Small Talk and Conversational Storytelling In Embodied Conversational Interface Agents

T. Bickmore, J. Cassell

Gesture and Narrative Language Group
MIT Media Laboratory
E15-315
20 Ames St, Cambridge, Massachusetts
+1 617 253 4899
{bickmore, justine}@media.mit.edu

Abstract

People engage in small talk and conversational storytelling to provide information in engaging ways and to serve interpersonal goals such as building rapport and credibility. Embodied conversational interface agents, that represent the computer in its interaction with human users by way of a graphical humanoid body on a screen, can also profitably engage in these behaviors. We describe the ongoing development of an embodied conversational interface agent that is capable of multimodal input understanding and output generation and operates in a limited application domain in which both social and task-oriented dialogue are important. We discuss our plans for adding small talk and conversational storytelling capabilities to the system, and the unique requirements these phenomena place on the understanding, discourse planning, and generation components of a real-time conversational interface.

Introduction

Interlocutors typically have multiple goals in conversation (Tracy and Coupland 1991). In task-oriented interactions, speakers may also have several interpersonal goals they are pursuing, such as developing a relationship (e.g., befriending, earning trust), establishing their reputations or expertise, or persuading their listeners to take some course of action. Even in casual small talk, where there are apparently no task goals being pursued, interlocutors are conscious of multiple goals related to conversation initiation, regulation and maintenance (Cegala, et al, 1988). In addition, there are many persistent goals which all speakers pursue to varying degrees, such as avoiding "face threats" (Goffman 1983) and being cooperative (Grice 1989).

We are currently developing an embodied conversational interface agent who performs the role of a real-estate salesperson. Our general interest in embodied conversational interface agents concerns the relationship between verbal and nonverbal behaviors in conversation, and between propositional and interactional goals in conversation. Real estate sales was selected as an application domain exactly because a significant portion of the interaction between the agent and a buyer is interpersonal in nature. Within this domain, we are currently modeling the initial buyer interview by studying videotapes of real agent-buyer interactions. In the initial buyer interview, the agent pursues several goals in parallel, including:

- The task goals of determining clients' housing preferences (size, location, style, etc.) and buying ability.
- The interpersonal goal of establishing rapport with clients (Hopkins 1982) and reducing their fear.
- The interpersonal goal of establishing the agent's expertise and credibility.

The interpersonal goals in this interaction are partly achieved through small talk and conversational storytelling. Small talk is non-task-oriented conversation about neutral topics whose primary functions are to mitigate face threats, provide an initial time interval in which the interlocutors can "size each other up," establish an interactional style, and establish some degree of mutual trust and rapport. Conversational storytelling is the spontaneous generation of a narrative embedded within face-to-face interaction and it is found both during periods of small talk and during task-oriented dialogue as a way of presenting task-oriented information in an engaging way, and/or meeting one or more of the goals stated above. From a functional perspective, small talk can best be thought of as a "frame" of interaction (Goffman 1981) which is used to satisfy the strategic goal of building rapport, while conversational storytelling can be characterized as a linguistic device which can be used to organize and relate information in a narrative form within any interactional frame.

We are interested in modeling these phenomena for several reasons. First, small talk and conversational storytelling occur in natural real-estate sales interactions and thus should be modeled in order to approximate the behavior of a real-estate salesperson in these situations. Second, building a generative model of small talk and storytelling as embedded within a service encounter helps us better understand these phenomena from a linguistic and computational linguistic perspective and will help to uncover deeper theoretical issues in constructing systems capable of real-time dialogue that is not strictly task-oriented. Finally, we feel that these behaviors do play a crucial role in the real-estate sales process and thus must be incorporated into a system that is to perform this role competently.

We begin this paper by presenting our development platform for building embodied conversational interface agents along with the current application domain of real-estate sales. We then describe small talk and conversational storytelling from a linguistic perspective, and discuss the requirements these capabilities place on a real-time conversational interface. Finally, we discuss some of the specific challenges we face in adding these capabilities to our system.

