Plans, Inference, and Indirect Speech Acts*

James F. Allen
Computer Science Department
University of Rochester
Rochester, NY 14627

C. Raymond Perrault
Computer Science Department
University of Toronto
Toronto, Canada M5S 1A7

Introduction

One of the central concerns of a theory of pragmatics is to explain what actions language users perform by making utterances. This concern is also relevant to the designers of conversational language understanding systems, especially those intended to cooperate with a user in the execution of some task (e.g., the Computer Consultant task discussed in Walker [1978]).

All actions have effects on the world, and may have preconditions which must obtain for them to be successfully executed. For actions whose execution causes the generation of linguistic utterances (or speech acts), the preconditions may include the speaker/writer holding certain beliefs about the world, and having certain intentions as to how it should change ([Austin, 1962], [Searle, 1969]).

In Cohen [1978] and Cohen and Perrault [1979] it is suggested that speech acts** be defined in the context of a planning system (e.g., STRIPS of Fikes and Nilsson [1971]) i.e., as a class of parameterized procedures called operators, whose execution can modify the world. Each operator is labelled with formulas stating its preconditions and effects.

The major problem of a theory of speech acts is relating the form of utterances to the acts which are performed by uttering them. Several syntactic devices can be used to indicate the speech act being performed: the most obvious are explicit performative verbs, mood, and intonation. But no combination of these provides a clear, single-valued function from form to illocutionary force. For example, (1.a)-(1.e) and even (1.f) can be requests to pass the salt.

(1.a) I want you to pass the salt.
(1.b) Do you have the salt?
(1.c) Is the salt near you?
(1.d) I want the salt.
(1.e) Can you pass the salt?
(1.f) John asked me to ask you to pass the salt.

Furthermore, all these utterances can also be intended literally in some contexts. For example, a parent leaving a child at the train station may ask "Do you know when the train leaves?" expecting a yes/no answer as a confirmation.

---

* This research was supported in part by the National Research Council of Canada under Operating Grant A9285.

** Unless otherwise indicated, we take "speech act" to be synonymous with "illocutionary act."

The object of this paper is to discuss, at an intuitive level, an extension to the work in Cohen [1978] to account for indirect speech acts. Because of space constraints, we will need to depend explicitly on the intuitive meanings of various terms such as plan, action, believe, and goal. Those interested in a more rigorous presentation should see [Allen, 1979] or [Perrault and Allen, forthcoming]. The solution proposed here is based on the following simple and independently motivated hypotheses:

(2.a) Language users are rational agents and thus speech acts are purposeful. In particular, they are a means by which one agent can alter the beliefs and goals of another.

(2.b) Rational agents are frequently capable of identifying actions being performed by others and goals being sought. An essential part of helpful behavior is the adoption by one agent of a goal of another, followed by an attempt to achieve it. For example, for a store clerk to reply "How many do you want?" to a customer who has asked "Where are the steaks?", the clerk must have inferred that the customer wants steaks, and then he must have decided to get them himself. This might have occurred even if the clerk knew that the customer had intended to get the steaks himself. Cooperative behavior must be accounted for independently of speech acts, for it often occurs without the use of language.

(2.c) In order for a speaker to successfully perform a speech act, he must intend that the hearer recognize his intention to achieve certain (perlocutionary) effects, and must believe it is likely that the hearer will be able to do so. This is the foundation the account of illocutionary acts proposed by Strawson [1964] and Searle [1969], based on Grice [1957].

(2.d) Language users know that others are capable of achieving goals, of recognizing actions, and of cooperative behavior. Furthermore, they know that others know they know, etc. Thus, a speaker may intend not only that his actions be recognized but also that his goals be inferred, and that the hearer be cooperative.

(2.e) Thus a speaker can perform one speech act A by performing another speech act B if he intends that the hearer recognize not only that B was performed but also that through cooperative behavior by the hearer, intended by the speaker, the effects of A should be achieved.

