

Initiative in Collaborative Interactions — Its Cues and Effects

Jennifer Chu-Carroll and Michael K. Brown

Multimedia Communications Research Lab

Bell Laboratories

600 Mountain Avenue

Murray Hill, NJ 07974, U.S.A.

E-mail: {jenc, mkb}@bell-labs.com

Abstract

In naturally-occurring collaborative dialogues, the participants often share the responsibility of taking initiative in the dialogue. Thus, in order for a system to interact with human agents in a natural and coherent fashion, it must be able to recognize and provide cues when shifts in initiative are intended, and to take initiative into account in interpreting user utterances and in generating responses to such utterances.

In this paper, we argue that it is necessary to distinguish between *task initiative*, which indicates the agent who is currently actively involved in planning toward accomplishing the agents' task, and *dialogue initiative*, which indicates the agent who has the lead in determining the current discourse focus. We identify a set of cues which may indicate changes in initiative that can be inferred based on linguistic and domain knowledge alone, and show how these cues can be used to determine the distribution of task and dialogue initiatives at the end of each dialogue turn. In addition, we show that the effects of task and dialogue initiatives on the dialogue system's response generation process are twofold: they determine which subcomponents of the dialogue system may be selected to address the problem of response generation, and they affect what the selected subcomponent will include in its response to the user's utterances.

Introduction

It is a well-known fact that naturally-occurring collaborative dialogues are very rarely, if ever, one-sided. Instead, initiative of the interaction shifts among the dialogue participants in a primarily principled fashion signaled by features such as linguistic cues, prosodic cues, and, in face-to-face interactions, eye gaze and gestures. It is thus important in developing a dialogue system to take into account the effects that these cues have on initiative shifts and to consider the effects that such shifts in initiative have on the utterance interpretation and response generation processes.

Previous work on modeling initiative in dialogue interactions has focused on tracking and allocating a single thread of control, the *conversational lead*, among participants (Novick 1988; Whittaker & Stenton 1988; Walker & Whittaker 1990; Kitano & Van Ess-Dykema 1991; Smith &

Hipp 1994; Guinn 1996). We argue that this view of conversational control fails to distinguish between *task initiative* and *dialogue initiative*, which together determine when an agent will address an issue during a dialogue, and how she will go about addressing it. Although physical cues, such as gestures and eye gaze, play an important role in coordinating shifts in initiative in face-to-face human-human interactions, we contend that a great deal of information regarding initiative shifts can be extracted from user utterances based on linguistic and domain knowledge alone. By utilizing a model for tracking shifts in task and dialogue initiatives, a dialogue system will then be able to better coordinate its interaction with the user and to generate responses to user utterances in a more cooperative and natural fashion.

In this paper, we will show how the distinction between task and dialogue initiatives accounts for different responses given by human agents in collaborative dialogues that previous models were unable to explain. We will present a set of cues that a dialogue system may adopt in modeling initiative based on linguistic and domain knowledge, as well as outline how these cues can be utilized to determine whether or not a shift in initiative should occur at the end of a dialogue turn. Finally, we show that by modeling initiative in dialogue interactions, the system is able to restrict its candidate set of response generation subcomponents and to select appropriate content to include in its response, thereby providing more cooperative and natural responses to user utterances.

Distinguishing Between Task and Dialogue Initiatives

Previous work on mixed initiative dialogue interactions has focused on modeling and allocating a single thread of control among dialogue participants based on features such as utterance types or priorities of agents' goals (Novick 1988; Whittaker & Stenton 1988; Walker & Whittaker 1990; Kitano & Van Ess-Dykema 1991; Smith & Hipp 1994; Guinn 1996). However, our analysis of naturally occurring dialogues (SRI Transcripts 1992; Gross, Allen, & Traum 1993; Heeman & Allen 1995) shows that this view of con-

versational initiative fails to explain an important feature of dialogue interactions as illustrated by the three alternative responses (utterances (3a), (3b), and (3c)), given by an advisor in response to a student's question in utterance (2) in the following dialogue:

(1) S: *I want to take NLP to satisfy my seminar course requirement.*

(2) *Who is teaching NLP?*

(3a) A: *Dr. Smith is teaching NLP.*

(3b) A: *You can't take NLP this semester because you haven't taken AI, which is a prerequisite for the course.*

(3c) A: *You can't take NLP this semester because you haven't taken AI, which is a prerequisite for the course.*

I would suggest that you take some other seminar course to satisfy your requirement, and sign up as a listener for NLP if you're really interested in it.

