

A Plan-Based Analysis of Indirect Speech Acts¹

C. Raymond Perrault

Department of Computer Science
University of Toronto
Toronto, Canada M5S 1A7

James F. Allen

Department of Computer Science
University of Rochester
Rochester, New York 14627

We propose an account of indirect forms of speech acts to request and inform based on the hypothesis that language users can recognize actions being performed by others, infer goals being sought, and cooperate in their achievement. This cooperative behaviour is independently motivated and may or may not be intended by speakers. If the hearer believes it is intended, he or she can recognize the speech act as indirect; otherwise it is interpreted directly. Heuristics are suggested to decide among the interpretations.

1. Introduction

Austin [1962] was one of the first to stress the distinction between the *action(s)* which a speaker performs by uttering a sentence (such as informing, requesting, or convincing) and the *truth conditions* of propositions contained in the sentence. Actions have effects on the world, and may have preconditions which must obtain for them to be felicitously performed. For actions whose execution involves the use of language (or *speech acts*), the preconditions may include the speaker holding certain beliefs about the world, and having certain intentions or wants as to how it should change.

As well as being important to the study of natural language semantics, speech acts are important to the designer of conversational natural language understanding systems. Such systems should be able to recognize what actions the user is performing. Conversely, if such a system is to acquire information or request assistance from its user, it should know how and when to ask questions and make requests. (See Bruce [1975] for an early attempt.)

Cohen and Perrault [1979] (hereafter referred to as CP) argue for the distinction between a *competence*

theory of speech acts, which characterizes what utterances an ideal speaker can make in performing what speech acts, and a *performance* theory which also accounts for how a particular utterance is chosen in given circumstances, or how it is recognized. We are only concerned here with a competence theory.

In Perrault, Allen, and Cohen [1978] we suggested that it is useful to consider speech acts in the context of a *planning system*. A planning system consists of a class of parameterized procedures called *operators*, whose execution can modify the world. Each operator is labelled with formulas stating its preconditions and effects. A plan construction algorithm is a procedure which, given a description of some initial state of the world and a goal state to be achieved, constructs a *plan*, or sequence of operators, to achieve it.

It is assumed there, and in all our subsequent work, that language users maintain a model of the world (their *beliefs*) and a set of goals (their *wants*). One person S's beliefs may include beliefs about another person A's beliefs and wants, including A's beliefs about S, etc. We do not concern ourselves with obligations, feelings, etc., which clearly can also be affected by speech acts.

CP discuss criteria for judging the correctness of the preconditions and effects of the operators corresponding to speech acts, and specifically those of the acts INFORM and REQUEST. However, the conditions on INFORM and REQUEST given in CP are at best necessary and certainly not sufficient. In particu-

¹ This research was supported in part by the National Research Council of Canada under Operating Grant A9285. Thanks to Phil Cohen, Michael McCord, Corot Reason, and John Searle for their comments. We assume the usual responsibility for remaining inaccuracies, misunderstandings, and downright errors.

lar they say nothing about the form of utterances used to perform the speech acts. Several syntactic devices can be used to indicate the speech act being performed: the most obvious are explicit performative verbs such as "I hereby request you to ...", and mood (indicative for assertions, imperative for requests to do, interrogative for requests to inform). But the mood of an utterance is well known to not completely specify its illocutionary force: 1.a-b can be requests to close the door, 1.c-e can be requests to tell the answer, and 1.f can be an assertion.

- (1.a) I want you to close the door.
- (1.b) Will you close the door?
- (1.c) Tell me the answer.
- (1.d) I want you to tell me the answer.
- (1.e) Do you know what the answer is?
- (1.f) Do you know that Jack is in town?

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 1.g expecting a yes/no answer as a confirmation.

- (1.g) Do you know when the train leaves?

The object of this paper is to extend the work in CP to account for indirect use of mood, loosely called *indirect speech acts*. The solution proposed here is based on the following intuitively simple and independently motivated hypotheses:

- (1) Language users are rational agents engaged in goal seeking behaviour. Among these goals are the modification of the beliefs and goals of other agents.
- (2) Rational agents are frequently capable of identifying actions being performed by others and goals being sought. An essential part of helpful or cooperative behaviour 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, then he must have decided to get them himself. This might have occurred even if the customer had intended to get the steaks him or herself. Cooperative behaviour must be accounted for independently of speech acts, for it often occurs without the use of language.
- (3) In order for a speaker to successfully perform a speech act, he must intend that the hearer recognize his intention to achieve the effects of the speech act, and must believe it is likely that the hearer will be able to do so. This is the foundation for the philosophical account of speech acts.

- (4) Language users know that others are capable of achieving goals, of recognizing actions, and of cooperative behaviour. Furthermore, they know that others know they know, etc. A speaker may intend not only that his actions be recognized but also that his goals be inferred.
- (5) 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 behaviour by the hearer, intended by the speaker, the effects of A should be achieved. The speaker must also believe that it is likely that the hearer can recognize this intention.

The process by which one agent can infer the plans of another is central to our account of speech acts. Schmidt et al [1978] and Genesereth [1978] present algorithms by which one agent can infer the goals of another, but assuming no interaction between the two. We describe the process in terms of a set of *plausible plan inference rules* directly related to the rules by which plans can be constructed. Let A and S be two agents and ACT an action. One example of a simple plan inference rule is:

"If S believes that A wants to do ACT then it is plausible that S believes that A wants to achieve the effects of ACT."

From simple rules like this can be derived more complex plan inference rules such as:

"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 complex 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". Throughout the paper we identify "want" and "intend".

We show that rules of the second type can account for S's recognition of many indirect speech acts by A, i.e. those in which S recognizes A's intention that S perform cooperative acts.

To distinguish the use of, say, the indicative mood, in an assertion from its use in, say, an indirect request, the speech act operators REQUEST and INFORM of CP are reformulated and two further acts S.REQUEST and S.INFORM are added. These *surface level acts* are realized literally as indicative and imperative utterances. An S.REQUEST to INFORM is realized as a question. The surface level acts can be recognized immediately as parts of the higher level (or *illocutionary level*) acts, to which the simple plan construction

and inference rules can apply. Alternatively, the complex rules can be applied to the effects of the surface acts, and the intended performance of one of the illocutionary acts inferred later.

For example, there are two ways an agent S could be led to tell A the secret after hearing A tell him "Can you tell me the secret?". Both start with S's recognition that A asked a yes/no question. In the first case, S assumes that A simply wanted to know whether S could tell the secret, then infers that A in fact wants to know the secret and, helpfully, decides to tell it. In the second case S recognizes that A *intends* S to infer that A wants to know the secret and that A *intends* S to tell A the secret, and thus that A has requested S to tell the secret.

In general, several of the plan inference rules could apply at any time, and none of them guarantees a valid consequence. The application of the rules is controlled by a set of heuristics which rate the plausibility of the outcomes.

Following a review of the relevant aspects of speech act theory in section 2, section 3 outlines our assumptions about beliefs, goals, actions, plans, and the plan inference process. Section 4 shows how the speech act definitions and the plan inference process can be used to relate literal to indirect meanings for REQUESTs and INFORMs. We show how utterances such as 1.h-l, and even 1.m can be used as requests to pass the salt, and what the origin of the several interpretations of 1.m is.

- (1.h) I want you to pass the salt.
- (1.i) Do you have the salt?
- (1.j) Is the salt near you?
- (1.k) I want the salt.
- (1.l) Can you pass the salt?
- (1.m) John asked me to ask you to pass the salt.

Similarly we show how 1.n can be used to inform while 1.o cannot. Section 5 relates this work to the literature, while section 6 suggests further problems and draws some conclusions.

- (1.n) Do you know that the train is late?
- (1.o) Do you believe that the train is late?

The speech act recognition process described here has been implemented as a computer program and tested by having it simulate an information clerk at a railway station. This domain is real, but sufficiently circumscribed so that interchanges between clerk and patrons are relatively short and are directed towards a limited set of goals. The program accepts as input simple English sentences, parses them using an ATN parser, and produces as output the speech act(s) it recognized and their associated propositional contents. It can handle all the examples discussed here. Details of the implementation can be found in Allen [1979].

2. Introduction to Speech Acts

2.1. Basic Definitions

Prior to Austin [1962], logicians considered the meaning of a sentence to be determined only by its truth value. However, Austin noted that some sentences cannot be classified as true or false; the utterance of one of these sentences constitutes the performance of an action, and hence he named them *performatives*. To quote Austin: "When I say, before the register or altar, etc., 'I do', I am not reporting on a marriage: I am indulging in it".

