Speech Acts as a Basis for Understanding Dialogue Coherence

by

C. Raymond Perrault and James F. Allen
Dept. of Computer Science
University of Toronto
Toronto Canada

and

Philip R. Cohen
Bolt Beranek and Newman
Cambridge Mass.

1. Introduction

Webster's dictionary defines "coherence" as "the quality of being logically integrated, consistent, and intelligible". If one were asked whether a sequence of physical acts being performed by an agent was coherent, a crucial factor in the decision would be whether the acts were perceived as contributing to the achievement of an overall goal. In that case they can frequently be described briefly, by naming the goal or the procedure executed to achieve it. Once the intended goal has been conjectured, the sequence can be described as a more or less correct, more or less optimal attempt at the achievement of the goal.

One of the mainstreams of AI research has been the study of problem solving behaviour in humans and its simulation by machines. This can be considered as the task of transforming an initial state of the world into a goal state by finding an appropriate sequence of applications of operators from a given set. Each operator has two modes of execution: in the first it changes the "real world", and in the second it changes a model of the real world. Sequences of these operators we call plans. They can be constructed, simulated, executed, optimized and debugged. Operators are usually thought of as achieving certain effects and of being applicable only when certain preconditions hold.

The effects of one agent executing his plans may be observable by other agents, who, assuming that these plans were produced by the first agent's plan construction algorithms, may try to infer the plan being executed from the observed changes to the world. The fact that this inferencing may be intended by the first agent underlies human communication.

Each agent maintains a model of the world, including a model of the models of other agents. Linguistic utterances are the result of the execution of operators whose effects are mainly on the models that the speaker and hearer maintain of each other. These effects are intended by the speaker to be produced partly by the hearer's recognition of the speaker's plan.

This view of the communication process is very close in spirit to the Austin-Grice-Strawson-Searle approach to illocutionary acts, and indeed was strongly influenced by it. We are working on a theory of speech acts based on the notions of plans, world models, plan construction and plan recognition. It is intended that this theory should answer questions such as:

(1) Under what circumstances can an observer believe that a speaker has sincerely and non-defectively performed a particular illocutionary act in producing utterance for a hearer? The observer could also be the hearer or speaker.

(2) What changes does the successful execution of a speech act make to the speaker's model of the hearer, and to the hearer's model of the speaker?

(3) How is the meaning (sense/reference) of an utterance x related to the acts that can be performed in uttering x?

A theory of speech acts based on plans must specify at least the following:

(1) A Planning System: a language for describing states of the world, a language for describing operators and algorithms for plan construction and plan inference. Semantics for the languages should also be given.

(2) Definitions of speech acts as operators in the planning system. What are their effects? When are they applicable? How can they be realized in words?

To make possible a first attempt at such a theory we have imposed several restrictions on the system to be modelled.

(1) Any agent Al's model of another agent A2 is defined in terms of "facts" that Al believes A2 believes, and goals that Al believes A2 is attempting to achieve. We are not attempting to model obligations, feelings etc.

(2) The only speech acts we try to model are some that appear to be definable in terms of beliefs and goals, namely REQUEST and INFORM. We have been taking these to be prototypical members of Searle's "directive" and "representative" classes (Searle (1976)). We represent questions as REQUESTs to INFORM. These acts are interesting for they have a wide range of syntactic realizations, and account for a large proportion of everyday utterances.

(3) We have limited ourselves so far to the study of so-called task-oriented dialogues which we interpret to be conversations between two agents cooperating in the achievement of a single high-level goal. These dialogues do not allow changes in the topic of discourse but still display a wide range of linguistic behaviour.

Much of our work so far has dealt with the problem of generating plans containing REQUEST and INFORM, as well as non-linguistic operators. Suppose that an agent is attempting to achieve some task, with incomplete knowledge of that task and of the methods to complete it, but with some knowledge of the abilities of another agent. How can the first agent make use of the abilities of the second? Under what circumstances can the first usefully produce utterances to transmit or acquire facts and goals? How can he initiate action on the part of the second?

