

The Identification of Users by Relational Agents

Daniel Schulman, Mayur Sharma, Timothy Bickmore
Northeastern University College of Computer and Information Science
360 Huntington Avenue, WVH 202, Boston, MA 02115
{schulman, mayur, bickmore}@ccs.neu.edu

ABSTRACT

Virtual agents designed to establish relationships with more than one user must be able to identify and distinguish among those users with high reliability. We describe an approach for relational agents in public spaces to identify repeat users based on two strategies: a biometric identification system based on hand geometry, and an identification dialogue that references previous conversations. The ability to re-identify visitors enables the use of persistent dialogue and relationship models, with which the agent can perform a range of behaviors to establish social bonds with users and enhance user engagement. The agent's dialogue encourages users towards repeat visits, and provides mechanisms of recovery from identification errors, as well as contextual information which may be used to improve the accuracy of the biometric identification. We have implemented and evaluated this identification system in a virtual guide agent for a science museum that is designed to conduct repeated and continuing interactions with visitors. We also present the results of a preliminary evaluation of the system, including user opinions of this technology, and of the effect of identification, both successful and unsuccessful, on acceptance and engagement of the agent.

Categories and Subject Descriptors

H.5.2 [Information Interfaces and Presentation]: User Interfaces – *graphical user interfaces, interaction style, natural language.*

General Terms

Design, Experimentation, Human Factors

Keywords

Relational agents, social interfaces, interaction installation, biometrics, user identification.

1. INTRODUCTION

Relational agents are computational artifacts designed to build and maintain long-term, social-emotional relationships with their users [2]. Such relationships are important in applications in which trust, rapport and working alliance have been shown to be important in human-human interactions, such as in education, sales and the helping professions. Regardless of which theoretical model is adopted to represent a user-agent

relationship, such models are always unique to a particular user-agent dyad. Thus, when these agents are placed in locations in which more than one user has access to them, user identification is crucial. In addition to the relational model, such a shared agent must remember its interaction and discourse history with each user so that prior conversations can be continued or referred to in the future. Such mechanisms will be important, for example, in conversational home robots that are purchased by households that have more than one occupant.

In terms of human social interaction among non-strangers, identification of one's interactant is a pre-requisite for any kind of relationship or shared history, and is the most fundamental form of acknowledging their "positive face", which is every person's need for social acceptance [7].

One solution to this user identification problem is to require users to "log in" to an agent at the start of each interaction. However, this is unnatural and inconvenient and continually violates user expectations that an agent will follow the rules of human social interaction. The use of ID cards (bar code or magnetic strip) has similar problems. RFID tags solve some of these problems, but require users to carry an artifact to make the interactions work.

Biometrics provides a solution to the user identification problem for relational agents and also, coincidentally, mimics the mechanism that people use to identify each other. Of course, these technologies can make mistakes, just as people can, and error correction mechanisms must be incorporated into these systems. People typically use dialogue for this error correction for both false positives (e.g., A: "Hi John.", U: "Excuse me?", A: "Oh, I'm very sorry, I couldn't see you clearly.") and false negatives (A: "Hello, do I know you?", U: "It's Mary.", A: "Oh, hi Mary.").

In this paper we present a hand geometry-based biometric user identification system for a relational agent that uses the identification function to resume relational and discourse models for repeat users, and conversational strategies for fallback error correction. In our approach we have intentionally sacrificed accuracy (especially the false positive rate so important in security applications) for the sake of naturalness and user satisfaction with the experience. The system is deployed as part of a virtual museum guide agent named "Tinker" who is currently installed in the ComputerPlace exhibit of the Boston Museum of Science.

1.1 Related Work

Here, we review work on agents designed to identify their users, as well as guide agents similar in function to our museum guide application. We then briefly review work in biometrics in general and hand geometry identification in particular.

