

Elements of a Plan-Based Theory of Speech Acts*

PHILIP R. COHEN

Bolt Beranek and Newman Inc.

AND

C. RAYMOND PERRAULT

University of Toronto

This paper explores the truism that people think about what they say. It proposes that, to satisfy their own goals, people often *plan* their speech acts to affect their listeners' beliefs, goals, and emotional states. Such language use can be modelled by viewing speech acts as operators in a planning system, thus allowing both physical and speech acts to be integrated into plans.

Methodological issues of how speech acts should be defined in a plan-based theory are illustrated by defining operators for requesting and informing. Plans containing those operators are presented and comparisons are drawn with Searle's formulation. The operators are shown to be inadequate since they cannot be composed to form questions (requests to inform) and multiparty requests (requests to request). By refining the operator definitions and by identifying some of the side effects of requesting, compositional adequacy is achieved. The solution leads to a metatheoretical principle for modelling speech acts as planning operators.

1. INTRODUCTION

The Sphinx once challenged a particularly tasty-looking student of language to solve the riddle: "How is saying 'My toe is turning blue,' as a request to get off my toe, similar to slamming a door in someone's face?" The poor student stammered that in both cases, when the agents are trying to communicate something, they have analogous intentions. "Yes indeed" countered the Sphinx, "but what are those intentions?" Hearing no reply, the monster promptly devoured the poor student and sat back smugly to wait for the next oral exam.

*The research described herein was supported primarily by the National Research Council of Canada, and also by the National Institute of Education under Contract US-NIE-C-400-76-0116, the Department of Computer Science of the University of Toronto, and by a summer graduate student associateship (1975) to Cohen from the International Business Machines Corporation.

Contemporary philosophers have been girding up for the next trek to Giza. According to Grice (1957)¹, the slamming of a door communicates the slammer's anger only when the intended observer of that act realizes that the slammer wanted both to slam the door in his face and for the observer to believe that to be his intention. That is, the slammer intended the observer to recognize his intentions. Slamming caused by an accidental shove or by natural means is not a communicative act. Similarly, saying "My toe is turning blue" only communicates that the hearer is to get off the speaker's toe when the hearer has understood the speaker's intention to use that utterance to produce that effect.

Austin (1962) has claimed that speakers do not simply produce sentences that are true or false, but rather perform speech actions such as requests, assertions, suggestions, warnings, etc. Searle (1969) has adapted Grice's (1957) recognition of intention analysis to his effort to specify the necessary and sufficient conditions on the successful performance of speech acts. Though Searle's landmark work has led to a resurgence of interest in the study of the pragmatics of language, the intentional basis of communicative acts requires further elaboration and formalization; one must state for any communicative act, precisely which intentions are involved and on what basis a speaker expects and intends those intentions to be recognized.

The Sphinx demands a competence theory of speech act communication—a theory that formally models the possible intentions underlying speech acts. This paper presents the beginnings of such a theory by treating intentions as plans and by showing how plans can link speech acts with nonlinguistic behavior. In addition, an adequacy test for plan-based speech act theories is proposed and applied.

1.1 A Plan-based Theory of Speech Acts

Problem solving involves pursuing a goal state by performing a sequence of actions from an initial state. A human problem-solver can be regarded as "executing" a *plan* that prespecifies the sequence of actions to be taken. People can construct, execute, simulate, and debug plans, and in addition, can sometimes infer the plans of other agents from their behavior. Such plans often involve the communication of beliefs, desires and emotional states for the purpose of influencing the mental states and actions of others. Furthermore, when trying to communicate, people expect and want others to recognize their plans and may attempt to facilitate that recognition.

Formal descriptions of plans typically treat actions as *operators*, which are defined in terms of applicability conditions, called *preconditions*, *effects* that will be obtained when the corresponding actions are executed, and *bodies* that describe the means by which the effects are achieved. Since operators are repre-

¹See also (Strawson, 1964; Schiffer, 1972)

sentations, their preconditions, effects, and bodies are evaluated relative to the problem-solver's model of the world. We hypothesize that people maintain, as part of their models of the world, symbolic descriptions of the world models of other people. Our plan-based approach will regard speech acts as operators whose effects are primarily on the models that speakers and hearers maintain of each other.²

Any account of speech acts should answer questions such as:

—Under what circumstances can an observer believe that a speaker has sincerely and successfully performed a particular speech act in producing an utterance for a hearer? (The observer could also be the hearer or speaker.)

—What changes does the successful performance of a speech act make to the speaker's model of the hearer, and to the hearer's model of the speaker?

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

To achieve these ends, a theory of speech acts based on plans should specify at least the following:

—A planning system: a formal language for describing states of the world, a language for describing operators, a set of plan construction inferences, a specification of legal plan structures. Semantics for the formal languages should also be given.

—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?

As an illustration of this approach, this paper presents a simple planning system, defines the speech acts of requesting and informing as operators within that system, and develops plans containing direct requests, informs and questions (which are requests to inform). We do not, however, discuss how those speech acts can be realized in words.

We argue that a plan-based theory, unlike other proposed theories of speech acts, provides formal adequacy criteria for speech act definitions: given an initial set of beliefs and goals, the speech act operator definitions and plan construction inferences should lead to the generation of plans for those speech acts that a person could issue appropriately under the same circumstances.³ This adequacy criterion should be used in judging whether speech act definitions pass certain tests, in particular, the test of compositionality. For instance, since a speaker can request that a hearer do some arbitrary action, the operator definitions should show how a speaker can request a hearer to perform a speech act. Similarly, since one can inform a hearer that an action was done, the definitions should capture a speaker's informing a hearer that a speech act was performed. We show how a number of previous formulations of requesting and informing are

compositionally inadequate, and then develop definitions of informing that can be composed into questions.

Another goal of this research is to develop metatheoretical principles that state how to formulate speech act definitions to pass these adequacy tests. This paper proposes such a principle and shows how its application leads to compositionally adequate definitions for multiparty requests (as in "Ask Tom to open the door").

To simplify our problems in the early stages of theory construction, several restrictions on the communication situation that we are trying to model have been imposed:

—Any agent's model of another will be defined in terms of "facts" that the first believes the second believes, and goals that the first believes the second is attempting to achieve. We are not attempting to model obligations, feelings, etc.

—The only speech acts we try to model are requests, informs, and questions since they appear to be definable solely in terms of beliefs and goals. Requesting and informing are prototypical members of Searle's (1976) "directive" and "representative" classes, respectively, and are interesting since they have a wide range of syntactic realizations, and account for a large proportion of everyday utterances.

—We have limited ourselves to studying "instrumental dialogues"—conversations in which it is reasonable to assume that the utterances are planned and that the topic of discourse remains fixed. Typically, such dialogues arise in situations in which the conversants are cooperating to achieve some task-related goal (Deutsch, 1974), for example, the purchasing of some item. The value of studying such conversations relative to the structure of a task is that the conversants' plans can be more easily formalized.

1.2 A Competence Theory of Speech Acts

At least two interdependent aspects of a plan-based theory should be examined—the plans themselves, and the methods by which a person could construct or recognize those plans. This paper will be concerned with theories of the first aspect, which we shall term *competence* theories, analogous to competence theories of grammar (Chomsky, 1965). A plan-based competence theory of speech acts describes the *set of possible plans* underlying the use of particular kinds of speech acts, and thus states the conditions under which speech acts of those types are appropriate. Such descriptions are presented here in the form of a set of operator definitions (akin to grammatical "productions") and a specification of the ways in which plans are created from those operators.

The study of the second aspect aims for a *process* theory, which concerns *how* an ideal speaker/hearer chooses one (or perhaps more than one) plan out of the set of possible plans. Such a theory would characterize how a speaker decides what speech act to perform and how a hearer identifies what speech act was performed by recognizing the plan(s) in which that utterance was to play a part.

By separating out these two kinds of theoretical endeavors we are not claiming that one can study speech act competence totally divorced from issues of processing. On the contrary, we believe that for a (careful) speaker to issue a particular speech act appropriately, she must determine that the hearer's speech

²This approach was inspired by Bruce and Schmidt (1974), Bruce (1975), and Schmidt (1975). This paper can be viewed as supplying methodological foundations for the analyses of speech acts and their patterned use that they present.

³Though this could perhaps be an empirical criterion, it will be used intuitively here.

act recognition process(es) will correctly classify her utterance. Thus, a competence theory would state the conditions under which a speaker can make that determination—conditions that involve the speaker's beliefs about the hearer's beliefs, goals, and inferential processes.

Our initial competence theory has been embodied in a computer program (Cohen, 1978) that can construct most of the plans presented here. Programs often point out weaknesses, inconsistencies, and incorrect assumptions in the statement of the competence theory, and can provide an operational base from which to propose process theories. However, we make no claims that computational models of plan construction and recognition are cognitive process theories; such claims would require empirical validation. Moreover, it is unclear whether there could be just one process theory of intentional behavior since each individual might use a different method. A more reasonable goal, then, is to construct computational models of speech act use for which one could argue that a person could employ such methods and converse successfully.

1.3 Outline of the Paper

The thread of the paper is the successive refinement of speech act definitions to meet the adequacy criteria. First, we introduce in sections 2 and 3 the tools needed to construct plans: the formal language for describing beliefs and goals, the form of operator definitions, and a set of plan construction inferences.

As background material, section 4 summarizes Austin's and Searle's accounts of speech acts. Then, Searle's definitions of the speech acts of requesting and informing are reformulated as planning operators in section 5 and plans linking those speech acts to beliefs and goals are given. These initial operator definitions are shown to be compositionally inadequate and hence are recast in section 6 to allow for the planning of questions. Section 7 shows how the definitions are again inadequate for modelling plans for composed requests. After both revising the preconditions of requests and identifying their side effects, compositional adequacy for multiparty requests is achieved. The solution leads to a metatheoretical "point of view" principle for use in formulating future speech act definitions within this planning system. Finally, section 8 discusses the limitations of the formalism and ways in which the approach might be extended to handle indirect speech acts.

