

Interpreting and Generating Indirect Answers

Nancy Green*
University of North Carolina at
Greensboro

Sandra Carberry†
University of Delaware

This paper presents an implemented computational model for interpreting and generating indirect answers to yes-no questions in English. Interpretation and generation are treated, respectively, as recognition of and construction of a responder's discourse plan for a full answer. An indirect answer is the result of the responder providing only part of the planned response, but intending for his discourse plan to be recognized by the questioner. Discourse plan construction and recognition make use of shared knowledge of discourse strategies, represented in the model by discourse plan operators. In the operators, coherence relations are used to characterize types of information that may accompany each type of answer. Recognizing a mutually plausible coherence relation obtaining between the actual response and a possible direct answer plays an important role in recognizing the responder's discourse plan. During generation, stimulus conditions model a speaker's motivation for selecting a satellite. Also during generation, the speaker uses his own interpretation capability to determine what parts of the plan are inferable by the hearer and thus do not need to be explicitly given. The model provides wider coverage than previous computational models for generating and interpreting indirect answers and extends the plan-based theory of implicature in several ways.

1. Introduction

In the following example,¹ Q asks a question in (1)i and R provides the requested information in (1)iii, although not explicitly giving (1)ii. (In this paper, we use square brackets as in (1)ii to indicate information which, in our judgment, the speaker intended to convey but did not explicitly state. For consistency, we refer to the questioner and responder as Q and R, respectively. For readability, we have standardized punctuation and capitalization and have omitted prosodic information from sources since it is not used in our model.)

- (1) i. Q: Actually you'll probably get a car won't you as soon as you get there?
ii. R: [No.]
iii. I can't drive.

Interpreting such responses, which we refer to as **indirect answers**, requires the hearer to derive a conversational implicature (Grice 1975). For example, the inference that R

* Department of Mathematical Sciences, Greensboro, NC 27412-5001

† Department of Computer and Information Sciences, Newark, DE 19716

1 Based on an example on page 220 in Stenström (1984). The reader may assume that any unattributed examples in the paper are constructed.

will not get a car on arrival, although licensed by R's use of (1)iii in some discourse contexts, is not a semantic consequence of the proposition that R cannot drive.

According to one study of spoken English (Stenström 1984) (described in Section 2), 13% of responses to certain *yes-no* questions were indirect answers. Thus, a robust dialogue system should be able to interpret indirect answers. Furthermore, there are good reasons for generating an indirect answer instead of just a *yes* or *no* answer. First, an indirect answer may be considered more polite than a direct answer (Brown and Levinson 1978). For example, in (1)i, Q has indicated (by the manner in which Q expressed the question) that Q believes it likely that R will get a car. By avoiding explicit disagreement with this belief, the response in (1)iii would be considered more polite than a direct answer of (1)ii. Second, an indirect answer may be more efficient than a direct answer. For example, even if (1)ii is given, including (1)iii in R's response contributes to efficiency by forestalling and answering a possible follow-up of *well, why not?* from Q, which can be anticipated since the form of Q's question suggests that Q may be surprised by a negative answer. Third, an indirect answer may be used to avoid misleading Q (Hirschberg 1985), as illustrated in (2).²

- (2) i. Q: Have you gotten the letters yet?
 ii. R: I've gotten the letter from X.

This example illustrates a case in which, provided that R had gotten some but not all of the letters in question, just *yes* would be untruthful and just *no* would be misleading (since Q might conclude from the latter that R had gotten *none* of them).

We have developed a computational model, implemented in Common LISP, for interpreting and generating indirect answers to *yes-no* questions in English (Green 1994). By a *yes-no* question we mean one or more utterances used as a request by Q that R convey R's evaluation of the truth of a proposition *p*. Consisting of one or more utterances, an indirect answer is used to convey, yet does not semantically entail, R's evaluation of the truth of *p*, i.e., that *p* is true, that *p* is false, that *p* might be true, that *p* might be false, or that *p* is partially true. In contrast, a **direct answer** entails R's evaluation of the truth of *p*. The model presupposes that Q and R mutually believe that Q's question has been understood by R as intended by Q, that Q's question is appropriate, and that R can provide one of the above answers. Furthermore, it is assumed that Q and R are engaged in a cooperative and polite task-oriented dialogue.³ The model is based upon examples of uses of direct and indirect answers found in transcripts of two-person telephone conversations between travel agents and their clients (SRI 1992), examples given in previous studies (Brown and Levinson 1978; Hirschberg 1985; Kiefer 1980; Levinson 1983; Stenström 1984) and constructed examples reflecting our judgments.

To give an overview of the model, generation and interpretation are treated, respectively, as construction of and recognition of the responder's **discourse plan specification** for a full answer. In general, a discourse plan specification (for the sake of brevity, hereafter referred to as discourse plan) explicitly relates a speaker's beliefs and discourse goals to his program of communicative actions (Pollack 1990). Discourse plan construction and recognition make use of the beliefs that are presumed

² (2) is Hirschberg's example (59).

³ We assume that it is worthwhile to model politeness-motivated language behavior for both generation and interpretation. For example in generation, it would seem to be a desirable trait for a software agent that interacts with humans. In interpretation, it would contribute to the robustness of the interpreter.

to be shared by the participants, as well as shared knowledge of discourse strategies, represented in the model by a set of **discourse plan operators** encoding generic programs of communicative actions for conveying **full answers**. A full answer consists of a direct answer, which we refer to as the **nucleus**, and “extra” appropriate information, which we refer to as the **satellite(s)**.⁴ In the operators, coherence relations are used to characterize types of satellites that may accompany each type of answer. **Stimulus conditions** are used to characterize the speaker’s motivation for including a satellite. An indirect answer is the result of the speaker (R) expressing only part of the planned response, i.e., omitting the direct answer (and possibly more), but intending for his discourse plan to be recognized by the hearer (Q). Furthermore, we argue that because of the role of **interpretation** in generation, Q’s belief that R intended for Q to recognize the answer is warranted by Q’s recognition of the plan.

The inputs to the interpretation component of the model (a model of Q’s interpretation of an indirect answer) are the semantic representation of the questioned proposition, the semantic representation of the utterances given by R during R’s turn, shared pragmatic knowledge, and Q’s beliefs, including those presumed by Q to be shared with R. (Beliefs presumed by an agent to be shared by another agent are hereafter referred to as **shared beliefs**, and those that are not presumed to be shared as **nonshared beliefs**).⁵ The output is a set of alternative discourse plans that might be ascribed to R by Q, ranked by plausibility. R’s inferred discourse plan provides the intended answer and possibly other information about R’s beliefs and intentions. The inputs to the generation component (a model of R’s construction of a response) are the semantic representation of the questioned proposition, shared pragmatic knowledge, and R’s beliefs (both shared and nonshared). The output of generation is R’s discourse plan for a full answer, including a specification of which parts of the plan do not need to be explicitly given by R, i.e., which parts should be inferable by Q from the rest of the answer.⁶

This paper describes the knowledge and processes provided in our model for interpreting and generating indirect answers. (The model is not intended as a cognitive model, i.e., we are not claiming that it reflects the participants’ cognitive states during the time course of comprehension and generation. Rather, its purpose is to compute the end products of comprehension and generation, and to contribute to a computational theory of conversational implicature.) As background, Section 2 describes some relevant generalizations about questions and answers in English. Section 3 describes the **reversible** knowledge in our model, i.e., knowledge used both in interpretation and generation of indirect answers. Sections 4 and 5 describe the interpretation and generation components, respectively. Section 5 includes a description of additional pragmatic knowledge required for generation. Section 6 provides an evaluation of the work. Finally, the last section discusses future research and provides a summary.

4 This terminology was adopted from Rhetorical Structure Theory (Mann and Thompson 1983, 1988), discussed in Section 2.

5 Our notion of shared belief is similar to the notion of one-sided mutual belief (Clark and Marshall 1981). However, following Thomason (1990), a shared belief is merely represented in the conversational record *as if* it were mutually believed, although each participant need not actually believe it.

6 However, our model does not address the interesting question of under what conditions a direct answer should be given explicitly even when it is inferable from other parts of the response. For some related work on the function of redundant information, see Walker (1993).

2. Background

This section begins with some results of a corpus-based study of questions and responses in English that provide the motivation for the notion of a full answer in our model. Next, we describe informally how **coherence relations** (similar to subject-matter relations of Rhetorical Structure Theory [Mann and Thompson 1983, 1988]) are used to characterize the possible types of indirect answers handled in our model.

2.1 Descriptive Study of Questions and Responses

Stenström (1984) describes characteristics of questions and responses in English, based on her study of a corpus of 25 conversations (face-to-face and telephone). She found that 13% of responses to polar questions (typically expressed as subject-auxiliary inverted questions) were indirect answers, and that 7% of responses to requests for confirmation (expressed as tag-questions and declaratives) were indirect.⁷ Furthermore, she points out the similarity in function of indirect answers to the extra information, referred to as *qualify* acts in her classification scheme, often accompanying direct answers (Stenström 1984).⁸ Stenström notes that both are used

- to answer an implicit *wh*-question, as in (3),⁹
 - (3) i. Q: Isn't your country seat there somewhere?
 - ii. R: [Yes/No].
 - iii. Stoke d'Abernon.
- for social reasons, as in (4),
 - (4) i. Q: Did you go to his lectures?
 - ii. R: [Yes.]
 - iii. Oh he had a really caustic sense of humour actually.
- to provide an explanation, as in (5),
 - (5) i. Q: And also did you find my blue and green striped tie?
 - ii. R: [No.]
 - iii. I haven't looked for it.
- or to provide clarification, as in (6).
 - (6) i. Q: I don't think you've been upstairs yet.
 - ii. R: [Yes, I have been upstairs.]
 - iii. Um only just to the loo.

In the above examples, coherence would not be affected by making the associated direct answer explicit. She suggests that the main distinction between *qualify* acts and indirect answers is the absence or presence of a direct answer.

⁷ Both of these types of requests are classified as *yes-no* questions in our model. Also, in Stenström's scheme, an utterance may be classified as performing more than one function. For example, an utterance may be classified as both a polar question and a request for identification (i.e., an implicit *wh*-question).

⁸ Other types of acts noted by Stenström as possibly accompanying direct answers, *amplify* and *expand*, are not relevant to the problem of modeling indirect answers.

⁹ (3), (4), (5), and (6) are based on Stenström's (65), (67), (68), and (142), respectively. In (3) either a *yes* or *no* could be conveyed, depending upon how *there* is interpreted and shared background knowledge about the location of Stoke d'Abernon.

Thus, in our model, the notion of a full answer is used to model both indirect answers and direct answers accompanied by qualify acts. A full answer consists of a direct answer, which we refer to as the nucleus, and possibly extra information of various types, which we refer to as satellites.¹⁰ Then, an indirect answer can be modeled as the result of R giving one or more satellites of the full answer, without giving the nucleus explicitly, but intending for the full answer to be recognized. A benefit of this approach is that it also can be used to model the generation of qualify acts accompanying direct answers. (That is, a qualify act would be a result of R providing the satellite(s) along with an explicit nucleus.) In the next section, we informally describe how different types of satellites of full answers (i.e., types of indirect answers) can be characterized.

2.2 Characterizing Types of Indirect Answers

Consider the constructed responses shown in (1) through (5) of Table 1, which are representative of the types of full answers handled in our model.¹¹ The (a) sentences are *yes-no* questions and each (b) sentence expresses a possible type of direct answer.¹² Each of the sentences labeled (c) through (e) could accompany the preceding (b) sentence in a full answer,¹³ or could be used without (b), i.e., as an indirect answer used to convey the answer given in (b). Also, to the right of each of the (c)–(e) sentences is a name intended to suggest the type of relation holding between that sentence and the associated (b) sentence. For example, (1c) provides a condition for the truth of (1b), (1d) elaborates upon (1b), and (1e) provides the agent's motivation for (1b). Many of these relations are similar to the subject-matter relations of Rhetorical Structure Theory (RST) (Mann and Thompson 1983, 1988), a general theory of discourse coherence. Thus, we refer to these as coherence relations. Other sentences providing the same type of information, i.e., satisfying the same coherence relation, could be substituted for each (c)–(e) sentence without destroying coherence. For example, another plausible condition could be substituted for (1c). Thus, as this table illustrates, a small set of coherence relations characterizes a wide range of possible indirect answers.¹⁴ Furthermore, as it illustrates, certain coherence relations are characteristic of only one or two types of answer, e.g., giving a cause instead of *yes*, or an obstacle instead of *no*.

To give a brief overview of Rhetorical Structure Theory as it relates to our model, one of the goals of RST is to provide a set of relations for describing the organization of coherent text. An RST relation is defined as a relation between two text spans, called the nucleus and satellite. The nucleus is the span which is "more essential to the writer's purpose [than the satellite is]" (Mann and Thompson 1988, 266). A relation definition provides a set of constraints on the nucleus and satellite, and an effect field. According to RST, implicit relational propositions are conveyed in discourse.

10 As noted earlier, this terminology is borrowed from Rhetorical Structure Theory, described below.

11 Constructed examples are used here to provide a concise means of demonstrating the classes of satellites.

12 Specifically, the possible types of direct answers handled in the model are: (1b) that *p* is true, (2b) that *p* is false, (3b) that there is some truth to *p*, (4b) that *p* may be true, or (5b) that *p* may be false, where *p* is the questioned proposition.

13 When more than one of the (c)–(e) sentences is used in the same response, coherence may be improved by use of discourse connectives.

14 However, we are not claiming that this set is exhaustive, i.e., that it characterizes all possible indirect answers.

Table 1
Examples of coherence relations in full answers.

1.	a.	Are you going shopping tonight?	
	b.	Yes.	
	c.	If I finish my homework.	Condition
	d.	I'm going to the mall.	Elaboration
	e.	I need new running shoes.	Cause
2.	a.	Aren't you going shopping tonight?	
	b.	No.	
	c.	I wouldn't have enough time to study.	Otherwise
	d.	My car's not running.	Obstacle
	e.	I'm going tomorrow night.	Contrast
3.	a.	Is dinner ready?	
	b.	To some extent.	
	c.	The pizza is ready.	Contrast
4.	a.	Is Lynn here?	
	b.	I think so.	
	c.	Her books are here.	Result
	d.	She's usually here all day.	Usually
	e.	I think she has a meeting here at 5.	Possible Cause
5.	a.	Is Lynn here?	
	b.	I don't think so.	
	c.	Her books are gone.	Result
	d.	She's not usually here this late.	Usually
	e.	I think she has a dentist appointment this afternoon.	Possible Obstacle

For example, (7) conveys, in addition to the propositional content of (7)i and (7)ii, the relational proposition that the 1899 Duryea is in the writer's collection of classic cars.¹⁵

- (7) i. I love to collect classic automobiles.
ii. My favorite car is my 1899 Duryea.

Such relational propositions are described in RST in a relation definition's effect field. The organization of (7) would be described in RST by the relation of Elaboration, where (7)i is the nucleus and (7)ii a satellite. To see the usefulness of RST for the analysis of full answers to *yes-no* questions, consider (8).

- (8) i. Q: Do you collect classic automobiles?
ii. R: Yes.
iii. I recently purchased an Austin-Healey 3000.