Cite as: The Identification of Users by Relational Agents, Schulman, Sharma, and Bickmore, *Proc. of 7th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2008)*, Padgham, Parkes, Müller and Parsons (eds.), May, 12-16., 2008, Estoril, Portugal, pp. 105-111.
Copyright © 2008, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

1.1.1 Agents that Incorporate User Identification

Shiomi, et al, report on an interactive robot installation at a science museum, in which four robots act as guides to the exhibit [15]. The robots identify visitors via RFID tags given to visitors at the exhibit entrance, can address visitors by name, play games (e.g., rock, paper, scissors) and take visitors to shows that are about to start. However, the robots only talk *at* visitors; they are not able to understand visitors or engage them in dialogue.

Also of relevance is Valerie the “Roboceptionist”, an animated and partially robotic receptionist [6]. Valerie can look up phone numbers and give directions (using typed text input and synthesized speech and synchronized animation output), as well as engage in social dialogue. The system identifies repeat users via a card reader that can read any of the magnetic stripe cards (e.g., credit cards) that a user carries; however, user identification is used only to greet return users by name. In addition, the system used a ‘chatbot’ rule-based interaction without a true dialogue model.

1.1.2 Virtual Guide Agents

There has been a significant amount of research on the development of interactive museum exhibits and mobile guide devices over the last decade (e.g., [13], [16]).

Kopp, et al, report on the development and fielding of an animated conversational character guide for the Heinz-Nixdorf-MuseumsForum computer museum in Paderborn, Germany [12]. The system features a visual system to track visitors, typed text input and synthesized speech and synchronized 3D animated character output. Initial evaluations characterizing visitor-agent dialogue indicated that people were willing to engage and converse with the agent. However, the system did not use any relational behavior, was not able to re-identify return visitors, and no outcome studies have been reported yet on the efficacy of

the system.

MACK was an animated guide to the MIT Media Laboratory, primarily designed to provide directions to visitors, and descriptions of laboratories. MACK featured a physical map that was placed between it and the visitor, and allowed the visitor (using a stylus) and/or MACK (using an overhead projector) to point at the map [3].

1.1.3 Biometrics

A wide range of different techniques for biometric identification have been investigated. Jain et. al. [10] give a review of the most widely used methods and the trade-offs involved. Focusing specifically on hand geometry recognition, the approach used here, prior work can be divided into earlier systems that require the use of guidance pegs to align the hand to a known location and position (e.g., [9]), and so-called “peg-free” systems which do not.

Peg-free systems can generally be divided into two approaches: the first, which we use here, is based on identifying specific geometric features of the hand (e.g., the length of a finger). Wong and Shi [18] give an approach, very similar to the approach we use here, which identifies key landmarks (e.g., the tips of the fingers) by analyzing the curvature of the hand contour. Amayeh et.al. [1] give an alternate method based on the use of complex image moments as features.

2. INSTALLATION CONCEPT

The motivation for using hand geometry-based identification grew out of requirements for the museum installation. The agent was going to appear as a human-sized embodied conversational agent that visitors could approach and talk to. Since the installation was going to be placed in a very crowded, noisy environment, the use of vision-based or voice-based user identification seemed to be ruled out. We also needed the identification process to be as natural and non-intrusive as possible, and needed to be able to accurately determine when a user was present and where they were standing so that the agent could appropriately initiate and termination conversations and gaze at its user.

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 and begin conversation initiation behaviors. The initial concept for the installation is shown in Figure 1.




Figure 1. Original Installation Concept




Figure 2. "Tinker" System Installation

3. IMPLEMENTATION

We first give a general description of the architecture and functionality of the Tinker system, followed by a more detailed examination of her user identification strategies and related behaviors.

3.1 The "Tinker" Museum Exhibit

Tinker appears as a six-foot tall anthropomorphic 3D cartoon robot, projected in front of visitors, and communicates with them using synthesized speech and synchronized nonverbal behavior. Tinker provides information on the exhibits within Computer Place, gives directions to points of interest in the museum, and discusses the theory and implementation underlying her own creation. Most importantly, Tinker is designed to engage users over repeated visits, applying a range of human relationship-building behaviors. Of particular relevance to this paper are behaviors intended to provide continuity between interactions. At the conclusion of an interaction, Tinker encourages the user to visit her again. When discussing particular exhibits, she will ask the user to go see the exhibit, then come back and talk to her about it. She will repeat this encouragement in her farewell ("Let me know how you like the exhibit").

