

Helper Agent: Designing an Assistant for Human-Human Interaction in a Virtual Meeting Space

Katherine Isbister¹, Hideyuki Nakanishi², Toru Ishida³, and Cliff Nass⁴
Kyoto University - NTT - Stanford University
Cross-Cultural Digital Environments Project

ABSTRACT

This paper introduces a new application area for agents in the computer interface: the support of human-human interaction. We discuss an interface agent prototype that is designed to support human-human communication in virtual environments. The prototype interacts with users strategically during conversation, spending most of its time listening. The prototype mimics a party host, trying to find a safe common topic for guests whose conversation has lagged. We performed an experimental evaluation of the prototype's ability to assist in cross-cultural conversations. We designed the prototype to introduce safe or unsafe topics to conversation pairs, through a series of questions and suggestions. The agent made positive contributions to participants' experience of the conversation, influenced their perception of each other and of each others' national group, and even seemed to effect their style of behavior. We discuss the implications of our research for the design of social agents to support human-human interaction.

Keywords

Social interface agents, human-human interaction, virtual meeting place, cross-cultural communication.

INTRODUCTION

Communication contexts are becoming an ever more prominent part of the computer interface. People spend a great deal of their computer time communicating with one another. They use a range of tools, from email, to chat, to elaborate 3-D meeting spaces (e.g. Mitsubishi Research Lab's Diamond Park [14], Electric Community's the Palace, and NTT's InterSpace [13]). People are not just talking to those they already know—virtual meeting places make casual meetings between strangers from across town,

or even across the world, easy. It is increasingly common for people to meet for the first time through a computer interface. This is a wonderful opportunity to build new social networks based more on common interests than on location, but it is also a tremendous challenge. Virtual meeting places usually provide very little socially meaningful context to use as a basis for finding common ground with each another. Because it is easy to arrive at a virtual meeting place from many entry points, it is often hard for visitors to assume much about one another's cultural backgrounds, group memberships, and other aspects of social identity. Psychologists have demonstrated that people need this sort of common context in order to build new relationships [2]. Some commercial chat rooms make use of human moderators to help fulfill this need. However, human moderators are a scarce resource.

We believe this is an appropriate new application domain for social interface agents. Social interface agents could provide ongoing, in-context help in forming social connections and building common ground between visitors to virtual environments. Our project is a first step toward exploring this new application space for interface agents.

Social Interface Agents—Related Work

HCI researchers have already discussed and demonstrated some benefits of interface agents in one-on-one task settings, such as taking an educational tutorial [8], going on a tour [6], or looking at real estate [1]. Although the CHI community has a range of opinions about when and where agents should be deployed in the interface, most would agree that there are some beneficial applications. Lester notes that the presence of an agent can lead to “a strong positive effect on students' perception of their learning experience” [p. 364, 8]. Cassell discusses the value of an

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, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CHI '2000 The Hague, Amsterdam
Copyright ACM 2000 1-58113-216-6/00/04...\$5.00

¹ NTT Open Lab, 2-4 Hikaridai, Seika-cho, Soraku-gun, Kyoto 619-0237 JAPAN, +81 774 93-1946, katherine@warui.com

² Department of Social Informatics, Kyoto University, Kyoto 606-8501, JAPAN, +81 75 753-5396, nuka@kuis.kyoto-u.ac.jp

³ Department of Social Informatics, Kyoto University, Kyoto 606-8501, JAPAN, +81 75 753-4821, ishida@kuis.kyoto-u.ac.jp

⁴ Communication Department, Stanford University, Stanford, CA 94305, USA +1 650 723-5499, nass@leland.stanford.edu

embodied conversation partner with the proper human verbal and nonverbal communication skills [1]. However, these findings concern task-support agents.

There are projects which have used text-tracking to create agent-based social support. Julia [3] engages in entertaining one-on-one conversation; the Extempo bartender agent converses with chat visitors, and is designed to enhance the social atmosphere [7]; and there are bots for websites that answer questions and direct visitors in a friendly way (e.g. <http://www.artificial-life.com>). However, these agents are designed to engage in one-on-one social interactions, rather than facilitating human-human interaction.