Examples like this, and his inability to rigorously distinguish performative sentences from those which purportedly have truth value (which he called *constatives*) led Austin to the view that all utterances could be described as actions, or *speech acts*. He classified speech acts into three classes, the *locutionary*, *illocutionary*, and *perlocutionary* acts.

A locutionary act is an act *of* saying something: it is the act of uttering sequences of words drawn from the vocabulary of a given language and conforming to its grammar.

An illocutionary act is one performed *in* making an utterance; "promise", "warn", "inform" and "request" are names of illocutionary acts. In general, any verb that can complete the sentence "I hereby <verb> you {that | to} ..." names an illocutionary act. An utterance has *illocutionary force* F if the speaker intends to perform the illocutionary act F by making that utterance. Verbs that name types of illocutionary acts are called *performative verbs*. From now on, we take *speech acts* to mean the illocutionary acts.

Perlocutionary acts are performed *by* making the utterance. For example, S may *scare* A by warning A, or *convince* A of something by informing A of it. The success of a perlocutionary act is typically beyond the control of the speaker. For example, S cannot convince A of something against A's will, S can only present A with sufficient evidence so that A will decide to believe it. Perlocutionary acts may or may not be intentional. For instance, S may or may not intend to scare A by warning A.

Searle [1969] suggests that illocutionary acts can be defined by providing, for each act, necessary and sufficient conditions for the successful performance of the act. Certain syntactic and semantic devices, such as mood and explicit performative verbs, are used to indicate illocutionary force.

One of the conditions included in Searle's account is that the speaker performs an illocutionary act only if he intends that the hearer recognize his intention to perform the act, and thereby recognize the illocutionary force. This is important for it links Austin's work

on speech acts with the work of Grice on meaning, and is discussed in the next section.

2.2. Communication and the Recognition of Intention

Many philosophers have noted the relationship between communication (or speaker meaning) and the recognition of intention (Grice [1957, 1968], Strawson [1964], Searle [1969], Schiffer [1972].) Grice presents informally his notion of a speaker *meaning* something as follows:

"S meant something by x' is (roughly) equivalent to 'S intended the utterance of x to produce some effect in an audience by means of the recognition of this intention'"

In other words, in order for S to communicate M by uttering x to A, S must get A to recognize that S intended to communicate M by uttering x. To use an example of Grice's, if I throw a coin out the window expecting a greedy person in my presence to run out and pick it up, I am not necessarily communicating to him that I want him to leave. For me to have successfully communicated, he must at least have recognized that I intended him to leave. The same arguments hold when discussing illocutionary acts. For example, the only way S can request A to do ACT is to get A to recognize S's intention to request A to do ACT.

2.3. The Indirect Speech Act Problem

The relation between speech acts and the devices used to indicate them is complicated by the fact that performative verbs are seldom present and the same device can be used to perform many illocutionary acts. The interrogative mood, for example, can be used to

request: "Can you pass the salt?"
 question: "Do you know the time?"
 inform: "Do you know that Sam got married?"
 warn: "Did you see the bear behind you?"
 promise: "Would I miss your party?"

As many authors have pointed out, an utterance conveys its indirect illocutionary force by virtue of its literal one (Searle [1975], Morgan [1977], Morgan [1978]). "It's cold here" can function as a request to, say, close the window, in part *because* it's an assertion that the temperature is low.

Most of the literature on the treatment of indirect speech acts within the theory of grammar stems from the work of Gordon and Lakoff [1975] (hereafter GL). They claim that direct and indirect instances of the same speech act have different "meanings", i.e. different logical forms, and they propose a set of "conversational postulates" by which literal forms "entail" indirect ones. The postulates for requests correspond to conditions that must obtain for a request to be sincere. For A to sincerely request B to do ACT, the following *sincerity conditions* must hold:

- (1) A wants ACT.
- (2) B can do ACT.
- (3) B is willing to do ACT.
- (4) B will not do ACT in the absence of the request.

They then propose that one can convey a request by asserting a speaker-based sincerity condition (condition 1), or querying a hearer-based sincerity condition (conditions 2-4).

The postulates for indirect requests given in GL do not account for the readings of 2.3a and 2.3b as requests, and although more rules could be added (and some should be weakened) we believe this solution to be misguided.

- (2.3a) Is the salt near you?
- (2.3b) John asked me to ask you to pass the salt.

GL's postulates directly relate the literal form of one speech act to the indirect form of another. Thus they do not *predict* why certain acts allow certain indirect forms. For example, the postulates do not account for why 2.3c-d can be requests while 2.3e-f cannot. But 2.3e is infelicitous as a (literal) question since there is no context where one can acquire information by querying one's own mental state. Utterance 2.3f is a reasonable question but even if the speaker found out the answer, it would not get him any closer to acquiring the salt (by having the hearer pass it). A theory of indirect speech acts should capture these facts; GL's does not (although they agree it should).

- (2.3c) I want the salt.
- (2.3d) Do you want to pass the salt?
- (2.3e) Do I want the salt?
- (2.3f) Does he want to pass the salt?

Similarly, GL's postulates fail to explain the relation between indirect forms of different speech acts. For example, 2.3g can be an assertion that P and 2.3h cannot, for the same reasons that 2.3i can be a request to do A and 2.3j cannot.

- (2.3g) I want you to know that P.
- (2.3h) Do I want you to know that P?
- (2.3i) I want you to A.
- (2.3j) Do I want you to A?

The hearer's knowing that P obtains is an intended perlocutionary effect of an informing act, just as the hearer's doing an act A is an intended effect of a request. A speaker can indirectly inform or request by informing the hearer that the speaker desires the perlocutionary effect of that act, and intending that the hearer recognize the speaker's intention that the perlocutionary effect should be achieved.

This paper shows that what GL achieve with their postulates can be derived from the five hypotheses given in the Introduction. Our proposal here is a de-

velopment of Searle [1975]. It requires separating the surface form conditions completely from the definitions of the illocutionary acts and introducing an intermediary level, the surface acts.

Our theory of indirection will however share with GL some problems brought up by Sadock [1970], Green [1975], and Brown [1980]. These are discussed further in section 4.5.

3. Plans, Plan Construction, and Plan Inference.

Our analysis of indirect REQUESTs and INFORMs relies on the inference by the hearer of some of the goals of the speaker and of some of the actions which the speaker is taking to achieve those goals. Section 3.1 outlines the form of the models of the world which language users are assumed to have, in particular their beliefs about the world (and about other agents), and their goals. In section 3.2 we define actions and how they affect the belief model. The rules for plan construction and inference are considered in sections 3.3 and 3.4. Because of space limitations, this section is very sketchy. More detail, motivation, and problems, are available in Allen [1979] and Allen and Perrault [1980].

3.1. Beliefs, Knowledge, and Goals

3.1.1. The Belief Model

We assume that every agent S has a set of beliefs about the world, which may include beliefs about other agents' beliefs. Agents can hold false beliefs. As Quine [1956] pointed out, belief creates a context where substitution of coreferential expressions need not preserve truth-value.

We add to a first-order language with equality the operator B, and B(A,P) (usually written $B_A(P)$) is to be read "A believes that P", for any formula P. The B operator is assumed to satisfy the following axiom schemas (inspired by Hintikka [1962]), where P and Q are schema variables ranging over propositions, and A ranges over agents:

(B.0) all theorems of First Order Predicate Calculus

(B.1) $B_A(P) \Rightarrow B_A(B_A(P))$

(B.2) $B_A(P) \wedge B_A(Q) \Rightarrow B_A(P \wedge Q)$

(B.3) $B_A(P) \vee B_A(Q) \Rightarrow B_A(P \vee Q)$

(B.4) $B_A(\sim P) \Rightarrow \sim B_A(P)$

(B.5) $(\exists x) B_A(P(x)) \Rightarrow B_A((\exists x)P(x))$

(B.6) $(B_A(P \Rightarrow Q) \wedge B_A(P)) \Rightarrow B_A(Q)$

The rules of inference are Modus Ponens and:

If T is a theorem, then $B_A(T)$ is a theorem, for every agent A.

i.e. every agent believes every valid consequence of the logical axioms.

The partial deduction system used in the implementation of Allen [1979] is based on Cohen [1978]. The foundations for a more elaborate system can be found in Moore [1979].