The installation is located at the entrance to Computer Place (Figure 2). Tinker is projected onto a 3' by 4' screen using a short-throw projector, and runs on two networked computers. Two pedestals are positioned approximately 40 inches high, and a user will stand between them. The left pedestal contains a hand recognition unit consisting of a glass plate on which a user rests his or her hand, and a camera beneath. In addition, a pressure sensor on the plate provides presence detection. The right pedestal supports a small touchscreen from which users select their responses during dialogue. The installation also contains a motion sensor which detects when visitors are nearby.

Tinker's dialogue is driven by scripts written using a hierarchical transition-network based dialogue model [2]. Users select their responses from a short list of choices using the touchscreen. Tinker's nonverbal behavior is mostly added to her scripted utterances both automatically, using BEAT [4], although behaviors can be added manually as needed (e.g., pointing gestures when giving directions).

Although the original concept for Tinker involved the use of two identical pedestals, after several design iterations we discovered that the lighting behind (over) the hand being recognized must be tightly controlled to provide a good silhouette, so an enclosure was designed over the left-hand pedestal with LED illumination. When a user places their hand on the reader plate and is sensed by the pressure sensor, the enclosure is illuminated and a message is sent to the interaction dialogue system to initiate a new conversation.

3.2 Biometric Identification

Tinker performs biometric identification of users based on hand geometry. When a user places his or her hand on the glass plate of the hand reader, the camera captures a grayscale image of the hand, backlit by the LEDs mounted above the glass plate. Once a hand image is captured, identification proceeds in two steps: A vector of features is extracted from the image, and this feature vector is compared to stored feature vectors from previous users in order to determine whether there is a match.

Feature extraction is implemented using the OpenCV image processing library¹. The hand is separated from the brightly lit background using a threshold, yielding a binary image, and the contour of the hand (a one-pixel-wide outline) is obtained. We analyze the curvature of each point along this contour. The points of highest curvature are assumed to be the tips of the fingers and the valleys between them – this technique is similar to that used by Wong and Shi [18]. These points are used as landmarks, from which we produce forty-nine features (Figure 3), including the length of each finger, the width of each finger at five points from the base to the tip, average width of each finger, aspect ratios of fingers, palm width, area of each finger, and perimeter of each finger.

The resulting feature vector is compared to stored feature vectors for the previous 20 users. The number of previous feature vectors to store is a trade-off between the false positive rate and the false negative rate. A larger number increases the probability of an incorrect match between a new visitor and a repeat visitor. However, a smaller number increases the probability that the feature vector corresponding to a repeat visitor will no longer be stored when they return. Preliminary observations at the museum indicated that 20 previous feature vectors would be sufficient to include nearly all repeat visitors, while still providing an acceptable false positive rate.

We use a support vector machine (SVM) [17] to perform the comparison. This was implemented using the libSVM library². Given a pair of feature vectors, we subtract them, obtaining a

¹ <http://www.intel.com/technology/computing/opencv>

² <http://www.csie.ntu.edu.tw/~cjlin/libsvm>



Figure 1. Hand Image, Hand Silhouette, and Features

vector of the differences on each individual feature. This vector of feature differences is used as input to an SVM, which classifies it either as a match (i.e. the two feature vectors were produced by the same hand) or no match.

In order to train the SVM, we collected samples of hand images from visitors to the museum. We obtained multiple hand images from each of 27 people, for a total of 98 images. Exact ages were not recorded, but based on visually estimated ages, visitors ranged from 5 to 50, and 29.6% were children age 12 or younger. With these images, we produced a training set consisting of the vector of feature differences between each possible pair of images.

10-fold cross-validation gives a predicted success rate of 88.7% when identifying a new visitor, and 83.1% when re-identifying a repeat visitor.

3.3 Incorporating Contextual Information in Biometric Identification

Virtual agents, and particularly conversational and relational agents, potentially have an advantage when applying user identification techniques: They engage the user in a larger interaction, of which identification is only a piece, and thus have access to a range of contextual information which can be used to assist identification.