REA: An Embodied conversational interface agent

In conversation one of the primary purposes of rapport-building behaviors is to establish that our interlocutor is, in some way, "just like us," that is, shares our basic values and beliefs, as well as sharing the ability to communicate (Laver 1981). One way for a software agent to overcome the credibility problem is to embody the agent in the interface, that is, provide it with a human appearance and lifelike behavior. It has been shown that people tend to treat computers as social agents anyway (Reeves and Nass, 1996), but providing an agent with an embodied presence provides a strong cue to users that the agent should be treated and interacted with as another person. To further support this image of computer-as-conversational-partner, we feel it is crucial that user and computer be able to use the face, hands, and melody of the voice in interaction to regulate the process of their conversation, as well as the ability to use verbal and nonverbal means to contribute content to the ongoing conversation (Cassell et al, 1999). In one study of personified interfaces, characters with realistic human faces were rated over characters with caricatures and animal faces as the "most intelligent," even though the behavior models underlying the different characters was the same (Koda and Maes, 1996).

The Rea project at the MIT Media Lab has as its goal the construction of an embodied, multi-modal real-time conversational interface agent. Rea implements the social, linguistic, and psychological conventions of conversation to make interactions as natural as face-to-face conversation with another person. Rea differs from other dialogue systems, and other interface agents in three ways:

- Rea has a human-like body, and uses her body in human-like ways during the conversation. That is, she uses eye gaze, body posture, hand gestures, and facial displays to organize and regulate the conversation.
- The underlying approach to conversational understanding and generation in Rea is based on discourse functions. Thus, each of the user's inputs are interpreted in terms of its conversational function and responses are generated according to the desired function to be fulfilled. (Cassell, et. al., 1999).
- Rea is able to respond to visual, audio and speech cues normally used in face to face conversation, such as speech, shifts in gaze, gesture, and non-speech audio (feedback sounds). She is also able to generate these cues, ensuring a full symmetry between input and output modalities.

Rea has a fully articulated graphical body, can sense the user passively through cameras and audio input, and is capable of speech with intonation, facial display, and gestural output. The system currently consists of a large projection screen on which Rea is displayed and which the user stands in front of (see Figure 1). Two cameras mounted on top of the projection screen track the user's head and hand positions in space. Users wear a microphone for capturing speech input. A single SGI Octane computer runs the graphics and conversation engine of Rea, while several other computers manage the speech recognition and generation and image processing.

Figure 1. User interacting with Rea

Rea is able to describe the features of the task domain while also responding to the users' verbal and non-verbal input. When the user makes cues typically associated with turn taking behavior such as gesturing, Rea allows herself to be interrupted, and then takes the turn again when she is able. She is able to initiate conversational repair when she misunderstands what the user says, and can generate combined voice and gestural output. Rea's responses are generated by an incremental natural language generation engine based on (Stone, 1998) that has been extended to synthesize redundant and complementary conversational gestures synchronized with speech output.

Rea's task domain is real estate sales (REA is an acronym for "Real Estate Agent"), and within this domain we are currently focused on modeling the initial interview with a prospective buyer. In a typical interview, we feel it is important for Rea to engage in some amount of small talk to build rapport with the user before going into the "sales" conversational frame. The rapport may convince users that Rea does in fact understand their communicative behaviors, as well as convincing them that she knows what she is talking about in the real estate domain.

Small Talk

Human social encounters between individuals who have never met or are unfamiliar with each other are usually initiated by small talk in which "light" conversation is made about neutral topics (e.g., weather, aspects of the interlocutor's physical environment) or in which personal experiences, preferences, and opinions are shared (Laver 1981). Even in business or sales meetings, it is customary to begin with some amount of small talk before "getting down to business" (at least in American culture). The purpose of this small talk is primarily to build rapport and trust among the interlocutors, provide time for them to "size each other up", establish an interactional style, and to allow them to establish their reputations (Dunbar 1996). Although small talk is most noticeable at the margins of conversational encounters, it can be used at various points in the interaction to continue to build rapport and trust (Cheepen 1988), and in real estate sales, a good agent will continue to focus on building rapport throughout the relationship with a buyer (Garros 1999).