3.1.2. Knowing

The word "know" is used in at least three different senses in English. One may *know that* a proposition P is true, *know whether* a proposition P is true or *know what* the referent of a description is.

We define "A knows that P", written $KNOW(A,P)$, as $P \wedge B_A(P)$. This is weaker than some definitions of "know" in the philosophical literature, where, among other things, "A knows that P" entails that A believes P for the "right reasons"; i.e. knowledge is true and *justified* belief (Ayer [1956], but see also Gettier [1963]). If S believes that A knows that P, S is committed to believing that P is true.

Unfortunately, the meaning of "A does not know that P" is not captured by $\sim(P \wedge B_A(P))$, but by the weaker $(P \wedge \sim B_A(P))$, i.e.

$$\sim KNOW(A,P) \equiv P \wedge \sim B_A(P)$$

In other words, if S believes A does not know P, then S must believe that P is true in addition to believing that A does not believe P is true. This problem is analogous to the wide/narrow scope distinction that Russell found in his account of definite descriptions (Russell [1919]). One solution to this problem is to consider KNOW as a "macro" whose expansion is sensitive to negation. Details may be found in Allen [1979].

A *knows whether* a proposition P is true if A KNOWs that P or A KNOWs that $\sim P$.

$$KNOWIF(A,P) \equiv KNOW(A,P) \vee KNOW(A,\sim P)$$

Knowing what the referent of a description is requires quantification into belief. One of its arguments is a formula with exactly one free variable.

$$\begin{aligned} KNOWREF(A,P(x)) \equiv \\ (\exists y) ((\forall z) P(z) \equiv y = z) \\ \wedge B_A((\forall z) P(z) \equiv y = z) \end{aligned}$$

A KNOWREF the departure time of TRAIN1 if TRAIN1 has a unique departure time y, and if A believes that y is TRAIN1's unique departure time.

3.1.3. Wanting

We let $W(A,P)$ (usually written $W_A(P)$) mean "agent A wants P to be true". P can be either a state or the execution of some action. In the latter case, if ACT is the name of an action, $W_A(ACT(b))$ means "A wants b to do ACT".

The logic of want is even more difficult than that of belief. It is necessary for us to accept the following:

- (W.1) $W_A(P) \equiv B_A(W_A(P)).$
 (W.2) $W_A(P \wedge Q) \Rightarrow W_A(P) \wedge W_A(Q).$

The most interesting interactions between the belief and want operators come from the models that agents have of each other's abilities to act and to recognize the actions of others. This will be further discussed in the following section.

3.2. Actions and Plans

Actions model ways of changing the world. As with the operators in STRIPS (Fikes and Nilsson [1971]), the actions can be grouped into families represented by *action schemas*, which can be viewed as parameterized procedure definitions. An action schema consists of a name, a set of parameters with constraints and a set of labelled formulas in the following classes:

Effects: Conditions that become true after the execution of the procedure.

Body: a set of partially ordered *goal states* that must be achieved in the course of executing the procedure. In the examples given here, there will never be more than one goal state in a body.

Preconditions: Conditions necessary to the successful execution of the procedure. We distinguish for voluntary actions a *want precondition*: the agent must want to perform the action, i.e. he must want the other preconditions to obtain, and the effects to become true through the achievement of the body.

The constraints on the parameters consist of type specifications, and necessary parameter interdependencies. Each action has at least one parameter, namely, the *agent* or instigator of the action. In the blocks world, for example, the action of putting one block on top of another could be defined as:

PUTON(a,b1,b2)
 constraints: AGENT(a) \wedge BLOCK(b1) \wedge BLOCK(b2)
 precondition: CLEAR(b1) \wedge CLEAR(b2)
 \wedge $W_a(\text{PUTON}(a,b1,b2))$
 effect: ON(b1,b2)

The preconditions, effects and body provide information to the plan construction and inference processes so that they can reason about the applicability and effect of performing the action in a given context. Finally, the body of the action specifies what steps must be achieved in the course of the execution of the action. *Primitive actions* have no bodies; their execution is specified by a non-examinable procedure.

All agents are assumed to believe that actions achieve their effects and require their preconditions. We need the following axioms:

For all agents a and b, and for all actions ACT, if PRE is the precondition of ACT and EFF its effect then:

- (ACT.1) $B_A(\text{ACT}(b) \Rightarrow \text{PRE}).$
 (ACT.2) $B_A(\text{ACT}(b) \Rightarrow \text{EFF}).$

Every predicate and modal operator in these axioms, and throughout the paper, should be indexed by a state or time. The resulting logic would be, accordingly, more complex. The issue is raised again in sect. 6.

3.3. Plan Construction

A *plan* to transform a world $W[0]$ (represented by a formula) into a world $W[n]$ is a sequence of actions A_1, \dots, A_n such that the preconditions of A_i are true in $W[i-1]$, and A_i transforms world $W[i-1]$ into $W[i]$.

An agent can achieve a goal by constructing and then executing a plan which transforms the current state of the world into one in which the goal obtains. This can be done by finding an operator which, if executed in some world, would achieve the goal. If its preconditions are satisfied in the initial world, the plan is complete. Otherwise, the planning process attempts to achieve the preconditions. This simple view of plan construction as a "backward chaining" process can be refined by assuming different levels of "detail" in the representation of the world and of the operators. This view (as developed in Sacerdoti [1973, 1975], for example) allows plans constructed at one level of detail to be expanded to a lower level through the bodies of their constituent acts.

As noted earlier, the agent of an action must believe that its precondition is true to believe that his executing the action will succeed. For agent A to plan that agent S should perform action ACT, A must achieve that S should believe that the precondition of ACT holds, and S's beliefs should not be inconsistent with A's, i.e. it must be true that $B_A(\text{KNOW}(S,P))$, where P is the precondition of ACT.

We assume that an agent cannot do an action without *wanting* to do that action. Thus a precondition of every action ACT by an agent A is that $W_A(\text{ACT}(A))$.

We are concerned with the model that agents have of each other's plan construction and inference process, and consider these two processes as consisting of *chains of plausible inferences* operating on goals and observed actions. The processes are specified in two parts: first as schemas of *rules* which conjecture that certain states or actions can be added to a plan being constructed. The plausibility of the plans containing the result of the inferences is then evaluated by *rating heuristics*. Thus the plan construction and inference rules are not to be interpreted as valid logical rules of inference.

The first three plan construction (PC) rules are:²

(PC.EA) [Effect-action rule] For any agent A, if Y is an effect of action X, then if A wants Y to hold, it is plausible that A will want action X to be done.

(PC.AP) [Action-precondition rule] For any agent A, if X is a precondition of action Y, and if A wants Y to be done, then it is plausible that S will want X to hold.

(PC.AB) [Action-body rule] For any agent A, if A wants an action Y to be done, and if X is a part of the body of Y then it is plausible that S will want X to be done.

If X and Y are systematically replaced by one of the pairs in Figure 1, then rules PC.EA, PC.AP, and PC.AB can all be written as

$$W_A(Y) =c=> W_A(X)$$

with $=c=>$ indicating that the rule is a construction rule.

We also need a rule based on KNOWIF:

(PC.KI) [KNOWIF rule] For any agent A, if A wants P to be true, then it is plausible that A should want to know whether P is true.

$$W_A(P) =c=> W_A(KNOWIF(A,P))$$

3.4. Plan Inference

For every plan construction rule

$$W_A(Y) =c=> W_A(X),$$

and every agent S, there is a corresponding *plan inference* (PI) rule which is written

$$B_S W_A(X) =i=> B_S W_A(Y).$$

The following rules correspond to the effect-action, action-precondition, and action-body rules of the previous section:

(PI.AE) [Action-effect rule] For all agents S and A, if Y is an effect of action X and if S believes that A wants X to be done, then it is plausible that S believes that A wants Y to obtain.

(PI.PA) [Precondition-action rule] For all agents S and A, if X is a precondition of action Y and if S believes A wants X to obtain, then it is plausible that S believes that A wants Y to be done.

(PI.BA) [Body-action rule] For all agents S and A, if X is part of the body of Y and if S believes that A wants X done, then it is plausible that S believes that A wants Y done.

There are two inverses to the KNOWIF rule: if A wants to know whether P is true, then A may want P to be true, or A may want P to be false.

² Throughout the rest of the paper agent A will usually denote the constructor/executor of plans, and S (or System) the recognizer of plans (usually constructed by A).