Based on initial observations of the system in the museum, we hypothesized that the majority of visitors who interacted with Tinker a second time would do so after only a short interval. To confirm this hypothesis, we observed visitors interacting with the system, noting repeat visits and the start and end times of all visits. During a five-hour period on a weekend day, we observed 83 interactions with Tinker, of which 13.3% were repeat visits. Of these, 54.5% were after an interval of less than 2 minutes from the first visit, and all but one were within 20 minutes. The average interval was 4 minutes. In the case of the very short intervals, the visitors appeared to be talking to Tinker, then

immediately returning in order to see if she would recognize them.

We incorporated this knowledge into Tinker’s hand recognition system using Bayesian methods. Using the SVM trained earlier, we train an additional sigmoid function which maps the output of the SVM to a probability that the two feature vectors under consideration are a match [14]. Given a pair of feature vectors, we can now estimate $P(f|match)$, the probability of the SVM output f given a match between the feature vectors. From the observations discussed above, we give an estimate of $P(i|match)$, the probability of time interval i given a match between the feature vectors. Assuming conditional independence between the SVM output and the time interval, we now have a situation equivalent to a naïve Bayesian classifier, and we can compute $P(match|f,i)$, the probability of a match given both the SVM output and the time interval since the user last visited.

By incorporating this information, we effectively make the hand recognition system more likely to match a stored feature vector if it belongs to a user who visited recently, rather than one who visited some time ago. We estimate that this results in a small (approximately 2%) increase in accuracy.

3.4 Dialogue-based Error Correction

Tinker also implements a secondary identification method based on conversational dialogue. This dialogue must consider two possible cases: Either the user has been matched by the biometric system, or they have not.

User’s hand not matched. When a user is not identified by the biometric system, Tinker introduces herself and asks the user to give his or her name. Supplying a name is optional – if the user chooses not to supply one, then Tinker will simply use slightly modified dialogue that eliminates any use of the user’s name. Tinker maintains a database of known names, and checks all names entered against it. This is intended to prevent users from entering expletives or other inappropriate words so that Tinker will speak them. Similar abusive behavior has been observed to

be commonplace in other museum guide agents (e.g., [12]). Tinker does not explicitly acknowledge that a name has been rejected, but simply omits the use of the name in her dialogue.

If the name entered by the user matches that given by a prior user in the last 10 minutes, then we consider it likely that the current user is a repeat visitor that the biometric system failed to identify. Tinker tries to confirm this situation by asking about the last topic discussed in the prior conversation: For example, “Hello Bob, were we just talking about the Theatre of Electricity?” The user is given a choice of either confirming (“Yes, that’s right.”) or denying (“No, we’ve never talked before.”). If the user confirms, Tinker acknowledges this, and continues with the conversation. Otherwise, she treats the user as a first-time visitor.

User’s hand is matched. If a user has been matched by the biometric system, Tinker greets them, using the name previously supplied, if any (e.g., “Hi Bob, great to see you!”). The user is given the option of acknowledging the greeting (“Hi Tinker”) or rejecting it (“We’ve never talked before.”). If the user acknowledges, providing grounding of their identity [5], then Tinker considers their identity to be confirmed. If they reject (either due to a false match of a new user or a mismatch of a prior user), then she apologizes, and asks for a name, continuing as if the biometric system had not identified them.

We view Tinker’s apology and recovery dialogue as a critical part of her behavior for two reasons: First, this type of affective behavior has shown promise for mitigating user frustration [11], and embodied conversational agents such as Tinker may be particularly effective [8]. Second, such a failure to acknowledge a user’s identity is a deep insult, causing an affront to their “positive face” [7]. Tinker’s apology and recover dialogue is an attempt to repair this insult and her relationship with the user.

Finally, we note that Tinker’s dialogue-based error correction capabilities rely on her relational capabilities, particularly her persistent discourse model. At the conclusion of an interaction, Tinker encourages the user to visit her again, asking them to come back and talk to her about particular exhibits that have been discussed. In addition to providing encouragement to return and continuity between dialogues, these utterances are stored in the discourse model, allowing us to reference the previous conversation during the identification dialogue in a natural and relevant manner.