Suppose we adopt a model that maintains a single thread of control, such as that in (Whittaker & Stenton 1988). In utterance (3a), A provides a direct response to S's question; thus the conversational lead remains with S. On the other hand, in utterances (3b) and (3c), A takes control of the dialogue by initiating a subdialogue to correct S's invalid proposal of taking NLP as a means of satisfying her seminar course requirement. However, existing initiative models are unable to explain the difference between (3b) and (3c), namely that in (3c), A becomes more *active* in the planning process by explicitly proposing actions that she believes should be incorporated into S's plan, while in (3b), A merely points out the invalidity in S's proposal and leaves it up to S to propose an alternative plan. Based on this observation, we argue that it is necessary to distinguish between *task initiative*, which tracks the lead in the development of the plan that the agents are constructing, and *dialogue initiative*, which tracks the lead in determining the current discourse focus. This distinction then allows us to model A's behavior: in utterance (3b), A responds to S's proposal of an invalid plan by taking over only the dialogue initiative, i.e., by merely informing S of the invalidity of the proposal. On the other hand, in utterance (3c), A responds to S's proposal by taking over both the task and dialogue initiatives, i.e., by both informing S of the invalidity and suggesting a possible remedy.

An agent has the *task initiative* if she is directing how the agents' task should be accomplished, i.e., if her utterances directly propose *actions* that the agents should perform. The agent's utterances may propose actions that contribute to achieving their domain goal, such as "*why don't we couple engine E2 to the boxcar that's at Elmira, and send it*

to Corning".¹ On the other hand, the agent's utterances may propose actions that contribute not directly to achieving their domain goal, but to how they would go about constructing a plan to achieve this goal,² such as "*Let's look at the first [problem] first, I think they are separate*". An agent has the *dialogue initiative* if he takes the conversational lead in order to establish mutual beliefs, such as mutual beliefs about a particular piece of domain knowledge or about the validity of a proposal, between the agents. Thus, when an agent takes over the task initiative, she also takes over the dialogue initiative, since the agent's proposal of actions can be viewed as an attempt to establish the mutual belief that a set of actions be adopted by the agents. On the other hand, in situations where an agent proposes other types of mutual beliefs, such as beliefs about particular pieces of domain knowledge, she may hold the dialogue initiative without holding the task initiative. For instance, in response to agent A's request to evaluate a proposed plan, agent B may take over the dialogue initiative (but not the task initiative) by giving a detailed analysis of the proposal based on his private knowledge.

Evidently, shifts in task and dialogue initiatives greatly affect *when* an agent will say something during the dialogue and *what* she will say during the dialogue. Furthermore, in order to ensure a smooth and coherent dialogue, such shifts in initiative should be signaled and the effects of the cues be taken into account by both dialogue participants. The rest of this paper discusses our model for tracking initiative and how this model affects the response generation process.

Modeling Shifts in Initiative

Cues for Shifts in Initiative

In order to identify cues for shifts in initiative, we annotated the TRAINS 91 corpus (Gross, Allen, & Traum 1993) to indicate the agents who hold the task and dialogue initiatives in each dialogue turn. We then analyzed the annotated dialogues and identified eight types of cues that may have contributed to the shift or lack of shift in initiative during the interactions. These eight cue types can be categorized into three classes based on the kind of knowledge needed to recognize them. Table 1 shows the three cue classes, the eight cue types, their subtypes if any, whether a cue may affect merely the dialogue initiative or both the task and dialogue initiatives, and the agent who is expected to hold the initiative in the turn after a cue is detected. For example, the last cue in Table 1, *ambiguity*, has two subclasses: ambiguities regarding proposed actions and ambiguities regarding proposed beliefs. Although cues in both subclasses are likely to result an initiative shift to the hearer, ambiguities regarding

¹Most sample utterances/dialogues in this paper are taken from (SRI Transcripts 1992; Gross, Allen, & Traum 1993).