Social Interface Agents to Support Human-Human Interaction—Design Concept for a Supporting Role

The agents described above were all designed to be communication partners with human users, and to be present and active at all times during the interaction. We do not believe that an interface agent should be an equal partner in the conversation, or equally active, when supporting human-human communication. This led us to make some important design decisions in creating our agent prototype.

We designed Helper Agent to be a supporting character rather than a central figure in the conversation. Helper's job is to pick up on contextual cues from the conversation, provide help, and then fade back into the background, allowing the central activities of the conversation environment to move forward. As we mentioned in the abstract, this behavior mimics the activities of a host at a party.

Target Application—Cross-cultural Conversations

For our first prototype, we focused on an extreme case of low social context in a virtual meeting space: strangers from different national cultures, meeting for the first time. Even when people can use a common language with reasonable fluency, they do not necessarily have a common context for their conversation. Different cultures have different notions of how to begin and develop conversations. What is a safe topic in one culture, may be very awkward in another culture. For example, in some cultures it is appropriate to ask about family members right away; whereas in other cultures this is private [5, 2].

We developed an agent prototype that could provide safe-topical suggestions, if the conversation was faltering. We focused on conversations between Japanese and Americans. These two national groups are known to have very different interaction styles and cultural norms [5], and so we felt this was a good test case.

In the rest of this paper, we will provide a more detailed overview of the design and features of the prototype, and discuss results of a cross-cultural test of the agent's usefulness. We will conclude with lessons learned for design of this new class of social interface agents, that we hope will be of use to the CHI community.

DESIGN OVERVIEW

The prototype Helper Agent was designed for a minimalist test of the effectiveness of our concept for a human-human communication assistant. The prototype tracks audio from a two-person conversation, looking for longer silences that will trigger its conversation aid. The agent basically acts in the same way a busy human party host does, looking for clues that the guests' conversations are going badly. (Pauses are a powerful cue for what is happening in a conversation [2].) Helper then directs a series of text-based, yes/no questions to both conversation partners in turn, and uses their answers to guide its suggestion for a new topic to talk about. Then the agent retreats until it is needed again.

Communication Environment

Our prototype works within an existing 3-D virtual meeting space called FreeWalk (see Figure 1), which was developed by us (Nakanishi and Ishida) [10]. Using FreeWalk allowed us to track and use audio silences, and supported the metaphor of the agent coming and going from the conversation, rather than becoming a conversation partner.



Figure 1. FreeWalk: Virtual Meeting Place Environment

Users are represented as three-dimensional pyramid objects, with their video image mapped onto one face of the pyramid (see Figure 1). In the lower right corner of the screen there is an overhead view of all avatars; however the user does not see his/her own avatar on the main screen (see Figure 2). The user does see a small video window of themselves in the lower left corner, to help them adjust their camera. Each user's voice is transmitted to others around them in the virtual space. The volume of other peoples' voices is proportional to how close they are to you in the space (farther away is fainter). Users can move around in the space, and rotate the orientation of the 'face' of their own pyramid, using the mouse or the arrow keys.

Helper Agent Features

Embodiment of the Agent in the Virtual Space

Helper Agent is presented on-screen the same way users are (see Figure 2). This allowed us to take advantage of nonverbal cues in designing the agent's behavior, such as turning to face users as it poses a question to them, and approaching and departing the conversation physically.

The agent is an animated dog, done in a style somewhere between typical Japanese and American cartoon dogs. We chose a dog because we wanted users to think of the agent as subservient, friendly, and reasonably socially intelligent. We chose stylized animation instead of more realistic, because we did not want the agent to be interpreted as a specific individual, but as more iconic and minor [9].

Nonverbal Communication Abilities

The dog has a set of animations of the proper nonverbal conversational moves for asking questions, reacting to affirmative or negative responses, and making suggestions. We crafted these animations as a supplement to the agent's speech [1], and focused on making them friendly and submissive in style [12].

The dog orients its face toward the user that it is addressing, and displays the proper animation for each phase: approach, first question, reaction, follow-up question, and finally topic suggestion. After concluding a suggestion cycle, the agent physically departs from the conversation zone, and meanders at a distance from the interaction, until it detects another awkward silence. This makes it clear to the

conversation pair that the agent need not be included in their discussion [4].

