seemed to be addressing an apparently-human or apparently-computer

partner. Overall, the coders were highly accurate in their judgments – 78.7% correct, well above chance. This suggests that there were important behavioral differences across conditions.

These two global findings – that people put more effort into their conversations with an apparently-human partner and that differences in behavior were detectable by naïve observers – suggested that people were behaving differently in some way depending on whether they believed they were interacting with a computer program or a person.

The Relationship Track

In casual preliminary observations, the task and communication tracks seemed to contain similar activity in the apparently-computer and apparently-human conditions – people simply communicated about the task at hand by explaining their reasons for each of their rankings. We surmised that the observed differences in effort and salient behaviors could be accounted for by the relationship track, so we focused on behavioral differences on this track.

Coding Scheme. A coding scheme was created to assess behavior on the relationship track. The coders began by making a list of all the types of behaviors they thought had informed their dichotomous guesses. We focused on those statement types that could be mapped to relationship space (see Figure 2). In total, eight categories of “relationship statements” were created that were considered connecting, influencing, and yielding. No behaviors were found to be explicitly disconnecting. We also created three additional categories that could be classified as hostile behaviors. Table 2 lists each of these categories along with examples and reliabilities.

The two coders, still blind to participant condition, recorded how many times each type of statement occurred in each transcript. Note that any one statement could meet criteria for one or more category and could thus be coded in multiple categories. Reliability, measured with a Cronbach’s alpha of the coders’ independent codings, satisfactorily ranged from 0.78 to 0.99 (see Table 2). For the data used in analysis, all disputes were settled by a group discussion with a third coder also blind to participant condition.

Relationship Statements by Factor. Each of the three factors was then examined to determine how it affected the participant’s relationship behavior.

Belief about Partner. This manipulation caused substantial differences in relationship behavior. In the apparently-human condition, the mean number of total relationship and hostile statements was 11.9 (SD=9.8), over four times more than the mean of 2.8 (SD=3.5) found in the apparently-computer condition ($t(127)=6.4$, $p<0.0001$). Figure 4 shows that participants in the apparently-human condition used significantly more connecting ($F(1,121)=42.3$, $p<0.0001$), influencing ($F(1,121)=15.0$, $p<0.0001$), and yielding ($F(1,121)=16.8$, $p<0.0001$) statements. This effect

Table 2: Relationship Statements and Hostility – categories, reliabilities, and examples

Connecting Relationship Statements

Positive feedback or compliment (alpha=0.92)

Excellent.

Thanks, you’ve boosted my confidence a tad.

Direct reference to agreement (alpha=0.99)

Agreed.

We both agree that the compass is an important item.

Reference to partner (alpha=0.89)

Well, I definitely would be thankful to have you by my side in this situation.

What do you think?

Self-disclosure (alpha=0.81)

I’m a vegetarian.

So I’m pretty much an idiot.

Influencing Relationship Statements

Giving advice (alpha=0.89)

Should be rated higher.

You should rank the water above the flashlight.

Direct command (alpha=0.81)

Rate it a little higher.

Give up on the food idea and stick to water.

Yielding Relationship Statements

Taking advice (alpha=0.91)

I should have put it a little higher, I know.

But I also understand the value of this book and am willing to be swayed by your argument.

Questions deferring to the partner (alpha=0.94)

How cold does this desert get at night?

Why did you pick this one fourth? I’m curious to hear your take.

Hostile Statements

Annoyance or hostility (alpha=0.91)

You rated the overcoat lower than water, come on.

It’s a frickin’ book. Your water argument can get flipped back on you.

Sarcasm (alpha=0.78)

I am not really sure what the salt tablets are for, so maybe you can enlighten me.

I have NO idea what you’re thinking, so just fill me in.

Negative feedback (alpha=0.89)

The bandage kit isn’t half as important as you make it.

You’re wrong.

was marginal for hostile statements ($F(1,121)=3.4$, $p=0.065$).

This shows that there was much more activity on the relationship track when participants believed they were interacting with a person. They were more likely to try to




Figure 4. Mean (\pm SE) number of relationship statements addressed to apparently-computer and apparently-human partners

connect, influence, yield to, and even become hostile toward their partners – everything it takes to negotiate the nature of an interpersonal relationship. When people believed they were interacting with a computer program they exhibited a factor of four fewer relationship behaviors.