We view the plan related aspects of language generation and recognition as indissociable, and strongly related to the process by which agents cooperate in the achievement of goals. For example, for agent2 to reply "It's closed" to agentl's query "Where's the nearest service station?" seems to require him to infer that agentl wants to make use of the service station which he could not do if it were closed. The reply "Two blocks east" would be seen as misleading if given alone, and unnecessary if given along with "It's closed". Thus part of cooperative behaviour is the detection by one agent of obstacles in the plans he believes the other agent holds, possibly followed by an attempt to overcome them. We claim that speakers expect (and intend) hearers to operate this way and therefore that any hearer can assume that inferences that he can draw based on knowledge that is shared with the speaker are in fact intended by the speaker. These processes underlie our

analysis of indirect speech acts (such as "Can you pass the salt?") - utterances which appear to result from one illocutionary act but can be used to perform another.

Section 2 of this paper outlines some requirements on the models which the various agents must have of each other. Section 3 describes the planning operators for REQUEST and INFORM, and how they can be used to generate plans which include assertions, imperatives, and several types of questions.

Section 4 discusses the relation between the operators of section 3 and the linguistic sentences which can realize them. We concentrate on the problem of identifying illocutionary force, in particular on indirect speech acts. A useful consequence of the illocutionary force identification process is that it provides a natural way to understand some elliptical utterances, and utterances whose purpose is to acknowledge, correct or clarify interpretations of previous utterances.

A critical part of communication is the process by which a speaker can construct descriptions of objects involved in his plans such that the hearer can identify the intended referent. Why can someone asking "Where's the screwdriver?" be answered with "In the drawer with the hammer" if it is assumed he knows where the hammer is, but maybe by "In the third drawer from the left" if he doesn't. How accurate must descriptive phrases be? Section 5 examines how the speaker and hearer's models of each other influence their references. Finally, section 6 contains some ideas on future research.

Most examples in the paper are drawn from a situation in which one participant is an information clerk at a train station, whose objective is to assist passengers in boarding and meeting trains. The domain is obviously limited, but still provides a natural setting for a wide range of utterances, both in form and in intention.

2. On models of others

In this section we present criteria that one agent's model of another ought to satisfy. For convenience we dub the agents SELF and OTHER. Our research has concentrated on modelling beliefs and goals. We claim that a theory of language need not be concerned with what is actually true in the real world: it should describe language processing in terms of a person's beliefs about the world. Accordingly, SELF's model of OTHER should be based on "believe" as described, for example, in Hintikka(1962) and not on "know" in its sense of "true belief".

Henceforth, all uses of the words "know" and "knowledge" are to be treated as synonyms for "believe" and "beliefs". We have neglected other aspects of a model of another, such as focus of attention (but see Grosz(1977)).

Belief

Clearly, SELF ought to be able to distinguish his beliefs about the world from what he believes other believes. SELF ought to have the possibility of believing a proposition P, of believing not-P, or of being ignorant of P. Whatever his stand on P, he should also be able to believe that OTHER can hold any of these positions on P. Notice that such disagreements cannot be represented if the representation is based on "know" as in Moore(1977).

SELF's belief representation ought to allow him to represent the fact that OTHER knows whether some proposition P is true, without SELF's having to know which of P or ~P he does believe. Such information can be represented as a disjunction of beliefs (e.g., OR(OTHER BELIEVE P, OTHER BELIEVE ~P)). Such disjunctions are essential to the planning of yes/no questions.

Finally, a belief representation must distinguish between situations like the following:

1. OTHER believes that the train leaves from gate 8.
2. OTHER believes that the train has a departure gate.
3. OTHER knows what the departure gate for the train is.