2. ON MODELS OF OTHERS

In this section, we present criteria that an account of one agent's (AGT1) model of another's (AGT2's) beliefs and goals ought to satisfy.⁴ A theory of speech acts need not be concerned with what is actually true in the real world; it should

⁴The representations used by Meehan (1976), and Schank and Abelson (1977) do not, in a principled way, maintain the distinctions mentioned here for belief or want.

describe language use in terms of a person's beliefs about the world. Accordingly, AGT1's model of AGT2 should be based on "believe" as described, for example, in Hintikka (1962; 1969). Various versions of the concept "know" can then be defined to be agreements between one person's beliefs and another's.

2.1 Belief

Apart from simply distinguishing AGT1's beliefs from his beliefs about AGT2's beliefs, AGT1's belief representation ought to allow him to represent the fact that AGT2 knows *whether* some proposition P is true, without AGT1's having to know which of P or $\sim P$ it is that AGT2 believes. A belief representation should also distinguish between situations like the following:

1. AGT2 believes that the train leaves from gate 8.
2. AGT2 believes that the train has a departure gate.
3. AGT2 knows what the departure gate is for the train.

Thus, case 3 allows AGT1 to believe *that* AGT2 knows what the departure gate is without AGT1's actually knowing which gate AGT2 thinks that is. This distinction will be useful for the planning of questions and will be discussed further in section 6.

Following Hintikka (1969), belief is interpreted as a model operator A BELIEVE(P), where A is the believing agent, and P the believed proposition.⁵ This allows for an elegant, albeit too strong, axiomatization and semantics for BELIEVE. We shall point out uses of various formal properties of BELIEVE as the need arises.

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, AGT1 needs two levels (AGT1 BELIEVE and

⁵The following axiom schemata will be assumed:

- B.1 aBELIEVE(all axioms of the predicate calculus)
- B.2 aBELIEVE(P) \Rightarrow aBELIEVE(aBELIEVE(P))
- B.3 aBELIEVE(P) OR aBELIEVE(Q) \Rightarrow aBELIEVE(P OR Q)
- B.4 aBELIEVE(P & Q) \Leftrightarrow aBELIEVE(P) & aBELIEVE(Q)
- B.5 aBELIEVE(P) $\Rightarrow \sim$ aBELIEVE($\sim P$)
- B.6 aBELIEVE(P \Rightarrow Q) \Rightarrow (aBELIEVE(P) \Rightarrow aBELIEVE(Q))
- B.7 $\exists x$ aBELIEVE(P(x)) \Rightarrow aBELIEVE($\exists x$ P(x))
- B.8 all agents believe that all agents believe B.1 to B.7

These axioms unfortunately characterize an idealized "believer" who can make all possible deductions from his beliefs, and doesn't maintain contradictory beliefs. Clearly, the logic should be weakened. However, we shall assume the usual possible worlds semantics of BELIEVE in which the axioms are satisfied in a model consisting of a *universe* U, a subset A of U of *agents*, a set of *possible worlds* W, and *initial world* WO in W, a *relation* R on the cross-product $A \times W \times W$, and for each world w and predicate P, a subset Pw of U called the *extension* of P in w. The truth functional connectives *and*, *or*, *not*, and \Rightarrow have their usual interpretations in all possible worlds. aBELIEVE(P) is true in world w if P is true in all worlds w1 such that R(a', w, w1), where a' is the interpretation of a in w. $\exists x$ P(x) is true in world w if there is some individual i in U such that P(x) is true in w when all free occurrences of x in P are interpreted as i. The resulting semantics for "quantifying into" BELIEVE is notoriously contentious.

AGT1 BELIEVE AGT2 BELIEVE). If AGT1 successfully lied to AGT2, he would have to be able to believe some proposition P, while believing that AGT2 believes that AGT1 believes P is false (i.e., AGT1 BELIEVE AGT2 BELIEVE AGT1 BELIEVE (\sim P)). Hence, AGT1 would need at least three levels. However, there does not seem to be any bound on the possible embeddings of BELIEVE. If AGT2 believes AGT1 has lied, he would need four levels. Furthermore, Lewis (1969) and Schiffer (1972) have shown the ubiquity of *mutual belief* in communication and face-to-face situations—a concept that requires an infinite conjunction of beliefs.⁶ Cohen (1978) shows how a computer program that plans speech acts can represent beliefs about mutual beliefs finitely.

2.2 Want

Any representation of AGT2's goals (wants) must distinguish such information from: AGT2's beliefs, AGT1's beliefs and goals, and (recursively) from AGT2'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 "AGT2 wants to take a train" as "There is a specific train that AGT2 wants to take" or as "AGT2 wants to take any train." Finally, it should allow arbitrary embeddings with BELIEVE. Wants of beliefs (as in "AGT1 WANTS AGT2 BELIEVE P") become the reasons for AGT1's telling P to AGT2, while beliefs of wants (i.e., AGT1 BELIEVES AGT1 WANTS P) will be the way to represent AGT1's goals P.⁷ In modelling planning behavior, we are not concerned with goals that the agent does not think he has, nor are we concerned with the subtleties of "wish," "hope," "desire," and "intend" as these words are used in English. The formal semantics of WANT, however, are problematic.

3. MODELS OF PLANS

In most models of planning (e.g., Fikes & Nilsson, 1971; Newell & Simon, 1963), real world actions are represented by *operators* that are organized into plans.⁸ To execute a plan, one performs the actions corresponding to the

⁶Lewis (1969) and Schiffer (1972) talk only about mutual or common knowledge, but the extension to mutual belief is obvious.

⁷This also allows a third place to vary quantifier scope, namely:

$\exists x$ aBELIEVE aWANT P(x)
aBELIEVE $\exists x$ aWANT P(x),
aBELIEVE aWANT $\exists x$ P(x)

⁸One usually generalizes operators to *operator schemata* in correspondence with *types* of actions; operator instances are then formed by giving values to the parameters of an operator schema. Since only operator instances are contained in plans we will not distinguish between the operator schema and its instances unless necessary. The same schema/instance, type/token distinction applies as well to speech acts modelled as planning operators.

operators in that plan. An operator will be regarded as transforming the planner's model of the world, the *propositions* that the planner believes, in correspondence with the changes to the real world made by the operator's associated action.⁹ An operator is *applicable* to a model of the world in which that operator's *preconditions* hold. Operators can be defined in terms of others, as stated in their *bodies* (Sacerdoti, 1975). The changes that an operator makes to the world model in which it is evaluated to produce a new world model are called that operator's *effects*.

We shall view plans for an arbitrary agent S to be constructed using (at least) the following heuristic principles of purposeful behavior:

At the time of S's planning:

1. S should not introduce in the plan actions whose effects S believes are (or will be) true at the time the action is initiated.
2. If E is a goal, an operator A that achieves E can be inserted into the plan.
3. If an operator is not applicable in the planner's belief model, all the preconditions of that operator that are not already true can be added to the plan.

The previous two inferences reflect an agent's reasoning "in order to do this I must achieve that."

4. If the planner needs to know the truth-value of some proposition, and does not, the planner can create a goal that it know whether that proposition is true or false.
5. If the planner needs to know the value of some description before planning can continue, the planner can create a goal that it find out what the value is.

The previous two inferences imply that the planner does not have to create an entire plan before executing part of it.

6. Everyone expects everyone else to act this way.

Since agents can sometimes recognize the plans and goals of others, and can adopt others' goals (or their negations) as their own, those agents can plan to facilitate or block someone else's plans. Bruce and Newman (1978) and Carbonell (1978) discuss these issues at length.

The process of planning to achieve a goal is essentially a search through this space of inferences to find a temporal sequence of operators such that the first operator in the sequence is applicable in the planner's current world model and the last produces a world model in which the goal is true. A new world model is obtained by the execution of each operator.

3.1 The Form of Operators

Early approaches to problem-solving based on first order logic (Green, 1969; McCarthy & Hayes, 1969) have emphasized the construction of provably correct

⁹We are bypassing the fact that people need to observe the success or failure of their actions before being able to accurately update their beliefs. The formalism thus only deals with operators and models of the world rather than actions and the real world. Operators names will be capitalized while their corresponding actions will be referred to in lower case.

plans. Such approaches formalize the changes an action makes to the state of the world model by treating an operator as a predicate of one whose arguments is a *state variable*, which ranges over states of the world model. Unfortunately, to be able to reason about what is true in the world after an action is executed, one must give axiom schemata that describe which aspects of the state of the world are *not* changed by each operator. For instance, calling someone on the telephone does not change the height of the Eiffel Tower. This thorny "frame problem" (McCarthy & Hayes, 1969) occurs because individual states of the world are not related to one another *a priori*.

To overcome this problem, Fikes and Nilsson (1971) in their STRIPS planning system assume that all aspects of the world stay constant except as described by the operator's effects and logical entailments of those effects. Such an assumption is not formalized in the reasoning system, making it difficult to prove the correctness of the resulting plans. Nevertheless, it has become the standard assumption upon which to build problem-solvers. We too will make it and thus shall describe an operator's effects by the propositions that are to be added to the model of the world.¹⁰

All operator schemata will have two kinds of preconditions—"cando" and "want" preconditions. The former, referred to as CANDO.PRs, indicate proposition schemata that, when instantiated with the parameter values of an operator instance, yield propositions that must be true in the world model for that operator instance to be applicable. We do not discuss how they can be proven true. The "want" precondition, henceforth WANT.PR, formalizes a principle of intentional behavior—the agent of an action has to want to do that action.

The following example serves to illustrate the form of such definitions.

MOVE(AGT, SOURCE, DESTINATION)	
CANDO.PR:	LOC(AGT, SOURCE)
WANT.PR:	AGT BELIEVE AGT WANT move-instance
EFFECT:	LOC(AGT, DESTINATION)

The parameters of an operator scheme are stated in the first line of the definitions and it is assumed that values of these parameters satisfy the appropriate selectional restrictions, (here, a person, and two locations, respectively). The WANT.PR uses a parameter "move-instance" that will be filled by any instance of the MOVE operator schema that is currently being planned, executed, or recognized. The CANDO.PR states that before an agent can move from the SOURCE location, he must be located there. The EFFECT of the MOVE indicates that the agent's new location is the DESTINATION.

S's plan to achieve goal G is pictured schematically in Figure 1 (P and Q are arbitrary agents, A1 and A2 are arbitrary actions). Instead of indicating the entire state of the planner's beliefs after each operator, those propositions that are effects of an operator and are preconditions of some other operator in the plan are presented.

¹⁰Those propositions that need to be deleted (or somehow made "invisible" in the *current* worldmodel) will not be discussed here.