Topic Knowledge

We gathered topics using an internet survey, that university students from Japan and the United States filled out. We used the collected pool of topics to select common safe and unsafe topics for people from both countries. From these topics, we crafted a set of questions that the agent could ask during interaction, drawing users into conversation. Safe topics included: movies, music, the weather, sports, and what you've been up to. Unsafe topics included: money, politics, and religion. A sample safe question: "Is the weather nice where you are right now?" A sample unsafe question: "So, do you think it is alright for a country to fish for and eat whales?"

Topic-Suggestion Mechanisms

Silence-detection: The agent decides there is silence when the sum of the voice volumes of both participants is below a fixed threshold value. When the agent detects a silence that lasts for more than a certain period of time, it decides the participants are in an awkward pause.




Figure 2. Conversation from both participant's point-of-view: (1) person A is asked the first question (2) and responds, (3) then the agent comments. (4) Next person B is asked a question. Note that the agent faces the person it is addressing.

Positioning: The agent decides how to position itself, based on the location and orientation of each participant. The agent turns toward the participant that it's currently addressing. If the participants move while the agent is talking, the agent adjusts its location and orientation. The agent tries to pick a place where it can be seen well by both people, but also tries to avoid blocking the view between them. If it's hard to find an optimal position, the agent will stand so that it can at least be seen by the participant to whom it is addressing the question.

State-transitions: The agent has three states—idling, approaching, and talking. When idling, the agent strolls at the corner of the virtual space, further away than the normal conversation zone [4]. When the agent detects an awkward pause in the participants' conversation, it begins an approach. Upon reaching the participants, the agent goes into the talking state. However, if the participants start talking again before the agent reaches them, it stops the approach and goes back to idling. (This behavior is strikingly similar to the actions of a hesitant subordinate trying to approach a superior, who is engaged in a conversation with another dominant person.)

The agent will also remain in idling state if the participants are standing far apart from each other (out of conversation range), or are not facing each other. If the participants turn away from each other during the agent's approach, or while it is talking, it will return to idling state, as well.

Conversation Model and Interface

The user interface for communicating with the agent is very easy to learn. The agent does not use voice—it presents questions to the participants in a text-balloon above its head. (We thought text was far less intrusive than audio.) The user indicates 'yes' or 'no' using the mouse to click on their answer. Both participants see all questions, but only the user addressed sees the Yes/No options. When the person answers the question, their answer is displayed in a text-balloon above their own avatar (see Figure 2).

Each topic has a tree structure, with nodes that are: first question for a participant, possible answers by participants, agent's reply to each answer, and flags indicating whether the agent will address its next question to the other person or to the same person. Topics were designed to draw participants into a dialogue, so each turn is tailored for this purpose. The cycle always concludes with a recommendation for how the participants could make use of the particular topic area, given their own answers to the agent.

When the agent approaches to start a cycle, it selects a topic from its repertoire of safe (or unsafe) topics randomly, out of those that have not yet been used. Then it randomly chooses one of the two participants as the target for the first question. Let's call this person A. When A answers, the agent replies to A's answer (see Figure 2). Based on what A answered, the agent then chooses a follow-up question. This

question might be directed at A or at B. If it is directed at B, the agent turns to B to pose the question. When B answers, the agent replies to B. Finally, the agent makes a general comment that is meant to guide the participants into using this topic. This general comment is selected based upon the previous answers from the participants, so that it makes sense given their replies. After making this comment, the agent departs. If at any time a user does not respond to the agent's question, the agent will wait for an interval, and then go back into idling mode, without trying to continue its question cycle

EVALUATION OF THE AGENT

Goals

We wanted to test the benefits of our prototype in a controlled setting. Our initial expectations were:

1. The safe-topic agent would create a more satisfying experience, than if there were no agent. Participants would feel they were more similar, would be happier with the interaction and partner, and would form more positive impressions of one another's nationality.
2. The unsafe topic agent would make people uncomfortable, but might lead to a more meaningful and interesting conversation than the safe topic agent.

Design

We designed a 3-condition experiment using pairs of students who were located in the United States and in Japan. Pairs either interacted one-on-one, or had the help of the safe-topic or unsafe-topic Helper Agent.