Case 1 can be represented by a proposition that contains no variables. Case 2 can be represented by a belief of a quantified proposition -- i.e.,

OTHER BELIEVE (
    $\exists$ x (the y : GATE(TRAIN,y) = x))

However, case 3 is represented by a quantified belief namely,

$\exists$ x OTHER BELIEVE
        (the y : GATE(TRAIN,y) = x)

The formal semantics such beliefs have been problematic for philosophers (cf. Quine (1956) and Hintikka (1962)). Our approach to them is discussed in Cohen (1978). In Section 3, we discuss how quantified beliefs are used during planning, and how they can be acquired during conversation.

Want

Any representation of OTHER's goals (wants) must distinguish such information from: OTHER's beliefs, SELF's beliefs and goals, and (recursively) from the other's model of someone else's beliefs and goals. The representation for WANT must also allow for different scopes of quantifiers. For example, it should distinguish between the readings of "John wants to take a train" as "There is a specific train which John wants to take" or as "John wants to take any train". Finally it should allow arbitrary embeddings with BELIEVE. Wants of beliefs (as in "SELF wants OTHER to believe P") become the reasons for telling P to OTHER, while beliefs of wants (e.g., SELF Believes SELF wants P) will be the way to represent SELF's goals P.

Levels of Embedding

A natural question to ask is how many levels of belief embedding are needed by an agent capable of participating in a dialogue. Obviously, to be able to deal with a disagreement, SELF needs two levels (SELF BELIEVE and SELF BELIEVE OTHER BELIEVE). If SELF were to lie to OTHER, he would have to be able to believe some proposition P (i.e. SELF BELIEVE (P)), while OTHER believes that SELF believes not P (i.e. SELF BELIEVE OTHER BELIEVE SELF BELIEVE (~P)), and hence he would need at least three levels.

We show in Cohen (1978) how one can represent, in a finite fashion, the unbounded number of beliefs created by any communication act or by face-to-face situations. The finite representation, which employs a circular data structure, formalizes the concept of mutual belief (cf. Schiffer (1972)). Typically, all these levels of belief embedding can be represented in three levels, but theoretically, any finite number are possible.

3. Using a Model of the Other to Decide What to Say

As an aid in evaluating speech act definitions, we have constructed a computer program, OSCAR, that plans a range of speech acts. The goal of the program is to characterize a speaker's capacity to issue speech acts by predicting, for specified situations, all and only those speech acts that would be appropriately issued by a person under the circumstances. In this section, we will make reference to prototypical speakers by way of the OSCAR program, and to hearers by way of the program's user.

Specifially, the program is able to:

- Plan REQUEST speech acts, for instance a speech act that could be realized by

"Please open the door", when its goal is to get the user to want to perform some action.

- Plan INFORM speech acts, such as one that could be realized by "The door is locked", when its goal is to get the user to believe some proposition.

- Combine the above to produce multiple speech acts in one plan, where one speech act may establish beliefs of the user that can then be employed in the planning of another speech act.

- Plan questions as requests that the user inform, when its goal is to believe something and when it believes that the user knows the answer.

- Plan speech acts incorporating third parties, as in "Ask Tom to tell you where the key is and then tell me."

To illustrate the planning of speech acts, consider first the following simplified definitions of REQUEST and INFORM as STRIPS-like operators (cf. Fikes and Nilsson (1971)). Let SP denote the speaker, H the hearer, ACT some action, and PROP some proposition. Due to space limitations, the intuitive English meanings of the formal terms appearing in these definitions will have to suffice as explanation.

REQUEST(SP,H,ACT)
preconditions:
SP BELIEVE H CANDO ACT
SP BELIEVE H BELIEVE H CANDO ACT
SP BELIEVE SP WANT TO REQUEST
effects:
H BELIEVE SP BELIEVE SP WANT H TO ACT

INFORM(SP,H,PROP)
preconditions:
SP BELIEVE PROP
SP BELIEVE SP WANT TO INFORM
effects:
H BELIEVE SP BELIEVE PROP

The program uses a simplistic backward-chaining algorithm that plans actions when their effects are wanted as subgoals that are not believed to be true. It is the testing of preconditions of the newly planned action before creating new subgoals that exercises the program's model of its user. We shall briefly sketch how to plan a REQUEST.