S BELIEVE S WANT:




Figure 1. A schematic of S's plan to achieve G.

This diagram illustrates the building block of plans—given goal G, S applies an inference of type 2 and selects operator A1, whose agent is Q as a producer of that effect. That operator is applicable when preconditions Ci and Cj hold and when agent Q wants to perform A1. Type 3 inferences allow each of the preconditions to be achieved by other actions (e.g., A2), which may be performed by another agent (e.g., P). This chaining of operators continues until all preconditions are satisfied. Plan diagrams are thus read from "top" to "bottom".

To indicate that this schematic is part of agent S's plan, the plan components are "embedded" in what S BELIEVE S WANTS. The truth or falsity of preconditions is evaluated with respect to S's beliefs. For example, verifying the WANT.PR of operator A1 (i.e., Q BELIEVE Q WANT Q do A1) would involve establishing that S BELIEVE Q BELIEVE Q WANT Q do A1. If Q is the same person as S (i.e., S is planning her own action A1) then this condition is trivially true since A1 is already part of S's plan, and since for all agents R, we assume that if R BELIEVE (P) then R BELIEVE R BELIEVE (P). However, if Q is not the same as S, the WANT.PR also needs to be achieved, leading, as we shall see, to S's planning a speech act.

4. SPEECH ACTS

4.1 Austin's Performatives

Austin (1962) notes a peculiar class of declarative utterances, which he termed *performatives*, that do not state facts but rather constitute the performance of an action. For instance saying, "I hereby suggest you leave" is an act of suggesting. Unlike the usual declaratives, such sentences are not true or false, but rather are subject to the same kinds of failures ("infelicities") as nonlinguistic actions—such as being applied in the wrong circumstances or being performed insincerely.

Generalizing further, Austin claims that in uttering any sentence, one performs three types of speech acts: the *locutionary*, *illocutionary*, and *perlocutionary*.

ary acts. A speaker performs a *locutionary* act by making noises that are the uttering of words in language satisfying its vocabulary and grammar, and by the uttering of sentences with definite meaning (though perhaps having more than one). Such acts are used in the performance of *illocutionary acts* which are those acts performed in making utterances. For instance, stating, requesting, warning, ordering, apologizing, are claimed to be different types of illocutionary acts, each of which is said to have a unique *illocutionary force* that somehow characterizes the nature of the act. Each illocutionary act contains *propositional content* that specifies what is being requested, warned about, ordered, etc.

New distinctions, however, bring new problems. Frequently, when performative verbs are not used, the utterance's illocutionary force is not directly interpretable from its content. For example, to understand the force of the utterance "The door," the hearer may need to use his beliefs that the door is currently closed, that the speaker has two arm-loads of groceries, and that he wants to be on the other side of the door in determining that the speaker has requested that the door be opened. Furthermore, a speaker may appear to be performing one illocutionary act, and actually may be trying to use it to do something else. Thus, "We have to get up early tomorrow" may simply be an assertion but when said at a party, may be intended as an excuse to the host for leaving, and may be intended as a request that the hearer leave. Such *indirect speech acts* (Gordon & Lakoff, 1971; Searle, 1975) are the touchstone of any theory of speech acts.

The last major kind of act identified by Austin is the *perlocutionary act*—the act performed by making an utterance. For instance, with the illocutionary act of asserting something, I may *convince* my audience of the truth of the corresponding proposition (or *insult* or *frighten* them). Perlocutionary acts produce *perlocutionary effects*: convincing produces belief and frightening produces fear. While a speaker often has performed illocutionary acts with the goal of achieving certain perlocutionary effects, the actual securing of those effects is beyond his control. Thus, it is entirely possible for a speaker to make an assertion, and for the audience to recognize the force of the utterance as an assertion and yet not be convinced.

4.2 Speech Acts à la Searle

Searle (1969) presents a formulation of the structure of illocutionary acts (henceforth referred to simply as speech acts) by suggesting a number of necessary and sufficient conditions on their successful performance. He goes on to state rules corresponding to these conditions, for a speaker's using any "indicator of illocutionary force" to perform a particular speech act.

As an example, let us consider Searle's conditions for a speaker S, in uttering T, to request that some hearer H do action A. The conditions are grouped as follows:

Normal Input/Output Conditions. These include such conditions as: H is not deaf and S is not mute, joking, or acting.

Propositional Content Conditions. Literal speech acts only use propositions of certain forms. The restrictions on these forms are stated in the *propositional content conditions*. For a request, the proposition must predicate a future act of H.

Preparatory Condition. A preparatory condition states what must be true in the world for a speaker to felicitously issue the speech act. For a request, the preparatory conditions include:

- H is able to do A.
- S believes H is able to do A.
- It is not obvious to S and H that H will do A in the normal course of events (the "non-obviousness" condition).

Searle claims the non-obviousness condition is not peculiar to illocutionary acts. This paper will support his claim by showing how the condition can be applied more generally to rational, intentional behavior.

Sincerity Condition. A *sincerity condition* distinguishes a sincere performance of the speech act from an insincere one. In the case of a request, S must want H to do A; for a promise, S must intend to do the promised action; for an assertion, S must believe what he is asserting.

Essential Condition. An *essential condition* specifies what S was trying to do. For a request, the act is an attempt to get H to do A.

Force Condition (our terminology). The purpose of the *force condition* is to require that the speaker utter a speech act only if he intends to communicate that he is performing that act. "Intending to communicate" involves having certain intentions regarding how the hearer will recognize the force of the utterance. The basic idea is that it is intended that the hearer recognize that the speaker is trying to bring about the satisfaction of the essential condition. For a request this amounts to the speaker's wanting the hearer to realize the speaker intends for him to do A.

5. A FIRST REFORMULATION OF SEARLE'S CONDITIONS

Searle (1969) unfortunately does not supply justifications for the adequacy of his definitions for various kinds of speech acts. A primary goal of this paper is to show how a plan-based theory provides the basis for such adequacy criteria by allowing one to see clearly how changes in speech act definitions affect the plans that can be generated.

A second, more specific point of this formulation exercise is to show which of Searle's conditions are better regarded as pertaining to more general aspects of intentional behavior than to particular speech acts. In this spirit, we show how the sincerity condition, which we shall argue is a misnomer, and the propositional content and "non-obviousness" conditions arise during the course of planning. Concerning the remaining conditions, we assume the "normal input/output conditions," but have chosen not to deal with the force condition until we have a better understanding of the plans for speech acts and how they can be recognized. The remaining conditions, the preparatory and essential conditions, will be mapped into the formalism as the preconditions and effects of speech act operators.

5.1 First Definition of REQUEST

Searle claims the preparatory conditions are required for the "happy" performance of the speech act—where "happy" is taken to be synonymous with Austin's use of "felicitous." Austin was careful to distinguish among infelicities, in particular, misapplications (performing the act in the wrong circumstances), and flaws (incorrectly performing the act). We take Searle's preparatory conditions as conditions guaranteeing applicability rather than successful performance, allowing them to be formalized as preconditions. Thus if an operator's preconditions are not satisfied when it is performed, then the operator was "misapplied." Before expressing preconditions in a formalism, a systematic "point of view" must be adopted. Since the applicability conditions affect the planning of that speech act, the preconditions are stated as conditions on the speaker's beliefs and goals. Correspondingly, the effects describe changes to the hearer's mental state.¹¹ We establish a *point-of-view principle*, that is intended to be a guideline for constructing speech act definitions in *this* planning system—namely: preconditions begin with "speaker believe" and effects with "hearer believe."

Let us consider Searle's preparatory conditions for a request: H is able to do ACT, and S believes H is able to do ACT. From our discussion of "belief," it should be clear what H can *in fact* do, i.e., what the real world is like is not essential to the success of a request. What may be relevant is that S and/or H thinks H can do ACT. To formalize "is able to do A," we propose a predicate CANDO(Q,ACT) that is true if the CANDO.PR's of ACT are true (with person Q bound to the agent role of ACT).¹²

The essential condition, which is modeled as the EFFECT of a REQUEST,

¹¹This does not violate our modelling just one person's view since a speaker, after having issued a speech act, will update his beliefs to include the effects of that speech act, which are defined in terms of the hearer's beliefs.

¹²This should be weakened to "... are true or are easily achievable"—i.e. if Q can plan to make them true.

is based on a separation of the illocutionary act from its perlocutionary effect. Speakers, we claim, cannot influence their hearers' beliefs and goals directly. The EFFECTs of REQUEST are modeled so that the hearer's actually wanting to do ACT is not essential to the successful completion of the speech act. Thus, the EFFECT is stated as the hearer's believing the speaker wants him to do the act. For important reasons, to be discussed in section 5.7, this formulation of the essential condition will prove to be a major stumbling block.

The operator REQUEST from SPEAKER to HEARER to do action ACT, which represents a literal request, can now be defined as:

REQUEST(SPEAKER,HEARER,ACT)

CANDO.PR:	SPEAKER BELIEVE HEARER CANDO ACT
	AND
	SPEAKER BELIEVE
	HEARER BELIEVE HEARER CANDO ACT
WANT.PR:	SPEAKER BELIEVE SPEAKER WANT request-instance
EFFECT:	HEARER BELIEVE
	SPEAKER BELIEVE SPEAKER WANT ACT

5.2 Mediating Acts and Perlocutionary Effects

To bridge the gap between REQUESTs and the perlocutionary effect for which they are planned, a mediating step named CAUSE-TO-WANT is posited, that models what it takes to get someone to want to do something. Our current analysis of this "act" trivializes the process it is intended to model by proposing that to get someone to want to do something, one need only get that person to know that you want them to do it.

The definition of an agent's (AGT1) causing another agent (AGT) to want to do ACT is:

CAUSE-TO-WANT (AGT1,AGT,ACT)

CANDO.PR:	AGT BELIEVE
	AGT1 BELIEVE AGT 1 WANT ACT
EFFECT:	AGT BELIEVE AGT WANT ACT

The plan for a REQUEST is now straightforward. REQUEST supplies the necessary precondition for CAUSE-TO-WANT (as will other act combinations). When the WANT.PR of some action that the speaker is planning for someone else to perform, is not believed to be true, the speaker plans a REQUEST. For example, assume a situation in which there are two agents, SYSTEM¹³(S) and JOHN, who are located inside a room (i.e., they are at location INROOM). Schematically, to get JOHN to leave the room by moving himself to location

¹³The agent who creates plans will often be referred to as "SYSTEM," which should be read as "planning system."