The Speech Act Model

In the spirit of Searle [1975], Gordon and Lakoff [1975], and Morgan [1978], we propose an account of speech acts with the following constituents:

(3.a) For each language user S, a model of the beliefs and plans of other language users A with which s/he is communicating, including a model of A's model of S's beliefs and plans, etc.

(3.b) Two sets of operators for speech acts: a set of surface level operators which are realized by utterances having specific syntactic and semantic features (e.g., mood), and a set of illocutionary level operators which are performed by performing surface level ones. The illocutionary acts model the intent of the speaker independent of the form of the utterance.

(3.c) A set of plausible inference rules with which language users construct and recognize plans. It is convenient to view the rules as either simple or augmented. A couple of examples of simple plan recognition rules are:

> [Action-Effect Inference]
>
> "If agent S believes that agent A wants to do action ACT then it is plausible that S believes that A wants to achieve the effects of ACT."
>
> [Know-Positive Inference]
>
> "If S believes A wants to know whether a proposition P is true, then it is plausible that S believes that A wants to achieve P."

Of course, given the conditions in the second inference above, S might also infer that A has a goal of achieving not P. This is another possible inference. Which applies in a given setting is determined by the rating heuristics (see 3.d below).

Simple rules can be augmented by adding the condition that the recognizer believes that the other agent intended him to perform the inference. An example of an augmented recognition rule is:

> "If S believes that A wants S to recognize A's intention to do ACT, then it is plausible that S believes that A wants S to recognize A's intention to achieve the effects of ACT."

Notice that the augmented rule is obtained by introducing "S believes A wants" in the antecedent and consequent of the simple rule, and by interpreting "S recognizes A's intention" as "S comes to believe that A wants." These rules can be constructed from the simple ones by assuming that language users share a model of the construction and recognition processes.

(3.d) A set of heuristics to guide plan recognition by rating the plausibility of the outcomes. One of the heuristics is: "Decrease the plausibility of an outcome in which an agent is believed to be executing an action whose effects he already believes to be true." Script-derived expectations also provide some of the control of the recognition process.

(3.e) A set of heuristics to identify the obstacles in the recognized plan. These are the goals that the speaker cannot easily achieve without assistance. If we assume that the hearer is cooperating with the speaker, the hearer will usually attempt to help achieve these goals in his response.

With these constituents, we have a model of helpful behavior: an agent S hears an utterance from some other agent A, and then identifies the surface speech act. From this, S applies the inference rules to reconstruct A's plan that produced the utterance. S can then examine this plan for obstacles and give a helpful response based on them. However, some of the inference rules may have been augmented by the recognition of intention condition. Thus, some obstacles may have been intended to be communicated by the speaker. These specify what illocutionary act the speaker performed.

An Example

This may become clearer if we consider an example. Consider the plan that must be deduced in order to answer (4.a) with (4.b):

(4.a) A: Do you know when the Windsor train leaves?

(4.b) S: Yes, at 3:15.

The goal deduced from the literal interpretation is that

(4.c) A wants to know whether S knows the departure time.

From this goal, S may infer that A in fact wants (4.d) by the Know-Positive Inference:

(4.d) A wants S to know the departure time

from which S may infer that

(4.e) A wants S to inform A of the departure time

by the precondition-action inference (not shown). S can then infer, using the action-effect inference, that

(4.f) A wants to know the departure time.

S's response (4.b) indicates that he believed that both (4.c) and (4.f) were obstacles that S could overcome in this response.

However, a sentence such as (4.a) could often be uttered in a context where the literal goal is not an obstacle. For instance, A might already know that S knows the departure time, yet still utter (4.a). In such cases, A's goals are the same as if he had uttered the request

(4.g) When does the Windsor train leave?

Hence (4.a) is often referred to as an indirect request.