Participant Assertiveness. This factor revealed important differences in relationship behavior. As Figure 5 shows, assertive participants used more of each type of relationship statement – but only in the apparently-human condition. In the apparently-human condition, the number of interpersonal statements was significantly greater for assertive participants (connecting: $F(1,60)=7.2, p<0.01$; influencing: $F(1,20)=6.1, p<0.05$; yielding: $F(1,60)=4.1, p<0.05$). In the apparently-computer condition, differences in behavior for assertive and nonassertive participants were not significant (connecting: $F(1,61)=0.7, p=0.4, n.s.$; influencing: $F(1,61)=0.0, p=1.0, n.s.$; yielding: $F(1,61)=1.7, p=0.2, n.s.$). Figure 5 shows the significant two-way interaction ($F(1,125)=7.5, p<0.01$) given the sum of all three of these types of statements.

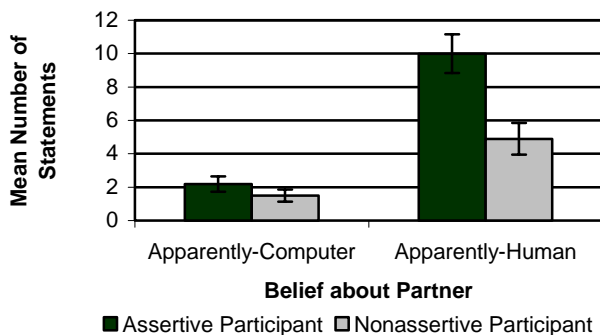


Figure 5. Mean (\pm SE) number of relationship statements by Participant Assertiveness and Belief about Partner. The dependent variable is the total of connecting, influencing, and yielding statements.

This suggests that assertive participants with an apparently-human partner were quite active on the relationship track. They tried to influence their partners using a full range of interpersonal behaviors – advice, flattery, connection, and

yielding. With apparently-computer partners, the very same kind of assertive individual did not become particularly engaged in the relationship track compared to a nonassertive participant.

Partner Assertiveness. This manipulation also caused important differences in relationship behavior. While there were no differences in connecting or yielding statements, there were marked differences in influencing statements. As Figure 6 shows, participants reacted to assertive behavior with influencing behavior in return – but only with an apparently-human partner. In the apparently-human condition, the number of influencing statements was significantly greater with an assertive partner ($F(1,60)=6.1, p<0.05$). Participants with an apparently-computer partner, however, did not react significantly differently to an assertive or nonassertive partner ($F(1,61)=0.9, p=0.3, n.s.$). The two-way interaction was significant ($F(1,125)=5.1, p<0.05$).




Figure 6. Mean (\pm SE) number of influencing statements by Partner Assertiveness and Belief about Partner

This suggests that participants responded on the relationship track when assertive behavior came from an apparently-human partner. The partner behaved in an assertive and dominating manner, so the participant took a cue to behave similarly in return. When the behavior came from an apparently-computer partner, however, the participant did not respond with similar relationship behavior.

Hostile Statements. Hostile behavior only occurred substantially in one condition. While participants in most conditions used means of fewer than one hostile statement, assertive participants with the apparently-human assertive partner used a mean of nearly three hostile statements. Figure 7 illustrates this significant three-way interaction ($F(1,121)=4.4, p<0.05$).

Thus assertive participants reacted with considerable hostility when the partner was assertive – but only when they thought the partner was a person. What we observe on this relationship track is an outright power struggle. Assertive individuals do not like to be pushed around, and they lashed back with annoyance, sarcasm, and negative feedback when the partner was dominating. Interestingly, when the very same behavior came from an apparently-

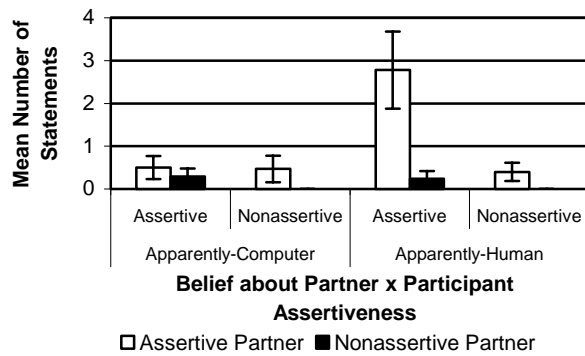


Figure 7. Mean (\pm SE) number of hostile statements across all conditions

computer partner, the assertive participant was not drawn into a power struggle at all. Just as calm as most other groups, these participants were not particularly engaged on the relationship track.