OUTROOM, the plan would be as in Figure 2. Notice that the WANT.PR of the REQUEST itself, namely

S BELIEVE
S WANT
REQUEST(S,JOHN,MOVE(JOHN,INROOM,OUTROOM))

is trivially true since that particular REQUEST is already part of S's plan. The CANDO.PR's of the REQUEST are true if S believes JOHN is located INROOM and if it believes JOHN thinks so too. Thus, once the planner chooses someone else, say H, to do some action that it believes H does not yet want to do, a directive act (REQUEST) may be planned.

5.3 Comparison with Searle's Conditions for a REQUEST

Searle's "non-obviousness" condition for the successful performance of a request stated that it should not be obvious to the speaker that the hearer is about to

do the action being requested, independently of the request. If that were obvious to the speaker, the request would be pointless. However, as Searle noted, the non-obviousness condition applies more generally to rational, intentional behavior than to speech acts alone. In our formalism, it is the WANT.PR of the act being requested (goal "++" in Figure 2). If the planning system believed the WANT.PR were already true, i.e., if it believed that John already wanted to leave the room, then the plan would proceed no further; no REQUEST would take place.

Searle's "sincerity" condition, stated that the speaker had to want the requested act to be performed. The sincerity condition in the plan of Figure 2 is the goal labeled "+." The speaker's wanting the hearer to move is the reason for planning a REQUEST.

Notice also that the propositional content of the REQUEST, a future act to be performed by the hearer, is determined by prior planning—i.e., by a combination of that act's WANT.PR, the mediating act CAUSE-TO-WANT, and by the EFFECT of a REQUEST. Searle's propositional content condition thus seems to be a function of the essential condition (which is approximated by the EFFECTs of the speech act operator), as Searle claimed. So far, we have factored out those aspects of a request that Searle suggested were eliminable. Future revisions will depart more significantly.

5.4 Definition of INFORM

The speech act of informing is represented by the operator INFORM, which is defined as a speaker's stating a proposition to a hearer for the purpose of getting the hearer to believe that the speaker believes that proposition to be true. Such acts will usually be planned on the basis of wanting the hearer to believe that proposition. For a SPEAKER to INFORM a HEARER that proposition PROP is true, we have:

<u>INFORM(SPEAKER, HEARER, PROP)</u>	
CANDO.PR:	SPEAKER BELIEVE PROP
WANT.PR:	SPEAKER BELIEVE SPEAKER WANT inform-instance
EFFECT:	HEARER BELIEVE SPEAKER BELIEVE PROP

The CANDO.PR simply states that the only applicability condition to INFORMing someone that proposition PROP is true is that the speaker believes PROP.¹⁴ The EFFECT of an INFORM is to communicate what the speaker believes. This allows for the hearer to refuse to believe the proposition without

¹⁴Other preconditions to the INFORM act could be added—for instance, to talk to someone one must have a communication link (Schank & Abelson, 1977); which may require telephoning or going to that person's location, etc. However, such preconditions would apply to any speech act, and hence probably belong on the locutionary act of making noises to someone.

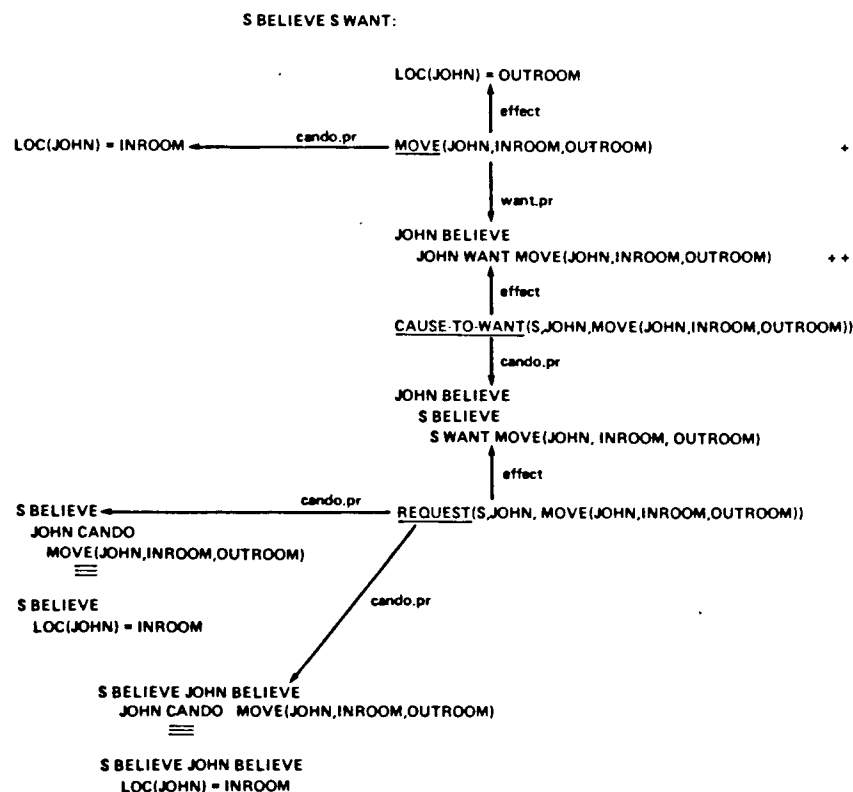


Figure 2. A plan for a REQUEST.

invalidating the speaker's action as an INFORM. Therefore, an intermediate "act," termed CONVINCE, is necessary to get the hearer to believe the proposition.

For a person AGT 1 to CONVINCE another person AGT that proposition PROP is true, we define:

<u>CONVINCE(AGT1, AGT, PROP)</u>	
CANDO.PR:	AGT BELIEVE
	AGT1 BELIEVE PROP
EFFECT:	AGT BELIEVE PROP

This operator says that for AGT 1 to convince AGT of the truth of PROP AGT need only believe that AGT1 thinks PROP is true. Though this may be a necessary prerequisite to getting someone to believe something, it is clearly not sufficient. For a more sophisticated precondition of CONVINCE, one might state that before AGT will be convinced, she needs to know the justifications for AGT1's belief, which may require that AGT believe (or be CONVINCE of) the justifications for believing those justifications, etc. Such a chain of reasons for believing might be terminated by mutual beliefs that people are expected to have or by a belief AGT believes AGT1 already has. Ideally, a good model of CONVINCE would allow one to plan persuasive arguments.¹⁵

5.5 Planning INFORM Speech Acts

The planning of INFORM speech acts now becomes a simple matter. For any proposition PROP, S's plan to achieve the goal H BELIEVE PROP would be that of Figure 3. Notice that it is unnecessary to state as a precondition to inform, that the hearer H does not already believe PROP. Again, this non-obviousness condition that can be eliminated by viewing speech acts in a planning context.

What would be Searle's sincerity condition for the INFORM above (S BELIEVE PROP) turns out to be a precondition for the speech act rather than a reason for planning the act as we had for REQUEST's sincerity condition, (i.e., SPEAKER BELIEVE SPEAKER WANT HEARER do ACT). If we were to use REQUEST as a model, the sincerity condition for an INFORM would be SPEAKER BELIEVE SPEAKER WANT HEARER BELIEVE PROP. One may then question whether Searle's sincerity condition is a consistent naming of distinctive features of various kinds of speech acts. Insincerity is a matter of falsely claiming to be in a psychological state, which for this model is either belief or want. By this definition, both conditions, SPEAKER BELIEVE PROP

¹⁵Without a specification of the justifications for a belief, this operator allows one to become convinced of the truth of one's own lie. That is, after speaker S lies to hearer H that P is true, and receives H's acknowledgment indicating H has been convinced, S can decide to believe P because he thinks H thinks so. Further research needs to be done on CONVINCE and BELIEVE to eliminate such bizarre behavior.

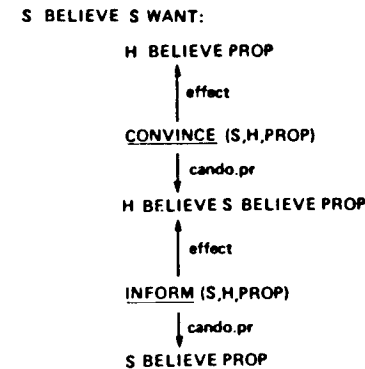


Figure 3. A plan for an INFORM.

and SPEAKER BELIEVE SPEAKER WANT HEARER BELIEVE PROP, are subject to insincerity.

5.6 Planning an INFORM of a WANT

As stated earlier, there are other ways to satisfy the precondition to CAUSE-TO-WANT. Since REQUEST was taken as a prototypical directive act, all members of that class share the same EFFECT (Searle's (1976) "illocutionary point"). However, issuing an INFORM of a WANT, as in "I want you to do X," also achieves it. Another plan to get John to move appears in Figure 4.

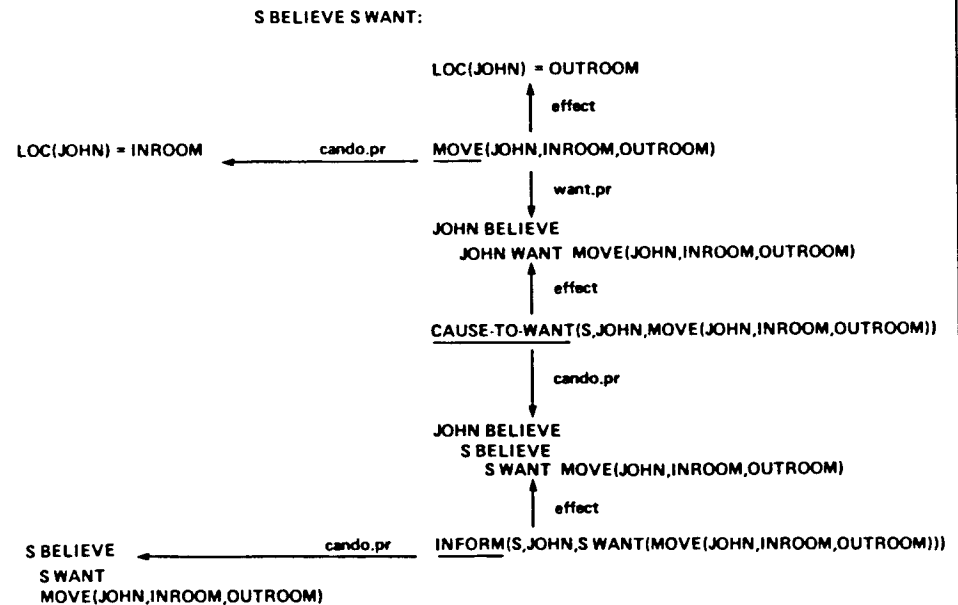


Figure 4. A plan for an INFORM of a WANT.

