
Public Displays of Affect: Deploying Relational Agents in Public Spaces

Timothy Bickmore
Laura Pfeifer
Daniel Schulman
Sepalika Perera
Chaamari Senanayake
Ishraque Nazmi

Northeastern University
College of Computer and
Information Science
360 Huntington Ave, WWH202
Boston, MA 02115
bickmore@ccs.neu.edu
laurap@ccs.neu.edu
schulman@ccs.neu.edu
sepalika@ccs.neu.edu
chaamari@ccs.neu.edu
nazmi.i@neu.edu

Abstract

Design principles for deploying agents designed for social and relational interactions with users in public spaces are discussed. These principles are applied to the development of a virtual science museum guide agent that uses human relationship-building behaviors to engage visitors. The agent appears in the form of a human-sized anthropomorphic robot, and uses nonverbal conversational behavior, empathy, social dialogue, reciprocal self-disclosure and other relational behavior to establish social bonds with users. The agent also uses a biometric identification system so that it can re-identify visitors it has already talked to. Results from a preliminary study indicate that most users enjoy the conversational and relational interaction with the agent.

Keywords

Relational agents, social interfaces, interactive installation

ACM Classification Keywords

H5.2 [Information Interfaces and Presentation]: User Interfaces—Graphical user interfaces, Interaction styles, Natural language, Voice I/O.

Introduction

Relational agents (RAs) are computational artifacts designed to build and maintain long-term social-emotional relationships with users [2]. These systems are often developed as anthropomorphic conversational interfaces in order to employ human verbal and nonverbal social and relational behavior such as proxemics and facial displays of attitude and affect. There are significant challenges in deploying these agents in public spaces in which issues such as user identification, user location, and bystanders must be addressed.

In this paper we present a set of design principles for relational agents, and review related work in deploying such agents in public spaces. We then describe the development of “Tinker”, a relational science museum guide agent, and how we implemented the principles and addressed the challenges outlined above.

Relational Agent Design Principles

Relational agents have a basic set of required components:

Design for multiple interactions. Relationships have temporal extent, so RAs must support multiple interactions with a user. Tinker was explicitly designed to support multiple conversations with each museum visitor, for example, by having follow-up dialogue about where the visitor went after their last conversation. Tinker promotes multiple interactions by asking visitors to return and talk to her at the end of each conversation.

Relational behavior. RAs must have a repertoire of behavior that can be used to increase bonding with users. Tinker uses social dialogue, verbal and nonverbal expressions of empathy, reciprocal self-disclosure, getting acquainted talk, calling the visitor

by name, explicit valuing of the relationship, explicit desire to continue the relationship, continuity behaviors (e.g., talking about what the visitor did while away), and increasing common ground [2].

Persistent relational model. An RA must maintain an assessment of the current status of its relationship with each user so that relational behavior can be used as needed, and so that other behavior (e.g., forms of address, discussing sensitive topics) are used in a manner that is appropriate to the current relationship. Tinker uses a one-dimensional scalar to represent social distance between her and each visitor, and she can report this value to visitors if they ask her about her relational model.

Persistent discourse model. An RA must record the propositional content of prior conversations, important not only for continuity in dialogue content, but for giving visitors the relational perception of sharing common ground. Tinker remembers topics previously discussed as well as many facts about each visitor, and takes these into account in future dialogues with them.

User identification. Obviously, relational and discourse models are unique to each human-agent dyad. Thus, if there is any possibility of multiple people interacting with a given agent, it is essential that the agent know who it is talking to. Tinker uses biometrics, together with an identification dialogue used for error resolution, to identify return visitors.

Natural conversational interaction. We believe that relational and other social behavior works best when it is presented in the context of a simulated human-human interaction (although it is certainly not a requirement). In addition, we also believe that natural language must be used in non-trivial interactions, and that an anthropomorphic character should be used to display human nonverbal behavior (hands for gesturing, etc.). For all of these reasons, we believe that embodied conversational agents [3] represent the best medium for relational agent

implementation. Tinker appears as an anthropomorphic robot, and can exhibit a wide range of human conversational and affective behavior in synchrony with her synthetic speech.

Related Work

Here we briefly review humanoid conversational agents (virtual and robotic) that are deployed in public spaces, two of which are installed in museums as guides (Kopp, et al [6] and Shiomi, et al [8]) and one which acts as a receptionist (Gockley, et al [5]). None of these agents use explicit models of the user-agent relationship, and they have a very limited repertoire of relational behavior (form of address and social dialogue). Two are able to identify visitors (Shiomi, based on RFID tags, and Gockley, based on magnetic strip ID cards), but only use this information to address users by name. The systems also only support very limited dialogue: Shiomi's robots can only talk at users (no dialogue support), while the other two systems use typed-text input and pattern-matching rules which support social chat but do not provide the deep dialogue models required for extended coherent conversation about a given topic.