GENERAL DISCUSSION

Summary of Findings

The goal of this study was to explore behavior on the relationship track during human-computer and (apparently) human-human interaction.

There were six key findings in the discourse analyses: (1) Participants used more words and spent more time in conversation when they believed the partner was a person; (2) Judges found it easy to distinguish discourse from apparently-computer and apparently-human conversations, and their intuitions were that the differences were on the relationship track; (3) Participants with an apparently-human partner used over four times more relationship statements – connecting, influencing, yielding, and hostile; (4) Assertive participants became more engaged on the relationship track, but only when they believed their partners were human; (5) When the partner’s scripted behavior was assertive, participants reacted with influencing behavior in return, but only when they believed the partner was human; (6) The partner’s scripted assertive behavior drew assertive participants into a hostile power struggle, but only when participants thought the partner was human.

In sum, this evidence suggests a much greater engagement on the relationship track for those who believed their partners were human compared to those who believed they were interacting with a computer program.

Implications for Design

These findings have several implications for design:

- The Media Equation should not automatically be applied to design. While people do exhibit some kinds of social reactions toward computers, clearly, there are some important differences as well. Research is still necessary to address the psychological mechanisms that might or might not lead people to react socially to technology.

- Sometimes there may be benefits to having human-computer interaction replace human-human interaction. Efficiency is one consideration. In this study, people put more time and effort into interactions with “people” than computers. Such behavior could be a hindrance to task goals, especially when time or efficiency is important. Emotion is another consideration. While relationships can bring about the most positive of emotions, they can also bring about the most negative, as demonstrated by the hostile assertive participants. When negative emotions such as anger, fear, anxiety, shame, or embarrassment could be the consequence of an interaction between human beings, a technological intervention might be more expedient.
- When designing an interface, it may be helpful to consider the user’s goals with respect to each of the three tracks of conversation – task, communication, and relationship. Sorting goals in this way may help to facilitate decisions about how to invest design resources (e.g., on social interfaces that can never be fully human or on making a solid task-oriented tool).
- Users’ beliefs or mental models about an interface have a critical impact on how they experience them or behave with them. The apparently-computer and apparently-human participants interacted with the very same interface – merely their beliefs about what they were doing made the difference.

Questions for Future Research: Mechanisms and Elicitors of Mechanisms

These results raise two fundamental theoretical questions. The first question is: why did participants exhibit so many more relationship behaviors when they believed the partner was a human? The second question is: what features might an interface need in order to evoke these kinds of reactions from a user? These questions are two sides of the same theoretical coin, intertwining issues in both human-computer interaction and psychology. We begin to explore these questions by considering the various mechanisms that might elicit relationship behavior.

The mechanism we have discussed here in greatest detail is conscious or unconscious relationship goals as understood within the framework of Interpersonal Theory. Perhaps relationship behavior comes about because the actor is motivated to establish the quality of the interpersonal relationship. In this study, we observed behaviors that indicated caring about the other’s feelings or impressions, a desire to share, a desire to influence, a desire to acknowledge, and even a desire to hurt. People use these types of behaviors to communicate to one another about how they want a relationship to operate, how strong they want the rapport to be and who they want to have the greater influence. This allows people to build functioning human relationships in which work can be accomplished.

There are other mechanisms that could evoke these types of behaviors. One such mechanism, in line with the SRCT

school of thought [e.g., 14], is a “mindless” habit of responding interpersonally when people receive social cues such as the conversational ones in this experiment. Another such mechanism is the possible inextricable link between the use of natural language and social interaction [4]. Perhaps relationship behaviors are simply difficult to filter out of communication and may arise as an artifact of using natural language in a conversational situation, no matter who the audience might be. Another potential mechanism is engagement in fantasy or play [4]. Relationship behaviors may arise out of a pretense that the person is interacting with a human being. Such a pretense may be for entertainment or to conform to what the actor believes to be the demands of the situation.