The initial stages of this plan are identical to that of Figure 2 through the CANDO.PR of CAUSE-TO-WANT. This precondition is achieved by an INFORM whose propositional content is S WANT MOVE (JOHN, INROOM, OUTROOM). In this instance, the planning system does not need to proceed through CONVINCE since an INFORM of a WANT produces the necessary effects. Testing the CANDO.PR of INFORM determines if the system believes this proposition, which it does since the MOVE by John is already one of its goals. The WANT.PR of INFORM is trivially true, as before, and thus the plan is complete.

5.7 REQUEST vs. INFORM of WANT

Searle claimed that the conditions he provided were necessary and jointly sufficient for the successful and nondefective performance of various illocutionary acts. Any behavior satisfying such a set of conditions was then said to be a particular illocutionary act. Thus, if two utterances have the same illocutionary force, they should be equivalent in terms of the conditions on their use. We believe that the two utterances "please open the door" and "I want you to open the door (please)" *can* have the same force as directives, differing only in their politeness. That is, they both *can be* planned for the same reasons. However, our treatment does not equate the literal speech acts that could realize them when they should be equated. The condition on REQUEST that distinguishes the two cases is the precondition SPEAKER BELIEVE HEARER BELIEVE HEARER CANDO ACT. Since there is no corresponding precondition in the plan for the INFORM of a WANT, there is no reason to check the hearer's beliefs.

In order to force an equivalence between a REQUEST and an INFORM of a WANT, various actions need to be redefined. We shall remove the above condition as a CANDO.PR from REQUEST and add it as a new CANDO.PR to CAUSE-TO-WANT. In other words, the new definition of CAUSE-TO-WANT would say that you can get a person to decide to want to do some action if she believes you want her to do it and if she believes she can do it. With these changes, both ways of getting someone to want to do some action would involve her believing she is able to do it. More formally, we now define:

REQUEST (SPEAKER, HEARER, ACT)

CANDO.PR: SPEAKER BELIEVE HEARER CANDO ACT
 WANT.PR: SPEAKER BELIEVE SPEAKER WANT request-instance
 EFFECT: HEARER BELIEVE
 SPEAKER BELIEVE SPEAKER WANT ACT

and

CAUSE-TO-WANT (AGT1, AGT, ACT)

CANDO.PR: AGT BELIEVE
 AGT1 BELIEVE AGT1 WANT ACT
 AND
 AGT BELIEVE AGT CANDO ACT
 EFFECT: AGT BELIEVE AGT WANT ACT

Though REQUEST and INFORM of a WANT can achieve the same effect, they are not interchangeable. A speaker (S), having previously said to a hearer (H) "I want you to do X," can deny having the intention to get H to want to do X by saying "I simply told you what I wanted, that's all." It appears to be much more difficult, however, after having requested H to do X, to deny the intention of H's wanting to do X by saying "I simply requested you to do X, that's all." S usually plans a request for the purpose of getting H to want to do some act X by means of getting H to believe that S wants H to do it. While maintaining the distinction between illocutionary acts and perlocutionary effects, thus allowing for the possibility that H could refuse to do X, we need to capture this distinction between REQUEST and INFORM of WANT. The solution (Allen, 1979; Per-rault & Allen, forthcoming) lies in formulating speech act bodies as plans achieving the perlocutionary effect—plans that a hearer is intended to recognize.

In the next two sections, we investigate the compositional adequacy of these operator definitions via the planning of REQUESTs that a hearer perform REQUEST or INFORM speech acts.

6. COMPOSITIONAL ADEQUACY: QUESTIONS

We are in agreement with many others, in proposing that questions be treated as requests for information. In terms of speech act operators, the questioner is performing a REQUEST that the hearer perform an INFORM. That is, the REQUEST leads to the satisfaction of INFORM's "want precondition." However, for a wh-question, the INFORM operator as defined earlier cannot be used since the questioner does not know the full proposition of which he is to be informed. If he did know what the proposition was there would be no need to ask; he need only decide to believe it.

Intuitively, one plans a wh-question to find out the value of some expression and a yes/no question to find out whether some proposition is true. Such questions are planned, respectively, on the basis of believing *that* the hearer knows what the value of that expression is or *that* the hearer knows whether the proposition is true, without the speaker's having to know what the hearer believes.

Earlier we stated that a person's (AGT1) belief representation should represent cases like the following distinctly:

1. AGT2 believes the Cannonball Express departs at 8 p.m.
2. AGT2 believes the Cannonball Express has a departure time.
3. AGT2 knows what the departure time for the Cannonball Express 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.,

AGT2 BELIEVE

$\exists x (\text{the } y : \text{DEPARTURE-TIME}(\text{CANNONBALL-EXPRESS}, y)) = x$

However, Case 3 can be approximated by a *quantified belief*, namely,

$\exists x$ AGT2 BELIEVE

$(\text{the } y : \text{DEPARTURE-TIME}(\text{CANNONBALL-EXPRESS}, y)) = x$,

where “the $y : P(y)$,” often written “ $iy P(y)$,” is the logical description operator read “the y which is P .” This formula is best paraphrased as “there is something which AGT2 believes to be the departure time for the Cannonball Express.”¹⁶ Typical circumstances in which AGT1 might acquire such quantified beliefs are by understanding a definite description uttered by AGT2 referentially (Donnellan, 1966). Thus, if AGT2 says “the pilot of TWA 461 on July 4,” AGT1 might infer that AGT2 knows who that pilot is.

Quantified beliefs often become goals when a planner needs to know the values of the parameters of an operator and when these parameters occur in that operator’s preconditions.¹⁷ We show how, when a quantified belief is a goal for AGT, AGT can plan a wh-question.

6.1 Planning Wh-Questions

First, a new operator, INFORMREF, and its associated mediating act CONVINCEREF, are needed.¹⁸

INFORMREF(SPEAKER, HEARER, $\lambda x D x$) (i.e., D is a predicate of one argument)

CANDO.PR: $\exists y$ SPEAKER BELIEVE $(ix D x) = y$

WANT.PR: SPEAKER BELIEVE SPEAKER WANT informref-instance

EFFECT: $\exists y$ HEARER BELIEVE SPEAKER BELIEVE $(ix D x) = y$

¹⁶Another conjunction can be added to the representation of (3) as suggested by Allen (1979) to refine our representations of “AGT2’s knowing what the value of the description is,” namely:

$$\exists x [(\text{the } y : D(y)) = x \ \& \ \text{AGT2 BELIEVE } ((\text{the } y : D(y)) = x)]$$

We shall, however, use the simpler quantified belief formulation.

¹⁷We would prefer to formalize declaratively that “the agent of an action must know the values of the parameters of the action.” One way of doing this is suggested by Moore (1979).

¹⁸In Cohen (1978) we achieved the same effect by parameterizing INFORM and CONVINCEREF so that different sets of preconditions and effects were used if the original goal was a quantified belief. In addition, Cohen (1978) did not use descriptions. We believe the formulation that follows, due to J. Allen, is clearer. The actual names for these acts were suggested by W. Woods.

Thus, before a speaker will inform a hearer of the value of some description, there must be some individual that the speaker believes is the value of the description, and the speaker must want to say what it is. The effect of performing this act is that there is then some individual that the hearer thinks the speaker believes to be the value of the description. As usual, we need a mediating act to model the hearer’s then believing that individual to be the value of the description. To this end, we define AGT1’s convincing AGT of the referent of the description as:

CONVINCEREF(AGT1, AGT, $\lambda x D x$)

CANDO.PR: $\exists y$ AGT BELIEVE AGT1 BELIEVE $(ix D x) = y$

EFFECT: $\exists y$ AGT BELIEVE $(ix D x) = y$

Using these operators, if the planning system wants to know where Mary is and believes that Joe knows where she is, it can create the plan underlying the question “Where is Mary?” as is shown in Figure 5. After the system plans for Joe to tell it Mary’s location, on the basis of believing that he knows where she is, it must get Joe to want to perform this act. In the usual fashion, this leads to a REQUEST and hence the construction of a question. The precondition to

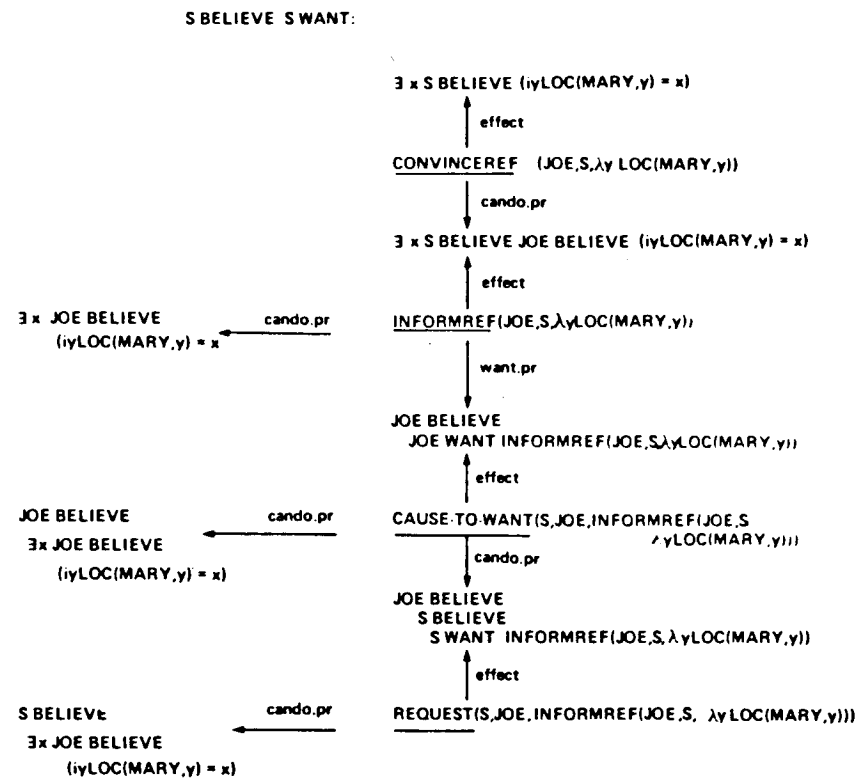


Figure 5. A plan for a wh-question.