Procedure

The study was a collaboration between the NTT Open Lab, Kyoto University's Department of Social Informatics, and Stanford University's Communication Department. We used a high-bandwidth (1.5 Mbps) dedicated line between the universities. The two research teams used chat software to communicate while running the study. We set up a PC with a small camera and microphone/headset at each location (see Figure 3), and installed FreeWalk and Helper Agent at both sites.



Figure 3. Set-up for the experiment (Stanford side).

We modified our prototype so that those in the agent conditions would all be exposed to the same number of

topics. We divided the conversation session into segments, and forced the agent to display a topic within each time segment. Thus, in the safe-agent condition, the agent introduced all 5 safe topics in random order. In the unsafe-agent condition, the agent introduced all 5 unsafe topics in random order.

Each research team recruited students for the study. The Stanford students were all part of an undergraduate class, which required study participation for credit. The Japanese students were undergraduates from Kyoto University and other nearby universities, who were paid for their participation. Because the study would be held in English, we screened Japanese students and selected those who scored at a reasonably high level on English proficiency tests. Both sets of students were screened for a high level of familiarity with one another's culture, and those with high experience were not asked to participate. In total we had 90 participating students. Due to some problems with equipment, we ended up with data from 45 Japanese students, and 43 American students, for our analysis. Students were assigned randomly to same-gender pairs. Each pair was randomly assigned to one of the three conditions.

Students were told that they would be testing out a new communication environment with a student from the other country. They were asked to talk about anything they liked, just "get to know each other a little bit". They were trained in how to use the system, then left alone to talk for 20 minutes. We made video recordings of all sessions, capturing what was on the screen on the Japan side onto videotape.

After their 20-minute conversation, participants filled out a web-based survey in their native language. (Questionnaire items were translated and then reverse-translated for accuracy.) The questionnaire included questions about the interaction, their conversation partner, the agent (in agent conditions), as well about the participant's own performance. We also asked them to make assessments of themselves, their partner, and the typical person of both participants' cultures on some commonly used stereotypic adjectives.

Results

Safe Agent versus No Agent

We got significant differences when comparing answers in the no agent and safe agent conditions. There was empirical confirmation that the concept of using a social interface agent to support human-human interaction has merit.

American Reaction

The safe agent had positive effects for American participants as we expected (see Figure 4; all items on an 8-point scale, 8 highest):

- **opinion of their own behavior higher**—they rated themselves as more confident, less domineering, and less restrained in the safe agent condition.

- **opinion of partner higher**—they rated their partner as significantly more trustworthy in the safe agent condition.
- **opinion of the typical Japanese person higher**—the safe agent condition had a positive effect on impression of typical Japanese people. Those in the safe agent condition rated the typical Japanese person as more creative and more friendly. However, no-agent condition participants rated the typical Japanese person as more emotionally expressive⁵.

variable	safe agent mean	no agent mean	t-value (df=24)
confident	6.46	5.54	-2.33**
domineering	4.00	4.92	2.03*
restrained	3.61	5.00	2.52**
partner trustworthy	6.54	5.91	-2.46**
Japanese creative	5.38	4.54	-2.06*
Japanese friendly	5.92	5.23	-2.08**
Japanese emotionally expressive	3.15	4.23	2.75***

Figure 4. Summary of t-test comparisons of American students' ratings, safe agent versus no agent (*p=.05, **p<.05, ***p=.01).

Japanese reaction

The Japanese participants had a different response to the safe agent's presence—it did not improve their experience. However, it did seem to make them think their partner was more like themselves, as expected (see Figure 5).

- **opinion of the experience lower**—Japanese in the safe agent condition rated the experience as less safe and more uncomfortable. They were less interested in continuing such a conversation, and were less satisfied afterward.
- **opinion of their own behavior lower**—they also rated themselves in a more negative light after the safe agent condition—safe agent condition participants rated themselves as more evasive and more quiet than no agent condition participants did.
- **opinion of their partner mixed**—their ratings of their American partners were mixed. The safe agent condition participants found their partners more talkative and more effusive, and less engaging⁶. Yet, they rated their partners as less typically American and more similar to themselves.
- **opinion of the typical American person lower**—the safe agent condition seemed to exacerbate negative views of Americans for Japanese participants. In the safe agent

⁵ Americans typically stereotype Japanese people as less creative, less friendly, and less emotionally expressive.