Development of the "Tinker" Agent

As an initial experiment in building a relational agent for use in a public space, we have developed a virtual museum guide agent named "Tinker" who is currently installed in the Computer Place exhibit at the Boston Museum of Science (Figure 1). Tinker appears as a six-foot-tall 3D cartoon robot, projected in front of visitors, and communicates with them using synthetic speech and synchronized nonverbal behavior. Tinker can provide visitors with information on and directions to a range of exhibits in the museum, as well as discuss the theory and implementation underlying her own




Figure 1. "Tinker" System Installation

creation. Most importantly, Tinker uses a range of human relationship-building behavior to engage users, along with a biometric sensor to re-identify return visitors so that the conversation, and relationship, can be continued. We hypothesize that museum exhibits that engage visitors in human bonding rituals could result in increased visitor satisfaction and engagement, potentially leading to increases in learning.

Tinker was developed over an eight-month period of time in close collaboration with the staff at Computer Place. This is a staffed area of the museum that provides visitors with explorations in computer science, communications, and robotics.

T: Hi, my name is Tinker. What is your name? You can type it on the screen to your right.

U: My name is Bob.

T: Hi Bob. Welcome to Computer Place!

T: What can I help you with?

U: Tell me about yourself.

T: What would you like to know about me?

U: What are you?

T: I am a relational agent, in the form of a virtual robot.

T: First, let me ask, how much do you know about computers?

U: Quite a bit.

T: That is great.

T: A relational agent is a computational artifact that is designed to build and maintain long-term, social-emotional relationships with users.

U: What do you mean by social-emotional?

...

T: So, what else can I help you with?

U: Nothing, thanks.

T: Come back and talke to me again before you leave today.

U: I will.

T: Have a great day at the museum, Bob. And let's see whether I can recognize you when you come back. Goodbye.

Figure 2. Sample Interaction

Dialogue Content

Development of Tinker's dialogue began by videotaping museum staff giving descriptions of exhibits and interacting with visitors, in order to characterize these conversations and the nonverbal behavior used. We then developed scripts using a hierarchical transition network-based dialogue model [2]. Tinker's nonverbal behavior was both specified explicitly (e.g., pointing gestures during direction giving) and automatically generated (using BEAT [4]). In addition, Computer Place staff felt that it was important that Tinker's dialogue about computers be tailored to each visitor's level of computer literacy. Consequently, Tinker establishes each visitor's computer literacy level through dialogue before discussing any technical content, and remembers this for future conversations. Relational dialogue, outlined above, was also added to the scripts, as well as tailoring of dialogue based on relationship status.

Installation: Relational Agents in Public Spaces

The concept for the installation went through several iterations. We felt that it was important that the character appear human-sized to facilitate naturalness of interaction, and decided to use projection technology so that as much of the character's body could be shown as possible without limiting its hand gesture space. We used multiple-choice touch screen input for user utterances, based on other work in developing a conversational agent for users who had no prior computer experience [1].

There are several significant challenges in deploying such relational agents in crowded settings such as museums. These include: user re-identification; user presence detection (for conversation initiation and

termination); and user location detection (so that the agent can appear to be looking directly at the visitor, required for human conversational turn-taking and grounding cues [3]). We solved all three of these problems by using a glass plate that visitors rest their hand on during their conversations with Tinker. Sensors on the plate provide presence detection, and a camera underneath provides hand shape-based user identification. In addition, with a visitor's left hand on this plate and their right hand using the touch screen, their location is fixed between the two, solving the agent gaze problem. We also use a motion sensor to determine if visitors are in Tinker's general area so that she can beckon them over to talk.

We added several other objects to Tinker's virtual environment to address other problems that are unique to public settings. A large scrolling text screen was placed behind Tinker, showing the content of the last several conversational turns. We felt this was important in order to support the involvement of bystanders who might be near Tinker once a conversation is underway, as well as supporting individuals with hearing problems or who have difficulty understanding the synthetic voice. We also placed a smaller sign behind Tinker to display system status information (e.g., indicating the system is down) as well as a demonstration animation sequence showing approaching visitors how to use the hand reader. Finally, a virtual hand recognition reader was placed in Tinker's environment so that she could demonstrate putting her hand in the reader when visitors approach.

The current installation is located at the entrance to Computer Place (Figure 1). Tinker is projected onto a screen using a short-throw projector, and runs on two

networked computers. Hand recognition is performed by extracting geometric features from hand images, and comparing them to those from prior visitors [7].

Preliminary Evaluation

We conducted an initial acceptance and usability study of the system as a pre-requisite to planned experimental evaluation studies. This evaluation was performed over an eight-hour period of time spanning two weekend days. Museum visitors were first observed using the system, after which semi-structured interviews were conducted to obtain subjective feedback. After this, they were asked to conduct a second interaction with Tinker, followed by a second interview.

Participants

During the period in question, 72 users initiated interactions with the system, with 50 successfully completing conversations with Tinker. Interviews were conducted with 34 of these visitors, with ages ranging from 5 to 55 (average 22.5), 68% of whom were male. Most of these visitors arrived in groups of two or more.