Small talk has received sporadic treatment in the linguistics literature, starting with the seminal work of Malinowski (1923) who defined "phatic communion" as "a type of speech in which ties of union are created by a mere exchange of words". Small talk is the language used in free, aimless social intercourse, which occurs when people are relaxing or when they are accompanying "some manual work by gossip quite unconnected with what they are doing." Jacobson (1960) also included a "phatic function" in his well-known conduit model of communication, which is focused on the regulation of the conduit itself (as opposed to the message, sender, or receiver). More recent work has further characterized small talk by describing the contexts in which it occurs, topics typically used, and even grammars which define its surface form in certain domains (Laver 1975, Cheepen 1988, Schneider 1988). In addition, degree of "phaticity" has been proposed as a persistent goal which governs the degree of politeness in all utterances a speaker makes, including task-oriented ones (Coupland and Coupland 1992).

Phatic communion is closely related to the notion of "face" (Goffman 1983); "positive face" is the desire of all speakers to be approved of by their listeners, while "negative face" is the desire of all speakers to be unobstructed in their autonomy. Small talk mitigates positive face threats by providing an interactional style in which it is very easy (even somewhat obligatory) for all interlocutors to carry on a conversation and thereby achieve some degree of cameraderie. It can also be used to mitigate negative face threats by establishing that one's interlocutors are non-hostile (e.g., as used to break uneasy silences in waiting rooms).

The effectiveness of small talk on the user's perception of a software agent has been demonstrated in text-only "chatterbot" systems such as ELIZA (Weizenbaum 1966). There are many documented cases in which users have established on-going relationships with such systems, even when they knew that the system did not understand what they were typing (Turkle 1995).

In the real-estate domain, we plan to utilize small talk to establish trust and rapport between the Rea and the user in such a way that the user relaxes enough to use normal speech and nonverbal behaviors, and also feels comfortable disclosing personal information (such as salary). Through the introduction of a metric of interpersonal distance which varies from "strangers" to "intimates", the discourse planner can actually plan to conduct small talk by adding a requisite level of interpersonal distance to the preconditions for asking particular questions. For example, before Rea asks the user invasive questions about financial status (a negative face threat) she should first establish some level of rapport, and this rapport might be achieved by conducting some amount of small talk.

Conversational Storytelling

It would not be appropriate for Rea to tell tall tales, or spin fantastic yarns. For this reason, we follow the literature on conversational storytelling, and consider stories to be "specific, affirmative, past time narratives which tell about a series of events which did take place at specific unique moments in a unique past time world" and which are told to others to make a point or transmit a message (Polanyi 1985). In casual conversation, stories are often used to relate interesting events or humor and thus contribute to the rapport-building function.

We believe conversational storytelling should be an integral part of an agent's small talk capability. (Cheepen 1988, actually characterizes most of the exchanges in small talk as story contributions.) One of the tricks that chatterbot programs use to convey the illusion of intelligence is to tell stories whenever possible in the course of the conversation (Mauldin 1994), a technique that is also espoused in several popular press books on conversational skills (e.g., RoAne 1997). There are two possible reasons why this is successful. First, people try to establish relevancy relationships among conversational contributions whenever possible and multi-utterance narratives may give listeners more material to work with. Second, a story contribution that is irrelevant to an on-going conversation may be tolerated if the entertainment or information content is high enough. This may be why most books on conversational skills suggest that speakers should always have several interesting stories ready to contribute to any social occasion.