4. EVALUATION

We conducted two preliminary evaluations of the system’s user identification capabilities as a pre-requisite to planned experimental evaluation studies. In the first study, museum visitors were simply observed using the system in order to evaluate its success rates and typical interaction pattern, while in the second, visitors were asked to use the system, then interviewed in order to determine user acceptance of the system.

4.1 Evaluation #1

The goal of the first evaluation was to identify typical usage patterns during normal usage of Tinker by museum visitors. We wished to observe how visitors reacted to her biometric systems and her dialogue-based error correction capabilities without the

potential bias introduced by prompting visitors to interact with Tinker.

The evaluation was performed during one weekend day, over a period of approximately 5 hours. Visitors were observed interacting with Tinker in the ComputerPlace exhibit. Visitors were not prompted to interact with Tinker, nor were they interrupted in any way during their interaction. We noted identifying characteristics of each visitor (e.g., clothing, hair color), thus making it possible to distinguish repeat visits.

4.1.1 Participants

During the period in question, 63 visits to Tinker were observed, of which 19% were identified as repeat visits. 43% of visitors were males, and 41.2% were estimated to be children below the age of teenager. 94% arrived with groups of two or more.

Of those visitors who interacted with Tinker at least twice, 50% were males, and 33% were estimated to be children. All arrived with a group of two or more.

4.1.2 Interactions

The duration of an interaction ranged from approximately 30 seconds to 7 minutes, with an average duration of approximately 1.5 minutes. Repeat visitors generally returned quickly; 83.3% returned after an interval of 2 minutes or less. The longest observed time between a first and second interaction was 47 minutes.

3.9% of first-time visitors were mistakenly identified as repeat visitors. 33.3% of repeat visitors were correctly identified by the biometric system. One visitor was identified through dialogue, giving a total success rate of 41.7%. The rest of the repeat visitors left immediately after the biometric system failed to identify them, before the recovery dialogue could proceed.

4.2 Evaluation #2

The goal of the second evaluation was to obtain information on user attitudes towards Tinker, and her user identification systems in particular, in the two cases of successful and unsuccessful identification.

The second evaluation was performed over the course of two weekdays, with a total time period of approximately 8 hours. Visitors were asked to interact with Tinker, after which a short semi-structured interview was conducted. Visitors were then asked to interact with Tinker a second time, followed by a longer semi-structured interview. In the case of children accompanied by a parent, both the child and parent were interviewed.

4.2.1 Participants

16 users interacted with Tinker and were interviewed. An additional 3 users did not complete a successful interaction due to software errors; these users all declined to be interviewed. Users ranged in age from 12 to 60, with an average age of 23. 62.5% of users were males. Most users (93.8%) arrived in a group of two or more.

4.2.2 Interactions

The duration of the first interaction ranged from 1 to 7 minutes (average 2.19). Second interactions were briefer than the first,

ranging from a few seconds (e.g., leaving immediately after the beginning of an interaction) to 3 minutes, with an average of 0.94 minutes.

56.3% of users were successfully re-identified by the biometric system. Most other users were identified through dialogue, giving a total success rate of 93.8%. This differs greatly from what was observed in the first evaluation, where few users were willing to interact with Tinker following a failure of the biometric system. We believe that this difference is due to visitors being asked to interact with Tinker by a researcher.

Those remaining left the interaction immediately after Tinker failed to identify them, before the recovery dialogue could proceed. In addition, two users had given names which were not in Tinker's database of names, and therefore their names were not used during dialogue. Consequently, both users did not realize that Tinker had re-identified them, although both were successfully identified by the biometric system.

4.2.3 Subjective Evaluation

Several users (31.3%) stated that they had liked being called by name when re-identified by Tinker. The most common reason was that it seemed more personal:

- "...not acting so much like a robot"
- "It was really nice... made it a little more personal"

One user gave a similar comment despite being identified by dialogue rather than biometrics. Only two users (28.6% of those recognized by dialogue) stated that they were disappointed that they had not been recognized automatically.