Every action has "want preconditions", which specify that before an agent does that action, he must want to do it. OSCAR plans REQUEST speech acts to achieve precisely this precondition of actions that it wants the user to perform. Similarly, the goal of the user's believing some proposition PROP becomes OSCAR's reason for planning to INFORM him of PROP.

Suppose, for example, that OSCAR is outside a room whose door is closed and that it believes that the user is inside. When planning to move itself into the room, it might REQUEST that the user open the door. However, it would only plan this speech act if it believed that the user did not already want to open the door and if it believed (and believed the user believed) that the preconditions to opening the door held. If that were not so, OSCAR could plan additional INFORM or REQUEST speech acts. For example, assume that to open a door one needs to have the key and OSCAR believes the user doesn't know where it is. Then OSCAR could plan "Please open the door. The key is in the closet". OSCAR thus employs its user model in telling him what it believes he needs to know.

Mediating Acts and Perlocutionary Effects

The effects of INFORM (and REQUEST) are modelled so that the hearer's believing P (or wanting to do ACT) is not essential to the successful completion of the speech act. Speakers, we claim, cannot influence their hearers' beliefs and goals directly. Thus, the perlocutionary effects of a speech act are not part of that act's definition. We propose, then, as a principle of communication that a speaker's purpose in sincere communication is to produce in the hearer an accurate model of his mental state.

To bridge the gap between the speech acts and their intended perlocutionary effects, we posit mediating acts, named CONVINCE and DECIDE, which model what it takes to get someone to believe something or want to do something. Our current analysis of these mediating acts trivializes the processes that they are intended to model by proposing that to convince someone of something, for example, one need only get that person to know that one believes it.

Using Quantified Beliefs -- Planning Questions

Notice that the precondition to OSCAR's getting the key -- knowing where it is -- is of the form:

∃ x OSCAR BELIEVE
     (the y : LOC(KEY,y) = x)

When such a quantified belief is a goal, it leads OSCAR to plan the question "Where is the key?" (i.e., REQUEST(OSCAR, USER, INFORM(USER, OSCAR, the y : LOC(KEY,y))). In creating this question, OSCAR first plans a CONVINCE and then plans the user's INFORM speech act, which it then tries to get him to perform by way of requesting.

The above definition of INFORM is inadequate for dealing with the quantified beliefs that arise in modelling someone else. This INFORM should be viewed as that version of the speech act that the planning agent (e.g., OSCAR) plans for itself to perform. A different view of INFORM, say INFORM-BY-OTHER, is necessary to represent acts of informing by agents other than the speaker. The difference between the two INFORMs is that for the first, the planner knows what he wants to say, but he obviously does not have such knowledge of the content of the second act.

The precondition for this new act is a quantified speaker-belief:

$\exists$ x USER BELIEVE
$\qquad$ (the y : LOC(KEY,y) = x)

where the user is to be the speaker. For the system to plan an INFORM-BY-OTHER act for the user, it must believe <u>that</u> the user knows where the key is, but <u>it</u> does not have to know that location! Similarly, the effects of the INFORM-BY-OTHER act is also a quantified belief, as in

$\exists$ x OSCAR BELIEVE
$\qquad$ USER BELIEVE
$\qquad\qquad$ (the y : LOC(KEY,y) = x)

Thus, OSCAR plans this INFORM-BY-OTHER act of the key's location in order to know where the user thinks the key is.

Such information has been lacking from all other formulations of ASK (or INFORM) that we have seen in the literature (e.g., Schank (1975), Mann et al. (1976), Searle (1969)). Cohen (1978) presents one approach to defining this new view of INFORM, and its associated mediating act CONVINCE.

