

# Mutual Beliefs of Multiple Conversants: A Computational Model of Collaboration in Air Traffic Control

David G. Novick and Karen Ward

Oregon Graduate Institute of Science & Technology  
19600 N.W. von Neumann Drive  
Beaverton, OR 97006-1999 USA  
novick@cse.ogi.edu, wardk@cse.ogi.edu

## Abstract

This work develops a computational model for representing and reasoning about dialogue in terms of the mutuality of belief of the conversants. We simulated cooperative dialogues at the speech act level and compared the simulations with actual dialogues between pilots and air traffic controllers engaged in real tasks. In the simulations, addressees and overhearers formed beliefs and took actions appropriate to their individual roles and contexts. The result is a computational model capable of representing the evolving context of complete real-world multiparty task-oriented conversations in the air traffic control domain.<sup>1</sup>

## Introduction

This paper addresses the question of how mutuality is maintained in conversation and specifically how mutuality can be usefully modeled in multiparty computational dialogue systems. The domain we studied and modeled is air traffic control (ATC). The problem we are solving is related to the distributed artificial intelligence research on ATC communications (e.g., Findler & Lo 1988), except that we are explicitly dealing with the mutuality aspects of interaction. While other ATC studies have developed domain models suitable for distributed processing via cooperating agents, we are interested in fundamental knowledge about how such cooperation is achieved through linguistic interaction. Interestingly, we find that the mutuality model by itself explains a great deal of the ATC communications that we observed.

This model and its associated representation of mutuality are not specific to ATC. While ATC served as a domain for the simulation, the mutuality maintenance mechanism presented here does depend on the domain's specifics. Thus, the representation should be useful for other multiparty interaction in, for example, applications of distributed artificial intelligence.

We define and validate a speech act model of ATC dialogue built around the mutuality of beliefs among conversants and overhearers. Complete actual dialogues between air traffic controllers and pilots were explicated in terms of task, belief, and event. The model was tested by computational simulation and was found to be successful in pre-

dicting and explaining the course of real-world conversation at the speech act level. In particular, we replicated a number of actual conversations using speech-act models of air traffic control and conversational mutuality. The simulations account for and produce the effects of belief formation in computational agents representing both conversants and overhearers.

## Motivation

Spoken language understanding systems require dialogue-level language models that are capable of representing and reasoning about the course of real-world task-oriented conversation. Speech act theory (Austin 1962; Searle 1969; Searle 1975; Searle & Vanderveken 1985) suggests that we can motivate dialogue and explain conversational coherence by modeling conversation in terms of the conversants' goals and their plans for reaching those goals. Language as action provides conversants with means for achieving mutual goals (Winograd & Flores, 1986). This approach accords well with findings that the structure of discourse about a particular task closely follows the structure of the task itself (e.g., Oviatt & Cohen 1988; Cohen & Perrault 1979; Grosz & Sidner 1986). In this view, language is just another tool to be used in accomplishing some goal, and utterance planning becomes incorporated into the larger task planning (e.g., Power 1979; Cohen 1984; Litman & Allen 1987).

Clark and his colleagues have proposed a theory of conversation as a collaborative process in which conversants collaborate in building a mutual model of the conversation (Clark & Marshall 1981; Clark & Wilkes-Gibbs 1986; Schober & Clark 1989; Clark & Schaefer 1989). The conversants' beliefs about the mutuality of their knowledge serves to motivate and explain the information that conversants exchange. The collaborative view of conversation offers an explanation for conversational coherence by viewing conversation as an ensemble work in which the conversants cooperatively build a model of shared belief (Suchman 1987; Clark & Schaefer 1989).

Our model is based on a synthesis of these principles, in that conversation is viewed as an attempt to establish and build upon mutual knowledge using speech acts. This synthesis was first proposed by Novick (1988) to explain conversational control acts. More recently, Traum and Hinkelman (1992) have developed a similar model of

1. This research was supported by National Science Foundation grant No. IRI-9110797.

mutuality maintenance as part of their theory of “conversation acts.” In this study, we use these principles to explain and motivate domain-level acts in a real-world task.

This article, then, explores the domain of air traffic control as a speech-act model, produces a computational representation of the beliefs, actions and inferences of the conversants, and validates the model and representation through detailed simulation of observed ATC dialogue. We pay particular attention to the maintenance of mutuality of belief in both conversants and overhearers. The next section of this article discusses the ATC domain and describes the dialogues we analyzed and modeled. We present the computational model, discussing the goals of the agents in the domain and presenting our representation of their beliefs. We next present inference rules for speech acts in the ATC domain, including detailed discussion of their mutuality effects. Based on the representation and inference rules, we then simulate observed dialogue in a multi-agent rule-based environment; we describe the simulation environment and procedure and discuss the results of the simulations.

### The Air Traffic Control Domain

Developing and validating a model of mutuality of belief in conversation, particularly for multiple party interaction, requires a domain in which there are naturally occurring cooperative tasks, relative ease of determining initial beliefs, and multiple conversants. For this study, we chose to examine air traffic control dialogue. Because we wanted the model to represent actual dialogue and not just encode the formalisms suggested by the Federal Aviation Administration (FAA 1989), we studied protocols of actual pilot-controller conversation (Ward, Novick, & Sousa 1990) rather than simply encoding the FAA’s rules for ATC communication.

In this study we examined four complete ATC dialogues, which ranged in size from 14 to 19 transmissions. Each of the four dialogues represent the entire conversation between an approach controller and the pilot of a commercial flight approaching the airport to land. Controller and pilot are cooperatively performing the task of guiding the aircraft through the controller’s airspace to the approach gate, a point from which the pilot can complete the approach and landing without further guidance from the controller. This task is referred to as an Instrument Landing System (ILS) approach procedure. For a detailed description of the ILS approach procedure from the controller’s standpoint, see *Air Traffic Control* (FAA 1989). The pilot’s view is described in *Airman’s Information Manual* (FAA 1991). Rosenbaum (1988) provides a good explanation for the non-pilot.

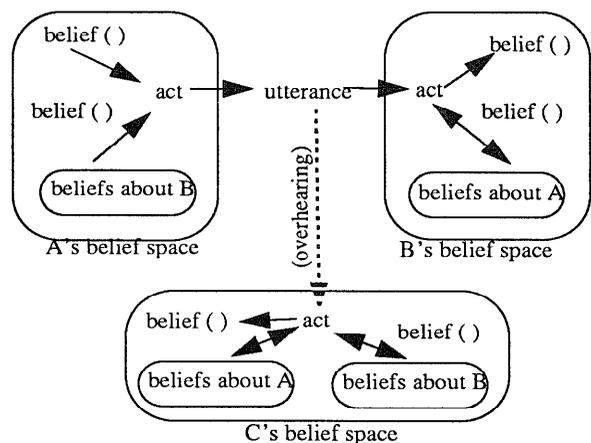
In ATC dialogue, the conversants often hold strong expectations about what the other knows and what the other will say. Pilot and controller are both trained in ILS approach procedures and each knows at least approximately what the other should do under normal circum-

stances. The system is not infallible, however, and circumstances are not always normal. The purpose of much of their dialogue, then, is to establish and confirm the mutuality of information that each thinks the other probably already knows or expects. Our goal, then, is to build a working model of mutuality that reflects this process, with representations that are usefully able to express differences among beliefs held by a single party, mutually held by some combination of parties, and perceived as held by the same or other combinations. The speech acts of the parties should reflect the natural consequences of their individual needs to achieve mutual understanding: Conversants with goals that require mutuality for their achievement should act appropriately; similarly capable conversants with goals that do not need mutuality should refrain from such actions.

### A Computational Model of ATC Dialogue

Our model of mutuality in multiparty dialogue encompasses relations among acts, utterances, and beliefs. Figure 1 summarizes this conceptual model. Conversants are modeled as autonomous agents, each having a separate belief space. An agent’s beliefs may include beliefs about another agent’s beliefs, depicted as smaller areas within an agent’s belief space. Agents communicate in a multi-step process:

- Agent A forms the intention to perform an act directed toward Agent B. This intention is based on A’s beliefs about B’s beliefs along with other beliefs that A holds.
- A’s intended act is expressed as an utterance. This utterance is transmitted to other agents in the system.
- B interprets this utterance as an act based on B’s own beliefs and B’s beliefs about A’s beliefs. Note that if B’s beliefs about A are in error or incomplete, B may infer a different act than A intended (misunderstanding).
- Agent B’s belief space is updated to reflect the effects that A’s act had on B’s beliefs.



**Figure 1. Conceptual Model of Agent Interaction**

Agent C represents an overhearer, an agent who hears A's utterance but is not the intended target (shown by a gray line from utterance to agent). Agent C interprets the overheard utterance in light of C's beliefs, including C's beliefs about the beliefs of A and B. Of course, C may arrive at an interpretation of the utterance that differs from that of either A or B; overhearers are at a disadvantage in following a conversation in that they do not actively participate in the collaborative process of reaching mutual belief (Schober & Clark 1989). However, overhearing plays a crucial role in ATC communications. Pilots maintain their situational awareness by monitoring the transmissions of others on the frequency and may respond or refer to dialogue that was not directed to them (Dunham 1991). The representation therefore explicitly supports multiple conversants by accounting for the effect that an utterance has on overhearers, including mutuality of belief judgements involving multiple agents.

Misunderstanding is modeled in terms of inconsistent beliefs; conversants' mental models of the conversation may have diverged without the conversants realizing it. A person may then take an action intended to embody a certain speech act, only to be surprised when the other's response indicates that a completely different speech act was understood. Similarly, disagreements can be modeled in terms of the conversants' beliefs about other conversants' differing beliefs. This model therefore supports both recognized and unrecognized belief inconsistency.

This conceptual model is similar in style to that proposed by Allen and Traum (Allen 1991; Traum 1991; Traum & Hinkelman 1992), in that a distinction is made between privately-held and shared knowledge in understanding a task-oriented conversation. Allen's TRAINS model is built around the negotiation and recognition of plans, however, and explicitly assumes that there are exactly two agents: system and person. The *ViewGen* algorithm (Ballim, Wilks, & Barden 1991) uses multiple environments to permit the system to ascribe and reason about conflicting beliefs among multiple agents. *ViewGen*'s emphasis is on understanding metaphor in text, however, and not on understanding task-oriented conversations. Here, we complement these approaches by providing a systematic representation for mutuality of belief.

## Representing Belief

An agent's understanding of the world is represented as a set of beliefs of the form

**belief** (*proposition*, *truth-value*, *belief-groups*).

The *proposition* clause represents the item about which a belief is held. An agent may hold beliefs about the state of the world, about acts that have **occurred**, and about acts that agents **intend**; these last two represent an agent's memory of past actions and expectations of future actions.

The *truth-value* clause represents the truth value that the belief-group assigns to the *proposition* (as understood by the agent in whose belief space the belief appears). This clause may take the values **true**, **false**, or **unknown**.

Clauses which do not appear in the agent's belief space are considered to have a truth-value of **unknown**. The *truth-value* does not indicate the agent's opinion about the "fact" represented in the *proposition*; rather, it reflects the agent's understanding of the beliefs of the belief group. This indirection allows an agent to ascribe beliefs to other agents.

A *belief-group* is a set of agents who mutually hold the belief. Each belief thus has one or more associated *belief-groups* that express the agent's view of the mutuality status of the proposition. For example, the belief

```
belief(altitude([sun512],5000),
      true,
      [[approach],[sun512],[approach,sun512]])
```

means the agent believes that the approach controller, and the pilot of Sundance 512 each believe (individually) that Sundance 512 is at 5000 feet; the agent also believes that controller and pilot have established that they mutually believe the pilot is at 5000 feet.

While we have paid special attention to mutuality, this representation clearly does not solve every problem associated with beliefs about actions in the world. Nevertheless, the representation is adequate for the purposes of demonstrating the effectiveness of our theory of conversational mutuality; its limitations mainly concern "domain" knowledge, particularly with respect to temporal and conditional actions. Thus the belief representation does not accommodate conditions that are in the process of changing from one state to another—that, for example, an aircraft is slowing from 190 knots to 170 knots. It does not capture notions of time; a more general representation would add a time interval argument to indicate when the agent believed the action would or did occur. It also does not capture events that take place without being triggered by the actions of some agent. As we have noted, though, these limitations are acceptable for this representation because our purpose is to model mutuality dialogue and not to model events in general.

## Core Speech Acts

In implementing and testing the belief model, we defined a small set of speech acts augmented with a preliminary set of domain acts and states sufficient to represent the dialogues being studied. These acts are: **acknowledge**, **authorize**, **contact**, **direct**, **report**, **request**. Acts are modeled at a high level, with preconditions and effects represented only to the degree necessary to understand the communication. The speech act representation is defined more fully in Ward (1992). Note that this is not a complete list of the acts needed to explain all ATC dialogue; it represents only the acts needed to explain typical dialogues seen in the ILS approach task.

## Dialogue Rules

ATC dialogue exhibits a strong task orientation that is reflected in the structure of the dialogue. Intentions to perform some act are assumed to be motivated by task

considerations, while the intentions to respond to others' acts are motivated primarily by mutuality considerations. The rules defined for the approach task dialogue are summarized in Table 1. Nine rules in four categories were needed to account for the four dialogues studied.

**Mutuality rules** establish and update agents' beliefs about the mutuality of knowledge. Included in this category are rules to translate an agent's perception of some act into a belief that the act occurred. These represent a rudimentary first step toward act recognition and interpretation.

**Expectation rules** capture expectations about an agent's response to another agent's acts. Like planning rules, they build expectations about agents' intended actions. Where planning rules seek to capture the procedures of complex domain tasks, however, expectation rules attempt to embody more basic rules of behavior in this domain: that agents acknowledge acts directed toward them; that agents follow instructions if able to do so; that agents confirm information that is reported to them.

**Planning rules** attempt to capture the procedural aspects of the domain's communications-oriented tasks, thus representing a rudimentary planning function.

**Performative rules** translate agents' intentions into acts. This category currently includes only a single, very general rule, **Perform Intentions**; it embodies the principle that agents carry out their intentions if possible. It locates an agent's unfulfilled intention to perform an act, checks that the preconditions for performing that act are satisfied, then does it.

### Simulating ATC Dialogue

To validate the mutuality model, we built a working implementation and tested it in saso (Novick 1990), a rule-based shell developed in Prolog as a tool for modeling multi-

agent problems involving simultaneous actions and subjective belief. The conversants are represented by computational agents (rules implemented as saso operators) that communicate using the acts defined in the model. Saso uses forward-chaining control to simulate the parallel execution of multiple rule-based agents with separate belief spaces. A user-defined set of extended STRIPS-style operators is used to specify the behavior of agents in terms of preconditions and effects. Agents communicate through acts; when an agent performs an act, clauses representing the action are posted to the memories of all agents in the system. The conflict resolution strategy was defined in terms of the rule types: Mutuality rules are applied in preference to expectation rules, expectation rules in preference to planning rules, and performative rules are triggered only when no other rules apply. This conflict resolution strategy reflects the idea that agents first update beliefs about their knowledge, then form expectations about the future behavior of agents, plan their own actions, and finally act.

For each dialogue, the inputs to the simulation were an initial state for each agent, a list of external non-conversation domain events that each agent would encounter, and a set of rules representing knowledge of mutuality and the domain task. The outputs were speech acts and changes in the agents' beliefs. The simulations were allowed to run to completion.

The simulations were validated by comparing their speech-act outputs on an act-for-act basis with a speech-act coding of the original ATC dialogues. To determine effects of overhearing, the internal belief states of the agents were compared with a manual analysis of mutuality of beliefs in the original dialogues.

Our first concern was to validate the basic model and rule set by simulating a single dialogue. The model's ability to track the belief states of overhearers was then tested by enlarging the simulation to include an agent who lis-

Table 1: Dialogue Rules for the Approach Task

	Rule	Purpose	Effect
<b>Mutuality Rules</b>	<ul style="list-style-type: none"> <li>• Act to Belief</li> <li>• Mark Belief Mutual</li> <li>• Mark Act Mutual</li> </ul>	<ul style="list-style-type: none"> <li>• Perceive acts.</li> <li>• Form beliefs about the mutuality of information.</li> <li>• Form beliefs about the mutuality of acts.</li> </ul>	<ul style="list-style-type: none"> <li>• Adds belief that act occurred</li> <li>• Updates mutuality of information</li> <li>• Updates mutuality of belief that act occurred</li> </ul>
<b>Expectation Rules</b>	<ul style="list-style-type: none"> <li>• Expect Acknowledge Act</li> <li>• Expect Follow Instructions</li> <li>• Expect Confirm Report</li> </ul>	<ul style="list-style-type: none"> <li>• Expect agents to acknowledge acts directed toward them.</li> <li>• Expect agents to follow instructions.</li> <li>• Expect agent to confirm information that is reported to them.</li> </ul>	<ul style="list-style-type: none"> <li>• Adds expectation that agent will acknowledge act.</li> <li>• Adds expectation that agent will follow instructions.</li> <li>• Adds expectation that agents will confirm report.</li> </ul>
<b>Planning Rules</b>	<ul style="list-style-type: none"> <li>• Plan Initial Contact</li> <li>• Plan Approach Clearance</li> </ul>	<ul style="list-style-type: none"> <li>• Pilot intends to contact controller.</li> <li>• Controller intends to authorize pilot to fly approach procedure</li> </ul>	<ul style="list-style-type: none"> <li>• Breaks task into a series of subgoals.</li> <li>• Breaks task into a series of subgoals.</li> </ul>
<b>Performative Rules</b>	<ul style="list-style-type: none"> <li>• Perform Intentions</li> </ul>	<ul style="list-style-type: none"> <li>• Agents attempt to carry out their intentions.</li> </ul>	<ul style="list-style-type: none"> <li>• Agent performs act.</li> </ul>

tened but did not participate in the conversation. Next, the model was tested against each of the other three dialogues from the protocol. Finally, the model was tested against the four dialogues as they were actually interleaved in the corpus.

## Results

The simulations ran to completion and correctly tracked the belief states of conversants and overhearers. With the exceptions noted below, the speech acts produced in the simulation corresponded on an act-for-act basis with the baseline codings. Agents' beliefs about the current state of all ongoing conversations were maintained, with agents responding appropriately to acts directed toward them and to events affecting them. Using the same rules as the active conversants, the overhearing agents "recognized" acts, formed expectations about the next actions of the conversants, and made judgments about the mutuality of knowledge and the evolving conversational state of the conversants.

Some aspects of the original dialogues were not reproduced in the simulations. As a consequence of the decision to avoid detailed event simulation, the model does not currently support directives that are not intended to be carried out immediately, such as an instruction to contact another controller when the aircraft reaches a certain point in its route. Such instructions were simulated as if they were to be carried out immediately. Also, the current rule set does not allow for an agent's inability to perform some action, e.g., a pilot's inability to see another aircraft; this aspect of the test dialogues was not reproduced. The model does not currently represent ranges of conditions, so ranges were instead represented as single values. Because saso simulates perfect understanding by conversants, correction sub-dialogues were not investigated in these simulations.

The model addresses only one aspect of the approach controller's responsibilities; it should be extended with substantial domain-task information to encompass a significant fraction of ATC dialogue. Also, there is clearly a need for many rules that weren't required for simulating the particular dialogues in this corpus, such as rules for reporting that an agent does not intend to comply with a direction. Another limitation to the generality of this model is the lack of a domain-specific planner. In the present study, this detailed planning function is simulated by introducing agents' intentions through saso's event queue. Although a detailed planner on the level of Weson's (1977) expert-level simulation of controller functions is probably not required for speech understanding, a certain level of knowledge of expected outcome of domain events is needed to motivate certain aspects of ATC dialogue.

## Conclusion

This work is a step toward developing a computational representation multiparty mutuality, which we have shown to be adequate for typical air traffic control dialogues. Com-

plete conversations were computationally simulated at the speech act level, demonstrating that real-world collaborative conversations can be motivated and represented in terms of belief and expectation, task and event. The goal of this study was to develop a computational model of air traffic control dialogue that would:

- Support reasoning about the beliefs and intentions of the conversants.
- Model agents' beliefs about other agents' beliefs.
- Capture sufficient domain and physical context to model the exchanges, particularly the evolving context of the dialogue itself.
- Support multiple (more than two) conversants, including overhearers.
- Permit different agents to hold different, possibly inconsistent, beliefs, and to support reasoning about the mutuality or inconsistency of conversants' beliefs.

We tested the model against these goals by simulation. Using a small set of rules, several actual dialogues were successfully simulated at the speech act level. The belief states of multiple conversational participants and overhearers were modeled through the course of several overlapping dialogues. Agents' beliefs were updated in response to the evolving conversational context as agents negotiated their tasks while maintaining a situational awareness of other agents' activities.

In modeling complete ATC dialogues in terms of the beliefs and goals of the conversants, typical patterns emerged. Conversants separately form similar beliefs based on separate inputs, then attempt to confirm the mutuality of the belief. As utterances are made and acknowledged, speaker and hearer form beliefs about what speech acts are intended by the utterance and form expectations that the act will be acknowledged and confirmed. These expectations form a strong context that permit the conversants to easily recognize the speech acts motivating non-standard responses that might otherwise be ambiguous.

A key contribution of this model is the representation of mutuality of belief in terms of belief sets. Belief sets capture both an agent's understanding of who believes a given piece of information and the mutuality that the agents holding that belief have established among themselves. This representation allows an agent to hold beliefs about other agents' possibly conflicting beliefs ("I believe A but he believes B.") as well as allowing agents to hold beliefs about the mutuality of their knowledge ("She and I have established that we both believe A; he should also believe A, but we have not yet confirmed that."). The belief set representation is flexible enough to represent and reason about mutuality combinations involving any number of agents, thus supporting the modeling of multi-agent conversations.

In this paper, then, we have proposed a representation for mutuality of belief that supports reasoning and action by multiple agents. We have shown the utility of the belief-groups representation through its use in a simulation of ATC dialogues. As part of the simulation, we developed a set of rules that maintain mutuality in this domain and that

are general enough to support straightforward extension to other domains involving cooperative tasks.

### Acknowledgments

The authors gratefully acknowledge the assistance of Larry Porter and Caroline Sousa.

### References

Allen, J. F. 1991. Discourse Structure in the TRAINS Project. In Proceedings of the Fourth DARPA Workshop on Speech and Natural Language.

Austin, J. L. 1962. *How To Do Things With Words*. London: Oxford University Press.

Ballim, A.; Wilks, Y.; and Barnden, J. 1991. Belief Ascription, Metaphor, and Intensional Identification. *Cognitive Science* 15:133-171.

Clark, H. H., and Marshall, C. R. 1981. Definite Reference and Mutual Knowledge. In I. A. Sag (Ed.), *Elements of discourse understanding*. Cambridge: Cambridge University Press.

Clark, H. H., and Wilkes-Gibbs, D. 1986. Referring as a Collaborative Process. *Cognition* 22:1-39.

Clark, H. H., and Schaefer, E. F. 1989. Contributing to Discourse. *Cognitive Science* 13:259-294.

Cohen, P. R., and Perrault, C. R. 1979. Elements of a Plan-based Theory of Speech Acts. *Cognitive Science* 3(3):177-212.

Cohen, P. R. 1984. The Pragmatics of Referring and the Modality of Communication. *Computational Linguistics* 10(2):97-146.

Dunham, S. 1991. Personal communication.

Federal Aviation Administration 1989. *Air Traffic Control* 7110.65f.

Federal Aviation Administration 1991. *Airman's Information Manual*.

Findler, N., and Lo, R. 1988. An Examination of Distributed Planning in the World of Air Traffic Control. In A. Bond and L. Gasser (eds.), *Readings In Distributed Artificial Intelligence*. San Mateo, Ca: Morgan Kaufman.

Grosz, B. J., and Sidner, C. L. 1986. Attention, Intentions, and the Structure of Discourse. *Computational Linguistics* 12(3):175-204.

Litman, D. J., and Allen, J. F. 1987. A Plan Recognition Model for Subdialogues in Conversations. *Cognitive Science* 11:163-200.

Novick, D. G. 1988. *Control Of Mixed-initiative Discourse Through Meta-locutionary Acts: A Computational Model*, Technical Report No. CIS-TR-88-18, Department of Computer and Information Science, University of Oregon.

Novick, D. G. 1990. Modeling Belief and Action in a Multi-agent System. In B. Zeigler & J. Rozenblit (eds.), *AI Simulation and Planning in High Autonomy Systems*. Los Alamitos, CA: IEEE Computer Society Press.

Oviatt, S. L., and Cohen, P. R. 1988. Discourse Structure and Performance Efficiency in Interactive and Noninteractive Spoken Modalities. Technical Note 454, SRI International.

Power, R. 1979. The Organization Of Purposeful Dialogues, *Linguistics* 17:107-152.

Rosenbaum, S. L. 1988. *A User's View Of The Air Traffic Control (ATC) System*. Internal Memorandum 46321-881130-01.IM, AT&T Bell Laboratories.

Schober, M. F., and Clark, H. H. 1989. Understanding by Addressees and Overhearers. *Cognitive Psychology* 21:211-232.

Searle, J. R. 1969. *Speech Acts: An Essay In The Philosophy Of Language*. Cambridge: Cambridge University Press.

Searle, J. R. 1975. Indirect Speech Acts. In J. L. Morgan (ed.) *Syntax and Semantics, Volume 3: Speech Acts*. New York: Academic Press.

Searle, J. R., and Vanderveken, D. 1985. *Foundations of Illocutionary Logic*. Cambridge: Cambridge University Press.

Suchman, L. A. 1987. *Plans and Situated Actions*. Cambridge: Cambridge University Press.

Traum, D. R. 1991. Towards a Computational Theory of Grounding in Natural Language Conversation, Technical Report No. 401, Computer Science Department, University of Rochester.

Traum, D., and Hinkelman, E. 1992. Conversation Acts in Task-Oriented Spoken Dialogue, Technical Report 425, Department of Computer Science, University of Rochester. (To appear in *Computational Intelligence*, 8(3)).

Ward, K.; Novick, D. G.; and Sousa, C. 1990. Air Traffic Control Communications at Portland International Airport, Technical Report, CS/E 90-025, Department of Computer Science and Engineering, Oregon Graduate Institute.

Ward, K. 1992. A Speech Act Model of Air Traffic Control Dialogue. M.S. thesis, Department of Computer Science and Engineering, Oregon Graduate Institute.

Wesson, R. B. 1977. Problem-solving with Simulation in the World of an Air Traffic Controller. Ph.D diss., University of Texas at Austin.

Winograd, T., and Flores, F. 1986. *Understanding Computers and Cognition*. Norwood, NJ: Ablex.