Although (8)ii is not semantically entailed by (8)iii, R could use (8)iii alone in response to (8)i to conversationally implicate (8)ii. Further, just as (7)ii provides an elaboration

¹⁵ This example is from Mann and Thompson (1983), page 81.

Table 2
Similar RST relations.

Coherence Relation	Similar RST Relation Name(s)
Cause	Non-Volitional Cause, Purpose, Volitional Cause
Condition	Condition
Contrast	Contrast
Elaboration	Elaboration
Obstacle	
Otherwise	Otherwise
Possible-cause	
Possible-obstacle	
Result	Non-Volitional Result, Volitional Result
Usually	

of (7)i, (8)iii provides an elaboration of (8)ii, whether (8)ii is given explicitly as an answer or not.¹⁶ Also, in giving just (8)iii as a response, R intends Q to recognize not only (8)ii but also this relation, i.e., that the car is part of R's collection.

Table 2 lists, for each of the coherence relations defined in our model (shown in the left-hand column), similar RST relations (shown in the right-hand column), if any. Although other RST relations can be used to describe other parts of a response (e.g., Restatement), only relations that contribute to the interpretation of indirect answers are included in our model. The formal representation of the coherence relations provided in our model is discussed in Section 3.

3. Reversible Knowledge

As shown informally in the previous section, coherence relations can be used to characterize various types of satellites of full answers. **Coherence rules**, described in Section 3.1, provide sufficient conditions for the mutual plausibility of a coherence relation. During generation, plausibility of a coherence relation is evaluated with respect to the beliefs that R presumes to be shared with Q. During interpretation, the same rules are evaluated with respect to the beliefs Q presumes to be shared with R. Thus, during generation R assumes that a coherence relation that is plausible with respect to his shared beliefs would be plausible to Q as well. That is, Q ought to be able to recognize the implicit relation between the nucleus and satellite.

However, the generation and interpretation of indirect answers requires additional knowledge. For example, for R's contribution to be recognized as an answer, there must be a discourse expectation (Levinson 1983; Reichman 1985) of an answer. Also, during interpretation, for a particular answer to be licensed by R, the attribution of R's intention to convey that answer must be consistent with Q's beliefs about R's intentions. For example, a putative implicature that *p* holds would not be licensed if R provides a disclaimer that it is not R's intention to convey that *p* holds. This and other types of knowledge about full answers is represented as discourse plan operators, described in Section 3.2. In our model, a discourse plan operator captures shared, domain-independent knowledge that is used, along with coherence rules, by

¹⁶ This may seem to conflict with the idea in RST that the nucleus, being more essential to the writer's purpose than a satellite, cannot be omitted. However, at least in the case of the coherence relations playing a role in our model, it appears that the nucleus need not be given *explicitly* when it is inferable in the discourse context.

It is mutually plausible to the agent that (*cr-obstacle q p*) holds,
 where q is the proposition that a state s_q does not hold during time period t_q ,
 and p is the proposition that an event e_p does not occur during time period t_p ,
 if the agent believes it to be mutually believed that s_q is a precondition
 of a typical plan for doing e_p ,
 and that t_q is before or includes t_p ,
 unless it is mutually believed that s_q does hold during t_q ,
 or that e_p does occur during t_p .

It is mutually plausible to the agent that (*cr-obstacle q p*) holds,
 where q is the proposition that a state s_q holds during time period t_q ,
 and p is the proposition that a state s_p does not hold during time period t_p ,
 if the agent believes it to be mutually believed that s_q typically prevents s_p ,
 and that t_q is before or includes t_p ,
 unless it is mutually believed that s_q does not hold during t_q ,
 or that s_p does hold during t_p .

Figure 1

Glosses of two coherence rules for *cr-obstacle*.

the generation component to construct a discourse plan for a full answer. Interpretation is modeled as inference of R's discourse plan from R's response using the same set of discourse plan operators and coherence rules. Inference of R's discourse plan can account for how Q derives an implicated answer, since a discourse plan explicitly represents the relationship of R's communicative acts to R's beliefs and intentions. Together, the coherence rules and discourse plan operators described in this section make up the reversible pragmatic knowledge, i.e., pragmatic knowledge used by both the generation and interpretation components, of the model. Other pragmatic knowledge, used only by the generation process to constrain content planning, is presented in Section 5.

3.1 Coherence Rules

Coherence rules specify sufficient conditions for the plausibility to an agent with respect to the agent's shared beliefs (which we hereafter refer to as the **mutual plausibility**) of a relational proposition (*CR q p*), where *CR* is a coherence relation and q and p are propositions. (Thus, if the relational proposition is plausible to R with respect to the beliefs that R presumes to be shared with Q, R assumes that it would be plausible to Q, too.) To give some examples, glosses of some rules for the coherence relation, which we refer to as *cr-obstacle* are given in Figure 1.¹⁷ The first rule characterizes a subclass of *cr-obstacle*, illustrated in (9), relating the nonoccurrence of an agent's volitional action (reported in (9)ii) to the failure of a precondition (reported in (9)iii) of a potential plan for doing the action.

- (9) i. Q: Are you going to campus tonight?
 ii. R: No.
 iii. My car's not running.

¹⁷ For readability, we have omitted the prefix *cr-* in Tables 1 and 2.

In other words, it is mutually plausible to an agent that the propositions conveyed in (9)iii and (9)ii are related by *cr-obstacle*, provided that the agent has a shared belief that a typical plan for R to go to campus has a precondition that R's car is running. The second rule in Figure 1 characterizes another subclass of *cr-obstacle*, illustrated in (10), relating the failure of one condition (reported in (10)i) to the satisfaction of another condition (reported in (10)ii).

- (10) i. R: My car's not running.
 ii. The timing belt is broken.

In other words, it is mutually plausible to an agent that the propositions conveyed in (10)ii and (10)i are related by *cr-obstacle*, provided that the agent has a shared belief that having a broken timing belt typically prevents a car from running.

Coherence rules are evaluated with respect to an agent's shared beliefs. Coherence rules and the agent's beliefs are encoded as Horn clauses in the implementation of our model. The sources of an agent's shared beliefs include:

- terminological knowledge: e.g., that driving a car is a type of action,
- domain knowledge, including
 - domain planning knowledge: e.g., that a subaction of a typical plan to go to campus is to drive to campus, and that a typical plan for driving a car has a precondition that the car is running,
 - other domain knowledge: e.g., that a broken timing belt typically prevents a car from running, and
- the discourse context: e.g., that R has asserted that R's car is not running.

3.2 Discourse Plan Operators

The discourse plan operators provided in the model encode generic programs for expressing full answers (and subcomponents of full answers).¹⁸ For example, the discourse plan operators for constructing full *yes* (*Answer-yes*) and full *no* (*Answer-no*) answers are shown in Figure 2.¹⁹

The first line of a discourse plan operator, its **header**, e.g., (*Answer-yes s h ?p*), gives the type of discourse action, the participants (*s* denotes the speaker and *h* the hearer), and a propositional variable. (Propositional variables are denoted by symbols prefixed with "?".) In top-level operators such as *Answer-yes* and *Answer-no*, the header variable would be instantiated with the questioned proposition. **Applicability conditions**, when instantiated, specify necessary conditions for appropriate use of a discourse plan operator.²⁰ For example, the first applicability condition of *Answer-yes* and *Answer-no* requires the speaker and hearer to share the discourse expectation that the speaker will inform the hearer of the speaker's evaluation of the truth of the questioned proposition *p*. Present in each of the five top-level answer operators, this particular applicability condition restricts the use of these operators to contexts where an answer is expected,

18 The particular formalism we adopted to encode the operators was chosen to provide a concise and perspicuous organization of the knowledge required for our interpretation and generation components. We make no further claims about the formalism itself.

19 There are three other "top-level" operators in the model for expressing the remaining types of full answers illustrated in Table 1.

20 In general, an applicability condition is a condition that must hold for a plan operator to be invoked, but that a planner will not attempt to bring about (Carberry 1990).

(Answer-yes s h ?p):	(Answer-no s h ?p)
Applicability conditions:	Applicability conditions:
(discourse-expectation	(discourse-expectation
(informif s h ?p))	(informif s h ?p)
(bel s ?p)	(bel s (not ?p))
Nucleus:	Nucleus:
(inform s h ?p)	(inform s h (not ?p))
Satellites:	Satellites:
(Use-condition s h ?p)	(Use-otherwise s h (not ?p))
(Use-elaboration s h ?p)	(Use-obstacle s h (not ?p))
(Use-cause s h ?p)	(Use-contrast s h (not ?p))
Primary goals:	Primary goals:
(BMB h s ?p)	(BMB h s (not ?p))

Figure 2
Discourse plan operators for *yes* and *no* answers.

and is needed to account for the hearer's attempt to interpret a response as an answer, even when it is not a direct answer.²¹ The second applicability condition of the top-level operators requires the speaker to hold the evaluation of *p* to be conveyed; e.g., in *Answer-no* it requires that the speaker believe that *p* is false. The primary goals of a discourse plan specify the discourse goals that the speaker intends for the hearer to recognize.²² For example, the primary goal of *Answer-yes* can be glossed as the goal that Q will accept the *yes* answer, at least for the purposes of the conversation.²³

The nucleus and satellites of a discourse plan describe primitive or nonprimitive acts to be performed to achieve the primary goals of the plan.²⁴ *Inform* is a primitive act that can be realized directly. The nonprimitive acts are defined by discourse plan operators themselves. (Thus, a discourse plan may have a hierarchical structure.) A full answer may contain zero, one, or more instances of each type of satellite, and the default (but not required) order of nucleus and satellites in a full answer is the order given in the corresponding operator.

Consider the *Use-elaboration* and *Use-obstacle* discourse plan operators, shown in Figure 3, describing possible satellites of *Answer-yes* and *Answer-no*, respectively. All satellite operators include a second propositional variable referred to as the **existential**

21 Without recourse to the notion of discourse expectation, it is difficult to account for the interpretation in (9)iii of *My car's not running* as *The speaker is not going to campus tonight*, while blocking interpretations such as *The speaker will rent a car*. Note that the latter interpretation may be licensed when the discourse expectation is that R will provide an answer to *Are you going to rent a car?* In general, discourse expectations provide a contextual constraint on what inferences are licensed by the speaker. (Similarly, it has been argued that scalar implicatures depend on the existence of a salient partially ordered set in the discourse context; see Section 4.3.) For a discussion of the overall role of discourse expectations in our model, see Section 4.2. One might argue that this type of applicability condition limits the generality of the operators and thus, could lead to a proliferation of context-specific operators, which would result in inefficient processing. First, we are not claiming that all discourse operators require this type of applicability condition, only those operators characterizing discourse-expectation-motivated units of discourse. Second, with an indexing scheme sensitive to discourse expectations, this would not necessarily lead to efficiency problems.

22 We refer to these as **primary** to distinguish them from other discourse goals the speaker may have but that he does not necessarily intend for the hearer to recognize.

23 During interpretation (see Section 4.1), in order for the implicature to be licensed, the applicability conditions and primary goals of any plan ascribed to R must be consistent with Q's beliefs about R's beliefs and goals. Thus, applicability conditions and primary goals play an important role in canceling spurious putative implicatures.

24 The discourse plan operators in our model are not intended to describe all acts that may accompany a direct answer. For example, the model does not address the generation of parts of the response, such as repetition or restatement, which entail the answer.

<pre>(Use-elaboration s h ?p): Existential variable: ?q Applicability conditions: (bel s (cr-elaboration ?q ?p)) (Plausible (cr-elaboration ?q ?p)) Nucleus: (inform s h ?q) Satellites: (Use-cause s h ?q) (Use-elaboration s h ?q) Primary goals: (BMB h s (cr-elaboration ?q ?p))</pre>	<pre>(Use-obstacle s h ?p): Existential variable: ?q Applicability conditions: (bel s (cr-obstacle ?q ?p)) (Plausible (cr-obstacle ?q ?p)) Nucleus: (inform s h ?q) Satellites: (Use-obstacle s h ?q) (Use-elaboration s h ?q) Primary goals: (BMB h s (cr-obstacle ?q ?p))</pre>
--	---

Figure 3 Two satellite discourse plan operators.

variable. For example, (9)ii–(9)iii could be described by a plan constructed from an *Answer-no* discourse plan operator

- whose header variable is instantiated with the proposition *p* that R is going to campus tonight, and
- which has a satellite constructed from a *Use-obstacle* discourse plan operator whose header variable is instantiated with (*not p*), the proposition that R is *not* going to campus tonight, and whose existential variable *q* is instantiated with the proposition that R’s car is not running.

In general, each satellite operator in our model has applicability conditions and primary goals analogous to those shown in Figure 3. (Each satellite operator has a name of the form, *Use-CR*, where *CR* is the name of a coherence relation.) The first applicability condition of a satellite operator, *Use-CR*, requires that the speaker believes that the relational proposition (*CR q p*) holds for propositions *q* and *p* instantiating the existential variable and header variable, respectively. The second applicability condition requires that, given the beliefs that the speaker presumes to be shared with the hearer, this relational proposition is plausible. (Mutual plausibility is evaluated using the coherence rules described in Section 3.1.) The primary goal of a satellite operator can be glossed as the goal that the hearer will accept the relational proposition.

4. Interpretation

This section describes the interpretation process. In our model, implicated answers are derived by an **answer recognizer**. Algorithms for the answer recognizer are described in Section 4.1. Of course, dialogue consists of more than questions and answers. Section 4.2 describes the role of the answer recognizer in a discourse-processing architecture. Finally, Section 4.3 discusses how this model relates to previous models of conversational implicature.

4.1 Answer Recognizer

4.1.1 Main Algorithm. The structure of the answer recognizer is shown in Figure 4. The inputs to the answer recognizer include:

- the set of discourse plan operators and coherence rules described in Section 3,

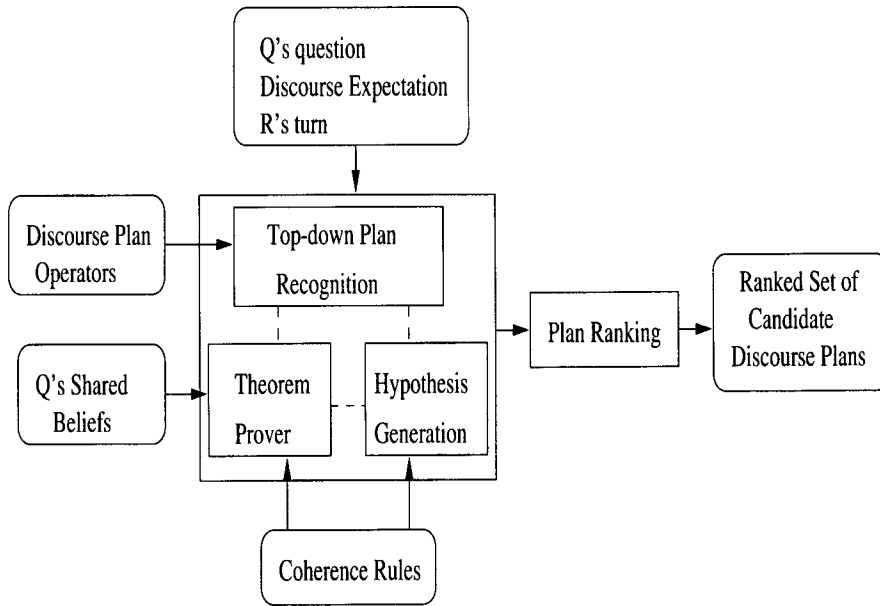


Figure 4
Structure of the answer recognizer.