⁶ Japanese tend to stereotype Americans as talkative and emotionally effusive.

condition, they rated the typical American as more competitive, more domineering, more selfish, and more effusive than those in the no agent condition⁷.

variable	safe agent mean	no agent mean	t-value (df=26)
unsafe	3.29	2.24	-2.05*
uncomfortable	5.14	2.71	-3.9*****
desire to continue	4.86	7.07	3.55*****
satisfying	4.79	6.14	2.32**
self evasive	5.86	4.71	-2.09**
self quieter	4.68	3.36	-2.08**
partner talkative	5.00	4.07	-2.06**
partner effusive	2.21	1.50	-2.06**
partner engaging	6.00	6.78	2.47**
partner typically American	5.22	6.14	2.26**
partner similar to self	5.28	4.21	-2.20**
Americans competitive	6.57	5.00	-2.40**
Americans domineering	6.07	3.36	-4.44*****
Americans selfish	6.14	4.93	-2.26**
Americans effusive	6.79	5.86	-2.21**

Figure 5. Summary of t-test comparisons of Japanese students' ratings, safe agent versus no agent (*p=.05, **p<.05, ***p=.01, ****p<.01, *****p=.001, *****p<.001).

We cannot be sure why the two groups had such different reactions. One reason may be that the agent's questions were implemented in English. It's possible that Japanese subjects felt it was a two-against-one situation. This might explain why they disliked the interaction, even though it seemed to make them rate their partner as more similar to themselves (our main goal!). We would need to test the system again, using a bilingual agent that address all questions to users with both languages displayed, to be sure. In any case, the positive American reaction was a strong support of our research concept.

Safe Agent versus Unsafe Agent

We got significant and interesting differences in participants' reactions to the safe and unsafe agents.

Awkward isn't necessarily bad

As we had expected, the unsafe agent made things more awkward, but also more interesting.

⁷ All stereotypical American traits, from the Japanese point of view. It might surprise American readers that 'effusive' is bad; in Japan one is expected to regulate one's emotional expressions, or risk appearing childlike and uncultured [5].

We counted awkward pauses, by observing the videotapes, and found a higher number of awkward pauses in the unsafe versus safe condition (Means = 4.34 and 3.09, $t(56)=-3.06$, $p < .01$).

Despite the higher level of awkwardness in these conversations, both Japanese and American participants found the conversation that included the unsafe topic agent more interesting. Americans rated the unsafe agent interaction more interesting; Japanese rated the unsafe agent experience more desirable to continue (see Figures 6 and 7). Japanese participants found the unsafe agent experience more comfortable as well.

variable	safe agent mean	unsafe agent mean	t-value (df=26)
desire to continue	4.86	6.21	-2.00^
uncomfortable	5.14	3.57	2.41**
partner similar to self	5.29	3.64	2.58**
partner considerate	7.31	5.93	3.02****
partner domineering	1.14	2.07	-2.43**
partner friendly	7.29	6.14	2.31**
partner talkative	5.00	3.79	2.14**
self evasive	5.86	4.64	2.03*
self restrained	5.43	3.79	2.37**
self self-abasing	5.54	4.07	2.33**
self team-oriented	4.00	2.64	2.34**
Americans domineering	6.07	5.00	2.26**
agent nice	3.43	5.29	-2.25**
agent competent	4.29	5.57	-2.04*
agent typically Japanese	4.50	3.43	2.71**
agent talkative	5.61	4.36	2.66***
agent nationalistic	1.43	3.29	-2.87***

Figure 6. Summary of t-test comparisons of Japanese students' ratings, safe agent versus unsafe agent conditions (*p=.05, **p<.05, ***p=.01, ****p<.01, ^p=.056).

American partner seemed better in the safe topic condition

Japanese participants rated their partner as less similar to themselves, less considerate, more domineering, less friendly, and less talkative in the unsafe condition. These rankings suggest that the safe agent led to more positive impressions of the partner, for Japanese participants.