²Such actions are referred to as *problem-solving actions* by Ramshaw (1991) and Lambert and Carberry (1991).

Class	Cue Type	Subtype	Effect	Initiative	Example
Explicit	Explicit requests	give up	both	hearer	"Any suggestions?"
		take over	both	speaker	"Summarize the plan up to this point" "Let me handle this one."
Discourse	End silence		both	hearer	
	No new info	repetitions	both	hearer	A: "Grab the tanker, pick up oranges, go to Elmira, make them into orange juice." B: "We go to Elmira, we make orange juice, okay." "Yeah", "Ok", "Right"
prompts		both	hearer		
	Questions	domain evaluation	DI DI	speaker hearer	"How far is it from Bath to Corning?" "What do you think about the plan so far?"
		Obligation fulfilled	task	both	hearer
Analytical	Invalidity	action	both	hearer	A: "Get the tanker car to Elmira and fill it with OJ." B: "Well, you need to get oranges to the OJ factory." C: "It's shorter to Bath from Avon." D: "It's shorter to Dansville." "The map is slightly misleading."
		belief	DI	hearer	
	Suboptimality		both	hearer	A: "Using Saudi on Thursday the eleventh." B: "It's sold out." A: "Is Friday open?" B: "Economy on Pan Am is open on Thursday."
	Ambiguity	action	both	hearer	A: "Take one of the engines from Corning." B: "Let's say engine E2." C: "We would get back to Corning at 4." D: "4PM? 4AM?"
		belief	DI	hearer	

Table 1: Cues for Shifts in Initiative

proposed actions will potentially lead to the hearer taking over both the task and dialogue initiatives, while ambiguities regarding proposed beliefs will likely lead to the hearer taking over only the dialogue initiative.

The first class, *explicit cues*, contains only explicit requests from the speaker to either give up or take over the initiative. For instance, in uttering “*any suggestions?*”, the speaker is explicitly requesting that the hearer propose actions for achieving their goal, thereby giving up both the task and dialogue initiatives. On the other hand, by saying “*Summarize the plan up to this point*”, the speaker is giving up only the dialogue initiative since the summarization process does not involve the addition of new actions to the agent's domain plan.

The second class, *discourse cues*, includes cues that can be inferred without the use of domain knowledge, e.g., from the surface form of an utterance or from the purpose an utterance is intended to serve. We have identified four types of discourse cues. The first type is perceptible silence observed at the end of an utterance, which suggests that the speaker does not have anything more to say in the current dialogue turn and may be an indication that the speaker intends to give up his initiative. The second type includes cases in which the speaker's utterances do not include any new information that has not been conveyed earlier in the dialogue. We further classify these utterances into two subtypes: *repetitions*, a subset of the *informationally redundant utterances* (Walker 1992), in which the speaker paraphrases a previous utterance by the hearer or repeats the utterance verbatim, and *prompts*, in which the speaker merely gives an acknowledgment for the hearer's previous utterance(s). Repetitions and prompts also suggest that the speaker has nothing more to say at the current point in the dialogue and are thus indications that the hearer should take over the initiative (Whittaker & Stenton 1988). The third type of discourse cues includes questions which, based on anticipated responses, are further divided into *domain questions* and *evaluation questions*. Domain questions are questions in which the speaker intends to obtain or verify a piece of domain knowledge. Such questions may take the form of a WH-question, such as “*where are there boxcars?*”, or the form of a YN-question, such as “*are there boxcars in Corning?*” These questions usually merely require a direct response and thus typically do not result in a shift in initiative. Evaluation questions, on the other hand, are questions in which the speaker intends to assess the quality of a proposed plan. They often require that the hearer analyze the proposal, and thus frequently cause a shift in dialogue initiative to take place. The final type of discourse cues includes situations in which the speaker's utterances satisfy an outstanding task or discourse obligation. Such obligations may have resulted from a prior request by the hearer, or from an interruption initiated by the speaker herself. In

either case, when the speaker fulfills the obligation of providing the requested information or satisfies the goal that resulted in the interruption, the initiative may be reverted back to the hearer who held the initiative prior to the request or interruption.