X	Y
ACT	effects of ACT
preconditions of ACT	ACT
body of ACT	ACT

Figure 1. Arguments for PC/PI rules.

(PI.KP) [Know positive] For all agents S and A,

$$B_S W_A(KNOWIF(A,P)) =i=> B_S W_A(P)$$

(PI.KN) [Know negative] For all agents S and A,

$$B_S W_A(KNOWIF(A,P)) =i=> B_S W_A(\sim P)$$

PI.W is the special case of the precondition-action rule where the precondition is the want precondition:

(PI.W) [Want rule] For all agents S, A, and C and for all actions ACT whose agent is C, it is plausible that

$$B_S W_A(W_C(ACT)) =i=> B_S W_A(ACT)$$

3.4.1. The Plan Inference Process

The plan inference rules generate formulas which the recognizing agent believes are possible. A separate mechanism is used to evaluate their plausibility. An agent S attempting to infer the plans of another agent A starts with an observed action of A and a (possibly empty) set of goals or *expectations* which S believes A may be trying to achieve. S attempts to construct a plan involving the action and preferably also including some of the expectations.

Plan inference is a search through a space of *partial plans* each consisting of two parts. One part is constructed using the plan inference rules from the observed action (and called the *alternative*); the other is constructed using the plan construction rules from an expected goal (and called the *expectation*).

The partial plans are manipulated by a set of *tasks* which decide what rules are to be applied, what "merges" between alternatives and expectations should be attempted, and when the process terminates. The partial plans and their associated tasks are *rated* by a set of heuristics, and the most highly rated task is executed first.

3.4.2. Rating Heuristics

The rating of a partial plan reflects how likely it is to be part of the "correct" plan, i.e. the plan the speaker is executing. If several incompatible inferences can be made from one point in the alternative, then its rating is divided among them. The heuristics described in this section are based on domain independent relations between actions, their bodies, preconditions, and effects. The need for more domain dependent measures is discussed later.

The heuristics are described here only in terms of increasing or decreasing ratings of partial plans.

Decrease the rating of a partial plan in which the preconditions of executing actions are currently false.

Decrease the rating of a partial plan containing a pending action ACT by an agent A if A is not able to do ACT.³

Decrease the rating of a partial plan in which the effects of a pending act already obtain or are not wanted by the planner.⁴

Other heuristics depending on how well the utterance fits with the expectations are not immediately relevant to understanding indirect speech acts and will not be discussed here. One further heuristic is added in section 4.3.

In general several rating heuristics are applicable to an partial plan. Their effects on the rating of the partial plan are cumulative.

3.4.3. Extending the Inference Rules

A hearer S identifies the illocutionary force of an utterance by recognizing that the speaker A has certain intentions, namely that S should recognize some intention P of A's. This can be represented by a formula of the form $B_S W_A(B_S W_A(P))$. To do the recognition, the simple plan construction and inference rules of sections 3.3 and 3.4 must be extended so that they can operate on these nested formulas. This can be done by assuming that every agent is aware that other agents construct and infer plans in the same way he can. In fact, both the simple inference and construction rules are necessary to derive the extended inference rules.

The extended rules are specified by "meta-rules" which show how to construct new PC/PI rules from old ones. The first extended construction rule (EC.1) is: A can achieve that S recognizes that A wants the effect of ACT by achieving that S recognizes that A wants ACT to be done, assuming that S would infer that the effects of ACT are also desired. The same rule applies if we replace "wants the effect of ACT" and "wants ACT to be done" by any pair of Y and X, as given in Figure 1. We assume all these substitu-

³ This definition is the same as Cohen's CANDO relation. Being able to do an action means that the action's preconditions are either presently true, achieved within the existing plan, or can be achieved by a "relatively simple plan", which we take to be a single action whose preconditions are presently true or achieved in the existing plan.

⁴ We have avoided the problem here of planning to do a task that requires one to deny a subgoal temporarily so that some action can execute, and then needing to re-achieve that (presently true) goal.

tions are possible in rules EC.1 - EC.3 and EI.1 - EI.3.

(EC.1) If $B_S W_A(X) = i \Rightarrow B_S W_A(Y)$ is a PI rule, then $W_A(B_S W_A(Y)) = c \Rightarrow W_A(B_S W_A(X))$ is a PC rule.

Similarly we can generate the corresponding PI rule:

(EI.1) If $B_S W_A(X) = i \Rightarrow B_S W_A(Y)$ is a PI rule, then $B_S W_A(B_S W_A(X)) = i \Rightarrow B_S W_A(B_S W_A(Y))$ is a PI rule.

EI.1 allows prefixing $B_S W_A$ to plan inference rules. Plan construction rules can also be embedded: if A wants S to want to do ACT, then A should be able to achieve this by achieving that S wants the effect of ACT, and by relying on S to *plan* ACT. In other words:

(EC.2) If $W_S(Y) = c \Rightarrow W_S(X)$ is a PC rule, then $W_A(W_S(X)) = c \Rightarrow W_A(W_S(Y))$ is a PC rule.

Correspondingly,

(EI.2) If $W_S(Y) = c \Rightarrow W_S(X)$ is a PC rule, then $B_S W_A(W_S(Y)) = i \Rightarrow B_S W_A(W_S(X))$ is a PI rule.

Finally, any agent A can plan for S to recognize A's intention that S plan, and for S to be able to recognize this intention in A. For example, A can plan for S to recognize A's intention that S want to close the door by planning for S to recognize A's intention that S want the door closed. These rules are obtained by using EI.2 as the PI rule which is "extended" by EC.1 and EI.1.

(EC.3) If $W_S(Y) = c \Rightarrow W_S(X)$ is a PC rule, then $W_A B_S(W_A W_S(X)) = c \Rightarrow W_A B_S(W_A W_S(Y))$ is a PC rule.

(EI.3) If $W_S(Y) = c \Rightarrow W_S(X)$ is a PC rule, then $B_S W_A(B_S W_A(W_S(Y))) = i \Rightarrow B_S W_A(B_S W_A(W_S(X)))$ is a PI rule.

Our "toolkit" is now sufficiently full to allow us to consider some speech acts and their recognition.

4. Plan Inference and Indirect Speech Acts

4.1. Speech Acts

The definitions of the speech acts REQUEST and INFORM used in this paper are slightly different from the ones in Cohen and Perrault [1979] in that they rely on the existence of speech act *bodies* to account for indirect forms. Plans including speech acts are now thought of as having two *levels*, the *illocutionary* level and the *surface* level. Acts at the illocutionary level model the intentions motivating an utterance independently of the syntactic forms used to indicate those intentions. Acts at the surface level are realized by utterances having specific illocutionary force indicators.

The first illocutionary level act is one by which a speaker informs a hearer *that* some proposition is true.

```
INFORM(speaker, hearer, P)
prec: KNOW(speaker,P) ^
      W(speaker,INFORM(speaker,hearer,P))
effect: KNOW(hearer,P)
body: B(hearer,W(speaker,KNOW(hearer,P)))
```

For A to sincerely inform S that P is true, A must believe A knows that P is true and want to inform S that P (the preconditions), and must intend to get S to know that P is true (the effect), which is done by constructing a plan that will achieve S's recognition of this intention (i.e. that $B_S(W_A(KNOW(S,P)))$). A then must depend on S to bring about the effect: S must *decide* to believe what A said. This is made explicit by introducing an admittedly simplistic DECIDE__TO__BELIEVE act:

```
DECIDE__TO__BELIEVE(agent, other, P)
prec: B(agent,W(other,KNOW(agent,P)))
effect: KNOW(agent,P)
```

Thus A can INFORM S of P by achieving $B_S W_A(KNOW(S,P))$ followed by DECIDE__TO__BELIEVE(S,A,P).