As with small talk, storytelling can be used to achieve multiple conversational goals simultaneously. In the real estate domain, it can be used to achieve any combination of the following goals:

- Phatic - Simply having the agent deliver a story meets the objectives of phatic communion by keeping the conversation going.
- Establishing expertise - As with the phatic goal, the successful delivery of any story helps to establish the agent's linguistic expertise and intelligence. In

addition, specific stories can help establish the agent's expertise in specific areas, such as: "I sold a house to someone just like you."; "I sold a house in the area you are interested in."; or "I sold 25 houses last month."

- Encouraging self-disclosure - It has been demonstrated that users are more willing to disclose personal information to a computer which has just disclosed similar information to them (Moon 1998). For example: "I just bought a wonderful home in Cambridge. Where are you looking to buy?"
- Persuasion/Problem solving - The agent can relate a story about a situation similar to the client's in which one of the client's problems is solved. For example: "I sold a house to someone who used Acme Mortgage Company and they closed escrow in only two weeks."
- Providing requested information - Stories give the agent the opportunity to answer a client's question while also achieving other goals. For example, in response to the question "Are there any new homes available in Boston?", the agent might respond "I just sold a brand new Foobar Development home to a couple out in Waltham."

Spontaneously generated, conversational stories differ from prepared narratives in several significant ways. Conversational stories are typically not delivered as monologues, but rather intimately involve the listeners in the production through their elicitation, feedback, requests for elaboration, and attempt to show relevance. Conversational stories must be *locally occasioned* and *recipient designed* (Sacks 1970-71). Speakers must take care to tell stories which are relevant to their listeners (locally occasioned) otherwise they suffer a loss of face due to their wasting the time of their audience (Labov, 1972). One general rule of story relevance presented by Polanyi (1982) is that whatever is "close" to the listener (in terms of space, time or relationship) is relevant. In the real-estate domain, for example, stories which relate to the area the client is looking to buy in, the style of home they are looking for, or the school district they are interested in would all be appropriate to tell. In addition to this strategic relevance requirement, stories must also be told at an appropriate point in a given social interaction, with the storyteller constructing the point of the story so that it relates directly to what is being discussed when it is introduced.

The problem of deciding which story to tell at any given time in a conversation has been addressed by several researchers. Gough (1990) analyzed the production of Xhosa folk narrative which is synthesized during a performance by weaving together several story fragments that the narrator has memorized. Gough identified two mechanisms which were used to index and modify story fragments based on relevancy relationships. Computational systems for relevant story indexing have also been developed using keywords (Bers & Cassell, 1998) and case-based retrieval mechanisms (Domeshek, 1992).

Stories must also be recipient designed, in that they need to be tailored for the specific audience they are delivered to. Describing the mechanics of escrow in an otherwise interesting story about creative financing would be inappropriate if told to a banker, but required for a first-time buyer new to home financing. Thus, the knowledge representation used to generate stories from must be hierarchical, with varying levels of detail, so that only the interesting and relevant portions to the listener can be selected and conveyed.

Once a story has been selected and tailored, it must be told in a "lifelike" manner by the embodied conversational character. The performance must be punctuated by emphasis at the appropriate points using relevance or information structure to determine placement and degree of prosody and gesture (Cassell 1995) to convey the emphasis to the listener. Appropriate linguistic devices must also be used to naturally introduce a story, mark its ending, and demonstrate relevance to the listener, if necessary (Jefferson 1978).

Implementation Requirements

Small talk and conversational storytelling impose several novel and theoretically interesting requirements on a real-time conversational interface which we feel have not been addressed by existent systems. First among these is that the discourse planner must be able to manage and pursue multiple conversational goals, some or all of which may be persistent or non-discrete (although see Stone, 1998). It is not sufficient that the planner work on one goal at a time, since a properly selected story, for example, can satisfy a task goal by providing information to the user while also advancing the interpersonal goals of the agent. In addition, many goals, such as phatic or face goals, are better represented by a utility model in which degrees of satisfaction can be planned for, rather than the discrete all-or-nothing goals typically addressed in AI planners (Hanks 1994).