Unsafe topics made Japanese act more American

Japanese rated themselves as less evasive, less restrained, less self-abasing, and less team-oriented in the unsafe

condition⁸. It seems they thought they acted more American than those in the safe agent condition. Americans rated their partner as more similar to themselves in the unsafe condition. This seems to corroborate the Japanese self-ratings.

variable	safe agent mean	unsafe agent mean	t-value (df=24)
interesting	5.85	6.77	-2.18**
partner similar to self	3.31	4.77	-2.55**
Japanese emotionally expressive	3.15	4.15	-2.16**
Japanese outgoing	4.08	4.77	-2.04*
Japanese talkative	3.77	4.85	-2.30**
Japanese evasive	3.85	4.85	-2.39**
Japanese quiet	6.00	4.38	2.82***
agent blunt	4.69	7.36	-3.84*****
agent domineering	3.38	5.25	-2.31**
agent restrained	3.15	1.92	2.52**
agent friendly	5.46	4.08	2.03*
agent typically American	6.62	4.92	2.40**

Figure 7. Summary of t-test comparisons of American students' ratings, safe agent versus unsafe agent (* $p=.05$, ** $p<.05$, *** $p=.01$, ***** $p=.001$).

Safe/unsafe topic choice affected stereotyping in contradictory ways

Japanese participants in the unsafe agent condition thought the typical American was less domineering. This conflicts with their ranking of their own partner's level of domineering-ness.

American participants rated the typical Japanese in conflicting ways: after the unsafe condition, they thought the typical Japanese person was more emotionally expressive, more outgoing, and more talkative; but also more evasive and quieter.

Safe/unsafe agents 'read' differently for Japanese and Americans

The two groups differed in their impressions of the safe and unsafe agents. The Americans formed the intended impression: they rated the unsafe agent's topics as less appropriate, thought it acted more blunt, more domineering, less restrained, and less friendly. They also said it was less typically American, distancing it from their own in-group's behavior. The Japanese thought that the unsafe agent was

nicer and more competent than the safe agent. They rated the unsafe agent as less typically Japanese, and as less talkative. They found the unsafe agent more nationalistic, probably because it brought up more political topics than the safe agent.

LESSONS FOR THE DEVELOPMENT OF AGENTS TO SUPPORT HUMAN-HUMAN COMMUNICATION

A social interface agent can help support human-human communication

Our evaluation demonstrated that a human-human communication assistant can have positive effects on perception of the experience, one's own qualities, one's partner, and even one's partner's cultural group.

Provocative help can be good

Our evaluation also suggested that a communication assistant can be helpful both when it offers safe topics to talk about, and when it steers the conversation in less safe directions. In fact, the Japanese participants seemed to prefer the unsafe topic agent, and both groups found it more interesting than the safe topic agent.

For overall conversational support purposes, both kinds of help may be desirable. We suspect that an agent with a model for offering both kinds of topics, depending upon the conversation flow, would be the most desirable.

User-adaptation would make the agent more effective

The two cultural groups had very different impressions of the same agent behaviors, and reacted in different ways. For example, behavior that was perceived as blunt and unfriendly by Americans was seen as nice and competent by Japanese. An effective agent for different types of people will probably need to adapt its behaviors to user subgroups, or perhaps to individuals' own interaction styles and preferences. We believe we created a more American identity for our agent by delivering its topic help in English. In future iterations, we plan to create an agent whose presentation is adapted to different user styles and preferences.

Agent behavior may shift user behavior

Both the Japanese and American participants noted that Japanese seemed to act more American in the unsafe agent condition. This result indicates that it may be possible to mold user behavior with the choices one makes about how the agent will behave, creating a different conversational environment by bringing different traits to the fore. This could have very interesting implications for those interested in setting a specific group conversational tone or style in a virtual meeting space.

Design Notes for creating a supporting character

We believe our prototype's success was partially due to design choices that made it a graceful supporting player in the interaction. We summarize key features here, as suggestions for CHI community members who may be interested in creating this kind of interface agent:

⁸ These are all stereotypically American traits, for Japanese.

- **Unobtrusive observation of users, and easily visible and controllable approaches.** The users could see when the agent was approaching, and could shoo it away simply by talking before it arrived.
- **Focused interactions with a limited duration.** The agent had a clear purpose in its approach, and the users could quickly grasp its interaction pattern, and knew it would leave after making a topic suggestion. They could see the agent in-between interactions, but did not need to include it in their conversation. This kept the focus on the task at hand: human-human interaction.
- **Easy to ignore.** If users ignored the agent, it simply gave up and went away! It did not hang about forever waiting for them to answer.