In many cases, agents reason about INFORM acts to be performed (by others or by themselves) where the information for the propositional content is not known at the time of plan construction. For example, A may plan for S to inform A whether P is true. A cannot plan for S to perform INFORM(S,A,P) since this assumes the truth of P. We get around this difficulty by defining INFORMIF, another view of the INFORM act.

```
INFORMIF(speaker, hearer, P)
prec: KNOWIF(speaker,P) ^
      W(speaker,INFORMIF(speaker,hearer,P))
effect: KNOWIF(hearer,P)
body: B(hearer,W(speaker,KNOWIF(hearer,P)))
```

Similarly, it must be possible for A to plan for S to tell A the referent of a description, without A knowing the referent. This is the role of the INFORMREF act.

```
INFORMREF(speaker, hearer, D(x))
prec: KNOWREF(speaker,D(x)) ^
      W(speaker,INFORMREF(speaker,
                          hearer,D(x)))
effect: KNOWREF(hearer,D(x))
body: B(hearer,W(speaker,KNOWREF(
                          hearer,D(x))))
```

Request is defined as:

```
REQUEST(speaker, hearer, action)
constraint: hearer is agent of action
prec: W(speaker,action(hearer))
effect: W(hearer,action(hearer))
body: B(hearer,W(speaker,action(hearer)))
```

The intention of a request is to get the hearer to want to do the action, and this is accomplished by getting the hearer to believe that the speaker wants the hearer to do the action and then depending on the hearer to decide to do it.⁵ To explicitly represent this decision process, a CAUSE__TO__WANT act defined along the lines of the DECIDE__TO__BELIEVE act above is necessary.

```
CAUSE__TO__WANT(agent, other, P)
prec: B(other,B(agent,W(agent,P)))
effect: W(other,P)
```

As examples of the use of speech acts, "Tell me whether the train is here" and "Is the train here?", intended literally, are both REQUESTs by A that S INFORMIF the train is here. "When does the train arrive?", intended literally, is a REQUEST by A that H INFORMREF of the departure time of the train.

Finally we define the two *surface* level acts: S.INFORM produces indicative mood utterances, and S.REQUEST produces imperative utterances, or interrogative utterances, if the requested act is an INFORM. These acts have no preconditions, and serve solely to signal the immediate intention of the speaker, the starting point for all the hearer's inferencing.

```
S.INFORM(speaker, hearer, P)
effect: B(hearer,W(speaker,KNOW(hearer,P)))

S.REQUEST(speaker, hearer, action)
effect: B(hearer,W(speaker,action(hearer)))
```

The effects of S.INFORM match the body of the INFORM act, reflecting the fact that it is a standard way of executing an INFORM. It is important, however, that S.INFORM is only one way of executing an INFORM. The same relationship holds between the S.REQUEST and REQUEST actions.

4.2. Recognizing Illocutionary Force

Given the speech act definitions of section 4.1, we say that *A performed an illocutionary act IA by uttering x to S* if A intends that S should recognize (and be able to recognize) that

- (1) x is an instance of a surface act SA, and
- (2) A intended S to infer (using the PI rules and associated heuristics) from A having performed SA that A wants to achieve the effects of IA.

This definition allows more than one illocutionary act to be performed by a single surface act. In this section we show how the hearer of an utterance can recognize the speaker's intention(s) indicated by a speech act, especially when these intentions are communicated indirectly.

⁵ See Cohen and Perrault [1979] for a discussion of why Searle's preparatory conditions "Speaker believes Hearer can do the action" need not be part of the preconditions on REQUEST.

All inferencing by S of A's plans starts from S's recognition that A intended to perform one of the surface acts, and that A in fact wanted to do the act. All inference chains will be shown as starting from a formula of the form $B_S W_A(A \text{ do the surface act})$. The object of the inferencing is to find what illocutionary level act(s) A intended to perform. The action-effect rule applied to the starting formula yields one of the form $B_S W_A(B_S W_A(P))$, i.e. S believes that A wants S to recognize A's intention that P. The inferencing process searches for plausible formulas of the form $B_S W_A(IA(A))$ where IA is an illocutionary level act.

-
- (1) S.REQUEST(A,S,PASS(S,A,SALT))
 PI.AE (2) $B_S W_A(PASS(S,A,SALT))$
 PI.BA (3) REQUEST(A,S,PASS(S,A,SALT))

Example 1. "Pass the salt."

Example 1 shows a direct request to pass the salt, where the surface request maps directly into the intended request interpretation.⁶ The actions relevant to the examples given here are:

- PASS(agent, beneficiary, object)
 prec: HAVE(agent, object)
 effect: HAVE(beneficiary, object)
- REACH(agent, object)
 prec: NEAR(agent, object)
 effect: HAVE(agent, object)

Let us also assume that S presently has the salt, i.e. HAVE(S,SALT) is true, and mutually believed by S and A.

The rating heuristics for the complex rules EI.1 to EI.3 are the same as for the PI rules but each heuristic may be applicable several times at different levels. For example, consider the frequently recurring inference chain:

- (1) $B_S W_A(ACT(S))$
 PI.BA (2) REQUEST(A,S,ACT(S))
 PI.AE (3) $W_S(ACT(S))$
 PI.PA (4) ACT(S)
 PI.AE (5) effects of ACT(S)

It shows the line of inference from the point where S recognizes that A requested S to do ACT (at step (2)) to the point where the effects of the requested action are inferred as part of A's plan. Of interest here is the evaluation of the plausibility of step (3). Two heuristics are applicable. The proposition " $W_S(ACT(S))$ " is

⁶ To improve readability of inference chains in the examples, we drop the prefix $B_S W_A$ from all propositions. The formula on line (n) follows from the one on line (n-1) by the rule at the beginning of line (n). Applications of EI.1 will be labelled "rule"/EI.1, where "rule" is a PI rule embedded by EI.1. Similarly, applications of EI.2 and EI.3 will be labelled "rule"/EI.2 and "rule"/EI.3, where "rule" is a PC rule name.

evaluated with respect to what S believes A believes. (Remember that $B_S W_A$ should appear as a prefix to all propositions in inference chains.) If $B_S B_A W_S(ACT(S))$ is true, the request interpretation is considered unlikely, by the effect-based heuristic. In addition, the preconditions of ACT(S) are considered with respect to what S believes A believes S believes. This step will only be reasonable if S can do the action, by a precondition-based heuristic.

To make more explicit the distinction between inferences in $B_S W_A$ and inferences in $B_S W_A B_S W_A$, let us consider two inference chains that demonstrate two interpretations of the utterance "Do you know the secret?". Lines 1-3 of Example 2 show the chain which leads S to believe that A asked a (literal) yes/no question; lines 1-6 of Example 3 show the interpretation as a request to S to inform A of the secret. Notice that in both interpretations S may be led to believe that A wants to know the secret. In the literal case, S infers A's goal from the literal interpretation, and may tell the secret simply by being helpful (lines 4-9). In the indirect case, S recognizes A's *intention* that S inform A of the secret (lines 1-6). Telling the secret is then conforming to A's intentions (lines 7-9).

There is in fact a third interpretation of this sentence. If A and S both know that A already knows the secret, then the utterance could be intended as

"If you don't know the secret, I will tell it to you."

This requires recognizing a conditional action and is beyond our present abilities.

4.3. The Level of Embedding Heuristic

Two sets of PI rules are applicable to formulas of the form $B_S W_A B_S W_A(P)$: the simple rules PI.1 to PI.6 operating "within" the prefix $B_S W_A$, and the rules generated by EI.1 and EI.3 which allow the simple rules to apply within the prefix $B_S W_A B_S W_A$. To reflect the underlying assumption in our model that intention will always be attributed if possible, the inferences at the most deeply nested level should be preferred.

Of course, if the inferences at the nested level lead to unlikely plans, the inferences at the "shallow" levels may be applied. In particular, if there are multiple mutually exclusive inferences at the nested level, then the "shallow" inferences will be preferred. This reflects the fact that the nested inferences model what the speaker intends the hearer to infer. If there are many inferences possible at the nested level, the speaker would not be able to ensure that the hearer would perform the correct (i.e., the intended) one.