Thus we have described two different interpretations of (4.a):

a) A said (4.a) merely expecting a yes/no answer, but S answered with the extra information in order to be helpful;

b) A said (4.a) intending that S deduce his plan and realize that A really wants to know the departure time.

Theoretically, these are very different: (a) describes a yes/no question, while (b) describes an (indirect) request for the departure time. But the distinction is also important for practical reasons. For instance, assume S is not able to tell A the departure time for some reason. With interpretation (a), S can simply answer the question, whereas with interpretation (b), S is obliged to give a reason for not answering with the departure time.

The distinction between these two cases is simply that in the latter, S believes that A intended S to make the inferences above and deduce the goal (4.f). Thus the inferences applied above were actually augmented inferences as described previously. In the former interpretation, S does not believe A intended S to make the inferences, but did anyway in order to be helpful.

Concluding Remarks

This speech act model was implemented as part of a program which plays the role of a clerk at a train station information booth [Allen, 1979]. The main results are the following:

(5.a) It accounts for a wide class of indirect forms of requests, assertions, and questions, including the examples in (1). This includes idiomatic forms such as (1.a) and non-idiomatic ones such as (1.f). It does so using only a few independently necessary mechanisms.

(5.b) It maintains a distinction between illocutionary and perlocutionary acts. In particular, it accounts for how a given response by one participant B to an utterance by A may be the result of different chains of inferences made by B: either B believed the response given was intended by A, or B believed that the response was helpful (i.e., non-intended). It also shows some ways in which the conversational context can favor some interpretations over others.

The main objective of our work is to simplify the syntactic and semantic components as much as possible by restricting their domain to literal meanings. The indirect meanings are then handled at the plan level.

There remain several open problems in a theory of speech acts which we believe to be largely independent of the issue of indirection, notably identifying the features of a text which determine literal illocutionary force, as well as constructing representations adequate to express the relation between several illocutionary force indicators which may be present in one sentence (see [Lakoff, 1974] and [Morgan, 1973]).

Bibliography

Allen, J.F. A Plan-Based Approach to Speech Act Recognition. Ph.D. thesis, Computer Science Department, University of Toronto, 1979.

Austin, J.L. How To Do Things With Words. New York, Oxford University Press, 1962.

Brown, G.P. An Approach to Processing Task-Oriented Dialogue, unpublished ms, MIT, 1978.

Cohen, P.R. On Knowing What to Say: Planning Speech Acts, TR 118, Computer Science Department, University of Toronto, January 1978.

Cohen, P.R. and Perrault, C.R. Elements of a Plan Based Theory of Speech Acts, forthcoming.

Cole, P. and Morgan, J.L. Syntax and Semantics, Vol 3: Speech Acts. New York, Academic Press, 1975.

Fikes, R.E. and Nilsson, N.J. STRIPS: A New Approach to the Application of Theorem Proving to Problem Solving. Artificial Intelligence 2, 189-205, 1971.

Gordon, D. and Lakoff, G. Conversational Postulates, in Cole and Morgan (eds), 1975.

Grice, H.H. Meaning. Phil. Rev. 66, 377-388, 1957.

Lakoff, G. Syntactic Amalgams. CLS 10, 321-344, 1974.

Morgan, J.L. Sentence Fragments and the Notion 'Sentence,' in B.B. Kachru et al. (eds), Issues in Linguistics. Urbana, University of Illinois Press, 1973.

Morgan, J.L. Towards a Rational Model of Discourse Comprehension, in Proceedings 2nd Conf. Theoretical Issues in Natural Language Processing, Champaign-Urbana, 1978.

Perrault, C.R. and Allen, J.F. A Plan-Based Analysis of Indirect Speech Acts, in preparation.

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

Searle, J.R. Indirect Speech Acts, in Cole and Morgan (eds), 1975.

Strawson, P.F. Intention and Convention in Speech Acts. Phil. Rev. 73, 4, 439-460, 1964.

Walker, D.E. Understanding Spoken Language. New York, North Holland, 1978.