Most users (62.5%) liked the biometric system. Comments from users who had been successfully recognized were generally more positive, with "good" (40%) and "cool" (30%) being the most common words used to describe it. One user, who was identified neither by hand nor by dialogue, stated that he disliked the hand recognition system. All remaining users gave no opinion.

A minority of users (25%) expressed concerns about the privacy issues related to the use of biometric identification, although most added that they were not concerned about this particular application:

- "This is okay, but if that was being used on a daily basis, I'd be very concerned about my fingerprints being taken."
- "I'm assuming you erase the tapes every night?"

One of Tinker's dialogues answers questions about privacy issues, explaining that she only stores a small amount of information and erases it after a short period. Although several users saw this dialogue, only one commented on it, stating that he "didn't believe a word it said there." One user was annoyed that the system had not explained it was using biometrics before her interaction.

5. DISCUSSION

The results of the second evaluation illustrate the importance of Tinker's dialogue and relational behaviors to her identification capabilities, and vice versa. Most visitors interviewed had positive opinions about Tinker, and her identification capabilities in particular, even after she had failed to identify them using

biometrics. For two visitors, Tinker's dialogue that was intended to provide confirmation and grounding after recognizing the user's identity failed when Tinker did not use their names. From the point of view of these two visitors, it was the same as if the biometrics themselves had failed. In the minds of visitors, successful identification does not occur until their identity has been grounded in dialogue.

However, in actual usage, these points were obscured by the fact that most visitors left immediately after Tinker failed to recognize them by biometrics, and thus never saw the recovery dialogue. Users in this instance are showing less politeness than would be expected in human-human interaction; while it would be normal to be insulted if someone forgot your name, it would be unusual to simply walk away. This behavior remains a major practical issue with the system.

We have not yet identified a definitive cause for the poor performance of the biometric system relative to the predicted performance, but we offer two possible hypotheses. First, the training set may have been too small, and not sufficiently representative of the population of visitors within the museum. In particular, the number of children in the training set (29.6%) is lower than that observed during the evaluation (41.2%).

Second, the discrepancy may be caused by differences in user hand placement between the training set and actual usage. Our approach is designed to accommodate differences in how far apart fingers are placed, but may fail if two fingers are touching – it cannot reliably identify the boundary between them. Also, the system is sensitive to the amount of pressure applied with the hand. If less pressure is applied, then more light may leak under the hand, obscuring the silhouette. Nearly all the images of the training set are quite clean, indicating that users may have placed their hands more deliberately and carefully than is typical during actual usage.

Finally, we note that although most visitors generally claimed to be accepting of biometric identification in this particular application, and the agent is programmed to discuss the biometric technology and privacy issues, several visitors voiced concerns about privacy. Such issues should be considered in any further work in this area.

5.1 Future Work

Our immediate future work will focus on improving the success rates of the biometric identification system. A possible area of investigation is incorporating additional contextual information, particularly information from the dialogue, such as topics discussed, length of interaction, and type of farewell address (e.g., did the user give a polite farewell, or simply leave?).

In addition, we plan to investigate different strategies for recovering from errors in biometric identification, possibly including an empirical study of the effect of varying the specific agent error recovery utterances used across dimensions such as politeness.

6. CONCLUSION

The incorporation of biometric and other seamless identification technologies represent a potential advance in the user of relational agents in public spaces. Users are unwilling to tolerate

inaccuracies in identification, and thus any deployed systems must be robust, and provide mechanisms which allow recovery from error when possible. We believe that the use of conversational dialogue as a mechanism for acknowledgement, grounding, and error recovery is a promising direction of research in this area.

7. ACKNOWLEDGMENTS

Thanks to the current and past members of the Relational Agents Group who worked on the design and implementation of the Tinker exhibit: Laura Pfeifer, Sepalika Perera, Chaamari Senanayake, Ishraque Nazmi, Erin Rapaki, Ben Braggins, Jeff Satterly, Brian Sweeney, and Thomas Brown.

Thanks to Dan Noren, Taleen Agulian and the staff of ComputerPlace.