- Q's beliefs (including the discourse expectation that R will provide an answer to the questioned proposition p),
- the semantic representation of p , and
- for each utterance performed by R during R's turn, the type of communicative act signaled by its form (e.g., to inform), and the semantic representation of its content.²⁵

Answer recognition is performed in two phases. The goal of the first phase is to derive a set of candidate discourse plans plausibly underlying R's response. The first phase makes use of two subcomponents: one that we refer to as the **hypothesis generation** component, and a theorem prover. The output of the first phase of answer recognition is a set of candidate discourse plans since there may be alternate interpretations of R's response. The goal of the second phase of answer recognition is to evaluate the relative plausibility of each candidate discourse plan. The final output of answer recognition consists of a partially ordered set of the candidates ranked by plausibility.

Plan recognition is primarily top-down, i.e., expectation-driven. More specifically, Q²⁶ attempts to interpret the response as having been generated from a discourse plan constructed from the discourse plan operators for full answers. The problem of reconstructing R's discourse plan has several aspects (to be described in more detail shortly):

- Instantiating discourse plan operators with the questioned proposition and appropriate propositions from R's response.

²⁵ The turn in question need not be the turn immediately following Q's asking of the question, as discussed in Section 4.2. Also, we make the simplifying assumption that R's answer is given within a single turn.

²⁶ For convenience, we refer to the answer recognizer component as Q, and to the answer generator as R.

- **Consistency checking:** determining whether the beliefs and goals that would be attributed to R by virtue of ascribing a particular discourse plan to R are consistent with Q's beliefs about R's beliefs and goals.
- **Coherence evaluation:** determining whether a putative satellite of a candidate plan is plausibly coherent, i.e., given a candidate plan's (or subplan's) nucleus proposition p , putative satellite proposition q , and the putative satellite's coherence relation CR , determining whether Q believes that $(CR\ q\ p)$ is mutually plausible. Coherence evaluation makes use of the coherence rules described in Section 3.1.
- **Hypothesis generation:** hypothesizing any "missing parts" of the response that are required in order to assimilate acts in R's response into a coherent candidate plan. Hypothesis generation also makes use of the coherence rules.

Initially, the header variable of each "top-level" answer discourse plan operator²⁷ is instantiated with the questioned proposition p , i.e., all occurrences of the header variable are replaced with p . Next, consistency checking is performed to eliminate any candidates whose applicability conditions or primary goals are not consistent with Q's beliefs about R's beliefs and goals. For all remaining candidates, the answer recognizer next attempts to recognize an act from R's turn as the nucleus of the plan, i.e., to check whether R gave a direct answer. If no acts in R's turn match the nucleus, then the nucleus is marked as hypothesized. For all remaining acts in R's turn, the answer recognizer attempts to recognize all possible satellites, as specified in each remaining candidate plan. In the model the discourse plan operators do not specify a required ordering of satellites.²⁸ The subprocedure of **satellite recognition** is described in more detail in Section 4.1.2.

4.1.2 Satellite Recognition. Satellite recognition is the (recursive) process of recognizing an instance of a satellite of a candidate plan. The inputs consist of:

- *sat-op*, a discourse plan operator for a possible satellite,
- the proposition p conveyed by the nucleus of the higher-level plan (i.e., the plan whose satellites are currently being recognized),
- *act-list*, a list of acts in R's turn that have not yet been assimilated into the candidate plan,
- *cur-act*, the current act (*inform s h q*) in *act-list*, where s is the speaker, h is the hearer, and q is the propositional content of the act.

The output is a set (possibly empty) of candidate instances of *sat-op*. To give a simplified, preliminary version of the algorithm, first, the header variable and existential variable of *sat-op* are instantiated with p and q , respectively. Then, coherence evaluation and consistency checking are performed. If successful, *cur-act* is recognized to

²⁷ Five of these are defined in our model, corresponding to the five types of answers illustrated in Table 1.

²⁸ The operators do specify a preferred order, however, which is used in generation. Also, our process model includes a structural constraint on satellite ordering. During interpretation, only instances satisfying this constraint are considered. That is, the constraint eliminates interpretations which, in our judgment, are not plausible due to incoherence. For a description of the constraint, see Green (1994). We expect that other such constraints may be incorporated into the process model.

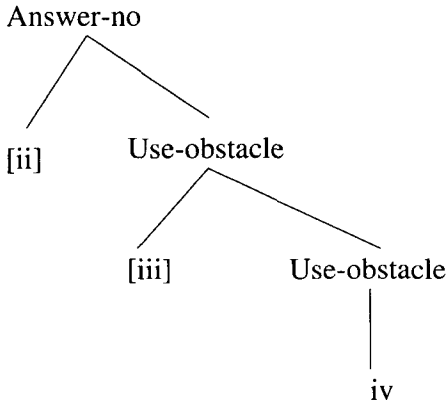


Figure 5
Candidate discourse plan with hypotheses.

be the nucleus of *sat-op*, and for each remaining act in *act-list*, satellite recognition is performed for each satellite of *sat-op*.

However, the satellite recognition algorithm as described so far would not be able to handle R’s response in (11), since there is no plausible coherence relation in the model directly relating (11)iv to (11)ii (or to any other direct answer that could be recognized in the model).

- (11) i. Q: Are you driving to campus tonight?
- ii. R: [No.]
- iii. [My car’s not running.]
- iv. My car has a broken timing belt.

Whenever the answer recognizer is unable to recognize *cur-act* as the nucleus of *sat-op*, a subprocedure we refer to as hypothesis generation is invoked. Hypothesis generation will be described in detail in the following section. It returns a set of alternative hypothesized propositions, each of which represents the content of a possible implicit *inform* act to be inserted at the current point of expanding the candidate plan.²⁹ In this example, the proposition conveyed in (11)iii would be returned as a hypothesized proposition, which is used to instantiate the existential variable of a *Use-obstacle* satellite, thereby enabling satellite recognition to proceed. Then, (11)iv can be recognized (without hypothesis generation being required) as a satellite of (11)iii. Ultimately, the plan shown in Figure 5 would be inferred. (Only the hierarchical structure and communicative acts are shown. By convention, the left-most child of a node is the nucleus and its siblings are the satellites. Labels of sentences in (11) that could realize a leaf node are used to label the node. Hypothesized nodes are indicated by square brackets.) The complete satellite recognition algorithm, employing hypothesis generation, is given in Figure 6.

²⁹ Thus, hypothesis generation may provide additional inferences, i.e., more than just the implicated answer. Hinkelman (1989) refers to such implicatures, licensed by attributing a plan to an agent, as plan-based implicatures.

- INPUT p: proposition from nucleus of higher-level plan
 cur-act: current act, (inform s h q), to be recognized
 act-list: list of remaining acts in R's turn
 op: discourse plan operator (Use-CR s h ?p)
- OUTPUT sat-cand-set: set of candidate instances of op
 underlying part of R's response
1. Instantiate header variable ?p of op with p.
 2. Instantiate existential variable ?q of op with q of cur-act.
 - a. Prove that it is plausible that q and p are related by CR.
 If not, go to step 2c.
 - b. Check consistency. If not consistent, then go to step 2c;
 else go to step 3a.
 - c. Try substituting each q returned by hypothesis generation for ?q:
 Check consistency and coherence as in steps 2a and 2b.
 For each q passing both checks, proceed with step 3b.
 If none pass, then fail.
 3. a. Mark cur-act as used. Go to step 4.
 b. Mark nucleus as hypothesized.
 4. For each unused act in act-list, attempt to recognize each satellite
 of op.

Figure 6
 Satellite recognition algorithm.

4.1.3 Hypothesis Generation. Based upon the assumption that the response is coherent, the goal of hypothesis generation is to fill in missing parts of a candidate plan in such a way that an utterance in R's turn can be recognized as part of the plan. The use of hypothesis generation broadens the coverage of our model to cases where more is missing from a full answer than just the nucleus of a top-level operator. (From the point of view of generation, it enables the construction of a more concise, though no less informative, response.) The hypothesis generation algorithm constructs chains of mutually plausible propositions, each beginning with the proposition (e.g., the proposition conveyed in (11)iv) to be related to a **goal** proposition in a candidate plan (e.g., the proposition conveyed in (11)ii), and ending with the goal proposition, where each pair of adjacent propositions in the chain is linked by a plausible coherence relation. The algorithm returns the proposition (e.g., the proposition conveyed in (11)iii) immediately preceding the goal proposition in each chain. Thus, when top-down recognition has reached an impasse, hypothesis generation (a type of bottom-up data-driven reasoning) provides a hypothesis that enables top-down recognition to continue another level of growth. An example of hypothesis generation is given in Section 4.1.5.

The algorithm for hypothesis generation, which is given in Figure 7, performs a breadth-first search subject to a processing constraint on the maximum depth of the search tree. Note that a chain may have a length greater than three, e.g., the chain may consist of propositions (p_0, p_1, p_2, p_3), where p_0 is the proposition to be related to the candidate plan, p_3 is the goal, and p_2 would be returned as a hypothesized proposition. In such a case, after p_2 has been assimilated into the candidate plan, if p_1 is not present in R's turn, then hypothesis generation is invoked again and p_1 would

INPUTS p_0 : initial proposition
 p_g : goal proposition
 GCR: goal coherence relation, i.e., coherence relation
 that must hold between hypothesized proposition and p_g
 S: set of coherence relations
 N: maximum search depth

OUTPUT hypoht-list: list of alternative hypothesized propositions

1. Initialize root of search tree with p_0 .
2. Expand nodes of tree in breadth-first order until either no more expansion is possible or maximum tree depth of N is reached, whichever happens first. To expand a node p_i :
 - a. Find all nodes p_{i+1} such that for some relation CR in S, (Plausible (CR p_i p_{i+1})) is provable.
 - b. Make each such p_{i+1} a child of p_i , linked by CR.
 - c. A goal state is reached whenever p_{i+1} is p_g and CR is identical to GCR.
 - d. Whenever a goal state is reached, add the parent of p_g to hypoht-list.

Figure 7

Hypothesis generation algorithm.

be hypothesized also.³⁰ Finally, the search for a proposition p_{i+1} in step 2a is performed in our implementation using a theorem prover.

4.1.4 Ranking Candidate Plans. Two heuristics are used to rank the relative plausibility of the set of candidate plans output by the first phase of answer recognition. First, plausibility decreases as the number of hypotheses in a candidate increases. (Assuming that all else is equal, it is safer to favor interpretations requiring fewer hypotheses.) Second, plausibility increases as the number of utterances in R's turn that are accounted for by the plan increases. (The more of R's turn accounted for, the more coherent the turn is likely to be, although not all of the utterances in R's turn are necessarily part of the full answer.) To give an example, consider the two candidate plans shown in Figure 8, corresponding to alternative interpretations of R's response in (12).³¹

- (12) i. Q: Are you going to campus tonight?
 ii. R: [No/Yes]
 iii. [My car's not running.]
 iv. My car has a broken timing belt,
 v. [so] I'm going to take the bus.
 vi. Do you know how much the fare is?

³⁰ The implementation saves the chains to avoid the expense of recomputing intermediate hypotheses.

³¹ This constructed example was designed to illustrate multiple aspects of the model. In our judgement, normally it would sound more coherent to give (12)v before (12)iv if (12)vi were not included. However, when (12)vi is included, (12)v not only elaborates upon the *yes*, but also serves as background for (12)vi. Another possible motivation for giving (12)v after (12)iv might be to delay giving dispreferred information (Levinson 1983), e.g., if the speaker believed that a *yes* was an unexpected or unwanted answer to (12)i.

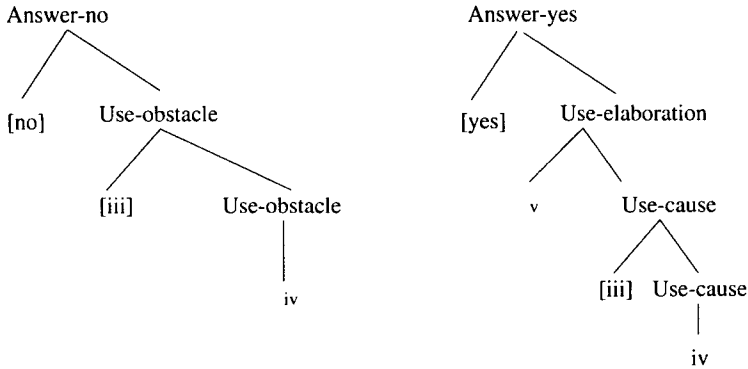


Figure 8
Ranking candidate plans.

By these heuristics, a *yes* answer would be the preferred interpretation, since the candidate *Answer-yes* plan uses the same number of hypotheses as the candidate *Answer-no* plan, and accounts for more of R's response. ((12)vi is not recognized as part of either answer.) The preference heuristics are intended to capture local coherence only. Since global information may play a role in selecting the correct interpretation, the higher-level discourse processor (described in section 4.2) must decide which plan to attribute to the speaker.

4.1.5 Answer Recognition Example. In this section, we illustrate the interpretation of indirect answers in the model by describing how the two candidate plans shown in Figure 8 would be derived from R's response of (12)iv through (12)vi. First, each of the five top-level answer discourse plan operators would be instantiated with the questioned proposition p , the proposition that R is going to campus tonight. Assuming that Q has no beliefs about R's beliefs and goals that are inconsistent with the applicability conditions and primary goals of these candidates, none of the candidates would be eliminated yet. Second, for each candidate the recognizer would check whether the communicative act specified in the nucleus was present in R's turn. In this example, since a direct answer was not explicitly provided by R, the recognizer would mark the nucleus of each candidate as hypothesized. The hypothesized nucleus of the candidate *Answer-no* and *Answer-yes* plans would be (*inform s h (not p)*) and (*inform s h p*), respectively. Next, the recognizer would try to recognize the acts expressed as (12)iv through (12)vi as satellites of each candidate plan. Assume that these acts are represented as (*inform s h p_{iv}*), (*inform s h p_v*), and (*inform s h p_{vi}*), respectively.

To recognize an instance of a satellite, first, a satellite discourse plan operator would be instantiated. The header variable would be instantiated by unifying the satellite plan header with the corresponding act in the higher-level plan. For example, the header variable of a *Use-obstacle* satellite of an *Answer-no* candidate would be instantiated with (*not p*) in this example. The existential variable would be instantiated with the proposition conveyed in some utterance to be recognized as a satellite, e.g., p_{iv} . However, before a candidate satellite may be attached to the higher-level candidate plan, the answer recognizer must verify that the candidate satellite passes the following two tests: First, the candidate satellite's applicability conditions and primary goals must be consistent with Q's beliefs about R's beliefs and goals. Second, the specified coherence relation must be plausible with respect to the beliefs that Q presumes to be shared with R, i.e., the satellite's instantiated applicability condition

of the form (*Plausible (CR q p)*) must be provable using the coherence rules described in Section 3. For example, given the beliefs that Q presumes to be shared with R and the coherence rules provided in the model, the act underlying (12)iv could not be the nucleus of a candidate *Use-obstacle* satellite of the *Answer-no* candidate, because the recognizer would not be able to prove that *cr-obstacle* is a plausible coherence relation holding between p_{iv} and (*not p*). On the other hand, the act underlying (12)v would be interpreted as the nucleus of a candidate *Use-elaboration* satellite of the *Answer-yes* candidate, since the above tests are satisfied, e.g., the recognizer could prove that *cr-elaboration* is a plausible coherence relation holding between p_v and p .