The discourse planner must also be very reactive, since the user's responses cannot be anticipated. The agent's goals and plans may be spontaneously achieved by the user (e.g., through volunteered information) or invalidated (e.g., by the user changing his/her mind) and the planner must be able to immediately accommodate these changes. This is true for conversational storytelling as well, since the user may modify a story being delivered by interrupting, asking for elaboration or providing or withholding feedback. Thus, all utterances, including story contributions which span multiple turns, must be incrementally generated. In contrast, most discourse planners developed to date have been designed to either plan multi-sentential monologues (e.g., Hovy 1991, Andre 1998) or plan dialogue as a sub-component of task planning (e.g., Rich and Sidner 1998), and none have been designed to pursue multiple, non-discrete goals in parallel in the manner described above.

The response time of the entire conversational system is especially crucial during the initial period of small talk, since part of the function of this behavior is to establish the capability and intelligence of the other party. A system with retarded responsiveness will likely not be held in high esteem. In some cultures the beginning of one turn actually overlaps with the end of the previous one (Tannen 1984). While we can't yet expect to mimic this behavior, some studies have shown that conversational system response time must be kept under 1.2 seconds for natural turn-taking behavior to be maintained (Jefferson 1989). We posit that responsiveness and reactivity in this situation (i.e., saying *something* reasonable within the required timeframe) are more important than soundness and completeness (Young, Moore and Pollack 1994), especially given that in general planning is NP-hard (Georgeff 1990).

Thus, we intend to use a reactive planning approach to discourse which focuses on the next move, only planning ahead enough to ensure that goals are being worked towards, and which is reactive to changes in the world (i.e., brought about by the user's contribution), in the spirit of such systems as Hap (Loyall and Bates 1991) and RAPs (Firby 1989), but extended to pursue multiple, and possibly non-discrete goals. Nonlinear planners used in most contemporary AI research systems seem particularly unsuited to this task since they are built around the idea that a plan's subgoal conflicts and interactions are common and dangerous, thus they should be noticed and dealt with as soon as possible, and this seems not to be the case with language (Hanks 1994). Within this planning framework, incremental story generation could be achieved using a mechanism which traverses a knowledge base of story representations using mechanism similar to that used in Salix (Sibun 1992).

The prospect of building a system capable of conducting "casual" conversation in real-time may seem to be an "AI complete" problem at first. However, there are several aspects of small talk which make this endeavor seem tractable. First, the structure of small talk is highly regular and the topics likely to be discussed are predictable (Cheepen 1988, Laver 1975, 1981), and in fact grammars have been developed which completely characterize small talk in particular contexts (Schneider 1988). Second, conversational contributions in small talk need not always be completely semantically relevant. As experience with chatterbot systems have shown, users work hard to establish the relevance of contributions made by a conversational partner (Mauldin 1994), and conversational stories especially depend on this. Also, since much of small talk is ritualized to some degree, the exact words spoken by the agent may not be as important as its style of delivery.[1] Third, it is natural for most real estate agents to maintain initiative in interactions with their clients, and keeping initiative is one of the tricks used by most chatterbot systems to maintain an illusion of intelligence by limiting what the user can say and minimizing the need to respond directly to their queries (Mauldin 1994). We contend that a properly designed system in a limited task domain such as real-estate sales can work on speech input using keyword spotting to detect topic shifts, feedback, and repair moves by the user, and a task-specific grammar to detect when the user has transitioned into task talk.

Conclusion

We have argued that small talk and conversational storytelling can be used by embodied conversational interface agents to achieve both task and interpersonal goals such as building trust and rapport with the user and demonstrating the expertise of the agent, and are thus an important part of an agent's social intelligence, especially in certain types of applications. We have laid out the demands of small talk and conversational storytelling, and have discussed the theoretical concerns that these demands bring to the fore. Developing systems capable of real-time, non-task-oriented dialogue provides numerous challenges and a fertile area of research.

Acknowledgments. Thanks to Hannes Vilhjálmsson, Hao Yan, Petra Chong and Jennifer Smith for their helpful comments on this paper.