The third class of cues, *analytical cues*, are cues that cannot be recognized unless the hearer performs an evaluation on the speaker's proposal using the hearer's private domain knowledge. After the evaluation, the hearer may find the speaker's proposal 1) *invalid*, if an action in the proposal is infeasible (Pollack 1986; Chu-Carroll & Carberry 1994), if the proposed plan is ill-formed (Pollack 1986; Chu-Carroll & Carberry 1994), if the hearer disagrees with a proposed mutual belief (Chu-Carroll & Carberry 1995), or if the speaker's utterances indicate that a miscommunication has occurred (McRoy & Hirst 1995; Traum & Dillenbourg 1996), 2) *suboptimal*, if the hearer detects a better alternative than the speaker's proposal (Joshi, Webber, & Weischedel 1984; van Beek, Cohen, & Schmidt 1993; Chu-Carroll & Carberry 1994), or 3) *ambiguous*, if the hearer cannot uniquely identify the speaker's intentions (van Beek, Cohen, & Schmidt 1993; Raskutti & Zukerman 1993; Logan *et al.* 1994). When such a problem is detected with respect to the speaker's proposal, the hearer may initiate a subdialogue to resolve the problem, resulting in a shift in task and/or dialogue initiatives.

Utilizing Cues in Modeling Initiative

Evidently, some of the cues in Table 1 provide stronger evidence for a shift in initiative than others. For instance, an explicit request for a hearer to take over the initiative will more likely result in an initiative shift than, say, a perceptible silence detected at the end of an utterance. Thus, in order to utilize these cues in modeling initiative, it is necessary to associate with each cue a probabilistic measure that estimates the likelihood that an initiative shift will occur when such a cue is observed. A companion paper describes how we utilize a training algorithm to obtain a set of basic probability assignments for the cues using the annotated TRAINS 91 corpus (Chu-Carroll & Brown Submitted). Furthermore, at the end of each dialogue turn, the system must be able to determine the accumulated effect of the cues observed during the turn and to determine the distribution of the task and dialogue initiatives in the next turn. This process is briefly outlined in this section and is discussed in further detail in (Chu-Carroll & Brown Submitted).

Our model keeps track of a *task initiative index* and a *dialogue initiative index* for each participant during the dialogue. The task or dialogue initiative index for an agent measures the amount of evidence available to support the agent holding the task or dialogue initiative, respectively. At the end of each turn, new initiative indices are calculated based on the current initiative indices and

Current initiative indices:

$$TI_{cur}(\{C\}) = .53 \quad TI_{cur}(\{T\}) = .47 \\ DI_{cur}(\{C\}) = .52 \quad DI_{cur}(\{T\}) = .48$$

Probability assignment for suboptimality:

$$TI_{sub}(\{T\}) = .25 \quad TI_{sub}(\{T, C\}) = .75 \\ DI_{sub}(\{T\}) = .25 \quad DI_{sub}(\{T, C\}) = .75$$

Probability assignment for domain-question:

$$DI_{dq}(\{C\}) = .03 \quad DI_{dq}(\{T\}) = .08 \quad DI_{dq}(\{T, C\}) = .89$$

New initiative indices:

$$TI_{new}(\{C\}) = .45 \quad TI_{new}(\{T\}) = .55 \\ DI_{new}(\{C\}) = .44 \quad DI_{new}(\{T\}) = .56$$

Figure 1: Calculation of New Initiative Indices

the effects of the cues observed during the current dialogue turn based on the Dempster-Shafer theory (Shafer 1976; Gordon & Shortliffe 1984). The new initiative indices then determine the distribution of the task and dialogue initiatives in the next dialogue turn, i.e., the agent with the higher task initiative index will hold the task initiative, and similarly for the dialogue initiative.