To return to consideration of the recognition of the *Answer-no* candidate, upon finding that the act underlying (12)iv cannot serve as a satellite, hypothesis generation would be attempted. Recall that the goal of hypothesis generation is to supply a hypothesized missing act of the plan so that top-down recognition can continue. Hypothesis generation would search for a chain of plausibly related propositions, beginning with the proposition (i.e., p_{iv}) to be related to the candidate *Answer-no* plan, and ending with the goal proposition (i.e., (*not p*)). As mentioned in Section 4.1.3, each pair of adjacent propositions in the chain must be linked by a plausible coherence relation. In this example, hypothesis generation would construct the chain (p_{iv} , p_{iii} , (*not p*)), where both pairs of adjacent propositions would be related by *cr-obstacle* and p_{iii} is the hypothesis that R's car is not running. Hypothesis generation would return the proposition immediately preceding the goal proposition in this chain, i.e., p_{iii} . Thus, p_{iii} would be used to instantiate the existential variable of a *Use-obstacle* satellite of the *Answer-no* candidate plan, and satellite recognition would proceed. (The nucleus of this satellite would be marked as hypothesized.) Then, the recognizer would recognize p_{iv} (without requiring hypothesis generation) as a *Use-obstacle* satellite of this *Use-obstacle* satellite. No remaining utterances in R's turn can be related to the candidate *Answer-no* plan, resulting in the candidate shown on the left in Figure 8.

Finally, to finish consideration of the recognition of the *Answer-yes* candidate, since neither the act underlying (12)iv nor the act underlying (12)vi can serve as a satellite of the *Answer-yes* candidate or its *Use-elaboration* satellite, hypothesis generation would again be invoked. Hypothesis generation would provide p_{iii} , the hypothesis that R's car is not running, as a plausible explanation for why R is going to take the bus. Thus, p_{iii} would be used to instantiate the existential variable of a *Use-cause* satellite of the *Use-elaboration* satellite of the *Answer-yes* candidate plan, and satellite recognition would proceed. (The nucleus of the *Use-cause* satellite would be marked as hypothesized.) Then, the recognizer would recognize p_{iv} (without requiring hypothesis generation) as a *Use-cause* satellite of this *Use-cause* satellite. No remaining utterances in R's turn can be related to the candidate *Answer-yes* plan, resulting in the candidate shown on the right in Figure 8.

Given the shared beliefs and the coherence rules provided in the model, none of the utterances in R's turn would be recognized as satellites of the other three top-level candidate answer plans. Candidates that do not account for any actual parts of the response are eliminated at the end of phase one. Thus the output of phase one of interpretation would be just the two candidates shown in Figure 8. Phase two would evaluate the *Answer-yes* candidate as more preferred than the *Answer-no* candidate, since the former interpretation requires the same number of hypotheses and also accounts for more of R's response.

4.2 Role of the Answer Recognizer in Discourse Processing

As discourse researchers have pointed out (e.g., Reichman 1985; Levinson 1983)) the asking of a *yes-no* question creates the expectation that R will provide the answer

(directly or indirectly), if possible. Other acceptable, though less preferred, responses include *I don't know* and replies that provide other helpful information. Furthermore, an answer need not be given in the turn immediately following the turn in which the question was asked. For example, in (13) the *yes-no* question in (13)i is not answered until (13)v, separated by a request for clarification in (13)ii and its answer in (13)iii.

- (13) i. Q: Is Dr. Smith teaching CS360 next semester?
 ii. R: Do you mean Dr. Smithson?
 iii. Q: Yes.
 iv. R: [No.]
 v. He will be on sabbatical next semester.

In Carberry's discourse-processing model for ellipsis interpretation (Carberry 1990), a mechanism is provided for updating the shared discourse expectations of dialogue participants throughout a conversation. Our answer recognizer would have the following role in such an architecture: The answer recognizer would be invoked whenever the current discourse expectation is that R will provide an answer. (If answer recognition were unsuccessful, then the discourse processor would invoke other types of recognizers for other types of responses.) The answer recognizer returns a partially ordered set (possibly empty) of answer discourse plans that it is plausible to ascribe to R as underlying (part or all of) the turn. The final choice of which discourse plan to ascribe to R should be made by the higher-level discourse processor, since it must select an interpretation consistent with the rest of the discourse.

4.3 Comparison to Previous Approaches to Conversational Implicature

Grice (1975) has proposed a theory of conversational implicature to account for certain types of conversational inferences. According to Grice, a speaker may convey more than the conventional meaning of an utterance by making use of the hearer's expectation that the speaker is adhering to general principles of cooperative conversation. Two necessary (but not sufficient) properties of conversational implicatures involve cancelability and speaker intention (Grice 1975; Hirschberg 1985). First, potential conversational implicatures may be canceled explicitly, i.e., disavowed by the speaker in the preceding or subsequent discourse context, or even canceled implicitly given a particular set of shared beliefs. In fact, potential implicatures may undergo a change in status from cancelable to noncancelable in the subsequent discourse (Gunji 1981). Second, conversational implicatures are part of the intended meaning of an utterance. Grice proposes several maxims of cooperative conversation that a hearer uses as justification for inferring conversational implicatures. However, Grice's theory is inadequate as the basis for a computational model of how conversational implicatures are derived. As frequently noted, Grice's maxims may support spurious or contradictory inferences. To date, few computational models have addressed the interpretation of conversational implicatures.

Hirschberg's model (Hirschberg 1985) addresses a class of conversational implicatures, **scalar implicatures**, which overlaps with the class of implicated answers addressed in our model. (That is, scalar implicatures arise in question-answer exchanges as well as in other contexts, and, not all types of implicated answers are scalar implicatures.) According to Hirschberg, a scalar implicature depends upon the existence of a partially ordered set of values that is salient in the discourse context. Her model provides licensing rules that specify, given such a set, which scalar implicatures are

It is mutually plausible to the agent that (*cr-contrast* $q p^*$) holds,
 where q is a proposition
 and p^* is the proposition that p is partly true,
 if the agent believes it to be mutually believed that q is less than p in a salient partial order,
 unless it is mutually believed that p is true or that q is not true.

It is mutually plausible to the agent that (*cr-contrast* $q (not p)$) holds,
 where q and p are propositions,
 if the agent believes it to be mutually believed that q is an alternate to p in a salient partial order,
 unless it is mutually believed that p is true or that q is not true.

Figure 9
 Glosses of two coherence rules for *cr-contrast*.

licensed in terms of values in the set that are lower than, alternate to, or higher than the value referred to in an utterance. For example, given a salient partially ordered set such that the value for *the letter from X* is lower than the value for *all of the letters in question*, in saying (2)ii (repeated below in (14)ii) R licenses the implicature that R has not gotten all of the letters in question.

- (14) i. Q: Have you gotten the letters yet?
- ii. R: I've gotten the letter from X.

In our model, the response in (14)ii would be analyzed as generated from an *Answer-hedge* discourse plan whose nucleus has not been explicitly given and which has a single *Use-contrast* satellite whose nucleus is expressed in (14)ii.³² The coherence rules for *cr-contrast*, which are based upon the notions elucidated by Hirschberg, are glossed in Figure 9.³³ However, the discourse plan operators in our model also characterize a variety of indirect answers that are not scalar implicatures, i.e., indirect answers based on the other coherence relations shown in Table 2.

A model such as Hirschberg's, which does not take the full response into account, faces certain problems in handling cancellation by the subsequent discourse context ("backwards" cancellation). For example, given a salient partially ordered set such that *going to campus* is ranked as an alternate to *going shopping*, Hirschberg's model would predict, correctly in the case of (15) and incorrectly in the case of (16), that R intended to convey a *no*.

- (15) i. Q: Are you going shopping?
- ii. R: [no]
- iii. I'm going to campus.
- iv. I have a night class.
- (16) i. Q: Are you going shopping?
- ii. R: [yes]

³² The nucleus of such a plan conveys that the questioned proposition is partly but not completely true.
³³ The uppermost rule in the figure is the one applying to this example. The other rule applies to *Use-contrast* satellites of *Answer-no* plans.

- iii. I'm going to campus.
- iv. The bookstore is having a sale.

In our model, (16) would be interpreted by recognizing an *Answer-yes* plan (with a *Use-elaboration* and a *Use-cause* satellite underlying (16)iii and (16)iv, respectively) as more plausible than an *Answer-no* plan, rather than by use of backwards cancellation.³⁴ In other words, in our model subsequent context can provide evidence for or against a particular interpretation, since a discourse plan may be expressed by multiple utterances.

Also, a model such as Hirschberg's provides no explanation for why potential implicatures may become noncancelable. Our model predicts that a potential implicature of an utterance becomes noncancelable after the point in the conversation when the full discourse plan accounting for that utterance has been attributed to the speaker. For example, imagine a situation in which Q and R mutually intend to discuss two job candidates, A and B. Also, suppose that they mutually believe that they should not discuss any candidate until two letters of recommendation have been received for the candidate, and further, that both letters for B have been received. Our model predicts that the scalar implicature potentially licensed in (17)ii (i.e., that R has not gotten both letters for A yet) is no longer cancelable after R's turn in (17)iv, since by that point, the participants apparently would share the belief that Q had succeeded in recognizing R's discourse plan underlying (17)ii.³⁵

- (17) i. Q: Have you gotten the letters for A yet?
 ii. R: I've gotten the letter from X.
 iii. Q: Then let's discuss B now.
 iv. R: O.K. I think we should interview B, don't you?

Inference of coherence relations has been used in modeling temporal (Lascarides and Asher 1991; Lascarides, Asher, and Oberlander 1992) and other defeasible discourse inferences (Hobbs 1978; Dahlgren 1989). Inference of plausible coherence relations is necessary but not sufficient for interpreting indirect answers. For example, Q also must believe that there is a shared discourse expectation of an answer to a particular question. In other words, in our model, discourse plans provide additional constraints on the beliefs and intentions of the speaker that a hearer uses in interpreting a response. Another limitation of the above approaches is that they provide no explanation for the phenomenon of loss of cancelability described above.

Plan recognition has been used to model the interpretation of indirect speech acts (Perrault and Allen 1980; Hinkelman 1989) and ellipsis (Carberry 1990; Litman 1986), discourse phenomena that share with conversational implicature the two necessary conditions described above, cancelability and speaker intention. However, these models are inadequate for interpreting indirect answers, i.e., for deriving an implicated answer *p* from an indirect answer *q*. In these models, for *p* to be derivable from *q*, it is necessary

34 Of course, in the case where R provides only *I'm going to campus*, both *yes* and *no* interpretations would be inferred as equally plausible in our model. Although prosodic information is not used in our model, it is an interesting question for future research whether it can help in recognizing the speaker's intentions in such cases.

35 In other words, it would sound as if R had changed his mind or was contradicting himself if he said *In fact I've gotten both letters for A* after saying (17)iv.

for the hearer to infer that the speaker is performing or at least constructing a domain plan relating p and q . However, q need not play such a role in the speaker's inferred or actual domain plans, as shown in (18).³⁶ That is, it is not necessary to infer that R has a domain plan involving the renting of a car by X in order to recognize R's intention to convey *no*.)

- (18) i. Q: X will be renting a car, won't he?
 ii. R: [No.]
 iii. He can't drive.

In other words, these models lack requisite knowledge encoded in our model in terms of possible satellites (based on coherence relations) of top-level discourse plan operators. Also, the above plan-based models face the same problems as Hirschberg's since they do not address multiutterance responses.

Philosophers (Thomason 1990; McCafferty 1987) have argued for a plan-based theory of implicature as an alternative to Grice's theory. Thomason proposes that implicatures are comprehended by a process of **accommodation** of the conversational record to fit the inferred plans of the speaker. According to McCafferty, "implicatures are things that the speaker plans that the hearer believe (and that the hearer can realize that the speaker plans that the hearer believe)" (p. 18). He claims that a theory based upon inferring the speaker's plan avoids the problem of predicting spurious implicatures, since the spurious implicature would not be part of the speaker's plan. Our model is consistent with this view of conversational implicature. McCafferty sketches a possible plan-based model to account for the implicated *yes* answer in (19).³⁷

- (19) i. Q: Has Smith been dating anyone?
 ii. R: [Yes.]
 iii. He's been flying to New York every weekend.

Although it was not McCafferty's intention to provide a computational model, but rather to show the plausibility of a plan-based theory of conversational implicature, some limitations of his suggestions for developing a computational model should be noted. First, his proposed rules cannot be used to derive an alternate, plausible interpretation of (19)iii, in which R scalar implicates a *no*.³⁸ Our model can account for both interpretations. The first interpretation would be accounted for by an inferred *Answer-yes* plan with a *Use-elaboration* satellite underlying (19)iii, while the latter would be accounted for by an inferred *Answer-no* plan with a *Use-contrast* satellite underlying (19)iii. More generally, his proposed rules cannot account for types of indirect answers described in our model by coherence relations whose definitions do not involve planning knowledge. Second, even if rules could be added to McCafferty's model to account for a speaker's plan to convey a *no* by use of (19)iii, his model does not provide a way of using information from other parts of the response, e.g., (20)iv, to help recognize the intended answer. As noted earlier, in our model such information can

36 (18) is based upon (1), modified for expository purposes.

37 (19) is from McCafferty (1987), page 67, and is similar to an example of Grice's. In a Gricean account, this implicature would be justified in terms of the Maxim of Relevance.

38 That is, an interpretation in which *flying to New York* is mutually believed to be an alternate to *dating someone* in a salient partially ordered set.

be used to provide evidence favoring one candidate discourse plan over another. (For example, (20)iv would be accounted for by the addition of a *Use-obstacle* satellite to the *Answer-no* candidate described above.)

- (20) i. Q: Has Smith been dating anyone?
 ii. R: [No.]
 iii. He's been flying to New York every weekend.
 iv. Besides, he's married.

4.4 Summary

This section closes with a summary of the argument for the adequacy of our model as a model of conversational implicature. As discussed earlier, two necessary conditions for conversational implicature are cancelability and speaker intention. We have demonstrated that our model can handle forward and backward cancellation, and provides an explanation for the "loss of cancelability" phenomenon. Regarding speaker intention, in our model a conversationally implicated answer is an answer that R planned that Q recognize (and that Q recognizes that R planned that Q recognize).³⁹ We have demonstrated how Q's recognition of R's discourse plan (in particular, the goal to provide an answer to the question) can be performed using the knowledge and algorithms in our model. Furthermore, we argue that Q's recognition of R's intention that Q recognize R's plan follows from the role of interpretation in generation, namely, Q and R mutually believe that R will not say what he does unless R believes that Q will be able to interpret the response as intended. In our model, during generation (to be described in Section 5), R constructs a model of Q's beliefs (using R's shared beliefs), and then simulates Q's interpretation of a trial pruned response. R's decision to use the pruned response depends upon whether R believes that Q would still be able to recognize the answer after the plan has been pruned. During interpretation, given the shared discourse expectation that R will provide an answer to Q's *yes-no* question, Q's use of (Q's) shared beliefs to interpret the response, and Q's belief that R expects that Q will be able to recognize the answer, Q's recognition of a discourse plan for an answer warrants Q's belief that R intended for Q to recognize this intention.