References

Abercrombie, D. 1956. *Problems and Principles*. London: Longmans. As cited in Laver 1975.

Andre, E., Rist, T., and Mueller, J. 1998. Integrating Reactive and Scripted Behaviors in a Life-Like Presentation Agent. In *Proceedings of Autonomous Agents 98*. ACM Press: Minneapolis/St. Paul.

Bers, M. and Cassell, J. 1998. Interactive Storytelling Systems for Children: Using Technology to Explore Language and Identity. *Journal of Interactive Learning Research* 9.

Cassell, J. 1995. The Role of Gesture in Stories as Multiple Participant Frameworks. AAAI Spring Symposium Series.

Cassell, J., Bickmore, T., Billinghurst, M., Campbell, L., Chang, K., Vilhjálmsson, H., Yan, H. 1999. Embodiment in Conversational Interfaces: Rea. CHI99, Pittsburgh, PA.

Cegala, D., Waldro, V., Ludlum, J., McCabe, B., Yost, S., and Teboul, B. 1988. A study of interactants' thoughts and feelings during conversation. Ninth Annual Conference on Discourse Analysis, Philadelphia, PA.

Cheepen, C. 1988. *The Predictability of Informal Conversation.* London: Pinter Publishers.

Coupland, J. and Coupland, N. 1992. "How are you?": Negotiating phatic communion. *Language in Society* 21: 207-230.

[1] Abercombie (1956) writes that: "The actual sense of the words used in phatic communion matters little," and goes on to recount the story of Dorothy Parker, alone and rather bored at a party [who was] asked "How are you? What have you been doing?" by a succession of distant acquaintenances. To each she replied "I've just killed my husband with an axe, and I feel fine." Her intonation and expression were appropriate to party small talk, and with a smile and a nod each acquaintance, unastonished, drifted on.

Domeshek, E. 1992. Do The Right Thing: A Component Theory for Indexing Stories as Social Advice. Technical Report #26, Institute for the Learning Sciences, Northwestern University.

Dunbar, R. 1996. *Grooming, Gossip, and the Evolution of Language*. Cambridge, MA: Harvard University Press.

Ervin-Tripp, S. and Kuntay, A. 1997. The Occasioning and Structure of Conversational Stories. In T. Givon, ed., *Conversation: Cognitive, communicative and social perspectives*: 133-166.

Firby, R. 1989. *Adaptive Execution in Complex Dynamic Worlds.* PhD thesis, Department of Computer Science, Yale University.

Garner, A. 1980. *Conversationally Speaking*. Los Angeles: Lowell House.

Garros, D. 1999. Personal communication. (Broker, Home and Hearth Realty, Cambridge, MA).

Georgeff, M. 1990. Planning. In J. Allen, J. Hendler, and A. Tate, eds. *Readings in Planning*: 5-25. San Mateo, CA: Morgan Kaufmann.

Goffman, E. 1981. *Forms of Talk*. University of Pennsylvania Press.

Gough, D. 1990. The Principle of Relevance and the Production of Discourse: Evidence from Xhosa Folk Narrative. in *Narrative Thought and Narrative Language*, B. Britton and A. Pellegrini, eds. Hillsdale, New Jersey: Lawrence Erlbaum Associates.

Grice, Paul. 1989. *Studies in the Way of Words*. Cambridge, MA: Harvard University Press.

Hanks, S. 1994. Discourse Planning: Technical Challenges for the Planning Community. AAAI Workshop on Planning for Inter-Agent Communication.

Hopkins, T. 1982. *How to Master the Art of Selling.* New York: Warner Books.

Hovy, E. 1991. Approaches to the planning of coherent text. In C. Paris, et al, eds, *Natural Language Generation in Artificial Intelligence and Computational Linguistics*: 83-102 Boston: Kluwer Academic Publishers.

Jakobson, R. 1960. Linguistics and poetics. In *Sebeok 1960*: 350-377.