CAUSE-TO-WANT, namely, JOE BELIEVE JOE CAN DO the INFORMREF is actually:

JOE BELIEVE
 $\exists y$ JOE BELIEVE
 $ixLOC(MARY, x) = y$

which is implied by

$\exists y$ JOE BELIEVE $ixLOC(MARY, x) = y$

that was asserted, for this example, to be one of the planning system's beliefs. Notice, that the planning of this question depends upon the system's having chosen Joe to tell it the answer, and upon its having chosen itself to get Joe to want to perform the INFORM. Section 7 discusses what happens when different decisions are made.

6.2 Plans for Yes/No Questions

To plan a yes/no question about some proposition P, one should think that the hearer knows whether P is true or false (or, at least "might know"). An approximate representation of AGT2's knowing whether P is true or false is OR (AGT2 BELIEVE P, AGT2 BELIEVE \sim P)).¹⁹ Such goals are often created, as modeled by our type 4 inference, when a planner does not know the truth-value of P. Typical circumstances in which an agent may acquire such disjunctive beliefs about another are telephone conversations, in which AGT1 believes that there are certain objects in AGT2's view. AGT1 then probably believes that AGT2 knows whether certain visually derivable (or easily computable) properties of those objects are true, such as whether object A is on top of object B.

To accommodate yes/no questions into the planning system, a third INFORM, called INFORMIF, and its associated mediating act CONVINCIF are defined as follows:

INFORMIF(SPEAKER, HEARER, P)

CANDO.PR: OR(SPEAKER BELIEVE P, SPEAKER BELIEVE \sim P)
 EFFECT: OR(HEARER BELIEVE SPEAKER BELIEVE P,
 HEARER BELIEVE SPEAKER BELIEVE \sim P)

WANT.PR: SPEAKER BELIEVE SPEAKER WANT informif-instance

CONVINCIF(AGT, AGT1, P)

CANDO.PR: OR(AGT BELIEVE AGT1 BELIEVE P,
 AGT BELIEVE AGT1 BELIEVE \sim P)
 EFFECT: OR(AGT BELIEVE P, AGT BELIEVE \sim P)

¹⁹Allen (1979) also points out that another conjunct can be added to the representation of "knowing whether" as a disjunctive belief, to obtain (P & AGT2 BELIEVE (P)) OR (\sim P & AGT2 BELIEVE (\sim P)).

The plan for a yes/no question to Joe is now parallel to that of a wh-question.²⁰ That is, in the course of planning some other act, if the system wants proposition P to be true or to be false, and if the truth-value of proposition P is unknown to it, it can create the goal OR(SYSTEM BELIEVE P, SYSTEM BELIEVE \sim P). For instance if P were LOC(MARY, INROOM), the illocutionary acts underlying the question to Joe "Is Mary in the room?" can be planned provided the planning system believes that Joe either believes P is true or he believes P is false. That disjunctive belief could be stated directly or could be inferred from a belief like $\exists y$ JOE BELIEVE($ixLOC(MARY, x) = y$)—i.e., there is something Joe believes is Mary's location. But if it had some idea where Joe thought Mary was, say OUTROOM, then it would not need to ask.

6.3 Summary

A plan for a question required the composition of REQUEST and INFORM and led to the development of two new kinds of informing speech acts, INFORMREF and INFORMIF, and their mediating acts. The INFORMREF acts lead to "what," "when," and "where" questions while INFORMIF results in a yes/no question.²¹ The reason for these new acts is that, in planning a REQUEST that someone else perform an INFORM act, one only has incomplete knowledge of their beliefs and goals; but an INFORM, as originally defined can only be planned when one knows what is to be said.

7. COMPOSITIONAL ADEQUACY AND THE POINT OF VIEW PRINCIPLE

Earlier, a guiding "Point of View Principle" (POVP) for defining speech acts as planning operators was proposed: the preconditions of the operator should be stated from the speaker's point of view, i.e., in terms of the speaker beliefs; the effects should be stated from the hearer's point of view. We now wish to judge the adequacy of speech act definitions formulated along these lines. The test case

²⁰Searle (1969) suggested there were different speech acts for real and teacher-student (or exam) questions, where in the latter case, the questioner just wants to know what the student thinks is the answer. Since teacher-student questions seem to have similar conditions on their appropriateness as real questions, save the questioner's intention to be convinced, we have good reason for factoring the mediating acts out of each of the three INFORM act types. This leaves the INFORM acts neutral with respect to what kind of question they are contained in. In general, if the perlocutionary effects of an INFORM were incorporated into the act's definition, then we would need two new primitive teacher-student question speech acts. For now, we opt for the former.

²¹The language for stating operators needs to be extended to account for "which," "how," and "why" questions. For instance, "why" and "how" questions involve quantifying over actions and/or plans.

will be the composing of REQUESTs, i.e., the planning of a REQUEST that some third party himself perform a REQUEST. For instance, the utterance "Ask Tom to tell you where the key is" is an example of such a third party request.

The current definitions of speech acts will be shown to be compositionally inadequate since they force speakers to have unnecessary knowledge about intermediaries' beliefs. Achieving compositional adequacy, however, requires more than a simple restatement of the point of view principle; the side effects of speech act operators also must be considered.

Our scrutiny will be focused upon the seemingly innocent precondition to REQUEST, SPEAKER BELIEVER HEARER CANDO ACT whose form depended on the POVP. The goal is to show how the POVP leads us astray and how a formulation of that precondition according to a new POVP that suggests a more neutral point of view for speech act definitions sets us back on course. From here on, the two versions of the precondition will be referred to as the "speaker-based" and "neutral" versions.

7.1 Plans for Multiparty Speech Acts

Multiparty speech acts can arise in conversations where communication is somehow restricted so as to pass through intermediaries.²² The planning system, since it is recursive, can generate plans for such speech acts using any number of intermediaries provided that appropriate decisions are made as to who will perform what action.

Let us suppose that the planning system wants to know where a particular key is and that it must communicate through John. We shall use the speaker-based precondition on REQUEST for this example, and for readability, the following abbreviations:

SYSTEM—S	TOM—T	JOHN—J
BELIEVE—B	WANT—W	LOC(KEY23,y)—D(y)

Figure 6 shows the plan for the specific three-party speech act underlying "Ask Tom to tell me where the key is."

S develops the plan in the following fashion: T is chosen to tell S the key's location since, we shall assume, he is believed to know where it is. Since T is not believed to already want to tell, and since S cannot communicate directly with T (but T can communicate with S), J is chosen to be the one to talk T into telling. Since J is not believed to already want to do that, S plans a REQUEST that J perform a REQUEST, namely REQUEST(S,J,REQUEST(J,T,INFORMREF(T,S,λyLOC(KEY23,y)))). J, then, is an intermediary who is just expected to do what he is asked; his status will be discussed soon.

²²For instance, in the Stanford Research Institute Computer-based Consultant research (Deutsch, 1974) communication between an expert and an apprentice was constrained in this way. The apprentice typically issued such speech acts, while the expert did not.

The preconditions that need to be satisfied in this plan are:

S BELIEVE:

- (P1) $\exists y T \text{ BELIEVE } [ixLOC(KEY23,x)=y]$
- (P2) T BELIEVE (P1) (implied by P1)
- (P3) J BELIEVE (P1)
- (P4) J BELIEVE J BELIEVE (P1) (implied by P3)
- (P5) S BELIEVE J BELIEVE (P1) (implied by P3)

S BELIEVE SWANT:

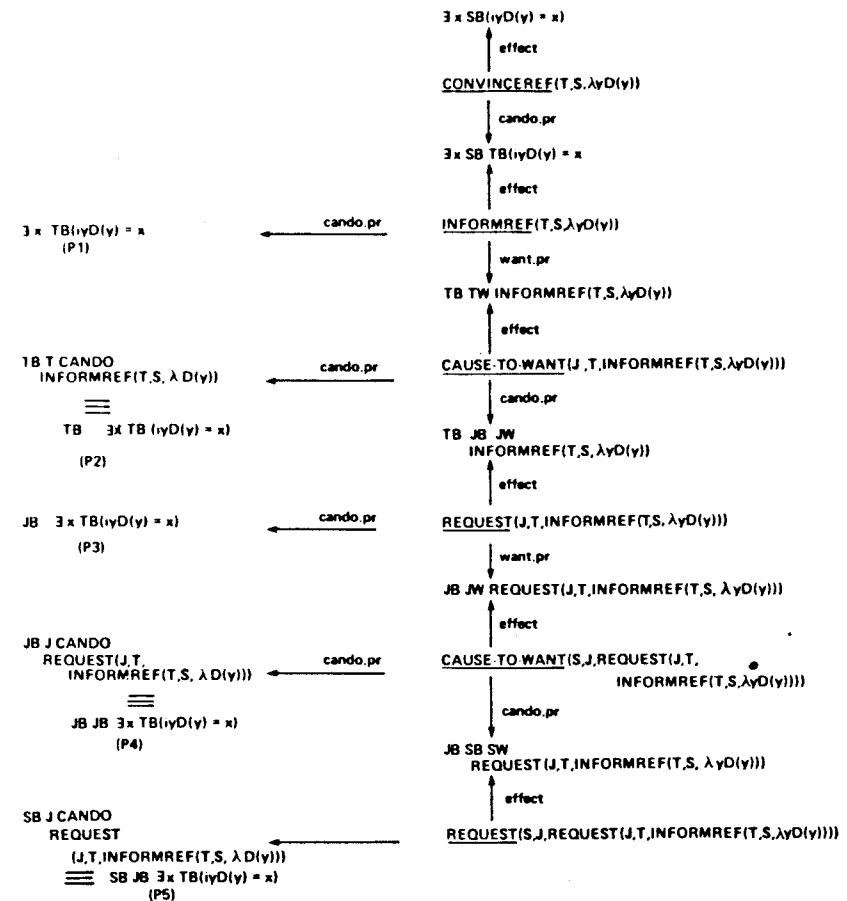


Figure 6. A plan for a third party REQUEST.

While the plan appears to be straightforward, precondition P3 is clearly unnecessary—S ought to be able to plan this particular speech act without having any *prior* knowledge of the intermediary's beliefs. This prior knowledge requirement comes about because precondition P5 is constructed by composing

REQUEST's precondition schema with precondition P3, and P3 is similarly constructed from P1.

The problem can be eliminated by reformulating REQUEST's precondition as HEARER CAN DO ACT. Consider a general plan for three-party REQUESTs, as in Figure 7. T's INFORMREF has been generalized to "ACT(T)" whose precondition is "P."

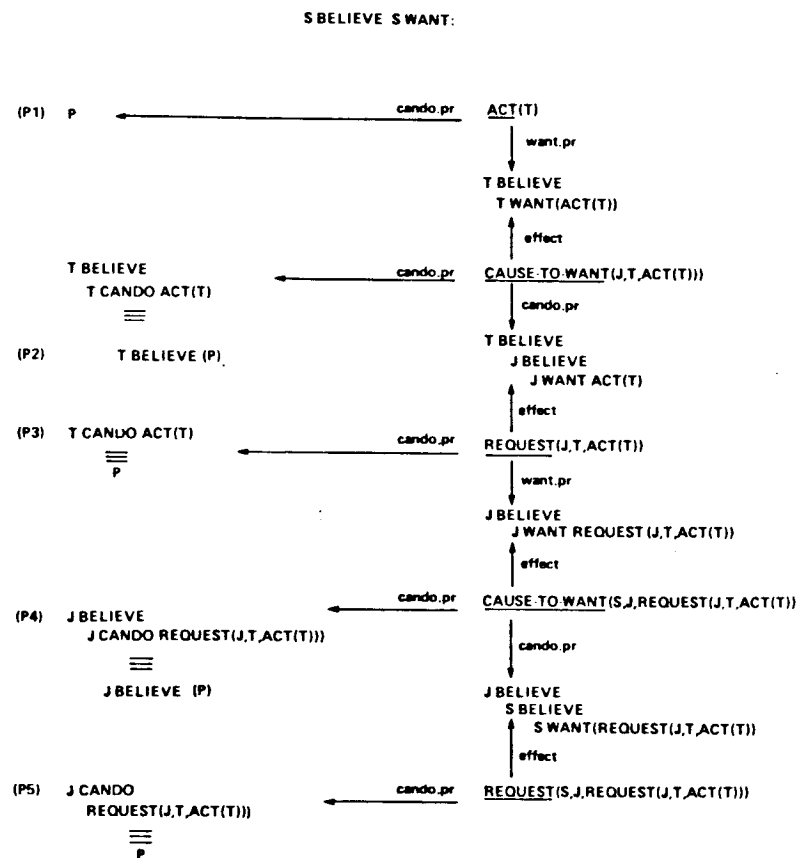


Figure 7. A third party REQUEST using the "neutral" precondition.

The preconditions that have to be satisfied in S's plan are:

- S BELIEVE:
- (P1) P (also P3 and P5)
 - (P2) T BELIEVE (P)
 - (P4) J BELIEVE (P)

Conditions P3 and P5 are the same as P1, and thus the preconditions to the REQUESTs in the plan, are independent of the speaker's beliefs; they depend only on the planner's beliefs. While the use of the neutral precondition eliminates prior knowledge requirements for REQUESTs *per se*, condition P4 still requires, as a precondition to CAUSE-TO-WANT, that the planner have some knowledge of the intermediary's beliefs. The next section shows why the planner need not have such beliefs at the time of plan construction.

7.2 Side Effects

The performance of a speech act has thus far been modeled as resulting in an EFFECT that is specific to each speech act type. But, by the very fact that a speaker has attempted to perform a particular speech act, a hearer learns more—on identifying which speech act was performed, a hearer learns that the speaker believed the various preconditions in the *plan* that led to that speech act held. The term *side effect* will be used to refer to the hearer's acquisition of such beliefs by way of the performance of a speech act. Since the plan the hearer infers for the

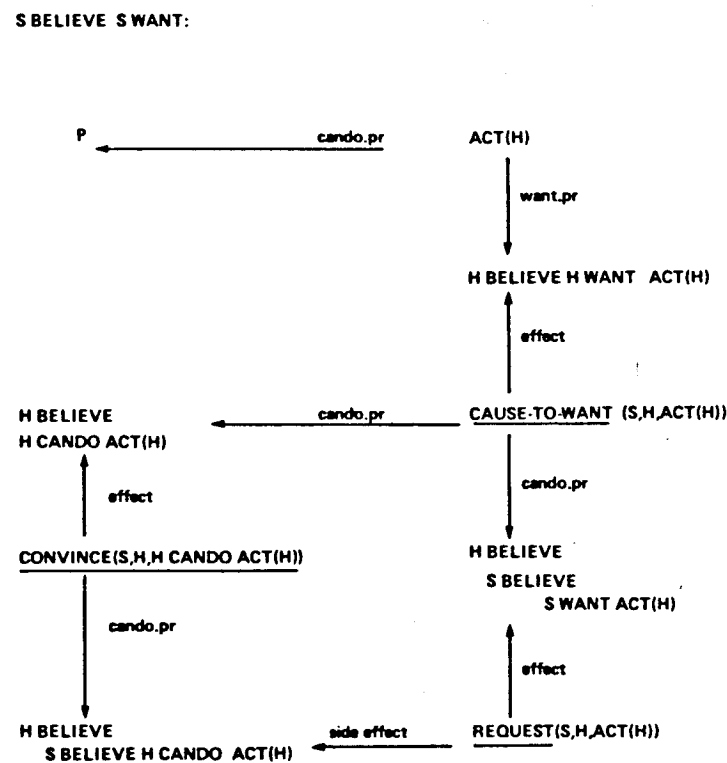


Figure 8. A REQUEST with side effects.

The *CANDO.PR*s of speech acts defined according to this principle not only resolve our difficulties with composite speech acts, but they also behave as desired for the usual noncomposite cases since preconditions now depend only on the *planner's* beliefs, and the planner is often the speaker. Thus speech act operator definitions are intimately bound to the form of the planning system.

The only result the new principle has on the form of the *EFFECT*s of speech acts is to make clear whose beliefs should be updated with those *EFFECT*s. After successfully executing a speech act to *H*, the speaker can update his model of *H* with the speech act's *EFFECT*s. But, for a composite speech act *ultimately* directed to *H*, the initial planner must observe or assume the success of the rest of the multiparty plan in order to conclude that the *EFFECT*s of the final speech act to *H* hold.

While the new principle guarantees that the *EFFECT*s of speech acts are independent of the use of intermediaries, hearers have every right to believe that the speakers of those speech acts believe that the preconditions hold. Because side effects are stated in terms of the hearer's beliefs about the speaker's beliefs, intermediaries are vulnerable to a charge of insincerity if they brazenly execute the speech acts they were requested to perform. It is to avoid such a charge, and thus make intermediaries "responsible for" the speech acts they execute, that we place the condition on *CAUSE-TO-WANT* stating that *AGT BELIEVE AGT CANDO ACT*.

Finally, to complete the reexamination of speech act definitions we point out that the *WANT.PR* also has a *SPEAKER BELIEVE* on it. One cannot, in the spirit of "housecleaning," remove the *SPEAKER BELIEVE SPEAKER WANT* from the *WANT.PR* of speech acts since a speaker's goal cannot be characterized independently of the speaker's beliefs, unless one is willing to model someone's "unconscious" goals. We are not.²⁷

7.4 New Definitions of REQUEST and INFORM

Using this principle, *REQUEST* is redefined as:

<u>REQUEST(SPEAKER, HEARER, ACT)</u>	
CANDO.PR:	HEARER CANDO ACT
WANT.PR:	SPEAKER BELIEVE SPEAKER WANT request-instance
EFFECT:	HEARER BELIEVE SPEAKER BELIEVE SPEAKER WANT ACT

The principle applied to the definition of the operator *INFORM* results in a *CANDO.PR* stated as *PROP* rather than as *SPEAKER BELIEVE PROP*.²⁸ Such a change allows one to plan to request an intermediary, say a child, to tell

someone else that the key is in the closet without the planner's having to believe, at the time of planning, that the child thinks so. The new definition of *INFORM* then becomes:

<u>INFORM(SPEAKER, HEARER, PROP)</u>	
CANDO.PR:	PROP
WANT.PR:	SPEAKER BELIEVE SPEAKER WANT inform-instance
EFFECT:	HEARER BELIEVE SPEAKER BELIEVE PROP

Regarding the other informing speech acts, the principle cannot be used to justify the deleting of the *SPEAKER BELIEVE* from the *CANDO.PR* of *INFORMREF* and *INFORMIF* since the highest elements of those conditions are "I" and "OR", respectively. Intuitively speaking, this is a sensible result since a speaker *SP* cannot plan for an intermediary, *INT*, to tell *H* whether *P* is true, or what the value of description *D* is unless *INT* is believed to have that information.

7.5 Summary

The appropriate planning of composite speech acts has turned out to be a powerful test of the adequacy of speech act definitions. To meet its demands on the planning of questions and multiparty speech acts, two new speech acts, *INFORMREF* and *INFORMIF* have been defined, and the preconditions to *REQUEST* and *INFORM* have been reformulated according to a point-of-view principle. Since these last two speech acts were taken to be prototypes of Searle's (1976) "directive" and "representative" classes, the principle will find wide application.

A side effect of direct requests was identified and used in planning multiparty speech acts. Side effects, however, cannot be calculated until the hearer has recognized the speaker's plan and thus has classified the observed utterance as a particular speech act type. Thus the minimal side effect formulation given here should be further justified on the basis of what a hearer needs to assume about the speaker's beliefs in order to identify an utterance's illocutionary force.

There may be other ways to meet compositional adequacy. For instance, one could state explicitly that an action's preconditions should be true at the time the action is to be done (Bruce, 1975). For our multiparty *REQUEST*s, such an approach (using a speaker-based precondition) produces preconditions like: *S* believes *J* will believe *P* will be true when *ACT* is to be done, which seems reasonable. However, the minimal side effect of *S*'s *REQUEST* then becomes: *J* now believes that (before that *REQUEST*) *S* expected *J* to believe that *P* would be true when *ACT* is done (where "now" is just after the *REQUEST* was made). As yet, we do not have an analogue of *CONVINCE* that would allow *J* to then come to believe that *P* would be true. Again, if *REQUEST* is defined using the neutral precondition, this problem does not arise.

²⁷The fact that a *WANT.PR* is found on every intentional act makes us suspect that it belongs on some single "element" that is present for every act.

²⁸Of course, what must be satisfied in any plan for *INFORM* is that the planner believe *PROP*.

8. CONCLUDING REMARKS

It has been argued that a theory of speech acts can be obtained by modelling them in a planning system as operators defined, at least, in terms of the speakers' and hearers' beliefs, and goals. Thus, speech acts are treated in the same way as physical acts, allowing both to be integrated into plans. Such an approach suggests new areas for application. It may provide a more systematic basis for studying real dialogues arising in the course of a task—a basis that would facilitate the tracking of conversants' beliefs and intentions as dialogue and task proceed. A similar analysis of characters' plans has also been shown (Bruce & Newman, 1978) to be essential to a satisfactory description of narrative. Finally, Allen (1979) and Cohen (1978) have suggested how computer conversants might plan their speech acts and recognize those of their users.

Given this range of application, the methodological issues of how speech acts should be modelled in a planning system become important. Specifically, a plan-based competence theory, given configurations of beliefs and goals, speech act operators, and plan construction inferences should generate plans for all and only those speech acts that are appropriate in those configurations. This paper developed tests that showed how various definitions of the speech acts of requesting and informing were inadequate, especially to the demand that they generate appropriate plans when composed with other speech acts to form questions and multiparty requests.

To resolve the difficulties, two "views" of INFORM to be used in constructing questions were defined, allowing the questioner to have incomplete knowledge of the hearer's beliefs. After revising both the form of speech act preconditions and identifying some speech act side effects, compositional adequacy for multiparty REQUESTS was achieved. The solution led to a metatheoretical "point-of-view" principle for use in defining future speech acts as operators within this planning system.

Our approach has both assumed certain idealized properties of speaker/hearers, and has been restricted in its scope. The preconditions and effects of our operators are stated in the language of logic, not because of any desire to perform logically valid inferences, but because the conditions in the plans should have well-defined semantics. While this has been partially realized through the adoption of the possible-worlds semantics for belief, the semantics is too strong to be a faithful model of human beliefs. For instance, it leads here to requiring a questioner to have very strong, though incomplete, knowledge of the hearer's beliefs. To reflect human beliefs more accurately, one needs to model (at least): degrees of belief, justifications, the failure to make deductions, inductive leaps, and knowing what/who/where something is. These refinements, though needed by a theory of speech acts, are outside its scope. Finally, the semantics for WANT and for actions are lacking (but see Moore (1979) for an interesting approach to the latter).

Only two kinds of speech acts, prototypes of Searle's (1976) directive and

representative classes, have been examined here, but the approach can be extended to other members of those classes (Bruce, 1975) and perhaps to the commissive class that includes promises. However, in order to model promises and warnings, a better understanding of the concepts of benefit and obligation is necessary.

Finally, we have so far discussed how a planning system can select illocutionary force and propositional content of a speech act, but not how utterances realizing it can be constructed nor how illocutionary acts can be identified from utterances. Extending the plan-based approach to the first area means investigating the extent of "pragmatic influence" of linguistic processing. An important subproblem here is the planning of referring expressions involved in performing illocutionary acts (Perrault & Cohen, forthcoming; Searle, 1969). Regarding speech act identification, the acid-test of a plan-based approach is its treatment of indirect speech acts (Searle, 1975). Gordon and Lakoff (1971) proposed "conversational postulates" to account for the relation between the direct or literal and the indirect illocutionary forces of an utterance. But, as Morgan (1977) notes, by calling them "postulates," one implies they cannot be explained by some other independently motivated analysis.

We suggest that the relation between direct and indirect readings can be largely accounted for by considering the relationship between actions, their preconditions, effects, and bodies, and by modelling how language users can recognize plans, which may include speech acts, being executed by others. The ability to recognize plans is seemingly required in order to be *helpful*, independent of the use of indirect speech acts. For instance, hearers often understand a speaker's utterance literally but go beyond it, inferring the speaker's plans and then performing acts that would enable the speaker's higher level goals to be fulfilled. Indirect speech acts arise because speakers can intend hearers to perform helpful inferential processing and they intend for hearers to know this. Allen (1979) and Perrault and Allen (forthcoming) formalize this process of intended plan-recognition (and thus Searle's force condition) extending our plan-based approach to the interpretation of indirect speech acts.

ACKNOWLEDGMENTS

We would like to thank Marilyn Adams, James Allen, Ron Brachman, Chip Bruce, Sharon Oviatt, Bill Woods and the referees for their comments, and Brenda Starr, Jill O'Brien, and Beverly Tobiason for their tireless assistance in the paper's preparation. Special thanks are extended to Brenda Starr for her invaluable editorial help.

REFERENCES

- Allen, J. A plan-based approach to speech act recognition. Ph.D. Thesis, Technical Report No. 131/79, Dept. of Computer Science, University of Toronto, January, 1979.
- Austin, J. L. *How to do things with words*. J. O. Urmson (Ed.), Oxford University Press, 1962.
- Bruce, B. Belief systems and language understanding. Report No. 2973, Bolt Beranek and Newman, Inc. January, 1975.
- Bruce, B., & Newman, D. Interacting plans. *Cognitive Science*, 1978, 2, 195-233.
- Bruce, B., & Schmidt, C. F. Episode understanding and belief guided parsing. Presented at the Association for Computational Linguistics Meeting at Amherst, Massachusetts (July 26-27, 1974).
- Carbonell, J. G. Jr. POLITICS: Automated idealogical reasoning. *Cognitive Science*, 1978, 2, 27-51.
- Chomsky, N. *Aspects of the theory of syntax*. Cambridge, Mass. MIT Press, 1965.
- Cohen, P. R. On knowing what to say: Planning speech acts. Ph.D. Thesis, Technical Report No. 118, Department of Computer Science, University of Toronto, January 1978.
- Deutsch, B. G. The structure of task-oriented dialogues. In L. D. Erman (Ed.), *Proceedings of the IEEE symposium on speech recognition*. Pittsburgh, PA: Carnegie-Mellon University, 1974.
- Donnellan, K. Reference and definite description. In *The Philosophical Review*, v. 75, 1960, 281-304. Reprinted in Steinberg & Jacobovits (Eds.), *Semantics*, Cambridge University Press, 1966.
- Fikes, R., & Nilsson, N. J. STRIPS: A new approach to the application of theorem proving to problem solving. *Artificial Intelligence*, 1971, 2, 189-208.
- Gordon, D., & Lakoff, G. Conversational postulates. *Papers from the Seventh Regional Meeting*, Chicago Linguistic Society, 1971, 63-84.
- Green, C. Application of theorem-proving techniques to problem-solving. In D. E. Walker & L. M. Norton (Eds.), *Proceedings of the international joint conference on artificial intelligence*. Washington, D.C., May 1969.
- Grice, H. P. Meaning. In *The Philosophical Review*, 1957, 66, 377-388. Reprinted in D. A. Steinberg & L. A. Jacobovits (Eds.), *Semantics: An interdisciplinary reader in philosophy, linguistics, and psychology*. New York: Cambridge University Press, 1971.
- Hintikka, J. *Knowledge and belief*. Ithaca: Cornell University Press, 1962.
- Hintikka, J. Semantics for propositional attitudes. In J. W. Davis et al. (Eds.), *Philosophical logic*. Dordrecht-Holland: D. Reidel Publishing Co., 1969. Reprinted in L. Linsky (Ed.), *Reference and modality*. New York: Oxford University Press, 1971.
- Lewis, D. K. *Convention: A philosophical study*. Cambridge, Mass: Harvard University Press, 1969.
- McCarthy, J., & Hayes, P. J. Some Philosophical Problems from the Standpoint of Artificial Intelligence. In B. Meltzer & D. Michie (Eds.) *Machine intelligence 4*, New York: American Elsevier, 1969.
- Meehan, J. R. Tale-spin, an interactive program that writes stories. In *Proceedings of the fifth international joint conference on artificial intelligence*, Cambridge, Mass., 91-98.
- Moore, R. C. Reasoning about knowledge and action. Ph.D. Thesis, Artificial Intelligence Laboratory, Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, February, 1979.
- Morgan, J. Conversational postulates revisited. *Language*, 1977, 277-284.
- Newell, A., & Simon, H. A. GPS, A program that simulates human thought. In E. A. Feigenbaum & J. Feldman (Eds.), *Computers and thought*. New York: McGraw Hill, 1963.
- Perrault, C. R., & Allen, J. F. A plan-based analysis of indirect speech acts. Forthcoming.
- Perrault, C. R., & Cohen, P. R. Inaccurate Reference. *Proceedings of the workshop on computational aspects of linguistic structure and discourse setting*, Joshi, A. K., Sag, I. A., & Webber, B. L. (Eds.), Cambridge University Press, forthcoming.
- Sacerdoti, E. D. A structure for plans and behavior. Ph.D. Thesis, Technical Note 109, Artificial Intelligence Center, Stanford Research Institute, Menlo Park, California, August 1975.
- Schank, R., & Abelson, R. *Scripts, plans, goals, and understanding*. Hillsdale, N.J.: Lawrence Erlbaum Associates, 1977.
- Schiffer, S. *Meaning*. Oxford: Oxford University Press, 1972.
- Searle, J. R. A taxonomy of illocutionary acts. In K. Gunderson (Ed.), *Language mind and knowledge*, University of Minnesota Press, 1976.
- Searle, J. R. Indirect speech acts. In P. Cole & J. L. Morgan (Eds.), *Syntax and semantics*, (Vol. 3), *Speech acts*. New York: Academic Press, 1975.
- Searle, J. R. *Speech acts: An essay in the philosophy of language*. Cambridge: Cambridge University Press, 1969.
- Strawson, P. F. Intention and convention in speech acts. In *The Philosophical Review*, v: lxxiii, 1964. Reprinted in *Logico-linguistic papers*, London: Methuen & Co., 1971.