5. Generation

This section describes our approach to the generation of indirect answers. Generation is modeled as a two-phase process of discourse plan construction. First, in the **content planning** phase, a discourse plan for a full answer is constructed. Second, the **plan pruning** phase uses the model's own interpretation capability to determine what information in the full response does not need to be stated explicitly. In appropriate discourse contexts, i.e., in contexts where the direct answer can be inferred by Q from other parts of the full answer, a plan for an indirect answer is thereby generated. When the direct answer must be given explicitly, the result is a plan for a direct answer accompanied by appropriate extra information. (According to the study mentioned in Section 2 [Stenström 1984], 85% of direct answers are accompanied by such information. Thus, it is important to model this type of response as well.)

While the pragmatic knowledge described in Section 3 is sufficient for interpretation, it is not sufficient for the problem of content planning during generation. Applica-

³⁹ Applying McCafferty's description of conversational implicature to indirect answers.

<pre>(Use-elaboration s h ?p): Existential variable: ?q Applicability conditions: (bel s (cr-elaboration ?q ?p)) (Plausible (cr-elaboration ?q ?p)) Stimulus conditions: (answer-ref-indicated s h ?p ?q) (clarify-concept-indicated s h ?p ?q) Nucleus: (inform s h ?q) Satellites: (Use-cause s h ?q) (Use-elaboration s h ?q) Primary goals: (BMB h s (cr-elaboration ?q ?p))</pre>	<pre>(Use-obstacle s h ?p): Existential variable: ?q Applicability conditions: (bel s (cr-obstacle ?q ?p)) (Plausible (cr-obstacle ?q ?p)) Stimulus conditions: (explanation-indicated s h ?p ?q) (excuse-indicated s h ?p ?q) Nucleus: (inform s h ?q) Satellites: (Use-obstacle s h ?q) (Use-elaboration s h ?q) Primary goals: (BMB h s (cr-obstacle ?q ?p))</pre>
--	---

Figure 10

Two satellite discourse plan operators with stimulus conditions.

bility conditions prevent inappropriate use of a discourse plan. However, they do not model a speaker's motivation for choosing to provide extra information. Consider (21).

- (21) i. Q: I need a ride to the mall.
 ii. Are you going shopping tonight?
 iii. R: [No.]
 iv. My car's not running.
 v. The timing belt is broken.

R's reason for providing the information in (21)iv might have been to give an excuse for not being able to offer Q a ride, and R's reason for providing the information in (21)v might have been to provide an explanation for news in (21)iv that may surprise Q. Furthermore, a full answer might be too verbose if every satellite whose applicability conditions held were included in the full answer. On the other hand, at the time when he is asked a question, R may not hold the primary goals of a potential satellite. (In our model the only goal R is assumed to have initially is the goal to provide the answer.) Thus, an approach to selecting satellites driven only by these satellite goals would fail.

To overcome these problems, we have augmented the satellite discourse plan operators, as described in Section 3, with one or more **stimulus conditions**. Two examples are shown in Figure 10. Stimulus conditions describe general types of situations in which a speaker is motivated to include a satellite during plan construction. They can be thought of as situational triggers, which give rise to new speaker goals (i.e., the primary goals of the satellite operator), and which are the compiled result of deeper planning based upon principles of cooperativity (Grice 1975) or politeness (Brown and Levinson 1978).⁴⁰ In order for a satellite to be included, all of its applicability conditions and at least one of its stimulus conditions must be true.

⁴⁰ It was beyond the scope of our research to model *recognition* of stimulus conditions. We argue in Section 5.3, however, that this does not compromise our approach as a model of conversational implicature.

Our methodology for identifying stimulus conditions was to survey linguistic studies, described in Section 5.1, as well as to analyze the possible motivation of the speaker in the examples in our corpus. The rules used in our model to evaluate stimulus conditions are given in Section 5.2. Section 5.3 presents our implemented generation algorithm, and Section 5.4 illustrates the algorithm with an example.

5.1 Linguistic Studies

In linguistic studies, the reasons given for including extra information⁴¹ in a response to a *yes-no* question can be categorized as:

- to provide implicitly requested information,
- to provide an explanation for an unexpected answer,
- to qualify a direct answer, or
- politeness-related.

5.1.1 Implicitly Requested Information. As mentioned in Section 2, Stenström claims that the typical reason for providing extra information is to answer an implicit *wh*-question. Kiefer (1980) observes that several types of *yes-no* questions, when used to perform indirect speech acts, have the property that one or both of the “binary” answers (i.e., *yes* or *no*) used alone is an inappropriate response to them. For example, in response to (22)i,⁴² when interpreted as (22)ii, an answer of (22)iii or (22)v⁴³ would be appropriate, but not (22)iv alone.

- (22) i. Q: Is John leaving for Stockholm TOMORROW?
 ii. Q: When is John leaving for Stockholm?
 iii. R: Yes.
 iv. R: No.
 v. R: No, John is going to leave the day after tomorrow.

Kiefer also provides examples of cases where the other binary answer alone is inappropriate, or where either alone is inappropriate.

Clark (1979) studied how different factors may influence the responder’s confidence that the literal meaning of a question was intended and confidence that a particular indirect meaning was intended. In one experiment, in which subjects responded to the question, *Do you accept credit cards?*, about half of the subjects provided information answering an indirect request of *What credit cards do you accept?* Clark speculates that the half who included information addressing the indirect request in their response had some, but not necessarily total, confidence that it was intended.

According to Levinson (1983), a *yes-no* question often may be interpreted as a **prerequisite** for another request, i.e., it may be used in the first position of the sequence

⁴¹ We are reporting only cases where the extra information may be used as an indirect answer.

⁴² In this example, Kiefer’s (11b), we follow Kiefer’s use of capitalization to indicate that *tomorrow* would be stressed in spoken English.

⁴³ Kiefer’s (21b).

$T1$ – $T4$, where the occurrence of $T3$ and $T4$ are conditional upon R 's answer in $T2$:

- $T1$: Q makes a prerequisite to determine if a precondition of an action to be requested by Q in $T3$ holds.
- $T2$: R gives an answer indicating whether the precondition holds
- $T3$: Q makes the request
- $T4$: R responds to the request in $T3$.

Levinson claims that prerequisites are used to check whether the planned request (in $T3$) is likely to succeed so that a dispreferred response to it can be avoided by Q . Another reason is that, since receiving an offer is preferred to making a request (Schegloff 1979), by making a prerequisite, Q gives R the opportunity to offer whatever Q would request in $T3$, i.e., the sequence would then consist of just $T1$ and $T4$. In analyses based on speech act theory, in a sequence consisting of just $T1$ and $T4$, the prerequisite would be referred to as an indirect speech act.

5.1.2 Explanation for Unexpected Answer. Stenström notes that a reason for providing extra information is to provide an explanation justifying a negative answer.⁴⁴ According to Levinson (1983), the presence of an explanation is a distinguishing feature of dispreferred responses to questions and other second parts of **adjacency pairs** (Schegloff 1972). In an adjacency pair, each member of the pair is produced by a different speaker, and the occurrence of the first part creates the expectation that the second part will appear, although not necessarily immediately following the first member. Levinson claims that dispreferred responses to first parts of adjacency pairs can be identified by structural features such as:

- use of pauses or displacement,
- prefacing with markers (e.g. *uh* or *well*), appreciations, apologies, or refusals,
- providing explanations, and
- declinations given in an indirect or mitigated manner.

For example in (23),⁴⁵ the marker *well* is used and an explanation is given.

- (23) i. Q : What about coming here on the way or doesn't that give you enough time?
 ii. R : Well no I'm supervising here

Although Levinson defines preference in terms of structural features, he notes that there is a correlation between preference and content. For example, unexpected answers to questions, refusals of requests and offers, and admissions of blame are typically marked with features from the above list.

⁴⁴ She found that 61% of negative direct answers but only 24% of positive direct answers were accompanied by qualify acts.

⁴⁵ Levinson's example (55).

5.1.3 Avoid Misunderstanding. Stenström notes that extra information may be given to qualify an answer. Hirschberg (1985) claims that speakers may give indirect answers to block potential unintended scalar implicatures of a *yes* or *no* alone. For example in (2), repeated below as (24), R's response is preferable to just *no*, since that would license the incorrect scalar implicature that R had not received any of the letters in question. However, by use of (24)ii in an appropriate discourse context, R is able to convey explicitly which letter has been received as well as to conversationally implicate that R has not gotten the other letters in question.

- (24) i. Q: Have you gotten the letters yet?
 ii. R: I've gotten the letter from X.

5.1.4 Politeness. Stenström claims that extra information may be given for social reasons. Kiefer notes that extra information may be given as an excuse when the answer indicates that the speaker has failed to fulfill a social obligation. Brown and Levinson (1978) claim that politeness strategies, which may at times conflict with Gricean maxims, account for many uses of language. According to Brown and Levinson, certain communicative acts are intrinsically **face-threatening acts** (FTAs). That is, doing an FTA is likely to injure some conversational participant's **face**, or public self-image. For example, orders and requests threaten the recipient's **negative face**, "the want . . . that his actions be unimpeded by others" (p. 67). On the other hand, disagreement or bearing "bad news" threatens the speaker's **positive face**, the want to be looked upon favorably by others. Further, they claim that politeness strategies can be ranked, and that the greater the threat associated with a face-threatening act, the more motivated a speaker is to use a higher-numbered strategy.

Brown and Levinson propose the following ranked set of strategies (listed in order from lower to higher rank):

1. perform the FTA. (Brown and Levinson claim that this amounts to following Gricean maxims.)
2. perform the FTA with **redressive action**, i.e., in a manner that indicates that no face threat is intended, using **positive politeness** strategies (strategies that increase the hearer's positive face). Such strategies include:
 - Strategy 1:* attending to the hearer's interests or needs
 - Strategy 6:* avoiding disagreement, e.g., by displacing an answer
3. perform the FTA with redressive action, using **negative politeness** strategies (strategies for increasing negative face). These include:
 - Strategy 6:* giving an excuse or an apology
4. perform the FTA **off-record**, i.e., by use of conversational implicature.

In the next section, we provide several stimulus conditions that reflect positive politeness strategy 1 and negative politeness strategy 6. However, although politeness considerations may motivate a speaker to convey an answer indirectly, it is beyond the scope of our generation model to choose between a direct and an indirect answer on this basis.

Table 3
Stimulus conditions of discourse plan operators.

Operator	Stimulus Conditions
Use-cause	explanation-indicated excuse-indicated
Use-condition	clarify-condition-indicated
Use-contrast	appeasement-indicated answer-ref-indicated clarify-extent-indicated substitute-indicated
Use-elaboration	answer-ref-indicated clarify-concept-indicated
Use-obstacle	explanation-indicated excuse-indicated
Use-otherwise	explanation-indicated excuse-indicated
Use-possible-cause	explanation-indicated excuse-indicated
Use-possible-obstacle	explanation-indicated excuse-indicated
Use-result	explanation-indicated
Use-usually	explanation-indicated

5.2 Stimulus Conditions

In this section we provide glosses of rules giving sufficient conditions for the stimulus conditions used in our model. (The rules are encoded as Horn clauses in our implementation of the model.) Table 3 summarizes which stimulus conditions appear in which discourse plan operators. As mentioned above, for an instance of a satellite operator to be selected during generation, all of its applicability conditions and at least one of its stimulus conditions must hold.

5.2.1 Explanation-indicated. This stimulus condition appears in all of the operators for providing causal explanations. For example in (1), repeated below as (25), R gives an explanation of why R won't get a car.

- (25) i. Q: Actually you'll probably get a car won't you as soon as you get there?
- ii. R: [No.]
- iii. I can't drive.

R's response may contribute to greater dialogue efficiency by anticipating a follow-up request for an explanation. The rule for this stimulus condition may be glossed as: (a speaker) *s* is motivated to give (a hearer) *h* an explanation for (*not p*), if *s* suspects that *h* suspects that (a proposition) *p* is true, unless it is the case that *s* has reason to believe that *h* will accept (*not p*) on *s*'s authority. *s* may acquire the suspicion that *h* suspects that *p* is true by means of syntactic clues from the *yes-no* question, e.g., from the form of the question in (25)i.

5.2.2 Excuse-indicated. Although this stimulus condition appears in some of the same causal operators as *explanation-indicated*, it represents a different kind of motivation.

A *yes-no* question may be interpreted as a prerequisite. Thus, a negative answer to a *yes-no* question used as a prerequisite may be interpreted as a refusal. To soften the refusal, i.e., in accordance with negative politeness strategy 6, the speaker may give an explanation of the negative answer, as illustrated in (21), repeated below in (26).

- (26) i. Q: I need a ride to the mall.
 ii. Are you going shopping tonight?
 iii. R: [No.]
 iv. My car's not running.
 v. The timing belt is broken.

The rule for this stimulus condition may be glossed as: *s* is motivated to give *h* an excuse for (*not p*), if *s* suspects that *h*'s request, (*informifshp*), is a prerequisite. Techniques for interpreting indirect speech acts (Perrault and Allen 1980; Hinkelman 1989) can be used to determine whether the rule's antecedent holds.

5.2.3 Answer-ref-indicated. This condition appears in *Use-elaboration*, illustrated by (27),⁴⁶ and in *Use-contrast*, illustrated by (28).⁴⁷

- (27) i. Q: Did you have a hotel in mind?
 ii. [What hotel did you have in mind?]
 iii. R: [Yes.]
 iv. There's a Holiday Inn right near where I'm working.
- (28) i. Q: You're on that?
 ii. [Who's on that?]
 iii. R: No no no.
 iv. Dave is.

In (27), R has interpreted the question in (27)i as a prerequisite for the *wh*-question shown in (27)ii. Thus, (27)iv not only answers the question in (27)i but also the anticipated *wh*-question in (27)ii. Similarly in (28), R may interpret the question in (28)i as a prerequisite for the *wh*-question in (28)ii, and so gives (28)iv to provide an answer to both (28)i and (28)ii. The rule for this stimulus condition may be glossed as: *s* is motivated to provide *h* with *q*, if *s* suspects that *h* wants to know the referent of a term *t* in *q*. As in *excuse-indicated*, techniques for interpreting indirect speech acts can be used to determine if the rule's antecedent holds.⁴⁸

46 From SRI Tapes (1992), tape 1.

47 Stenström's (102). A *no* answer may be conversationally implicated by use of (28)iv alone.

48 However, following Clark (1979), the rule does not require that R be certain that Q was making an indirect request.

5.2.4 Substitute-indicated. This condition appears in *Use-contrast*, illustrated by (29).

- (29) i. Q: Do you have Verdi's *Otello* or *Aida*?
 ii. R: [No.]
 iii. We have *Rigoletto*.