4. <u>Recognizing</u> <u>Speech</u> <u>Acts</u>

In the previous section we discussed the structure of plans that include instances of the operators REQUEST and INFORM without explaining the relation between these speech acts and sentences used to perform them. This section sketches our first steps in exploring this relation. We have been particularly concerned with the problem of recognizing illocutionary force and propositional content of the utterances of a speaker. Detailed algorithms which handle the examples given in this section have been designed by J. Allen and are being implemented by him. Further details can be found in (Allen and Perrault 1978) and Allen's forthcoming Ph.D. dissertation.

Certain syntactic clues in an utterance such as its mood and the use of explicit performatives indicate what act

the speaker intends to perform, but, as is well known, utterances which taken literally would indicate one illocutionary force can be used to indicate another. Thus "Can you close the door?" can be a request as well as a question. These so-called indirect speech acts are the acid test of a theory of speech acts. We claim that a plan-based theory gives some insight into this phenomenon.

Searle(1975) correctly suggests that "In cases where these sentences <indirect forms of requests> are uttered as requests, they still have their literal meaning and are uttered with and as having that literal meaning". How then can they also have their indirect meaning?

Our answer relies in part on the fact that an agent participating in a cooperative dialogue must have processes to:

(1) Achieve goals based on what he believes.
(2) Adopt goals of other agents as his own.
(3) Infer goals of other agents.
(4) Predict future behaviour of other agents.

These processes would be necessary even if all speech acts were literal to account for exchanges where the response indicates a knowledge of the speaker's plan. For example

Passenger: "When does the next train to Montreal leave?"
Clerk   : "At 6:15 at Gate 7"
or
Clerk   : "There won't be one until tomorrow."

Speakers expect hearers to be executing these processes and they expect hearers to know this. Inferences that a hearer can draw by executing these processes based on information he thinks the speaker believes can be taken by the hearer to be intended by the speaker. This accounts for many of the standard examples of indirect speech acts such as "Can you close the door?" and "It's cold here". For instance, even if "It's cold here" is intended literally and is recognized as such, the helpful hearer may still close the window. When the sentence is uttered as a request, the speaker intends the hearer to recognize the speaker's intention that the hearer should perform the helpful behaviour.

If indirect speech acts are to be explained in terms of inferences speakers can expect of hearers, then a theory of speech acts must concern itself with how such inferences are controlled. Some heuristics are particularly helpful. If a chain of inference by the hearer has the speaker planning an action whose effects

are true before the action is executed, then the chain is likely to be wrong, or else must be continued further. This accounts for "Can you pass the salt?" as a request for the salt, not a question about salt-passing prowess. As Searle(1975) points out, a crucial part of understanding indirect speech acts is being able to recognize that they are not to be interpreted literally.

A second heuristic is that a chain of inference that leads to an action whose preconditions are known to be not easily achievable is likely to be wrong.

Inferencing can also be controlled through the use of expectations about the speaker's goals. Priority can be given to inferences which relate an observed speech act to an expected goal. Expectations enable inferencing to work top-down as well as bottom-up.

The use of expected goals to guide the inferencing has another advantage: it allows for the recognition of illocutionary force in elliptical utterances such as "The 3:15 train to Windsor?", without requiring that the syntactic and semantic analysis "reconstitute" a complete semantic representation such as "Where does the 3:15 train to Windsor leave?". For example, let the clerk assume that passengers want to either meet incoming trains or board departing ones. Then the utterance "The 3:15 train to Windsor?" is first interpreted as a REQUEST about a train to Windsor with 3:15 as either arrival or departure time. Only departing trains have destinations different from Toronto and this leads to believing that the passenger wants to board a 3:15 train to Windsor. Attempting to identify obstacles in the passenger's plan leads to finding that the passenger knows the time but probably not the place of departure. Finally, overcoming the obstacle then leads to an INFORM like "Gate 8".