Next Steps

Though silence-sensing produced strong results, we would like to incorporate content recognition, to make Helper Agent a more powerful assistant. Also, our first prototype supported only two people. With more participants, we can experiment with additional features, such as recognizing active conversation groups and leading a newcomer to an ongoing conversation. We are also interested in supporting groups with the same base language, but different sub-cultural memberships.

It would of course be interesting to run follow-up studies using pairs from other cultures, and to continue to refine and deepen our understanding of how cultural differences should affect an interface agent's behavior. Our work is really only a beginning in exploring this terrain.

CONCLUSION

We built a social agent prototype, that was designed to facilitate human-human interaction. A cross-cultural evaluation of the prototype demonstrated its effectiveness, and raised interesting considerations for further development of this class of interface agent.

We feel the support of human-human interaction in virtual meeting places is an exciting and useful new domain for interface agents. Given the proliferation of online spaces, and the interest in community formation that far exceeds the industry's ability to staff communities with human hosts, this kind of agent may become a familiar part of the virtual landscape. We hope that CHI community members who are called upon to think about and design these kinds of interface agents will find our prototype design and evaluation results useful for their work.

ACKNOWLEDGMENTS

Thanks to GEMNet for providing the broadband line, and to Eva Jettmar for her assistance in the evaluation.

REFERENCES

1. Cassell, J., Bickmore, T., Billingham, M., Campbell, L., Chang, K., Vilhjalmsson, H., and Yan, H., Embodiment in Conversational Interfaces: Rea, *International Conference on Human Factors in Computing Systems (CHI-99)*, pp.520-527, 1999.
2. Clark, H.H., *Using Language*, Cambridge University Press, 1996.
3. Foner, L., Entertaining Agents: A Sociological Case Study, in *Proceedings of the First International Conference on Autonomous Agents*, Marina del Rey, CA, 1997.
4. Hall, E.T., *The Hidden Dimension*, Anchor Books/Doubleday, 1982 (1966).
5. Hall, E.T., and Hall, M.R., *Hidden Differences: Doing Business with the Japanese*, Anchor Books, 1990 reprint.
6. Isbister, K., and Doyle, P. Touring Machines: Guide Agents for Sharing Stories about Digital Places, in *Proceedings of the Workshop on Narrative and Artificial Intelligence*, AAAI Fall Symposium Series, 1999.
7. Isbister, K., and Hayes-Roth, B. Social Implications of Using Synthetic Characters, in *Animated Interface Agents: Making Them Intelligent* (a workshop in IJCAI-97, Nagoya, JAPAN, August 1997), 19-20.
8. Lester, J.C., Barlow, S.T., Converse, S.A., Stone, B.A., Kahler, S.E., and Bhogal, R.S. The Persona Effect: Affective Impact of Animated Pedagogical Agents. *Proceedings of CHI '97* (Atlanta GA, March 1997), ACM Press, 359-366.
9. McCloud, S., *Understanding Comics: The Invisible Art*, Harper Perennial, 1993.
10. Nakanishi, H., Yoshida, C., Nishimura, T., and Ishida, T., FreeWalk: A 3D Virtual Space for Casual Meetings, *IEEE MultiMedia*, 6(2), pp.20-28, 1999.
11. Parise, S., Kiesler, S., Sproull, L., and Waters, K. My Partner is a Real Dog: Cooperation with Social Agents, in *Proceedings of Computer Supported Cooperative Work '96* (Cambridge, MA), ACM Press, 399-408.
12. Reeves, B., and Nass, C., *The Media Equation: How People Treat Computers, Television, and New Media Like Real People and Places*, Cambridge University Press, 1996.
13. Sugawara, S., Suzuki, G., Nagashima, Y., Matsuura, M., Tanigawa, H., and Moriuchi, M., InterSpace: Networked Virtual World for Visual Communication, *IEICE (The Institute of Electronics, Information and Communication Engineers) Transactions on Information and Systems*, E77-D(12), pp.1344-1349, 1994.
14. Waters, R.C., and Barrus, J.W., The Rise of Shared Virtual Environments, *IEEE Spectrum*, 34(3), pp.20-25, 1997.