Although Q may not have intended to use (29)i as a prerequisite for the question *What Verdi operas do you have?*, R suspects that the answer to this *wh*-question might be helpful to Q, and so provides it (in accordance with positive politeness strategy 1). The rule for this stimulus condition may be glossed as: *s* is motivated to provide *h* with *q*, if *s* suspects that it would be *helpful* for *h* to know the referent of a term *t* in *q*. The rule's antecedent would hold whenever obstacle detection techniques (Allen and Perrault 1980) determine that *h*'s not knowing the referent of *t* is an obstacle to an inferred plan of *h*'s. However, not all helpful responses, in the sense described in Allen and Perrault (1980), can be used as indirect answers. For example, even if the clerk (R) at the music store believes that Q's not knowing the closing time could be an obstacle to Q's buying a recording, a response of (30) alone would not convey *no* since it cannot be coherently related to an *Answer-no* plan.

- (30) i. R: We close at 5:30 tonight.

5.2.5 Clarify-concept-indicated. This stimulus condition appears in *Use-elaboration*, as illustrated in (31).⁴⁹

- (31) i. Q: Do you have a pet?
 ii. R: We have a turtle.

In (31), R was motivated to elaborate on the type of pet R has since turtles are not prototypical pets. The rule for this stimulus condition may be glossed as: *s* is motivated to clarify *p* to *h* with *q*, if *p* contains a concept *c*, and *q* provides an atypical instance of *c*. Stereotypical knowledge would be used to evaluate the rule's antecedent.

5.2.6 Clarify-condition-indicated. This stimulus condition appears in the operator *Use-condition*, as illustrated by (32).⁵⁰

- (32) i. Q: Um let me can I make the reservation and change it by tomorrow?
 ii. R: [Yes.]
 iii. If it's still available.

In (32), a truthful *yes* answer depends on the truth of (32)iii. The rules for this stimulus condition may be glossed as: *s* is motivated to clarify a condition *q* for *p* to *h* if 1) *s* doesn't know if *q* holds, or 2) *s* suspects that *q* does not hold.

49 Example (177) from Hirschberg (1985).

50 From SRI Tapes (1992), tape 10ab.

5.2.7 Clarify-extent-indicated. This stimulus condition appears in *Use-contrast*, as illustrated by (2), repeated below as (33).

- (33) i. Q: Have you gotten the letters yet?
 ii. R: I've gotten the letter from X.

On the strict interpretation of (33)i, Q is asking whether R has gotten *all* of the letters, but on a looser interpretation, Q is asking if R has gotten *any* of the letters. Then, if R has gotten some but not all of the letters, just *yes* would be untruthful. However, if Q is speaking loosely, then just *no* might lead Q to erroneously conclude that R has not gotten any of the letters. R's answer circumvents this problem, by conveying the extent to which the questioned proposition (on the strict interpretation) is true.

The rules for this stimulus condition may be glossed as: *s* is motivated to clarify to *h* the extent *q* to which *p* is true, or the alternative *q* to *p* which is true, if *s* suspects that *h* does not know if *q* holds, and *s* believes that *q* is the highest **expression alternative** to *p* that does hold. According to Hirschberg (1985) (following Gazdar), sentences *p_i* and *p_j* (representing the propositional content of two utterances) are expression alternatives if they are the same except for having comparable components *e_i* and *e_j*, respectively. As mentioned earlier, Hirschberg claims that in a discourse context there may be a partial ordering of values that the discourse participants mutually believe to be salient. She claims that the ranking of *e_i* and *e_j* in this ordering can be used to describe the ranking of *p_i* and *p_j*. In the above example, (33)ii is a realization of the highest true expression alternative to the questioned proposition, *p*, i.e., the proposition that R has gotten all the letters.⁵¹

5.2.8 Appeasement-indicated. This stimulus condition appears in *Use-contrast*, as illustrated by (34).⁵²

- (34) i. Q: Did you manage to read that section I gave you?
 ii. R: I've read the first couple of pages.

In (34), R conveys that there is some (though not much) truth to the questioned proposition in an effort to soften his answer (in accordance with positive politeness strategy 1). More than one stimulus condition may motivate R to include the same satellite. For example, in (34), R may have been motivated also by *clarify-extent-indicated*, which was described above. However, it is possible to provide a context for (35) where *appeasement-indicated* holds but *clarify-extent-indicated* does not, or a context for (34) where the converse is true.

- (35) i. Q: Did you wash the dishes?
 ii. R: I brought you some flowers.

The rules for this stimulus condition may be glossed as: *s* is motivated to appease *h* with *q* for *p* not holding or only being partly true, if *s* suspects that (*not p*) is undesirable

⁵¹ Recall that additional constraints on *p* and *q* arise from the applicability conditions of operators containing this stimulus condition, namely *Use-contrast* in this case. Thus, another constraint is that it is plausible that *cr-contrast* holds. The coherence rule for *cr-contrast* was described in Section 4.3.

⁵² Example (56) from Hirschberg (1985).

Table 4
General principles underlying stimulus conditions.

-
1. Efficiency: explanation-indicated, answer-ref-indicated
 2. Accuracy: clarify-concept-indicated, clarify-extent-indicated, clarify-condition-indicated
 3. Politeness: excuse-indicated, appeasement-indicated, substitute-indicated
-

to h but that q is desirable to h . The antecedents to this rule would be evaluated using heuristic rules and stereotypical and specific knowledge about h 's desires. For example, two heuristics of rational agency that might lead to beliefs about h 's desires are 1) if an agent wants you to perform an action A , then your failure to perform A may be undesirable to the agent, and 2) if an agent wants you to do A , then it is desirable to the agent that you perform a part of A .

5.2.9 Summary. In summary, the stimulus conditions in our model can be classified according to three general principles, as shown in Table 4. The first category, efficiency, includes the motivation to provide implicitly requested information as well as to provide an explanation for unexpected information. In other words, giving this type of extra information contributes to the efficiency of the conversation by eliminating the need for follow-up *wh*-questions or follow-up *why?* or *why not?* questions, respectively. In the category of accuracy, in addition to the reason cited by Hirschberg (which is represented in our model as *clarify-extent-indicated*), we have identified two other reasons for giving extra information, which contribute to accuracy. The category of politeness includes reasons for redressing face-threatening acts using positive and negative politeness strategies.

5.3 Generation Algorithm

The inputs to generation consist of:

- the set of discourse plan operators (described in Section 3) augmented with stimulus conditions,
- the set of coherence rules (also described in Section 3),
- the set of stimulus condition rules (described in Section 5.2),
- R 's beliefs (including the discourse expectation that R will provide an answer to some questioned proposition p), and
- the semantic representation of p .

The model presupposes that when answer generation begins, the speaker's (R 's) only goal is to satisfy the above discourse expectation. R 's nonshared beliefs (including beliefs whose strength is not necessarily certainty) about Q 's beliefs, intentions, and preferences are used in generation to evaluate whether a stimulus condition holds. The output of the generation algorithm is a discourse plan that can be realized by a tactical generation component (McKeown 1985).⁵³

⁵³ The plan that is output specifies an ordering of discourse acts based upon the ordering of coherence relations specified in the discourse plan operators. However, reordering may be required, e.g., to model a speaker who has multiple goals.

The answer generation algorithm has two phases. In the first phase, content planning, the generator creates a discourse plan for a full answer, i.e., a direct answer and extra appropriate information. In the second phase, plan pruning, the generator determines which propositions of the planned full answer do not need to be explicitly stated. For example, given an appropriate model of R's beliefs, the system would generate a plan for asserting only the proposition conveyed in (36)v and (36)vi as an answer to (36)i.⁵⁴

- (36) i. Q: Is Mark here [at the office]?
 ii. R: [No.]
 iii. [He's at home.]
 iv. [He is caring for his daughter.]
 v. His daughter has the measles,
 vi. but he's logged on.

An advantage of this approach is that, even when it is not possible to omit the direct answer, a full answer is generated.

5.3.1 Content Planning. Content planning is performed by top-down expansion of an answer discourse plan operator. First, each top-level answer discourse plan operator is instantiated with the questioned proposition until one is found such that its applicability conditions hold.⁵⁵ Next, the satellites of this operator are expanded (recursively). The algorithm for expanding a satellite adds each instance of a satellite such that all of its applicability conditions and at least one of its stimulus conditions hold. Thus, different instantiations of the same type of satellite may be included in a plan for different reasons. For example, (36)iii and (36)vi both realize *Use-contrast* satellites, the former included due to the *answer-ref-indicated* stimulus condition, and the latter due to the *substitute-indicated* stimulus condition.

For each stimulus condition of a satellite, our implementation of the algorithm uses a theorem prover to search the set of R's beliefs (encoded as Horn clauses) for a proposition satisfying a formula consisting of a conjunction of the applicability conditions and that stimulus condition. A proposition satisfying each such formula is used to instantiate the existential variable of the satellite operator. For example, to generate the response in (36), the following formula would be constructed from the *Use-contrast* operator's applicability conditions and one of its stimulus conditions, (*answer-ref-indicated*), where p is the proposition that Mark is at the office, and $?q$ is the existential variable to be instantiated.

```
((and (bel s (cr-contrast ?q (not p)))
      (Plausible (cr-contrast ?q (not p)))
      (answer-ref-indicated s h ?q)))
```

The result of the search is to instantiate $?q$ with the proposition that Mark is at home, due to the speaker's belief that the hearer might have been using (36)i as a prerequisite

⁵⁴ Parts (36)i–(36)v were overheard by one of the authors in a naturally occurring dialogue, and (36)vi was added for expository purposes.

⁵⁵ It is assumed that exactly one is appropriate.

for the question, *Where is Mark?* For a more complete description of how R's response in (36) is generated, see Section 5.4.

We have employed a simple approach to planning because the focus of our research was on the use of the response as an indirect answer, i.e., on aspects of response generation that play a role in interpreting the implicature. In a more sophisticated discourse planning formalism, such as argued for in Young, Moore, and Pollack (1994), it would be possible to represent and reason about other intended effects of the response. (In our model, the effects or primary goals are used in interpretation but their only role in generation is in simulated interpretation. However, their role in interpretation is important; they constrain what discourse plans can be ascribed to the speaker.) While we believe that use of more sophisticated planning formalisms is well motivated for discourse generation in general, we leave the problem of generating indirect answers in such formalisms for future research. The use of stimulus conditions to motivate the selection of optional satellite operators is sufficient for our current goals.

5.3.2 Plan Pruning. The output of the content planning phase, an expanded discourse plan representing a full answer, is the input to the plan pruning phase of generation. The expanded plan is represented as a tree of discourse acts. The goal of this phase is to make the response more concise, i.e., to determine which of the planned acts can be omitted while still allowing Q to infer the full discourse plan.⁵⁶ To do this, the generator considers each of the acts in the frontier of the tree from right to left. (This ensures that a satellite is considered before its related nucleus.) The generator creates a trial plan consisting of the original plan minus the nodes pruned so far and minus the current node. Then, using the answer recognizer, the generator simulates Q's interpretation of a response containing the information that would be given explicitly according to the trial plan. In the simulation, R's shared beliefs are used to model Q's shared beliefs. If Q could infer the full plan (as the most preferred interpretation), then the current node can be pruned. Otherwise, it is left in the plan and the next node is considered.

For example, consider Figure 11 as we illustrate the possible effect of pruning on a full discourse plan. The leaf nodes, representing discourse acts, are numbered 1 through 8. Arcs labeled *N* and *S* lead to a nucleus or satellite, respectively. Node 8 corresponds to the direct answer. Plan pruning would process the nodes in order from 1 to 8. The maximal set of nodes that could be pruned in Figure 11 is the set containing 2, 3, 4, 7, and 8. That is, nodes 2 through 4 might be inferable from 1, node 7 from 5 or 6, and node 8 from 4 or 7, but nodes 1, 5, and 6 cannot be pruned since they are not inferable from other nodes.⁵⁷ In the event that it is determined that no node can be pruned, the full plan would be output. The interpretation algorithm (described in Section 4) would use hypothesis generation to recognize missing propositions other than the direct answer, i.e., the propositions at nodes 2, 3, 4, and 7.

To comment on the role of interpretation in generation, it is a key component of our claim to have provided an adequate model of conversational implicature. Given the shared discourse expectation that R will provide an answer to Q's *yes-no* question, Q's use of (Q's) shared beliefs to interpret the response, and Q's belief that R expects that Q will be able to recognize the answer, Q's recognition of a discourse plan for an answer warrants Q's belief that R intended for Q to recognize this intention. In particular, R would not have pruned the direct answer unless, given the beliefs that R presumes

⁵⁶ Conciseness is not the only possible motive for omitting the direct answer. As mentioned earlier, an indirect answer may be used to avoid performing a face-threatening act. However, it is beyond the scope of our model to determine whether to omit the direct answer on grounds of politeness.

⁵⁷ In fact, leaves that have no satellites of their own cannot be pruned.

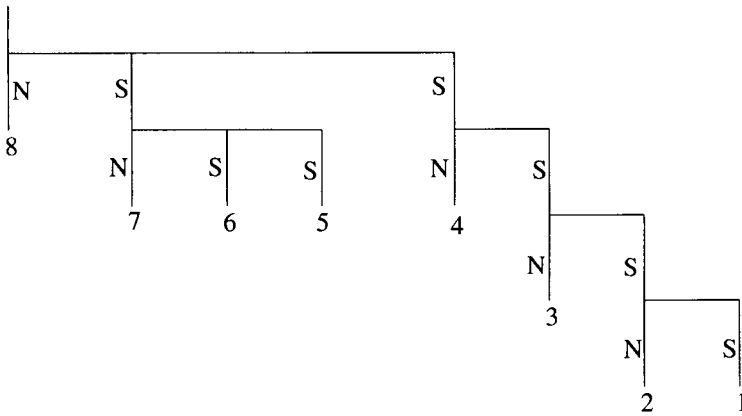


Figure 11
Example of full discourse plan before pruning.

- Q: Is Mark here?
 R: [No.]
 a. [He's at home.]
 b. [He is caring for his daughter.]
 c. His daughter has the measles.
 d. but he is logged on.

Figure 12
Generation example: exchange.

to be shared with Q, R believes that Q will be able to recognize a chain of mutually plausible coherence relations from the actual response to the intended answer, and thus be able to recognize R's plan. Note that although stimulus conditions are not recognized during interpretation in our approach, the model does account for the recognition of those parts of the plan concerning the answer. For example, although Q may not know whether R was motivated by *excuse-indicated* or *explanation-indicated* to provide (21)iv in response to (21)ii, Q can recognize R's intention to convey a *no* by Q's recognition of (21)iv as the nucleus of a *Use-Obstacle* satellite of *Answer-No*. Thus, Q can thereby attribute to R the primary goal of the *Answer-No* plan to convey a *no*.

5.4 Generation Example

This example models R's generation of the response in the exchange shown in Figure 12, which repeats (36). The discourse plan constructed by the algorithm is depicted in Figure 13, where (a) through (d) refer to communicative acts that could be performed by saying the sentences with corresponding labels in Figure 12. Square brackets in the plan indicate acts that have been pruned, i.e., that are not explicitly included in the response.