For example, consider the following dialogue segment, taken from a transcript of airline reservation dialogues (SRI Transcripts 1992):

- (4) T: *You're looking for Saudi Arabian Airlines.*
(5) C: *Yeah, using Saudi on the night of Thursday the eleventh.*
(6) T: *Right, that's what I'm looking at. It's sold out.*
(7) C: *Is Friday open?*
(8) T: *Let me check here. I'm showing economy on Pan Am is open on the eleventh.*

In utterance (7), the travel agent detects two cues. The first cue is *domain-question*, observed from the form of the utterance as well as the discourse purpose it is intended to serve. The second cue is *suboptimality*, which is detected when the travel agent evaluates the customer's (implicit) proposal of traveling on Friday instead of Thursday and finds a better alternative to this proposal, i.e., to change the airline instead of the date of travel. The agents then need to determine the distribution of the task and dialogue initiatives for the next turn based on the current initiative indices and the basic probability assignments for each of the two observed cues using the Dempster-Shafer theory, as shown in Figure 1.³ The new initiative indices indicate that the

³For details on utilizing Dempster's combination rule to determine the accumulated effect of multiple basic probability assignments, see (Shafer 1976; Gordon & Shortliffe 1984).

travel agent should hold both the task and dialogue initiatives in the next turn, which agrees with the continuation of the actual dialogue in utterance (8).

We are currently working on incorporating our model for tracking initiative into a dialogue manager based on (Abella, Brown, & Buntschuh 1996). As discussed earlier, the distribution of task and dialogue initiatives between the agents affects the system's responses to user utterances. Thus using our initiative model, we are able to determine, at each dialogue turn, whether or not the system should hold the task and/or dialogue initiative, and to tailor the system's response to user utterances accordingly. In the next section, we discuss the effects of the distribution of initiative on the response generation process.

Response Generation in Mixed Initiative Dialogues

Abella, Brown, and Buntschuh (1996) proposed a set of motivations for conducting a dialogue, which they called *dialogue motivators*. They then incorporated a subset of these dialogue motivators into a dialogue manager which accepts the semantic representation of an utterance, determines which dialogue motivator to apply under the given circumstance, and subsequently generates an appropriate response to the given utterance. However, their dialogues are entirely system-driven and user utterances are only prompted by requests for information by the system.

Our goal, on the other hand, is to develop a dialogue manager that interacts with the user in a manner that more closely resembles human-human dialogue interactions. In particular, we intend to allow both the task and dialogue initiatives to shift between the agents in a coherent and natural fashion, thereby allowing mixed-initiative in the dialogue interactions. The dialogue manager's task can thus be outlined as follows. Given a user utterance, the dialogue manager infers the user's intentions from the utterance, and extracts the explicit, discourse, and analytical cues from the utterance itself and the inferred intentions. It then invokes the initiative model to determine the task and dialogue initiative holders for the next dialogue turn. Finally, it invokes the appropriate dialogue motivator⁴ based on the inferred user intentions and the distribution of the initiative, and generates an appropriate response to the user utterance.

The distribution of initiative affects two aspects of the response generation process. First, it affects the candidate set of dialogue motivators. For instance, the dialogue motivator *suggest alternative* is invoked when the system detects a better alternative to the user's proposal (Chu-Carroll & Carberry 1994). Activation of this dialogue motivator will

⁴We have expanded the list of dialogue motivators proposed in (Abella, Brown, & Buntschuh 1996) to include extra features such as *correction of invalid plans/beliefs* and *suggestion of better alternatives* (Chu-Carroll & Carberry 1994; 1995).

result in the system suggesting the alternative plan to the user, thereby taking over both the task and dialogue initiatives. Thus if the system does not have task initiative in the next turn, *suggest alternative* will not be considered as a candidate dialogue motivator for response generation. Second, the distribution of initiative affects the content of the system's response. In order to illustrate this effect, we return to the dialogue shown earlier where the student proposes taking NLP to satisfy her seminar course requirement and is interested in finding out who the instructor of NLP is. If at the end of utterance (2), the initiative model suggests that the system does not have either the task or the dialogue initiative (perhaps because the system is instructed to act as a database query system and therefore does not perform any analysis on the validity of the student's proposed plan), then the system will provide a direct, although arguably unhelpful, response to the student's question, as in (3a). If the initiative model indicates that the system does not have the task initiative but does have the dialogue initiative, then the system will inform the student of the invalidity of her plan, but will not propose an alternative plan, as in (3b). Finally, if the initiative model decides that the system should have both the task and dialogue initiatives, the system will not only point out the invalidity of the proposal, but will also propose what it believes to be a valid plan as an alternative, as in (3c).