- | | |
|-------|--|
| | (1) S.REQUEST(A,S,INFORMIF(S,A,KNOWREF(S,SECRET(x)))) |
| PI.AE | (2) B _S W _A (INFORMIF(S,A,KNOWREF(S,SECRET(x)))) |
| PI.BA | (3) REQUEST(A,S,INFORMIF(S,A,KNOWREF(S,SECRET(x)))) |
| PI.AE | (4) W _S (INFORMIF(S,A,KNOWREF(S,SECRET(x)))) |
| PI.W | (5) INFORMIF(S,A,KNOWREF(S,SECRET(x))) |
| PI.AE | (6) KNOWIF(A,KNOWREF(S,SECRET(x))) |
| PI.KP | (7) KNOWREF(S,SECRET(x)) |
| PI.PA | (8) INFORMREF(S,A,SECRET(x)) |
| PI.AE | (9) KNOWREF(A,SECRET(x)) |

Example 2. "Do you know the secret?" (yes/no question)

- | | |
|------------|--|
| | (1) S.REQUEST(A,S,INFORMIF(S,A,KNOWREF(S,SECRET(x)))) |
| PI.AE | (2) B _S W _A (INFORMIF(S,A,KNOWREF(S,SECRET(x)))) |
| PI.AE/EI.1 | (3) B _S W _A (KNOWIF(A,KNOWREF(S,SECRET(x)))) |
| PI.KP/EI.1 | (4) B _S W _A (KNOWREF(S,SECRET(x))) |
| PI.PA/EI.1 | (5) B _S W _A (INFORMREF(S,A,SECRET(x))) |
| PI.BA | (6) REQUEST(A,S,INFORMREF(S,A,SECRET(x))) |
| PI.AE | (7) W _S (INFORMREF(S,A,SECRET(x))) |
| PI.W | (8) INFORMREF(S,A,SECRET(x)) |
| PI.AE | (9) KNOWREF(A,SECRET(x)) |

Example 3. "Do you know the secret?" (indirect request)

- | | |
|------------|--|
| | (1) S.INFORM(A,S,W _A (PASS(S,A,SALT))) |
| PI.AE | (2) B _S W _A (B _S (W _A (PASS(S,A,SALT)))) |
| PI.PA/EI.1 | (3) B _S W _A (CAUSE_TO_WANT(A,S,PASS(S,A,SALT))) |
| PI.AE/EI.1 | (4) B _S W _A (W _S (PASS(S,A,SALT))) |
| PI.W/EI.1 | (5) B _S W _A (PASS(S,A,SALT)) |
| PI.BA | (6) REQUEST(A,S,PASS(S,A,SALT)) |

Example 4. "I want you to pass the salt."

4.4. More Indirect Requests

Example 4 shows the interpretation of "I want you to pass the salt" as a request. Taking the utterance literally, S infers that A wants him to know that A wants him to pass the salt. This yields proposition (2) which leads through the next three inferences to the intention that would be recognized from a request act, i.e. that A wants S to pass the salt (5). Notice that an application of the body-action rule to step (2) yields:

INFORM(A, S, W_A(PASS(S, A, SALT))),

for, in fact, the speaker may be performing both speech acts. The level of inferencing heuristic favours the indirect form.

The key step in Example 5 is the application of the know-positive rule from line (3) to line (4). Since, given the context, S assumes that A knows whether S has the salt, the literal interpretation (from (2)) would not produce a reasonable goal for A. This supports

the nested know-positive inference, and attributes further intention to the speaker (4). Once this is done, it is easy to infer that A wants S to pass him the salt (5), hence the request interpretation.

"Can you pass the salt?" and "Do you want to pass the salt?" are treated similarly, for they inquire about the preconditions on PASS(S, A, SALT).

Example 6 begins like Example 5, leading to the inference that A wants S to be able to reach the salt (4).⁷ Since being able to reach the salt is a precondition to reaching the salt (5), which then enables passing the salt (6), S can infer that he is being requested to pass the salt. "Is the salt near you?" can be treated in the same way, as being near the salt is a precondition on reaching the salt.

⁷ Let CANDO(S,ACT) be true if S believes the preconditions of ACT are true.

- (1) S.REQUEST(A,S,INFORMIF(S,A,HAVE(S,SALT)))
 PI.AE (2) B_SW_A(INFORMIF(S,A,HAVE(S,SALT)))
 PI.AE/EI.1 (3) B_SW_A(KNOWIF(A,HAVE(S,SALT)))
 PI.KP/EI.1 (4) B_SW_A(HAVE(S,SALT))
 PI.PA/EI.1 (5) B_SW_A(PASS(S,A,SALT))
 PI.BA (6) REQUEST(A,S,PASS(S,A,SALT))

Example 5. "Do you have the salt?"

- (1) S.REQUEST(A,S,INFORMIF(S,A,CANDO(S,REACH(S,SALT)))
 PI.AE (2) B_SW_A(INFORMIF(S,A,CANDO(S,REACH(S,SALT)))
 PI.AE/EI.1 (3) B_SW_A(KNOWIF(A,CANDO(S,REACH(S,SALT)))
 PI.KP/EI.1 (4) B_SW_A(CANDO(S,REACH(S,SALT)))
 PI.PA/EI.1 (5) B_SW_A(REACH(S,SALT))
 PI.AE/EI.1 (6) B_SW_A(HAVE(S,SALT))
 PI.PA/EI.1 (7) B_SW_A(PASS(S,A,SALT))
 PI.BA (8) REQUEST(A,S,PASS(S,A,SALT))

Example 6. "Can you reach the salt?"

- (1) S.INFORM(A,S,W_A(HAVE(A,SALT)))
 PI.AE (2) B_SW_A(B_S(W_A(HAVE(A,SALT))))
 PI.PA/EI.1 (3) B_SW_A(CAUSE_TO_WANT(A,S,HAVE(S,A,SALT)))
 PI.AE/EI.1 (4) B_SW_A(W_S(HAVE(A,SALT)))
 PI.AE/EI.3 (5) B_SW_A(W_S(PASS(S,A,SALT)))
 PI.W/EI.1 (6) B_SW_A(PASS(S,A,SALT))
 PI.BA (7) REQUEST(A,S,PASS(S,A,SALT))

Example 7. "I want the salt." (= "I want to have the salt.")

Example 7 includes in the step from (3) to (4), an application, through EI.3, of the effect-action rule. A informs S of A's goal of having the salt (2) and then depends on S's planning on that goal to infer the PASS action. Because the action is the "obvious" way of achieving the goal, S believes that A intended him to infer it.

Since questions are treated as requests to inform, most of them are handled in a similar manner to the requests above. 4.4a-h can all be understood as questions about the departure time of some train.

- (4.4a) When does the train leave?
 (4.4b) I want you to tell me when ...
 (4.4c) I want to know when ...
 (4.4d) Tell me when ...
 (4.4e) Can you tell me when ...
 (4.4f) Do you know when ...
 (4.4g) Do you want to tell me when ...
 (4.4h) Will you tell me when ...

4.5. An Example of an Indirect INFORM

An interesting example of an indirect INFORM is 4.5a for it is very similar to 4.5b-c which both seem to only be requests. The interpretation of 4.5a as an indirect INFORM follows from the fact that inference chains which would make it a REQUEST are all inhibited by the heuristics.

- (4.5a) Do you know that the RAPIDO is late?
 (4.5b) Do you believe that the RAPIDO is late?
 (4.5c) Do you know whether the RAPIDO is late?

In Example 8, the possible body-action inference from (2) to

REQUEST(A,S,INFORMIF(S,A,KNOW(S,P)))

is downgraded because the embedded inference to (3) is possible. The interesting case is the embedded know-negative inference which is also possible from (3). It implies that B_SW_A(~KNOW(S,P)), or equivalently

(4.5d) B_SW_A(P ∧ ~B_S(P))

	(1) S.REQUEST(A,S,INFORMIF(S,A,KNOW(S,P)))
PI.AE	(2) $B_S W_A$ (INFORMIF(S,A,KNOW(S,P)))
PI.AE/EI.1	(3) $B_S W_A$ (KNOWIF(A,KNOW(S,P)))
PI.KP/EI.1	(4) $B_S W_A$ (KNOW(S,P))
PI.BA	(5) INFORM(A,S,P)

Example 8. "Do you know that P?"

But such a goal is highly unlikely. A is attempting to achieve the goal $\sim B_S(P)$ by having S recognize that A wants P to be true! As a result, no speech act interpretation is possible from this step. For instance, the bodies of the acts INFORM(A, S, P) and INFORM(A, S, $\sim P$) are $B_S W_A(P \wedge B_S(P))$, and $B_S W_A(\sim P \wedge B_S(\sim P))$, respectively. Both of these are contradicted by part of 4.5d. Thus the know-negative possibility can be eliminated. This allows the know-positive inference to be recognized as intended, and hence leads to the indirect interpretation as an INFORM(A, S, P).

4.5b has only a literal interpretation since both the know-positive and know-negative rules are applicable at the nested level; without a reason to favour either, the literal

REQUEST(A,S,INFORMIF(S,A, $B_S(P)$))

is preferred. The interpretations of 4.5c are similar to those of Examples 2 and 3.

4.6. Using Knowledge of Deduction

All the examples of indirect speech acts so far have been explained in terms of rules PI.1-PI.6, and complex inference rules derived from them. In this section, we give one more example relying on somewhat more specific rules. A full investigation of how many such specific rules are necessary to account for common forms of indirect REQUESTs and INFORMs remains to be done.

This example shows how a completely non-standard form can be intended indirectly. Suppose that A tells S

(4.6a) "John asked me to ask you to leave"

This has at least three possible interpretations:

(4.6b) A is asking S to leave, and giving a reason.

(4.6c) A wants to simply report the fact to S that John did the action of asking S to leave.

(4.6d) A wants to inform S that John wants him to leave.

Interpretations c and d can hold even if S decides that A actually does want him to leave. However, in these cases, he would not say that A intended to communicate the intent that he leave, i.e. he would not say the utterance was a REQUEST.

Both interpretations rely on axioms ACT.1 and ACT.2 (of section 3.2) which state that if some agent A believes that agent S executed some action ACT, then A may believe that the preconditions of ACT obtained before, and the effects of ACT obtained after, the execution of ACT.

They also require a new PC/PI rule: if A wants S to believe some proposition P, then A may get S to believe some proposition Q, as long as A believes that S believes that Q implies P.

(PC.I) $W_A(B_S(P)) = c \Rightarrow W_A(B_S(Q))$,
if $B_A B_S(Q \Rightarrow P)$.

(PI.I) $B_S W_A(B_S(Q)) = i \Rightarrow B_S W_A(B_S(P))$,
if $B_S B_A B_S(Q \Rightarrow P)$.

In Example 9, S recognizes that A asked him to leave. The interpretation depends on S concluding that John performed his REQUEST successfully (through PI.I and ACT.2), and hence that A wants to request S to leave. It is then an easy step to infer that A wants S to leave, which leads to the request interpretation. Interpretation (c), a simple report of some previous action, follows from (2) by PI.BA.

In Example 10, S recognizes that A intended to tell him that John wants him to leave. This depends on the fact that S concludes that John wanted to perform the REQUEST that A reported. Most of the needed inferences call for the use of EI.1 to embed simple inference rules twice. Note that an INFORM act could have been inferred at each of the four previous steps; for example, from (5) the body inference would produce

INFORM(A, S, W_J (REQUEST(A, S, LEAVE(S)))).

But the inferences at the " $B_S W_A B_S W_J$ " level were so direct that they were continued.

5. Gordon and Lakoff Revisited

The examples of the previous section show how our plan inference rules account for the indirect interpretations of the requests which GL's postulates were designed for, as well as several others. Our approach differs from GL's in that an utterance may carry both a literal and an indirect interpretation, and of course in that its inference rules are language independent.

	(1) S.INFORM(A,S,REQUEST(J,A,REQUEST(A,S,LEAVE(S))))
PI.AE	(2) B _S W _A (B _S (REQUEST(J,A,REQUEST(A,S,LEAVE(S))))))
PI.I/EI.1	(3) B _S W _A (B _S W _A (REQUEST(A,S,LEAVE(S))))
(PI.AE/EI.1)/EI.1	(4) B _S W _A (B _S W _A (W _S (LEAVE(S))))
(PI.W/EI.1)/EI.1	(5) B _S W _A (B _S W _A (LEAVE(S)))
(PI.BA/EI.1)	(6) B _S W _A (REQUEST(LEAVE(S)))
PI.AE/EI.1	(7) B _S W _A (W _S (LEAVE(S)))
PI.W/EI.1	(8) B _S W _A (LEAVE(S))
PI.BA	(9) REQUEST(A,S,LEAVE(S))

Example 9. "John asked me to ask you to leave." (Interpretation b)

	(1) S.INFORM(A,S,REQUEST(J,A,REQUEST(A,S,LEAVE(S))))
PI.AE	(2) B _S W _A (B _S (REQUEST(J,A,REQUEST(A,S,LEAVE(S))))))
PI.I/EI.1	(3) B _S W _A (B _S W _J (REQUEST(J,A,REQUEST(A,S,LEAVE(S))))))
(PI.AE/EI.1)/EI.1	(4) B _S W _A (B _S W _J (W _A (REQUEST(A,S,LEAVE(S))))))
(PI.W/EI.1)/EI.1	(5) B _S W _A (B _S W _J (REQUEST(A,S,LEAVE(S))))
(PI.AE/EI.1)/EI.1	(6) B _S W _A (B _S W _J (W _S (LEAVE(S))))
(PI.W/EI.1)/EI.1	(7) B _S W _A (B _S W _J (LEAVE(S)))
(PI.BA/EI.1)/EI.1	(8) INFORM(A,S,W _J (LEAVE(S)))

Example 10. "John asked me to ask you to leave." (Interpretation d)

However, in some ways both solutions are too strong.

Consider, for example, the following:

- (5.a) Can you reach the salt?
- (5.b) Are you able to reach the salt?
- (5.c) I hereby ask you to tell me whether you are able to reach the salt.

Although 5.a-c are all literally questions about the hearer's ability, only 5.a normally conveys a request.

Sadock [1974] suggests that forms such as 5.a differ from 5.b in that the former is an idiom which is directly a request while 5.b is primarily a yes/no question. However, as Brown [1980] points out, this fails to account for responses to 5.a which follow from its literal form. One can answer "Yes" to 5.a and then go on to pass the salt.

Brown proposes what she calls "frozen ISA forms" which directly relate surface form and indirect illocutionary force, bypassing the literal force. Frozen forms differ from normal rules mapping illocutionary forces to illocutionary forces in that they point to the relevant normal rule which provides the information necessary to the generation of responses to the surface forms.

The speaker of 5.b or 5.c may in fact want the hearer to reach the salt, as does the speaker of 5.a, but he does not want his intention to be recognized by the hearer. Thus it appears that from the hearer's point of

view the chain of inferences at the intended level should get turned off, soon after the recognition of the literal act. It seems that in this case (Example 6 of section 4.4) the plausibility of the inferences after step 3 should be strongly decreased. Unfortunately it is not obvious that this can be done without making the rating heuristics sensitive to syntax.

The indirect interpretation can also be downgraded in the presence of stronger expectations. If a speaker entered a room full of aspiring candidates for employment and said: "I want to know how many people here can write a sort/merge program" and then turning to each individually asked "Can you write a sort/merge?" the question would not be intended as a request to write a program, and would not be recognized as such by a PI algorithm which rated highly an illocutionary act which fits well in an expectation.

In several of the earlier examples of questions intended as indirect requests, the literal interpretation is blocked because it leads to acts whose effects were true before the utterance. The literal interpretation of 5.d gets blocked because the reminding gets done as part of the understanding of the literal act. Thus only an indirect interpretation is possible.

- (5.d) May I remind you to take out the garbage?

Sadock [1970] points out that some co-occurrence rules depend on conveyed rather than literal illocutionary force. The morpheme *please* can occur initially only in sentences which convey a request.

(5.e) Please, can you close the window?

(5.f) Please, it's cold in here.

(5.g) *Please, do you know that
James is in town?

But it can occur in final position only in sentences which both convey a request and are literally requests:

(5.h) Can you close the window, please?

(5.i) *It's cold in here, please.

(5.j) *Do you know that John is in town, please?

These remain problematic for Brown and for us.

6. Conclusion

We have given evidence in this paper for an account of indirect speech acts based on rationality (plan construction), imputing rationality to others (plan inference), surface speech act definitions relating form to "literal" intentions, and illocutionary acts allowing a variety of realizing forms for the same intentions.

The reader may object that we are suggesting a complex solution to what appears to be a simple problem. It is important to distinguish here the general explanation of indirect speech acts (which is presented here partly through an algorithm) from the implementation of such an algorithm in a practical natural language understanding system. We claim that the elements necessary for a theoretically satisfying account of indirect speech acts are independently motivated. It is almost certain that a computationally efficient solution to the indirect speech act problem would short-cut many of the inference chains suggested here, although we doubt that all searching can be eliminated in the case of the less standard forms such as 4.6a. The implementation in Brachman et al [1980] does just that. However, the more fundamental account is necessary to evaluate the correctness of the implementations.

Many problems remain. Other syntactic forms that have significance with respect to illocutionary force determination should be considered. For example, tag questions such as

"John is coming to the party tonight, isn't he?"
have not been analysed here (but see Brown [1980]). Furthermore, no "why" or "how" questions have been examined.

Besides the incorporation of more syntactic information, another critical area that needs work concerns the control of inferencing. To allow the use of specialized inferences, a capability that is obviously required by the general theory, much research needs to

be done outlining methods of selecting and restricting such inferences.