First, each top-level answer operator is instantiated with the questioned proposition, *p*, the proposition that Mark is at the office. An (*Answer-no s h p*) plan would be selected for expansion since its applicability conditions can be proven. To expand this plan, the algorithm attempts to expand each of its satellites as described in Section 5.3.1. The generation algorithm searches for (at most) one instance of a satellite for each possible motivation of a satellite. In this example, two satellites of (*Use-contrast s h (not p)*) are added to the plan. In one, motivated by the *Answer-ref* stimulus condi-

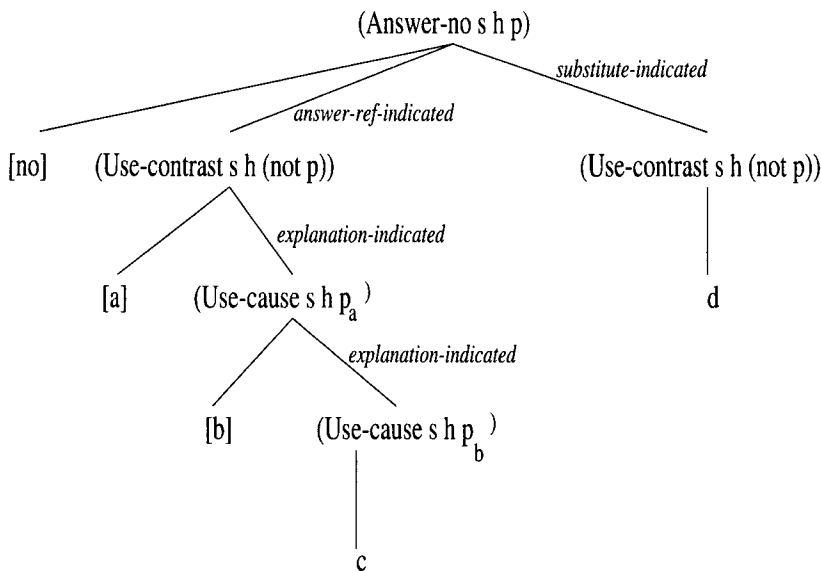


Figure 13
Final plan.

tion, the existential proposition is instantiated with p_a , the proposition that Mark is at home. In other words, proposition, p_a was found to satisfy $?q$ in formula 1 below.

1. ((and (bel s (cr-contrast ?q (not p)))
(Plausible (cr-contrast ?q (not p)))
(answer-ref-indicated s h ?q)))

In the other (*Use-contrast s h (not p)*) satellite, motivated by the *substitute-indicated* stimulus condition, the existential proposition is instantiated with p_a , the proposition that Mark is logged on. That is, p_a was found to satisfy $?q$ in formula 2 below.

2. ((and (bel s (cr-contrast ?q (not p)))
(Plausible (cr-contrast ?q (not p)))
(substitute-indicated s h ?q)))

The former (*Use-contrast s h (not p)*) satellite (i.e., the one constructed using p_a) can be expanded by adding a (*Use-cause s h p_a*) satellite to it. This satellite's existential variable is instantiated with p_b , the proposition that Mark is caring for his daughter, which was found to satisfy $?q$ in formula 3 below.

3. ((and (bel s (cr-cause ?q p_a))
(Plausible (cr-cause ?q p_a))
(explanation-indicated s h p_a)))

Finally, this satellite is expanded using p_c , the proposition that Mark's daughter has the measles, which was found to satisfy $?q$ in formula 4 below.

4. ((and (bel s (cr-cause ?q p_b))
(Plausible (cr-cause ?q p_b))
(explanation-indicated s h p_b)))

The output of phase one is a discourse plan for a full answer, as shown in Figure 14. The second phase of generation, plan pruning, will walk the tree bottom-up. The root

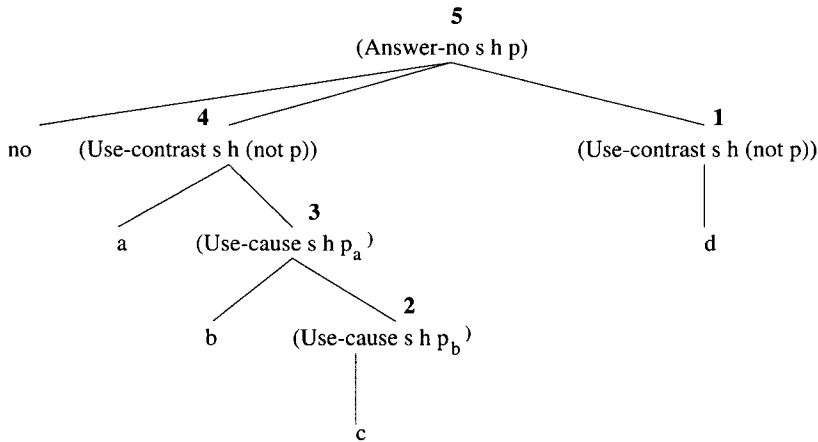


Figure 14
Plan for full answer (before pruning).

of each subtree has been annotated with a sequence number to show the order in which a subtree is visited in the bottom-up traversal of the tree, i.e., 1 through 5. Since subtree 1 has no satellites (only a nucleus), the traversal moves to subtree 2. For the same reason, the traversal moves to subtree 3. Next, the nucleus of subtree 3 is tentatively pruned, i.e., a trial response consisting of the direct answer plus (a), (c), and (d) is created. Simulated interpretation of this trial response results in the inference of a discourse plan identical to the full plan as the most preferred (in fact, the only) interpretation of the trial response. Thus, (b) can be pruned, and subtree 4 is considered next. By a similar process, (a) is also pruned. Last, the tree with root labeled 5 is examined, and it is determined that the direct answer (*no*) can also be pruned. The final result of the traversal is that the direct answer, (a), and (b) are marked as pruned, and a response consisting of just acts (c) and (d) is returned by the generator.

5.5 Related Work in Generation

This work differs from most previous work in cooperative response generation in that the information given in an indirect answer conversationally implicates the direct answer. Hirschberg (1985) implemented a system that determines whether a *yes* or *no* alone licenses any unwanted scalar implicatures, and if so, proposes alternative true scalar responses that do not. In our model, that type of response is generated by constructing a response from an *Answer-no* or *Answer-hedge* operator having a single *Use-contrast* satellite, motivated by *clarify-extent-indicated*, as illustrated in Section 5.2.7.⁵⁸ However, Hirschberg's model does not account for other types of indirect answers, which can be constructed using the other operators (or other combinations of the above operators) in our model, nor for other motives for selecting *Use-contrast* such as *answer-ref-indicated* and *appeasement-indicated*.

Rhetorical or coherence relations (Grimes 1975; Halliday 1976; Mann and Thompson 1988) have been used in several text-generation systems to aid in ordering parts of a text (e.g., Hovy 1988) as well as in content planning (e.g., McKeown 1985; Moore and Paris 1993). The discourse plan operators based on coherence relations in our model

⁵⁸ As mentioned earlier, the coherence rules for *cr-contrast* as well as the rules for *clarify-extent-indicated* make use of notions elucidated by Hirschberg (1985).

(i.e., the operators used as satellites of top-level operators) play a similar role in content planning. However, none of the above approaches model the speaker's motivation for selecting optional satellites. Stimulus conditions provide principled discourse-level knowledge (based upon principles of efficiency, accuracy, and politeness) for choice of an appropriate discourse strategy. Also, stimulus conditions enable content selection to be sensitive not only to the current discourse context, but also to the anticipated effect of a part of the planned response. Finally, none of the above systems incorporate a model of discourse plan recognition into the generation process, which enables indirect answers to be generated in our model.

Moore and Pollack (1992) show the need to distinguish the intentional and informational structure of discourse, where the latter is characterized by the sort of relations classified as subject-matter relations in RST. In our model, the operators used as satellites of top-level answer discourse plan operators are based on relations similar to RST's subject-matter relations. The primary goals of these operators are similar to the effect fields of the corresponding RST relation definitions. However, our model does distinguish the two types of knowledge. In our model stimulus conditions reflect, though they do not directly encode, communicative subgoals leading to the adoption of informational subgoals. For example, the *explanation-indicated* stimulus condition may be triggered in situations when the responder's communicative subgoal would lead R to select a *Use-cause* satellite of *Answer-yes*, rather than a *Use-elaboration* satellite.

Moore and Paris (1993) argue that it is necessary for generation systems to represent not only the speaker's top-level goal, but also the communicative subgoals that a speaker hoped to achieve by use of an informational relation so that, if that subgoal is not achieved, then an alternative rhetorical means can be tried. Although stimulus conditions do reflect the speaker's motivation for including satellites in a plan, it was beyond the scope of our work to address the problem of failure to achieve a subgoal of the original response. Therefore, our system does not record which stimulus condition motivated a satellite; if the stimulus condition was recorded in the final plan then our system would have access to information about the speaker's motivation for the satellite. In our current approach, if a follow-up question is asked then a response to the follow-up question is planned independently of the previous response. However, if R's beliefs have changed since the original question was asked by Q (e.g., as a result of information about Q's beliefs obtained from Q's follow-up question), then it is possible in our approach for R's response to contain different information. Furthermore, in our approach the *original* response may provide the information that a questioner would have to elicit by follow-up questions in a system that can provide only direct answers.

Finally, our use of interpretation during plan pruning has precursors in previous work. In Horacek's approach to generating concise explanations (Horacek 1991), a set of propositions representing the full explanation is pruned by eliminating propositions that can be derived from the remaining ones by a set of contextual rules. Jameson and Wahlster (1982) use an anticipation feedback loop algorithm to generate elliptical utterances.

6. Implementation and Evaluation

We have implemented a prototype of the model in Common LISP. The implemented system can interpret and generate the types of examples discussed in Sections 4 and 5 and the specific examples tested in the experiments described below. The overall coverage of the implemented system can be defined as all (direct and indirect) responses that can be composed from the 5 top-level operators and 10 satellite operators




Figure 15
Blocks world picture used in Experiment 1.

(for 8 stimulus conditions and 24 coherence relation rules) provided in the model. The performance of the system running on a UNIX workstation depends mainly on the amount of hypothesis generation performed (which can be controlled by setting a parameter limiting the depth of the breadth-first search during hypothesis generation).

We have evaluated the system with two experiments. The purpose of the first experiment was to determine whether users' interpretations of indirect answers would agree with the system's interpretations. The purpose of the second experiment was to see how users would evaluate the unrequested information selected for an indirect answer by the system. The system was run to verify that it could actually interpret or generate the responses that were evaluated in the first and second experiments, respectively. Each experiment was conducted by means of a questionnaire given to 10 adult subjects who were not familiar with this research work. At the beginning of each questionnaire, subjects were given a brief textual and pictorial description of the setting in which the questions and responses supposedly had occurred. (A black-and-white version of the picture shown to subjects for the first experiment is given in Figure 15. Since the picture used in the experiment was in color, we have annotated the objects in the figure to indicate their color. A similar picture was used for the second experiment.) The fictional setting was described as a laboratory inhabited by a talking robot and a mouse; outside of the laboratory is a manager who cannot see inside the laboratory, but the robot and manager can communicate with each other. In both experiments, the questionnaire consisted of questions supposedly posed by the manager to the robot, and the robot's possible or actual responses. This setting was selected because it could easily be presented to subjects with a minimum of description about the beliefs that might motivate the robot's responses. Also, a new domain was used so that our experience with the other domains would not influence the evaluation.

6.1 Experiment 1

6.1.1 Experiment. The first experiment addressed whether the subjects' interpretations of indirect answers would agree with the system's interpretations. The subjects were given 19 *yes-no* question-response exchanges. Each response consisted of from 1 to 3 sentences without an explicit *yes* or *no*, e.g., as in (37) (item 3 in the questionnaire for Experiment 1).

- (37) i. Q: Is the yellow ball on the table?
 ii. R: The yellow ball is on the floor.

(For more examples, see the appendix.) Fourteen of the responses were indirect, i.e., our system would interpret them as generated from *Answer-No* or *Answer-Yes*.⁵⁹ (For example, (37)ii was interpreted as a *no* generated by *Answer-No*.) These 14 responses made use of all of the possible satellites of *Answer-No* and *Answer-Yes* in the model. Several responses made use of multiple satellites. For example, the response in item 19 of the questionnaire was similar to the response shown on the right-hand side of Figure 8. The other 5 responses we characterize as **bogus**, i.e., would not be interpreted as answers by our system, e.g., (38) (item 2 in the questionnaire).

- (38) i. Q: Can you pick up the ball?
 ii. R: A red block is on the table.

The purpose of the so-called bogus responses was to make certain that the subjects were not just interpreting every response as saying *yes* or *no*. For each response, the subjects were asked to select one of the following interpretations:

1. *Yes* (glossed as *I would interpret this as yes*)
2. *Yes-?* (glossed as *I could interpret this as yes but I am uncertain*),
3. *No* (glossed as *I would interpret this as no*),
4. *No-?* (glossed as *I could interpret this as no but I am uncertain*), or
5. *Other* (glossed as *I would not interpret this as yes or no*).

6.1.2 Results. The results are shown in Table 5. The rows of the table present the results for each question-response pair. The second column gives the system's interpretation of the response for cases where the system would interpret the response as an indirect *yes* or *no*, or indicates *bogus* for a bogus response. The third column gives the number of subjects who selected *yes* or *no* in agreement with the system's interpretation. The fourth column gives the number of subjects who selected *yes* or *no* in agreement with the system's interpretation but with some uncertainty (i.e., a *Yes-?* or *No-?*). The last column gives the number of subjects who judged the response as *Other*.

In the overwhelming majority of cases the subjects' interpretations agreed with the systems' interpretation as a *yes* or *no*, though occasionally with some uncertainty. None of the subjects selected the opposite interpretation, i.e., a *Yes/Yes-?* for a *no*, or a *No/No-?* for a *yes*. Of the 14 questions interpreted as a *yes* or *no* by the system,

⁵⁹ We have not yet evaluated indirect answers that would be generated by the other three top-level operators in our model.

Table 5
Results of experiment on interpretation of indirect answers.

Example Number	Indirect Answer	Indirect Interpretation	Indirect Interpretation (?)	Other
1	Yes	6	3	1
2	(bogus)	0	1	9
3	No	9	0	1
4	No	5	5	0
5	No	8	2	0
6	Yes	7	3	0
7	(bogus)	0	1	9
8	No	9	1	0
9	Yes	9	1	0
10	No	8	2	0
11	Yes	3	7	0
12	(bogus)	0	0	10
13	Yes	8	2	0
14	No	8	2	0
15	No	7	3	0
16	Yes	4	6	0
17	(bogus)	0	0	10
18	(bogus)	0	0	10
19	Yes	10	0	0

only 2 items were interpreted by subjects as *Other* (each by a different subject). In 28% of the instances where the subjects interpreted a response as saying *yes* or *no*, they noted some degree of uncertainty. During debriefing, the subjects who tended to express uncertainty said that while they might interpret the response as *yes* or *no*, one generally had some uncertainty when the direct answer was omitted. Only one subject interpreted a bogus question as answering *yes* or *no*.

To test the statistical significance of the pattern of responses shown in Table 5, we took a very conservative approach. We grouped *Indirect Interpretation (?)* with *Other* in Table 5 for question-response instances where the system interpreted the response as an indirect answer, and we grouped *Indirect Interpretation (?)* with *Indirect Interpretation* for instances where the system did *not* interpret the response as an indirect answer. Thus *Indirect Interpretation (?)* responses by the subjects were treated as disagreeing with the system's interpretation of the example. We then applied Cochran's Q test (Cochran 1950) to the resulting two columns of data. The result shows that the pattern of responses is statistically significant (not the result of random chance) at better than the level $p < .005$. To determine whether the subjects differentiated between responses that the system interpreted as indirect answers and those that it did not, we applied the Mann Whitney U statistic (Siegel 1956), which showed no score overlap at the level $p < .005$.