Interactions

The duration of the first interaction ranged from 1.5 to 10.5 minutes (average 4.6). Thirteen participants agreed to interact with the system for a second time. The second interactions were briefer compared to the first, lasting an average of 2.6 minutes. Many of these participants seemed to be excited to find out whether Tinker would recognize them when they returned. The system correctly identified 77% of the return users, although only 31% were recognized based on handprint, with the rest being recognized through dialogue.

Usability

Twenty-one visitors (62%) had some problem using the system, with the majority of these (13) having some initial problem using the hand reader. The majority of these corrected themselves after the system provided feedback (e.g., if they remove their hand, Tinker interrupts the conversation and tells them they must keep their hand in the reader for the conversation to continue). A few visitors also had confusion about turn-taking cues (not knowing it was their turn to say something, 6%) or had problems using the touch screen (3%).

We also observed that the system was used collaboratively when participants were in a group. Adults or older children accompanying a young child often helped the child enter inputs on the touch screen. Often, when a participant made a mistake or was unable to proceed, they were prompted on the correct use of the system by other members of the group.

Subjective Evaluation

Most visitors (62%) enjoyed using the system ("cool" was by far the most frequent term used to describe it, 29%). Only two (6%) did not like it ("weird", "boring"). As some visitors put it: "Great. An introduction to science in simple words."; "Interesting, a whole new way of interacting with the museum."

Visitors were amazed when Tinker recognized them, and several children smiled when Tinker called them by their name, even if they had just typed it in. They were also excited to see that Tinker not only knew their names but also what they talked about earlier. We heard responses like: "Wow, she remembers me! This is so cool."; "It is unusual that it recognized [me]. It's a

lot of fun.”, and “I liked that Tinker remembered what she talked about last time.” Conversely, visitors did not like it when Tinker mis-recognized them. Three visitors terminated their first interaction immediately after they were misidentified by the system and called by the wrong name.

When asked whether they would have rather talked to a museum staff member than Tinker, most participants (56% of respondents) said they would prefer Tinker (31% said they would prefer a person, 13% were unsure):

“I can keep asking questions without bothering her.”

“Good Idea. I can avoid having to talk to a person”

“It’s more appropriate for a science museum.”

“It’s easier to pick what you want than try to explain it to a person.”

Many visitors enjoyed the social and relational dialogue more than the science content, and wanted to chat more with Tinker about their background and hers. Visitors also liked Tinker’s discussion of her relational model, and their current relationship ‘score’: “How it finds the friendship through numbers is cool.”

Conclusion

Overall, visitors enjoyed talking to Tinker, and many were entranced when she recognized them and continued earlier conversations. Many also enjoyed the social dialogue and other relational behavior.

Future Work

Based on this evaluation, our next development iteration will focus on improving the hand recognition rate, making the hand recognition unit more intuitive to

use, and providing more turn-taking cues, such as having Tinker glance at the input menu when it is the visitor’s turn to say something. We also plan to conduct a series of experiments to evaluate our hypotheses regarding engagement of museum visitors.

Acknowledgements

Thanks to Dan Noren, Taleen Agulian and the staff at Cahner’s Computer Place. This work was supported by NSF CAREER grant IIS-0545932.

References

- [1] Bickmore, T., Caruso, L., Clough-Gorr, K., and Heeren, T. “It’s just like you talk to a friend” - Relational Agents for Older Adults. *Interacting with Computers*, 17, 6, (2005) 711-735.
- [2] Bickmore, T. and Picard, R. Establishing and Maintaining Long-Term Human-Computer Relationships. *ACM Transactions on Computer Human Interaction*, 12, 2, (2005) 293-327.
- [3] Cassell, J., Sullivan, J., Prevost, S., and Churchill, E., Eds., *Embodied Conversational Agents*, The MIT Press, Cambridge, MA, 2000.
- [4] Cassell, J., Vilhjálmsón, H., and Bickmore, T., BEAT: The Behavior Expression Animation Toolkit, *SIGGRAPH ’01*, 2001, pp. 477-486.
- [5] Gockley, R., Bruce, A., Forlizzi, J., Michalowski, M., Mundell, A., Rosenthal, S., Sellner, B., Simmons, R., Snipes, K., Schultz, A. C., and Wang, J., Designing Robots for Long-Term Social Interaction, *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2005.
- [6] Kopp, S., Gesellensetter, L., Krämer, N., and Wachsmuth, I., A conversational agent as museum guide - design and evaluation of a real-world application. , *Intelligent Virtual Agents*, 2005, pp. 329-343.
- [7] Schulman, D., Sharma, M., and Bickmore, T., The Identification of Users by Relational Agents, *Autonomous Agents and Multi Agent Systems (AAMAS)*, Estoril, Portugal, 2008.
- [8] Shiomi, M., Kanda, T., Ishiguro, H., and Hagita, N., Interactive Humanoid Robots for a Science Museum, *HRI’06*, 2006.