Conclusions and Future Work

In this paper, we have shown that distinguishing between task and dialogue initiatives allows us to better model the interaction between participants in naturally occurring collaborative dialogues. We presented a set of explicit, discourse, and analytical cues that a collaborative system can adopt in modeling initiative based on linguistic and domain knowledge alone. By utilizing these cues in modeling initiative during dialogue interactions, we are able to determine the distribution of the task and dialogue initiatives for each turn in the dialogue. Furthermore, we discussed how this model for tracking initiative can be incorporated into a dialogue manager to affect the system's selection of the dialogue motivator and the content of its response to the user's utterances, thereby allowing the system to interact with the user in a more natural and coherent manner.

Our current work involves constructing the dialogue manager to illustrate the system's mixed-initiative behavior. In the future, we intend to address the following three main issues. First, we intend to generalize our model for tracking initiative in terms of its application domains. We have performed a preliminary evaluation of our model by comparing the system's predictions of the distribution of task and dialogue initiatives with the distributions in actual dialogues. This evaluation is carried out on several types of collaborative dialogues: collaborative planning dialogues

in the TRAINS domain (Gross, Allen, & Traum 1993; Heeman & Allen 1995), airline reservation dialogues (SRI Transcripts 1992), instruction-giving dialogues (Map Task Dialogues 1996), and non-task-oriented dialogues (Switchboard Credit Card Corpus 1992). We found that the system's performances on collaborative planning dialogues, i.e., the TRAINS dialogues and the airline reservation dialogues, most closely resemble one another (Chu-Carroll & Brown Submitted). We believe that the difference in the system's performance in various application environments can be explained by the differences in the way people interact in different collaborative environments and hence how they interpret cues. We intend to develop a method for dynamically adjusting the basic probability assignments for the cues to reflect the differences in the application environments as well as in individual users. Second, we intend to generalize our model in terms of the number of dialogue participants it can handle. We believe that our basic framework of utilizing the Dempster-Shafer theory to track initiative is directly applicable to multi-agent interactions. However, in multi-agent interactions, we must further refine the cues to distinguish between cases in which the speaker intends to pass the initiative to one particular agent, those in which the speaker intends to pass the initiative to a particular subset of the agents, and those in which the speaker intends to pass the initiative to any other agent. Finally, we intend to investigate other areas for cues that may signal shifts in dialogue initiative. In particular, we are interested in studying whether or not incorporating prosodic information into our analysis will improve the system's performance in predicting the dialogue initiative holder.

References

- Abella, A.; Brown, M. K.; and Buntschuh, B. M. 1996. Development principles for dialog-based interfaces. In *Proceedings of the ECAI96 workshop on Dialogue Processing in Spoken Language Systems*, 1-7.
- Chu-Carroll, J., and Brown, M. K. Submitted. Tracking initiative in collaborative dialogue interactions. Submitted for publication.
- Chu-Carroll, J., and Carberry, S. 1994. A plan-based model for response generation in collaborative task-oriented dialogues. In *Proceedings of the Twelfth National Conference on Artificial Intelligence*, 799-805.
- Chu-Carroll, J., and Carberry, S. 1995. Response generation in collaborative negotiation. In *Proceedings of the 33rd Annual Meeting of the Association for Computational Linguistics*, 136-143.
- Gordon, J., and Shortliffe, E. H. 1984. The Dempster-Shafer theory of evidence. In Buchanan, B., and Shortliffe, E., eds., *Rule-Based Expert Systems: The MYCIN Ex-*

- periments of the Stanford Heuristic Programming Project*. Addison-Wesley. chapter 13, 272–292.
- Gross, D.; Allen, J. F.; and Traum, D. R. 1993. The TRAINS 91 dialogues. Technical Report TN92-1, Department of Computer Science, University of Rochester.
- Guinn, C. I. 1996. Mechanisms for mixed-initiative human-computer collaborative discourse. In *Proceedings of the 34th Annual Meeting of the Association for Computational Linguistics*, 278–285.
- Heeman, P. A., and Allen, J. F. 1995. The TRAINS 93 dialogues. Technical Report TN94-2, Department of Computer Science, University of Rochester.
- Joshi, A.; Webber, B.; and Weischedel, R. M. 1984. Living up to expectations: Computing expert responses. In *Proceedings of the Fourth National Conference on Artificial Intelligence*, 169–175.
- Kitano, H., and Van Ess-Dykema, C. 1991. Toward a plan-based understanding model for mixed-initiative dialogues. In *Proceedings of the 29th Annual Meeting of the Association for Computational Linguistics*, 25–32.
- Lambert, L., and Carberry, S. 1991. A tripartite plan-based model of dialogue. In *Proceedings of the 29th Annual Meeting of the Association for Computational Linguistics*, 47–54.
- Logan, B.; Reece, S.; Cawsey, A.; Galliers, J.; and Sparck Jones, K. 1994. Belief revision and dialogue management in information retrieval. Technical Report 339, University of Cambridge, Computer Laboratory.
- Map Task Dialogues. 1996. Transcripts of DCIEM Sleep Deprivation Study, conducted by Defense and Civil Institute of Environmental Medicine, Canada, and Human Communication Research Centre, University of Edinburgh and University of Glasgow, UK. Distributed by HCRC and LDC.
- McRoy, S. W., and Hirst, G. 1995. The repair of speech act misunderstanding by abductive inference. *Computational Linguistics* 21(4):435–478.
- Novick, D. G. 1988. *Control of Mixed-Initiative Discourse Through Meta-Locutionary Acts: A Computational Model*. Ph.D. Dissertation, University of Oregon.
- Pollack, M. E. 1986. A model of plan inference that distinguishes between the beliefs of actors and observers. In *Proceedings of the 24th Annual Meeting of the Association for Computational Linguistics*, 207–214.
- Ramshaw, L. A. 1991. A three-level model for plan exploration. In *Proceedings of the 29th Annual Meeting of the Association for Computational Linguistics*, 36–46.
- Raskutti, B., and Zukerman, I. 1993. Eliciting additional information during cooperative consultations. In *Proceedings of the 15th Annual Meeting of the Cognitive Science Society*.
- Shafer, G. 1976. *A Mathematical Theory of Evidence*. Princeton University Press.
- Smith, R. W., and Hipp, D. R. 1994. *Spoken Natural Language Dialog Systems — A Practical Approach*. Oxford University Press.
- SRI Transcripts. 1992. Transcripts derived from audiotape conversations made at SRI International, Menlo Park, CA. Prepared by Jacqueline Kowtko under the direction of Patti Price.
- Switchboard Credit Card Corpus. 1992. Transcripts of telephone conversations on the topic of credit card use, collected at Texas Instruments. Produced by NIST, available through LDC.
- Traum, D. R., and Dillenbourg, P. 1996. Miscommunication in multi-modal collaboration. In *Proceedings of the AAAI-96 Workshop on Detecting, Repairing, and Preventing Human-Machine Miscommunication*.
- van Beek, P.; Cohen, R.; and Schmidt, K. 1993. From plan critiquing to clarification dialogue for cooperative response generation. *Computational Intelligence* 9(2):132–154.
- Walker, M., and Whittaker, S. 1990. Mixed initiative in dialogue: An investigation into discourse segmentation. In *Proceedings of the 28th Annual Meeting of the Association for Computational Linguistics*, 70–78.
- Walker, M. A. 1992. Redundancy in collaborative dialogue. In *Proceedings of the 15th International Conference on Computational Linguistics*, 345–351.
- Whittaker, S., and Stenton, P. 1988. Cues and control in expert-client dialogues. In *Proceedings of the 26th Annual Meeting of the Association for Computational Linguistics*, 123–130.