6.2 Experiment 2

6.2.1 Experiment. Although the linguistic studies discussed in Section 5 show that in human-human dialogue people often include additional unrequested information in their responses to *yes-no* questions, we conducted a second experiment to determine how users would evaluate responses consisting of the kinds of extra unrequested information produced by our system. The subjects were given 11 *yes-no* questions (some preceded by 1 or 2 sentences to establish some additional context), each with a

set of 4 possible responses. The subjects were told to suppose that all of the responses in a set were true, and were asked to select the best response in each set. For each question, the response choices included:

- a direct response of *Yes* or *No* (depending on the correct answer to the question),
- the direct response with further emphasis (such as *No, I can't*),
- 2 extended responses, containing the direct answer with extra unrequested information,

e.g., (39)iii–(39)vi, respectively. ((39) is item 1 in the questionnaire for Experiment 2.)

- (39) i. M: I am looking for the blue ball.
 ii. Is it on the table?
 iii. R: No.
 iv. R: No. It's not.
 v. R: No. It's on the floor.
 vi. R: No. It cost \$5.

In 9 of the 11 sets, 1 of the extended responses was motivated by our stimulus conditions (e.g., (39)v), and 1 was not (e.g., (39)vi). In the other 2 sets (questionnaire items 3 and 7), neither of the extended responses was motivated by any stimulus condition. The purpose of these 2 so-called bogus examples was to make certain that the subjects were not inclined to always select responses with extra information.

6.2.2 Results. The results are shown in Table 6. The rows of the table present the results for each question. The second column lists the stimulus condition, if any, that our system used to trigger one of the extended responses to the question. Items 3 and 7 contained bogus responses, i.e., none of the responses was motivated by a stimulus condition. The next three columns indicate respectively the number of subjects who selected the response motivated by the listed stimulus condition, the number who selected the direct answer alone or the direct answer with emphasis but no additional information, and the number who selected an extended response not motivated by a stimulus condition. Note that none of the subjects selected a response with extra information for the two bogus questions, indicating that they were not merely inclined to select responses with extra information.

Items 8 and 10 warrant some discussion. Question 8 was problematic. The original question given to the first four subjects asked whether the robot could tell the lab manager the time. The response "*No. There is no clock in here.*" was motivated by the stimulus condition *excuse-indicated*. However, two of the four subjects selected just *No* as the best response, and explained during debriefing that if the robot could tell time, then certainly he had an internal clock that he could use (since all computers have internal clocks) and thus the absence of a clock in the room was not relevant. Since the prior beliefs of these subjects conflicted with the beliefs that were intended as the context for interpreting the robot's response, we altered the question for the remainder of the study to circumvent this problem. In item 10, the extra information in the system's response was motivated by the *appeasement-indicated* stimulus condition. In that

Table 6
Results of experiment on including extra information.

Example Number	Stimulus Condition SC	Response SC	Direct Answer Only	Other Extended Response
1	answer-ref-indicated	10	0	0
2	substitute-indicated	10	0	0
3	none	0	10	0
4	explanation-indicated	9	1	0
5	clarify-extent-indicated	10	0	0
6	clarify-condition-indicated	10	0	0
7	none	0	10	0
8	excuse-indicated	8	2	0
9	clarify-concept-indicated	10	0	0
10	appeasement-indicated	4	5	1
11	explanation-indicated	9	1	0

response, the robot answers *No* (that he has not yet done the requested task) and then attempts to appease the questioner by describing another task that he has completed. Since only 4 of the 10 subjects selected this response, it is possible that the subjects did not view appeasement as an appropriate stimulus condition in human-machine dialogue, despite the fact that it does occur in human-human dialogue. Alternatively, the subjects did not have enough information to recognize the response as attempted appeasement.

To test the statistical significance of the pattern of responses in Table 6, we again took a conservative approach and grouped *Other Extended Response* (which was selected only once by a subject) with *Direct Answer Only* so that it was treated as disagreeing with the system response. Once again we applied Cochran's Q test (Cochran 1950) and the Mann Whitney U statistic (Siegel 1956). Cochran's Q test showed that the pattern of responses in Table 6 is statistically significant at the level $p < .005$, and the Mann Whitney U statistic showed that there is no score overlap at the level $p < .0253$.

6.3 Summary

The first experiment suggests that our system's interpretations of indirect answers agree with the judgments of human interpreters. The second experiment suggests that our stimulus conditions result in the construction of responses containing extra information that users will view favorably. However, we have not addressed the question of when other stylistic considerations might limit the amount of extra information that is included, nor the question of choosing between a direct and an indirect response when both are possible. These issues will be addressed in future research.

7. Conclusions

In summary, we have proposed and implemented a computational model for interpreting and generating indirect answers to *yes-no* questions in English. This paper describes the knowledge and processes provided by the model. Generation and interpretation are treated, respectively, as construction of and recognition of a responder's discourse plan for a full answer. A discourse plan explicitly relates a speaker's beliefs and discourse goals to his program of communicative actions. An indirect answer is the result of the responder providing only part of the planned response, but intending

for his discourse plan to be recognized by the questioner. Discourse plan construction and recognition make use of the beliefs that are presumed to be shared by the participants, as well as shared knowledge of discourse strategies, represented in the model by shared discourse plan operators. In the operators, coherence relations are used to characterize types of satellites that may accompany each type of answer. Recognizing a mutually plausible coherence relation obtaining between the actual response and a possible direct answer plays an important role in recognizing the responder's discourse plan. The use of hypothesis generation in interpretation broadens the coverage of the model to cases where more is missing from a full answer than just the nucleus of a top-level operator. (From the point of view of generation, it enables the construction of a more concise, though no less informative, response.) Stimulus conditions model a speaker's motivation for selecting a satellite. During generation, the speaker uses his own interpretation capability to determine what parts of the plan are inferable by the hearer in the current discourse context and thus do not need to be explicitly given. We argue that because of the role of interpretation in generation, Q's belief that R intended for Q to recognize the answer is warranted by Q's successful recognition of the plan.

Although it was not our goal to develop a cognitive model of how implicatures are produced and comprehended, certain aspects of the model might be incorporated into a cognitive model. To a large extent the model is *recognitional* rather than *inferential*. Of course, we make no claims about the cognitive plausibility of the particular coherence relations and discourse plan operators used in our model, which were encoded solely on the basis of their descriptive and computational utility. We await further cognitive studies on coherence relations as begun in Sanders, Spooren, and Noordman (1992), Knott and Dale (1994), and Knott (1995). Since the work reported in this paper was performed, the first author has investigated the automatic compilation of discourse plan operators in a computational cognitive architecture (SOAR) (Green and Lehman 1998).

In conclusion, our model provides wider coverage than previous computational models for generating and interpreting answers. Specifically, it covers both direct and indirect answers, multiple-sentence responses, a variety of types of indirect answer (i.e., characterized in terms of multiple coherence relations), and multiple types of speaker motivation for deciding to provide extra information (i.e. characterized in terms of different stimulus conditions). In addition, it appears that this approach could be extended to other discourse-expectation-based types of conversational implicature. As a computational model of conversational implicature, it extends current plan-based theories of implicature in several ways. First, it demonstrates the role of shared discourse expectations and pragmatic knowledge. Second, it makes predictions about cancelability in terms of intentional structure of discourse. Lastly, it treats generation as a process drawing upon the speaker's own interpretation mechanism.

Appendix: Questionnaire for Experiment 1

1. Q: Can you make a stack of 3 blocks?
R: I can put the green block on the blue block.
2. Q: Can you pick up the ball?
R: A red block is on the table.
3. Q: Is the yellow ball on the table?
R: The yellow ball is on the floor.

4. Q: Can you pick up the green block?
R: The green block is glued to the table.
I glued it yesterday.
5. Q: Can you build a stack with the green block on top of the cone?
R: The cone does not have a flat top.
6. Q: Are you going to move the green block?
R: The green block is very heavy.
I'm going to use the forklift.
7. Q: Can the mouse climb up the red block?
R: The yellow ball is on the floor.
8. Q: Can you build a stack with the green block on top of the cone?
R: The green block would fall off the cone.
9. Q: Can you pick up the red block?
R: If I move the blue block off of it.
10. Q: Can you pick up the blue block?
R: My arms won't move.
The human forgot to oil me.
11. Q: Is there a yellow ball on the floor?
R: The mouse pushed the yellow ball off the table.
12. Q: Are there any blocks on the table?
R: The table is about three feet tall.
13. Q: Is something on the red block?
R: I put the blue block on the red block.
14. Q: Can you put the mouse on the green block?
R: He runs too fast.
15. Q: Is the blue block on the table surface?
R: I put the blue block on the red block.
16. Q: Did you move the cone off the green block?
R: I wanted to pick up the green block.
17. Q: Is the blue block on the red block?
R: The mouse squeaks a lot.
18. Q: Did you pick up the cone?
R: There are three blocks, one cone, a yellow ball, and a mouse.
19. Q: Are you going to pick up the blue block?
R: The blue block is sticky.
The mouse poured honey on it.
I am going to use a pair of tongs.

Acknowledgments

This describes the first author's dissertation research at the University of Delaware. We would like to thank Dan Chester of the University of Delaware for providing us with an implementation of a Horn clause theorem prover (Chester 1980) to use in this work, and Fred Masterson also of the University of Delaware for help with the statistical analysis of the experiments. Also, we wish to thank the journal referees for their helpful comments.

References

- Allen, James F. and C. Raymond Perrault. 1980. Analyzing intention in utterances. *Artificial Intelligence* 15:143–178.
- Brown, Penelope and Stephen C. Levinson. 1978. Universals in language usage: Politeness phenomena. In Esther N. Goody, editor, *Questions and Politeness: Strategies in Social Interaction*. Cambridge University Press, pages 56–289.
- Carberry, Sandra. 1990. *Plan Recognition in Natural Language Dialogue*. MIT Press, Cambridge, MA.
- Chester, Daniel. 1980. HCPRVR: An interpreter for logic programs. In *Proceedings of the First Annual National Conference on Artificial Intelligence*, pages 93–95.
- Clark, Herbert H. 1979. Responding to indirect speech acts. *Cognitive Psychology* 11:430–477.
- Clark, Herbert and C. Marshall. 1981. Definite reference and mutual knowledge. In Aravind K. Joshi, Bonnie Webber, and Ivan Sag, editors, *Elements of Discourse Understanding*. Cambridge University Press.
- Cochran, W. G. 1950. The comparison of percentages in matched samples. *Biometrika* 37:256–266.
- Dahlgren, Kathleen. 1989. Coherence relation assignment. In *Proceedings of the Eleventh Annual Meeting of the Cognitive Science Society*, pages 588–596.
- Green, Nancy L. 1994. *A Computational Model for Generating and Interpreting Indirect Answers*. Ph.D. thesis, University of Delaware.
- Green, Nancy and Jill F. Lehman. 1998. An application of explanation-based learning to discourse generation and interpretation. In *Papers from the 1998 AAAI Spring Symposium on Applying Machine Learning to Discourse Processing*, pages 33–39.
- Grice, H. Paul. 1975. Logic and conversation. In Peter Cole and Jerry L. Morgan, editors, *Syntax and Semantics III: Speech Acts*. Academic Press, pages 41–58.
- Grimes, Joseph E. 1975. *The Thread of Discourse*. Mouton, The Hague.
- Gunji, Takao. 1981. *Toward a Computational Theory of Pragmatics—Discourse, Presupposition, and Implicature*. Ph.D. thesis, Ohio State University.
- Halliday, M. A. K. and Ruqaiya Hasan. 1976. *Cohesion in English*. Longman, New York.
- Hinkelman, Elizabeth A. 1989. *Linguistic and Pragmatic Constraints on Utterance Interpretation*. Ph.D. thesis, University of Rochester.
- Hirschberg, Julia B. 1985. *A Theory of Scalar Implicature*. Ph.D. thesis, University of Pennsylvania.
- Hobbs, Jerry R. 1978. Resolving pronoun references. *Lingua*, 44:311–338.
- Horacek, Helmut. 1991. Exploiting conversational implicature for generating concise explanations. In *Proceedings of the European Association for Computational Linguistics*, pages 191–193.
- Hovy, Eduard H. 1988. Planning coherent multisentential text. In *Proceedings of the 26th Annual Meeting*, pages 163–169. Association for Computational Linguistics.
- Jameson, Anthony and Wolfgang Wahlster. 1982. User modelling in anaphora generation: Ellipsis and definite description. In *Proceedings of the European Conference on Artificial Intelligence*.
- Kiefer, Ferenc. 1980. Yes-no questions as wh-questions. In John Searle, Ferenc Kiefer, and Manfred Bierwisch, editors, *Speech Act Theory and Pragmatics*. Reidel, Dordrecht, Holland, pages 48–68.
- Knott, Alistair. 1995. *A Data-Driven Methodology for Motivating a Set of Coherence Relations*. Ph.D. thesis, University of Edinburgh.
- Knott, Alistair and Robert Dale. 1994. Using linguistic phenomena to motivate a set of coherence relations. *Discourse Processes*, 35–62.
- Lascarides, Alex and Nicholas Asher. 1991. Discourse relations and defeasible knowledge. In *Proceedings of the 29th Annual Meeting*, pages 55–62. Association for Computational Linguistics.
- Lascarides, Alex, Nicholas Asher, and Jon Oberlander. 1992. Inferring discourse relations in context. In *Proceedings of the 30th Annual Meeting*, pages 1–8. Association for Computational Linguistics.

- Levinson, Stephen C. 1983. *Pragmatics*. Cambridge University Press, Cambridge.
- Litman, Diane J. 1986. Understanding plan ellipsis. In *Proceedings of the Fifth National Conference on Artificial Intelligence*, pages 619–624.
- Mann, William C. and Sandra A. Thompson. 1983. Relational propositions in discourse. Technical Report ISI/RR-83-115, Information Sciences Institute, University of Southern California, Marina del Rey, CA.
- Mann, William C. and Sandra A. Thompson. 1988. Rhetorical Structure Theory: Toward a functional theory of text organization. *Text* 8(3):167–182.
- McCafferty, Andrew S. 1987. *Reasoning about Implicature: A Plan-Based Approach*. Ph.D. thesis, University of Pittsburgh, Pittsburgh.
- McKeown, Kathleen R. 1985. *Text Generation*. Cambridge University Press, Cambridge.
- Moore, Johanna D. and Cecile Paris. 1993. Planning text for advisory dialogues: Capturing intentional and rhetorical information. *Computational Linguistics* 19(4):651–694.
- Moore, Johanna D. and Martha E. Pollack. 1992. A problem for RST: The need for multi-level discourse analysis. *Computational Linguistics* 18(4):537–544.
- Perrault, Raymond and James Allen. 1980. A plan-based analysis of indirect speech acts. *American Journal of Computational Linguistics* 6(3-4):167–182.
- Pollack, Martha. 1990. Plans as complex mental attitudes. In Philip R. Cohen, Jerry Morgan, and Martha Pollack, editors, *Intentions in Communication*. MIT Press, Cambridge, MA.
- Reichman, Rachel. 1985. *Getting Computers To Talk Like You And Me*. MIT Press, Cambridge, MA.
- Sanders, Ted J., