

Engagement Rules for Human-Robot Collaborative Interaction

Candace L. Sidner, Christopher Lee, Neal Lesh

TR2003-50 May 2003

Abstract

This paper reports on research on developing the ability for robots to engage with humans in a collaborative conversation for hosting activities. It defines the engagement process in collaborative conversation and reports on our progress in creating a robot to perform hosting activities. The paper then presents the analysis of a study of human-human look tracking and discusses rules that will allow a robot to track humans so that engagement is maintained.

Proceedings of the IEEE Conference on Systems, Man and Cybernetics, 2003

This work may not be copied or reproduced in whole or in part for any commercial purpose. Permission to copy in whole or in part without payment of fee is granted for nonprofit educational and research purposes provided that all such whole or partial copies include the following: a notice that such copying is by permission of Mitsubishi Electric Research Laboratories, Inc.; an acknowledgment of the authors and individual contributions to the work; and all applicable portions of the copyright notice. Copying, reproduction, or republishing for any other purpose shall require a license with payment of fee to Mitsubishi Electric Research Laboratories, Inc. All rights reserved.

Engagement Rules for Human-Robot Collaborative Interactions

Candace L. Sidner

Christopher Lee

Neal Lesh

Mitsubishi Electric Research Laboratories

201 Broadway

Cambridge, MA 02139

{[@merl.com](mailto:Sidner, Lee, Lesh)}

Abstract - *This paper reports on research on developing the ability for robots to engage with humans in a collaborative conversation for hosting activities. It defines the engagement process in collaborative conversation, and reports on our progress in creating a robot to perform hosting activities.. The paper then presents the analysis of a study of human-human “look tracking” and discusses rules that will allow a robot to track humans so that engagement is maintained.*

Keywords: Human-robot interaction, engagement, conversation, collaboration, collaborative interface agents, gestures in conversation.

1. Introduction

This paper reports on our research toward developing the ability for robots to engage with humans in a collaborative conversation for hosting activities. Engagement is the process by which two (or more) participants establish, maintain and end their perceived connection during interactions they jointly undertake. Engagement is supported by the use of conversation (that is, spoken linguistic behavior), ability to collaborate on a task (that is, collaborative behavior), and gestural behavior that conveys connection between the participants. While it might seem that conversational utterances alone are enough to convey connectedness (as is the case on the telephone), gestural behavior in face-to-face conversation conveys much about the connection between the participants [1].

Collaborative conversations cover a vast range of interactions activities from call centers to auto repair to physician-patient dialogues. In order to narrow our research efforts, we have focused on hosting activities. Hosting activities are a class of collaborative activity in which an agent provides guidance in the form of information, entertainment, education or other services in the user's environment and may also request that the user undertake actions to support the fulfillment of those services. Hosting activities are situated or embedded activities, because they depend on the surrounding environment as well as the participants involved. They are social activities because, when undertaken by humans, they depend upon the social roles that people play to determine the choice of the next actions, timing of those actions, and negotiation about the choice of actions. In our research, physical robots, who serve as guides, are the hosts of the environment.

Conversational gestures generally concern gaze at/away from the conversational partner, pointing behaviors, (bodily) addressing the conversational participant and other persons/objects in the environment, all in appropriate synchronization with the conversational, collaborative behavior. These engagement gestures are culturally determined, but every culture has some set of behaviors to accomplish the engagement task. These gestures are not about where the eye camera gets its input from, but what the eyes, hands, and body tell conversational participants about their interaction.

Not only must the robot produce these behaviors, but it must interpret similar behaviors from its collaborative partner (hereafter CP). Proper gestures by the robot and correct interpretation of human gestures dramatically affect the success of conversation and collaboration. Inappropriate behaviors can cause humans and robots to misinterpret each other's intentions. For example, a robot might look away for an extended period of time from the human, a signal to the human that it wishes to disengage from the conversation and could thereby terminate the collaboration unnecessarily. Incorrect recognition of the human's behaviors can lead the robot to press on with a conversation in which the human no longer wants to participate.

While other researchers in robotics are exploring aspects of gesture (for example, [2], [3]), none of them have attempted to model human-robot interaction to the degree that involves the numerous aspects of engagement and collaborative conversation that we have set out above. Robotics researchers interested in collaboration and dialogue [4] have not based their work

on extensive theoretical research on collaboration and conversation, as we will detail later. Our work is also not focused on emotive interactions, in contrast to Breazeal among others. For 2D conversational agents, researchers (notably, [5],[6]) have explored agents that produce gestures in conversation. However, they have not tried to incorporate recognition as well as production of these gestures, nor have they focused on the full range of these behaviors to accomplish the maintenance of engagement in conversation.