One possible explanation of the present data is that these latter mechanisms elicited some relationship behavior in both the apparently-computer and apparently-human conditions, while the motive of interpersonal relationship building brought about the (four times) greater frequency of relationship behavior in the apparently-human condition. However, this study alone does not provide enough information to verify this inference, it simply points the way for more work. Future research must seek to understand each of these mechanisms more deeply and how and when they influence individuals in both human-computer and human-human interaction.

It is also important to investigate how changes in the specific circumstances of these interactions might affect behavior. Many characteristics of a human-human interaction (e.g., individual preferences, extent of acquaintance, work versus personal setting, work versus play) should change the degree and quality of relationship behavior. Likewise, many characteristics of a human-computer interaction should affect relationship behavior. Future research should address what features of an interface (e.g., natural language processing, facial expressions, voice, physical movement, the story the user is told about what the interface is) or characteristics of a user (e.g., age, experience with computers, mental model of the technology, working versus playing with the technology) might allow an interface to effect interpersonal mechanisms and evoke different patterns of relationship behaviors.

ACKNOWLEDGEMENTS

The authors would like to thank the CHI reviewers and the many colleagues at Stanford and SRI who gave valuable feedback on earlier versions of this paper.

REFERENCES

- Berzins, J. I. (1977). Therapist-patient matching. In A. S. Gurman (Ed.), *Effective psychotherapy* (pp. 221-251). New York: Pergamon.
- Blatt, S. J. (1990). Interpersonal relatedness and self-definition: Two personality configurations and their implications for psychopathology and psychotherapy. In J. L. Singer (Ed.), *Repression and dissociation*. Chicago: University of Chicago Press.
- Carson, R. (1969). *Interaction concepts of personality*. Chicago, IL: Aldine Publishing Company.
- Clark, H. (1996). *Using language*. Cambridge: Cambridge University Press.
- Dryer, D. C., & Horowitz, L. M. (1997). When do opposites attract? Interpersonal complementarity versus similarity. *Journal of Personality & Social Psychology*, 72, 592-603.
- Hobbs, J. R., & Evans, D. A. (1980). Conversation as planned behavior. *Cognitive Science*, 4, 349-377.
- Horowitz, L. (in press). *An interpersonal approach to psychopathology*. Washington D.C.: APA Books.
- Kiesler, D. (1996). *Contemporary interpersonal theory and research*. New York: Wiley.
- Lafferty, J., & Eady, P. (1974). *The desert survival problem*. Plymouth, MI: Experimental Learning Methods.
- LaForge, R., & Suczek, R. F. (1955). The interpersonal domain of personality: An interpersonal checklist. *Journal of Personality*, 24.
- Leary, T. (1957). *Interpersonal diagnosis of personality*. New York: Ronald Press Co.
- Moon, Y. (2000). Intimate exchanges: Using computers to elicit self-disclosure from consumers. *Journal of Consumer Research*.
- Morkes, J., Kernal, H. K., & Nass, C. (1999). Effects of humor in task-oriented human-computer interactions and computer-mediated communication: A direct test of SRCT theory. *Human-Computer Interaction*, 14, 395-435.
- Nass, C., & Moon, Y. (2000). Machines and mindlessness: Social responses to computers. *Journal of Social Issues*, 56, 81-103.
- Nass, C., Isbister, K., & Lee, E-J. (2000). Truth is beauty: Research conversational agents. In J. Castells (Ed.), *Embodied conversational agents*. Cambridge, MA: MIT Press.
- Nass, C., Fogg, B. J., & Moon, Y. (1996). Can computers be teammates? *International Journal of Human-Computer Studies*, 45, 669-678.
- Nass, C., Moon, Y., & Carney, P. (1999). Are people polite to computers? Responses to computer-based interviewing systems. *Journal of Applied Social Psychology*, 29, 1093-1110.
- Nass, C., Moon, Y., Fogg, B. J., Reeves, B., & Dryer, D. C. (1995). Can computer personalities be human personalities? *International Journal of Human-Computer Studies*, 43, 223-239.
- Person, E. (2000). *The effect of polite language on the assertive person's reaction to advice*. Stanford: Stanford University.
- Reeves, B., & Nass, C. (1996). *The media equation: how people treat computers, televisions, and new media like real people and places*. New York: Cambridge University Press.
- Sullivan, H. S. (1953). *The interpersonal theory of psychiatry*. New York: Norton.