Our analysis of elliptical utterances raises two questions. First, what information does the illocutionary force recognition module expect from the syntax and semantics? Our approach here has been to require from the syntax and semantics a hypothesis about the literal illocutionary force and a predicate calculus-like representation of the propositional content, but where undetermined predicates and objects could be replaced by patterns on which certain restrictions can be imposed. As part of the plan inferencing process these patterns become further specified.

The second question is: what should the hearer do if more than one path between the observed utterance and the expectations is possible? He may suspend plan deduction and start planning to achieve a goal which would allow plan deduction to continue. Consider the following example.

Passenger : When is the Windsor train?
Clerk : The train to Windsor?
Passenger : Yes.
Clerk : 3:15.

After the first sentence the clerk cannot distinguish between the expectations "Passenger travel by train to Windsor" and "Passenger meets train from Windsor", so he sets up a goal : (clerk believes passenger wants to travel) or (clerk believes passenger wants to meet train). The planning for this goal produces a plan that involves asking the passenger if he wants one of the alternatives, and receiving back the answer. The execution of this plan produces the clerk response "The train to Windsor?" and recognizes the response "Yes". Once the passenger's goal is known, the clerk can continue the original deduction process with the "travel to Windsor" alternative favoured. This plan is accepted and the clerk produces the response "3:15" to overcome the obstacle "passenger knows departure time".

5. Reference and the Model of the Other

We have shown that quantified beliefs are needed in deciding to ask someone a question. They are also involved, we claim, in the representation of singular definite noun phrases and hence any natural language system will need them. According to our analysis, a hearer should represent the referring phrase in a speaker's statement "The pilot of TWA 510 is drunk" by:

$$\exists \; x \; \text{SPEAKER BELIEVE}$$
$$(\text{the } y \; : \; \text{PILOT}(y,\text{TWA510}) = x \; \& $$
$$\text{DRUNK}(x))$$

This is the reading whereby the speaker is believed to "know who the pilot of TWA 510 is" (at least partially accounting for Donnellan's (1966) referential reading). This is to be contrasted with the reading of whoever is piloting that plane is drunk (Donnellan's attributive noun phrases). In this latter case, the existential quantifier would be inside the scope of the belief.

These existential presuppositions of definite referential noun phrases give one important way for hearers to acquire quantified speaker-beliefs. Such beliefs, we have seen, can be used as the basis for planning further clarification questions.

We agree with Strawson (1950) (and many others) that hearers understand referring phrases based on what they believe speakers intend to refer to.

Undoubtedly, a hearer will understand a speaker's (reference) intentions by using a model of that speaker's beliefs. Speakers, of course, know of these interpretation strategies and thus <u>plan</u> their referring phrases to take the appropriate referent within the hearer's model of them. A speaker cannot use private descriptions, nor descriptions that he thinks the hearer thinks are private, for communication.

For instance, consider the following variant of an example of Donnellan's (1966): At a party, a woman is holding a martini glass which Jones believes contains water, but of which he is certain everyone else believes (and believes he believes) contains a martini. Jones would understand that Smith, via question (1), but not via question (2) is referring to this woman.

(1) Who is the woman holding the martini?
(2) Who is the woman holding the water?

since Jones does not believe Smith knows about the water in her glass.

Conversely, if Jones wanted to refer to the woman in an utterance intended for Smith, he could do so using (1) but not (2) since in the latter case he would not think the hearer could pick out his intended referent.

Thus it appears that for a speaker to plan a successful singular definite referential expression requires that the speaker believe the expression he finally chooses have the right referent in the hearer's model of the speaker. Our concept of mutual belief can be used (as in Cohen (1978)) to ensure that the expression denotes appropriately in all further embedded belief models. This example is problematic for any approach to reference where a communicating party assumes that its reality is the only reality. Speakers and hearers can be "wrong" or "ignorant" and yet communication can still be meaningful and successful.