In this paper we discuss our research agenda for creating a robot with collaborative conversational abilities, including gestural capabilities in the area of hosting activities. We will also discuss the results of a study of human-human hosting and how we are using the results of that study to determine rules and associated algorithm for the engagement process in hosting activities. We will also critique our current engagement rules, and discuss how our study results might improve our robot's future behavior.

2. Communicative capabilities for collaborative robots

To create a robot that can converse, collaborate, and engage with a human interactor, a number of different communicative capabilities must be included in the robot's repertoire. Most of these capabilities are linguistic, but some make use of physical gestures as well. These capabilities are:

- Engagement behaviors: initiate, maintain or disengage in interaction;
- Conversation management: turn taking [7], interpreting the intentions of the conversational participants, establishing the relations between intentions and goals of the participants and relating utterances to the attentional state [8] of the conversation.
- Collaboration behavior: choosing what to say or do next in the conversation, to foster the shared collaborative goals of the human and robot, as well as how to interpret the human's contribution (either spoken acts or physical ones) to the conversation and the collaboration.

Turn taking gestures also serve to indicate engagement because the overall choice to take the turn is indicative of engagement, and because turn taking involves gaze/glance gestures. There are also gestures in the conversation (such as beat gestures, which are used to indicate old and new information [9,10]) that CPs produce and observe in their partners. These capabilities are significant to robotic participation in conversation because they allow the robot to communicate using the same strategies and techniques that are normal for humans, so that humans can quickly perceive the robot's communication.

Humans do not turn off their own conversational and engagement capabilities when talking to robots. Hence robots can and must make use of these capabilities to successfully communicate with humans. To use human linguistic and gestural behavior, a robot must fuse data gathered from its visual and auditory sensors to determine the human gestures, and infer what the human intends to convey with these gestures.

In our current work, we have developed (1) a model of engagement for human-human interaction and adapted it for human-robot interaction, and (2) an architecture for a collaborative non-mobile conversational robot.

Our engagement model describes an engagement process in three parts, (1) initiating a collaborative interaction with another, (2) maintaining the interaction through conversation, gestures, and, sometimes, physical activities, and (3) disengaging, either by abruptly ending the interaction or by more gradual activities upon completion of the goals of the interaction. The rules of engagement, which operate within this model, provide choices to a decision-making algorithm for our robot about what gestures and utterances to produce.

Our robot, which looks like a penguin, as shown in Figure 1, uses its head, wings and beak for gestures that help manage the conversation and also express engagement with its human interlocutor (3 DOF in head/beak, 2 in wings). While the robot can only converse with one person at a time, it performs gestures to acknowledge the onlookers to the conversation. Gaze for our robot is determined by the position of its head, since its eyes do not move. Since our robot cannot turn its whole body, it does not make use of rules we have already created concerning addressing with the body. Because bodily addressing (in US culture) is a strong signal for whom a CP considers the main other CP, change of body position is a significant engagement signal. However, we will be mobilizing our robot in the near future and expect to test these rules following that addition.

To create an architecture for collaborative interactions, we use several different systems, largely developed at MERL. The conversational and collaborative capabilities of our robot are provided by the CollagenTM middleware for collaborative

agents [8,11,12,13] and commercially available speech recognition software (IBM ViaVoice). We use a face detection algorithm [14], a sound location algorithm, and an object recognition algorithm [15] and fuse the sensory data before passing results to the Collagen system. The robot's motor control algorithms use the engagement rule decisions and the conversational state to decide what gestures to perform. Further details about the architecture and current implementation can be found in [16].



Figure 1: Mel, the penguin robot

3. Current engagement capabilities

The greatest challenge in our work on engagement is determining rules governing the maintenance of engagement. Our first set of rules, which we have tested in scenarios as described in [16], are a small and relatively simple set. The test scenarios do not involve pointing to or manipulating objects. Rather they are focused on engagement in more basic conversation. These rules direct the robot to initiate engagement with gaze and conversational greetings. For maintaining engagement, gaze at the speaking CP signals engagement, when speaking, gaze at both the human interlocutor and onlookers maintains engagement while gaze away for the purpose of taking a turn does not signal disengagement. Disengagement from the interaction is understood as occurring when a CP fails to take an expected turn together with loss of the face of the human. In those cases where the human stays engaged until the robot has run out of things to say, the robot closes the conversation using known rules of conversational closing [17,18].

While this is a fairly small repertoire of engagement behaviors, it was sufficient to test the robot's behavior in a number of scenarios involving a single CP and robot, with and without onlookers. Much of the robot's behavior is quite natural. However, we have observed oddities in its gaze at the end of its turn (it will incorrectly look at an onlooker instead of its CP when ending its turn) as well as confusion about where to look when the CP leaves the scene.

These conversations have one drawback: they are of the how-are-you-and-welcome-to-our-lab format. However, our goal is hosting conversations, which involve much more activity. While other researchers have made considerable progress on the navigation involved for a robot to host visitors [e.g. 19] and gestures needed to begin conversation [e.g. 20], many aspects of conversation are open for investigation. These include extended explanations, pointing at objects, manipulating them (on the part of the humans or robots), moving around in a physical environment to access objects and interact with them. This extended repertoire of tasks requires many more gestures than our current set. In addition, some of these gestures could alternatively be taken to indicate that the human participant is disengaging from the conversation (e.g. looking away from the human speaker). So many behaviors appear to be contextually sensitive.

To explore these conversations, we have provided our robot with some additional gestural rules and new recipes for action (in the Collagen framework), so that our penguin robot now undertake a hosting conversation to demo an invention created at MERL. Mel greets a visitor (other visitors can also be present), convinces the visitor to participate in a demo, and proceeds to show the visitor the invention. Mel points to demo objects (a kind of electronic cup sitting on a table), talks (with speech) the visitor through the use of the cup, asks the visitor questions, and interprets the spoken answers, and includes the onlookers in its comments. The robot also expects the visitor to say and do certain activities, as well as gaze at objects, and will await or insist on such gestures if they are not performed. The entire interaction lasts about five minutes (a video will

be available). Not all of Mel's behaviors in this interaction appear acceptable to us. For example, Mel often looks away for too long (at the cup and table) when explaining them, it (Mel is "it" since it is not human) fails to make sure it's looking at the visitor when it calls the visitor by name, and it sometimes fails to look for a long enough when it turns to look at objects. To make Mel perform more effectively, we are investigating gesture in human-human interactions.

4. Evidence for engagement in human behavior

Much of the available literature on gestures in conversation (e.g. [7], [21]) provides a base for determining what gestures to consider, but does not provide enough detail about how gestures are used to maintain conversational engagement, that is, to signal that the participants are interested in what the other has to say and in keeping the interaction going.

The significance of gestures for human-robot interaction can be understood by considering the choices that the robot has at every point in the conversation for its head movement, its gaze, and its use of pointing. The robot must also determine whether the CP has changed its head position or gaze and what objects the CP points to or manipulates. Head position and gaze are indicators of engagement. Looking at the speaking CP is evidence of engagement, while looking around that room, for more than very brief moments, is evidence of disinterest in the interaction and possibly the desire to disengage. However, looking at objects relevant to the conversation is not evidence of disengagement. Furthermore, the robot needs to know that the visitor has or has not paid attention to what it points at or looks at. If visitor fails to look at what the robot looks at, the visitor might miss something crucial to the interaction.

A simple hypothesis for maintaining engagement (for each CP) is: Do what the speaking CP does: look wherever the CP looks, look at him if he looks at you, and look at whatever objects are relevant to the discussion when he does. When the robot is the speaking CP, this hypothesis means it will expect perfect tracking of its looking by the human interlocutor. The hypothesis does not constrain the CP's decision choices for what to look and point at.

Note that there is evidence that the type of utterances that occur in conversation affect the gaze of the non-speaking CP. [22] provides evidence that in direction giving tasks, the non-speaking CP will gaze more often at the speaking CP when the speaking CP's utterance pairs are assertion followed by elaboration, and more often at a map when the utterance pairs are assertion followed by the next map direction.

To evaluate the simple hypothesis for engagement, we have been evaluating interactions in videotapes of human-human hosting activities, which were recorded in our laboratory. In these interactions, a member of our lab hosted a visitor who was shown various new inventions and computer software systems. The host and visitor were followed by video camera as they experienced a typical tour of our lab and its demos. The host and visitor were not given instructions to do anything except to give/take the lab tour. We have obtained about 3.5 hours of video, with three visitors, each on separate occasions being given a tour by the same host. We have transcribed portions of the video for all the utterances made and the gestures (head, hands, body position, body addressing) that occur during portions of the video. We have not transcribed facial gestures. We report here on our results in observing gestural data (and its corresponding linguistic data) for just over five minutes of one of the host-visitor pairs.

The subject of the investigated portion of the video is a demonstration of an "Iglassware" cup which P (a male) demos and explains to C (a female). P produces a gestures, with his hands, face, and body, gazes at C and other objects. He also point to the cup, holds it and interacts with a table to which the cup transfers data. He uses his hands to produce iconic and metaphorical gestures [10], and he uses the cup as a presentation device as well. We do not discuss iconic, metaphorical or presentation gestures, in large part because our robot does not have hands with which to perform similar actions.

We report here on C's tracking of P's gaze gestures (since P speaks the overwhelming majority of the utterances in their interaction). Gaze in this analysis is expressed in terms of head movements (looking). We did not code eye movements due to video quality.

There are 81 occasions on which P changes his gaze by moving his head with respect to C. Seven additional gaze changes occurred that are not counted in this analysis because it is unclear to where P changed his gaze. Of the counted gaze changes, C tracks 45 of them (55%). The remaining failures to track gaze (36, or 45% of all gazes) can be subclassed into 3 groups: quick looks, nods (glances followed by gestural or verbal feedback), and uncategorized failures (see Table 1).

These tracking failures indicate that our simple hypothesis for maintaining engagement is incorrect. Of these tracking failures, the *nod failures* can be explained because they occur when P looks at C even though C is looking at something else (usually the cup or the table, depending on which is relevant to the discussion). However, in all these cases, C immediately nods and often articulates with "Mm-hm," "Wow" or other phrases to indicate that she is following her conversational part-

ner. In grounding terms [23], P is attempting to ascertain by looking at C that she is following his utterances and actions. When C cannot look, she provides feedback by nods and comments. She grounds his comments and thereby indicates that she is still engaged. In the nod cases, she also is not just looking around the room, but paying attention to an object that is highly relevant to the demo.

	Count	% of tracking failures	% of total tracking
Quick looks	11	31	14
Nods	13	36	16
Uncategorized	12	33	15

Table 1: Failures to Track Gaze Change

The quick looks and uncategorized failures represent 64% of the failures and 28% of the gaze changes in total. Closer study of these cases reveals significant information for our robot vision detection algorithms.

In the *quick look* cases, P looks quickly (including moving his head) at something without C tracking his behavior. Is there a reason to think that C is not paying attention or has lost interest? P never stops to check for her feedback in these cases. Has C lost interest or is she merely not required to follow? In all of these cases, C does not track P for the simple reason that P's changes in looking (usually to the cup or the table, but occasionally to something else he is commenting on) happen very quickly. It's a quick up-and-down or down-and-up head movement. There simply is insufficient time for C to do so. In every case C is paying attention to P, so C is able to see what's happening. It appears she judges on the velocity of the movement not to track it.

Of the *uncategorized failures*, the majority (8 instances) occur when C has other actions or goals to undertake. For example, C may be finishing a nod and not be able to track P while she's nodding. One uncategorized failure is similar to a quick look but longer in duration. Of the remaining 3 tracking failures, each occurs for seemingly good reasons to us as observers, but may not be known at the time of occurrence. For example, one failure occurs at the start of the demo when C is looking at the new (to her) object that P displays and does not track P when he looks up at her.

This data clearly indicates that our rules for maintaining engagement must be more complex than the "do whatever the speaking CP does" hypothesis. In fact, tracking is critical to the interaction. It is not completely essential because very quick glances away can be disregarded, and one can avoid tracking the interlocutor by providing verbal feedback. In those cases where arguably one should track the speaking CP, failure to do so may lead the CP to restate an utterance or to just go on because the lost information is not critical to the interaction. The data in fact suggest that for our robot engagement rules, tracking the speaking CP in general is the best move, but when another goal interferes, providing verbal feedback maintains engagement. Furthermore, the robot as tracker can ignore head movements of brief duration, as these quick looks do not need tracking.

5. Future directions

An expanded set of rules for engagement requires a means to evaluate them. We want to use the standard training and testing set paradigm common in computational linguistics and speech processing, but training and test sets are hard to come by for interaction with robots. Our solution has been to approach the process with a mix of techniques. We have developed a graphical display of simulated, animated robots running our engagement rules. We plan to observe the interactions in the graphic display for a number of scenarios to check and tune our engagement rules. Following that, we will undertake something akin to a testing set by testing the robot's interactions with human visitors as the robot demonstrates the Iglassware cup.

6. Summary

This paper has discussed the nature of engagement in human-robot interaction, and outlined our methods for investigating rules for engagement for the robot. We report on analysis of human-human look tracking where the humans do not always track the changes in looks by their conversational interlocutors. We conclude that such tracking failures indicate both the default behavior for a robot and when it can fail to track without its human conversational partner inferring that it wishes to disengage from the interaction.

7. Acknowledgements

The authors wish to acknowledge the work of Charles Rich on aspects of Collagen critical to this effort.

8. References

1. D. McNeill. *Hand and Mind: What Gestures Reveal about Thought*. University of Chicago Press, Chicago, 1992.
2. C. Breazeal "Affective interaction between humans and robots", *Proceedings of the 2001 European Conference on Artificial Life (ECAL2001)*. Prague, Czech Republic, (2001).
3. T. Kanda, H. Ishiguro, M. Imai, T. Ono, and K. Mase. "A constructive approach for developing interactive humanoid robots. *Proceedings of IROS 2002*, IEEE Press, NY, 2002.
4. T. Fong, C. Thorpe, and C. Baur, Collaboration, Dialogue and Human-Robot Interaction, *10th International Symposium of Robotics Research*, Lorne, Victoria, Australia, November, 2001.
5. J. Cassell, J. Sullivan, S. Prevost and E. Churchill, *Embodied Conversational Agents*. MIT Press, Cambridge, MA, 2000.
6. W.L. Johnson, J.W. Rickel, and J.C. Lester. "Animated Pedagogical Agents: Face-to-Face Interaction in Interactive Learning Environments," *International Journal of Artificial Intelligence in Education*, 11: 47-78, 2000.
7. S. Duncan. Some signals and rules for taking speaking turns in conversation. in *Nonverbal Communication*, S. Weitz (ed.), New York: Oxford University Press, 1974
8. B. J. Grosz and C.L. Sidner. Attention, intentions, and the structure of discourse. *Computational Linguistics*, 12(3): 175--204, 1986.
9. M.A.K. Halliday. *Explorations in the Functions of Language*, London: Edward Arnold, 1973.
10. J. Cassell. Nudge nudge wink wink: Elements of face-to-face conversation for embodied conversational agents. in *Embodied Conversational Agents*, J. Cassell, J. Sullivan, S. Prevost, and E. Churchill (eds.), Cambridge, MA: MIT Press, 2000.
11. C. Rich, C.L. Sidner, and N. Lesh. "COLLAGEN: Applying Collaborative Discourse Theory to Human-Computer Interaction," *AI Magazine, Special Issue on Intelligent User Interfaces*, AAAI Press, Menlo Park, CA, Vol. 22: 4: 15-25, 2001.
12. C. Rich and C.L. Sidner. "COLLAGEN: A Collaboration Manager for Software Interface Agents," *User Modeling and User-Adapted Interaction*, Vol. 8, No. 3/4, 1998, pp. 315-350, 1998.
13. K.E. Lochbaum. A Collaborative Planning Model of Intentional Structure. *Computational Linguistics*, 24(4): 525-572, 1998.
14. Viola, P. and Jones, M. Rapid Object Detection Using a Boosted Cascade of Simple Features, *IEEE Conference on Computer Vision and Pattern Recognition*, Hawaii, pp. 905-910, 2001.
15. Beardsley, P.A. *Piecode Detection*, Mitsubishi Electric Research Labs TR2003-11, Cambridge, MA, February, 2003.
16. C.L. Sidner and C. Lee, *An Architecture for Engagement in Collaborative Conversations between a Robot and Humans*, Mitsubishi Electric Research Labs TR2003-13, Cambridge, MA, April, 2003, also in submission.
17. E. Schegeloff, & H. Sacks. Opening up closing. *Semiotica*, 7(4): 289-327, 1973.
18. H.H. Luger. "Some Aspects of Ritual Communication," *Journal of Pragmatics*. Vol. 7: 695-711, 1983.
19. Burgard, W., Cremes, A. B., Fox, D., Haehnel, D., Lakemeyer, G., Schulz, D., Steiner, W. & Thrun, S. "The Interactive Museum Tour Guide Robot," *Proceedings of AAAI-98*, 11-18, AAAI Press, Menlo Park, CA, 1998.
20. A. Bruce, I. Nourbakhsh, R. Simmons. "The Role of Expressiveness and Attention in Human Robot Interaction," In *Proceedings of the IEEE International Conference on Robotics and Automation*, Washington DC, May 2002.
21. A. Kendon. Some functions of gaze direction in social interaction. *Acta Psychologica*, 26: 22-63, 1967.
22. Y. Nikano, G. Reinstein, T. Stocky, J. Cassell. *Towards a Model of Face-to-Face Grounding*, accepted for ACL, 2003.
23. H.H. Clark, *Using Language*, Cambridge University Press, Cambridge, 1996.