Jefferson, G. 1978. Sequential Aspects of Storytelling in Conversation. In J. Schenkein, ed., *Studies in the Organization of Conversational Interaction*: 219-248. New York: Academic Press.

Jefferson, G. 1989. Notes on a possible metric for a "standard maximum" silence of approximately one second in conversation. In D. Roger and P. Bull eds. *Conversation: an Interdisciplinary Perspective*: 166-196. Philadelphia: Multilingual Matters.

Koda, T. and Maes, P. 1996. Agents with Faces: The Effects of Personification of Agents. HCI'96: 98-103. London, UK, August 20-23.

Labov, W. and Waletzky, J. 1967. Narrative Analysis: Oral versions of personal experience. In J. Helm, ed., *Essays on the verbal and visual arts*. Seattle: University of Washington Press.

Laver, J. 1975. Communicative functions of phatic communion. In A. Kendon, R. Harris and M. Key, eds. *The organization of behavior in face-to-face interactio:* 215-238. The Hague: Mouton.

Laver, J. 1981. Linguistic routines and politeness in greeting and parting. In F. Coulmas, ed., *Conversational routine:* 289-304. The Hague: Mouton.

Loyall, A. and Bates, J. 1991. *Hap: A Reactive, Adaptive Architecture for Agents.* Pittsburgh, PA: Carnegie Mellon University technical report CMU-CS-91-147.

Malinowski, B. 1923. The problem of meaning in primitive languages. Supplement to C. Ogden and I. Richards, *The meaning of meaning:* 146-152. London: Routledge and Kega Paul.

Mauldin, M. 1994. Chatterbots, Tinymuds, and the Turing Test: Entering the Loebner Prize Competition. AAAI 94.

Moon, Y. 1998. *Intimate Self-Disclosure Exchanges: Using Computers to Build Reciprocal Relationships with Consumers*. Cambridge, MA: Harvard Business School Working Paper 99-059.

Polanyi, L. 1982. Linguistic and Social Constraints on Storytelling. *Journal of Pragmatics* 6: 509-524.

Polanyi, L. 1985. *Telling the American Story: A Structural and Cultural Analysis of Conversational Storytelling*. Norwood, New Jersey: Ablex Publishers.

Reeves, B. and Nass, C. 1996. *The Media Equation: How People Treat Computers, Television, and New Media Like Real People and Places*. Cambridge University Press.

RoAne, S. 1997. *What do I Say Next? Talking your way to business and social success*. New York: Warner Books.

Rich, C. and Sidner, C. 1998. *Collagen: A Collaboration Manager for Software Interface Agents.* Mitsubishi Electric Research Laboratory technical report TR-97-21a.

Sacks, H. 1970-71. Unblished lecture notes: conversational storytelling. University of California at Irvine. As cited in Polanyi, L. 1982.

Schneider, K. 1988. *Small Talk: Analysing Phatic Discourse*. Marburg Germany: Hitzeroth.

Sibun, P. 1992. Generating Text Without Trees. *Computational Intelligence: Special Issue on Natural Language Generation.* 8(1): 102-122.

Stone, M. 1998. *Modality in Dialogue: Planning, Pragmatics, and Computation*. PhD Thesis, University of Pennsylvania.

Tannen, D. 1984. *Conversational Style: Analyzing talk among friends*. Ablex Publishing, Norwood, New Jersey.

Tracy, K. and Coupland, N. eds. 1991. *Multiple goals in discourse*. Clevedon: Multilingual Matters.

Turkle, S. 1995. *Life on the Screen: Identity in the Age of the Internet*. New York: Touchstone.

Weizenbaum, J. 1966. Eliza -a computer program for the study of natural language communication between man and machine. *Communications of the ACM* 9: 26-45.

Young, M., Moore, J. and Pollack, M. 1994. Towards a Principled Representation of Discourse Plans. *Proceedings of the Cognitive Science Society*.