This work was supported by NSF CAREER IIS-0545932.

8. REFERENCES

- [1] Amayeh, G., Bebis, G., Erol, A., and Nicolescu, M. 2006. Peg-Free Hand Shape Verification Using High Order Zernike Moments. Proceedings of the 2006 Conference on Computer Vision and Pattern Recognition Workshop, pp. 40.
- [2] Bickmore, T. and Picard, R. 2005. Establishing and Maintaining Long-Term Human-Computer Relationships. ACM Transactions on Computer Human Interaction. 12, 2, 293-327.
- [3] Cassell, J., Stocky, T., Bickmore, T., Gao, Y., Nakano, Y., Ryokai, K., Tversky, D., Vaucelle, C., and Vilhjálmsón, H. 2002. MACK: Media lab Autonomous Conversational Kiosk. Imagina '02.
- [4] Cassell, J., Vilhjálmsón, H., and Bickmore, T. 2001. BEAT: The Behavior Expression Animation Toolkit. SIGGRAPH '01, pp. 477-486.
- [5] Clark, H. H. and Brennan, S. E. 1991. Grounding in Communication. In L. B. Resnick, J. M. Levine, and S. D. Teasley, Eds. Perspectives on Socially Shared Cognition. American Psychological Association, Washington, pp. 127-149.
- [6] Gockley, R., Bruce, A., Forlizzi, J., Michalowski, M., Mundell, A., Rosenthal, S., Sellner, B., Simmons, R., Snipes, K., Schultz, A. C., and Wang, J. 2005. Designing Robots for Long-Term Social Interaction. IEEE/RSJ International Conference on Intelligent Robots and Systems.
- [7] Goffman, I. 1967. On face-work. In Interaction Ritual: Essays on Face-to-Face Behavior. Pantheon, New York, pp. 5-46.
- [8] Hone, K. 2006. Empathic agents to reduce user frustration: The effects of varying agent characteristics. Interacting with Computers. 18, 2, 227-245.
- [9] Jain, A. K., Ross, A., and Pankanti, S. 1999. A prototype hand geometry-based verification system. Proc. of 2nd Int'l Conference on Audio- and Video-based Biometric Person Authentication (AVBPA), pp. 166-171.
- [10] Jain, A. K., Ross, A., and Prabhakar, S. 2004. An introduction to biometric recognition. Circuits and Systems for Video Technology, IEEE Transactions on. 14, 1, 4-20.
- [11] Klein, J., Moon, Y., and Picard, R. 2002. This Computer Responds to User Frustration: Theory, Design, Results, and Implications. Interacting with Computers. 14, 119-140.
- [12] Kopp, S., Gesellensetter, L., Krämer, N. C., and Wachsmuth, I. 2005. A conversational agent as museum guide: design and evaluation of a real-world application. In Intelligent Virtual Agents. pp. 329-343.
- [13] Lehm, D. v., Hindmarsh, J., Luff, P., and Heath, C. 2007. Engaging constable: revealing art with new technology. Proceedings of the SIGCHI conference on Human factors in computing systems, pp. 1485-1494.
- [14] Platt, J. C. 1999. Probabilistic outputs for support vector machines and comparison to regularized likelihood methods. In A. J. Smola, Bartlett, P., Schoelkopf, B., Schuurmans, D., Ed., Advances in Large Margin Classifiers. MIT Press, Cambridge, MA, pp. 61-74.
- [15] Shiomi, M., Kanda, T., Ishiguro, H., and Hagita, N. 2006. Interactive Humanoid Robots for a Science Museum. HRI'06.
- [16] Thom-Santelli, J., Boehner, K., Gay, G., and Hembrooke, H. 2006. Beyond just the facts: transforming the museum learning experience. CHI '06 extended abstracts on Human factors in computing systems, pp. 1433-1438.
- [17] Vapnik, V. N. 1998. Statistical Learning Theory. Wiley-Interscience, New York.
- [18] Wong, A. L. and Shi, P. 2002. Peg-free hand geometry recognition using hierarchical geometry and shape matching. IAPR Workshop on Machine Vision Applications, pp. 281-284.