This paper has concentrated on recognition. Allen [1979] shows how the construction algorithms would have to be modified to allow the generation of surface acts, including indirect forms. McDonald [1980] discusses the planning of low-level syntactic form.

According to the definition of INFORM of section 4.1, any utterance that causes S to infer that A has a plan to achieve $KNOW(S,P)$ by achieving $B_S W_A(KNOW(S,P))$ is considered by S to be an INFORM. Strawson [1964] argues that one level of recognition of intention is not sufficient for the definition of a speech act. Schiffer [1972] gives a series of counterexamples to show that no finite number of conditions of the form $B_S W_A(B_S W_A(\dots(KNOW(S,P))))$ is sufficient either. The solution he proposes is that the recognition of intention must be *mutually believed* between the speaker and the hearer. Cohen and Levesque [1980] and Allen [forthcoming] show how the speech act definitions given here can be extended in this direction.

We have only considered acts to request and inform because many of their interesting properties can be based on belief and want. At least primitive accounts of the logics of these propositional attitudes are available. Clearly there is room for much work here. Extending the analysis to other speech acts, such as promises, will require a study of other underlying logics such as that of obligation.

There also remain many problems with the formalization of actions. We believe this work shows that the concepts of preconditions, effects, and action bodies are fruitful in discussing plan recognition. The operator definitions for speech acts used here are intended to facilitate the statement of the plan construction and inference rules. However, their expressive power is insufficient to handle complex actions involving sequencing, conditionals, disjunctions, iterations, parallelism, discontinuity, and a fortiori requests and promises to do such acts. They are also inadequate, as Moore [1979] points out, to express what the agent of an action knows (and does not know) after the success or failure of an act. Moore's logic of action includes sequencing, conditionals, and iterations, and is being applied to speech acts by Appelt [1980]. Much remains to be done to extend it to parallel and discontinuous actions typical of multiple agent situations.

These difficulties notwithstanding, we hope that we have helped show that the interaction of logic, philosophy of language, linguistics and artificial intelligence is productive and that the whole will shed light on each of the parts.

References

- Allen, J.F. *A Plan-Based Approach to Speech Act Recognition*. Ph. D. Thesis, Dept. of Computer Science, University of Toronto, 1979.
- Allen, J.F. Recognizing Intentions from Natural Language Utterances, forthcoming.
- Allen, J.F. and Perrault, C.R., Analyzing Intention in Utterances, *Artificial Intelligence*, 15, 3, pp. 143-178, 1980.
- Appelt, D.E., Problem Solving Applied to Language Generation, *Proc. of 18th Annual Meeting of Assoc. for Comp. Ling.*, pp. 59-63, 1980.
- Austin, J.L. *How To Do Things With Words*. New York: Oxford University Press, 1962.
- Ayer, A.J. *The Problem of Knowledge*. London: Penguin Books, 1956.
- Brachman, R.J., Bobrow, R.J., Cohen, P.R., Klovstad, J.W., Webber, B.L., and Woods, W.A., Research in Natural Language Understanding, Report 4274, Bolt, Beranek, and Newman, Cambridge, Mass., 1980.
- Brown, G.P. Characterizing Indirect Speech Acts :A Computational Approach, *Am. J. of Comp. Ling.*, this issue, pp. 150-166, 1980.
- Bruce, B.C. Belief Systems and Language Understanding, Technical Report 2973, Bolt Beranek and Newman, Cambridge Mass, 1975.
- Cohen, P.R. *On Knowing What to Say: Planning Speech Acts*, Technical Report no. 118, Dept. of Computer Science, University of Toronto January, 1978.
- Cohen, P.R., and Levesque, H.J., Speech Acts and the Recognition of Shared Plans, in *Proc. of 3rd Conf. of Can. Soc. for Comp. Studies of Intelligence*, Victoria, B.C., pp. 263-271, 1980.
- Cohen, P.R. and Perrault, C.R. Elements of a Plan Based Theory of Speech Acts, *Cognitive Science*, 3, 3, pp. 177-212, 1979.
- 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, pp. 189-205, 1971.
- Genesereth, M.R. *Automated Consultation for Complex Computer Systems*, Ph.D. Thesis, Harvard University, 1978.
- Gettier, E.L. Is Justified True Belief Knowledge?. *Analysis* 23, pp. 121-123, 1963. Reprinted in M.D. Roth and L. Galis (eds.) *Knowing*, New York: Random House, 1970.
- Gordon, D. and Lakoff, G. Conversational Postulates. in Cole and Morgan (eds), 1975.
- Green, G. M. How to Get People to Do Things with Words: the Whimperative Question, in Cole and Morgan (eds.), 1975.
- Grice, H.P. Meaning. *Phil. Rev.* 66, pp. 377-388, 1957.
- Grice, H.P., Utterer's Meaning, Speaker Meaning, and Word Meaning, *Found. of Language*, 4, pp. 225-242, 1968.
- Hintikka, J. *Knowledge and Belief*, Ithaca, New York: Cornell University Press, 1962.
- McDonald, D.D., *Natural Language Production as a Process of Decision Making under Constraint*, Ph.D. Dissertation, Dept. of E.E. and Comp. Sc., MIT, 1980.
- Moore, R.C. *Reasoning about Knowledge and Action*, Ph.D. Thesis, Massachusetts Institute of Technology, 1979.
- Morgan, J.L. Conversational Postulates Revisited, *Language* 53, 2, 1977.
- Morgan, J.L. Towards a Rational Model of Discourse Comprehension, in *Proc. of Second Conf. Theoretical Issues in Natural Language Processing*, Champaign-Urbana, 1978.
- Morgan, J.L. Two Types of Convention in Indirect Speech Acts, in P. Cole (ed.) *Syntax and Semantics (vol. 9): Pragmatics*, New York: Academic Press, 1978.
- Perrault, C.R., Allen, J.F. and Cohen, P.R. Speech Acts as a Basis for Understanding Dialogue Coherence, in *Proc. of Second Conf. Theoretical Issues in Natural Language Processing*, Champaign-Urbana, 1978.
- Perrault, C.R. and Cohen, P.R. It's for Your Own Good: a Note on Inaccurate Reference, to appear in *Elements of Discourse Understanding*, A.K. Joshi, I.A. Sag, and B.L. Webber (eds.), New York: Cambridge University Press.
- Quine, W.V.O. Quantifiers and Propositional Attitudes. *Journal of Philosophy* 53, pp177-187, 1956. Reprinted in L. Linsky (ed.) *Reference and Modality*, Oxford: Clarendon Press, 1971.
- Russell, B. *Introduction to Mathematical Philosophy*. London: George Allen and Unwin, 1919.
- Sacerdoti, E.D. Planning in a Hierarchy of Abstraction Spaces. *Proc. of Third Int. Joint Conf. on Artificial Intelligence*, Stanford CA, 1973.
- Sacerdoti, E.D. The Nonlinear Nature of Plans. *Proc. of Fourth Int. Joint Conf. on Artificial Intelligence*, Tbilisi USSR, 1975.
- Sadock, J.M. Whimperatives, in J.M. Sadock and A. Vanek *Studies Presented to R.B. Lees by his Students*, pp. 223-38, Edmonton, 1970.
- Sadock, J. M. *Towards a Linguistic Theory of Speech Acts*, New York: Academic Press, 1974.
- Schiffer, S.R. *Meaning*. London: Oxford University Press, 1972.
- Schmidt, C.F., Shridharan N.S. and Goodson J.L., The Plan Recognition Problem: An Intersection of Artificial Intelligence and Psychology, *Artificial Intelligence* 11, pp. 45-83, 1978.
- Searle, J.R. *Speech Acts*, New York: Cambridge University Press, 1969.
- Searle, J.R. Indirect Speech Acts. in Cole and Morgan (eds), 1975.
- Stampe, D.W. Meaning and Truth in the Theory of Speech Acts. in Cole and Morgan (eds), 1975.
- Strawson, P.F. Intention and Convention in Speech Acts. *Phil. Rev.*, 73,4, pp. 439-460, 1964.

C. Raymond Perrault is Associate Professor in the Department of Computer Science at the University of Toronto. He received the Ph.D. degree in computer and communication sciences from the University of Michigan in 1975.

James F. Allen is Assistant Professor in the Department of Computer Science at the University of Rochester. He received the Ph.D. degree in computer science from the University of Toronto in 1979.