6. <u>Further Research</u>

We believe that speech acts provide an excellent way of explaining the relations between utterances in a dialogue, as well as relating linguistic to non-linguistic activity. Until we better understand the mechanisms by which conversants change the topic and goals of the conversation it will be difficult to extend this analysis beyond exchanges of a few utterances, in particular to non-task oriented dialogues. Fuller justification of our approach also requires its application to a much broader range of speech acts. Here the problem is mainly representational: how can we handle promises without first dealing with obligations, or warnings without the notions of danger and undesirability? We are currently considering an extension of the approach to understanding stories which report simple dialogue.

Much remains to be done on the representation of the abilities of another agent. A simple setting suggests a number of problems. Let one agent H be seated in a room in front of a table with a collection of blocks. Let another agent S be outside the room but communicating by telephone. If S believes that there is a green block on the table and wants it cleared, but knows nothing about any other blocks except that H can see them, then how can S ask H to clear the green block? The blocks S wants removed are those which are in fact there, perhaps those which he could perceive to be there if he were in the room. The goal seems to be of the form

S BELIEVE
$\forall$ x (x on the green block => S WANT
(x removed from green block))

but our planning machinery and definition of REQUEST are inadequate for generating "I request you to clear the green block".

We have not yet spent much time investigating the process of giving answers to How and Why questions, or to WH questions requiring an event description as an answer. We conjecture that because of the speech act approach answers to "What did he say?" should be found in much the same way as answers to "What did he do?" and that this parallelism should extend to other question types. The natural extension of our analysis would suggest representing "How did AGT achieve goal G?" as a REQUEST by the speaker that the hearer inform him of a plan by which AGT achieved G. We have not yet investigated the repercussions of this extension on the representation language.

Finally consider the following dialogue. Assume that S is a shady businessman, A his secretary.

A : IRS is on the phone.

S : I'm not here.

How is A to understand S's utterance? Although its propositional content is literally false, maybe even nonsensical, the utterance's intention is unmistakable. How tolerant does the understanding system have to be to infer its way to a correct interpretation? Must "I'm not here" be treated idiomatically?

Bibliography

Allen, J.F. and Perrault, C.R., "Participating in Dialogue: Understanding via Plan Deduction", 2nd National Conference of the Canadian Society for Studies in Computational Intelligence, Toronto, July, 1978.

Cohen, P.R., "On Knowing What to Say: Planning Speech Acts", TR118 Dept. of Computer Science, University of Toronto, 1978.

Donnellan, K., "Reference and Definite Description", The Philosophical Review, vol. 75, 1960, pp280-304. Reprinted in Semantics, Steinberg and Jacobovits, eds., Cambridge University Press, 1970.

Fikes, R. E. and Nilsson, N. J., 1970, "STRIPS: A new approach to the application of theorem proving", Artificial Intelligence 2, 1970.

Grosz, B. J., "The Representation and Use of Focus in Natural Language Dialogues", 5IJCAI, 1977.

Hintikka, K.J., Knowledge and Belief, Cornell University Press, 1962.

Mann, W.C., Moore, J.A., Levin, J.A.; "A Comprehension Model for Human Dialogue", 5IJCAI, 1977.

Moore, R.C.; "Reasoning about Knowledge and Action", 5IJCAI, 1977.

Quine, W.V., "Quantifiers and Propositional Attitudes", The Journal of Philosophy 53, (1956), 177-187.

Schiffer, S., Meaning, Oxford University Press, 1972.

Schank, R. and Abelson, R., "Scripts, Plans and Knowledge", 4IJCAI, 1975.

Searle, J. R., Speech Acts, Cambridge University Press, 1969.

Searle, J. R.; "Indirect Speech Acts" in Syntax and Semantics, Vol. 3: Speech Acts, Cole and Morgan (eds), Academic Press, 1975.

Searle, J. R., "A Taxonomy of Illocutionary Acts", Language, Mind and Knowledge, K. Gunderson (ed.), University of Minnesota Press, 1976.

Strawson, P. F., "On Referring", Mind, 1950.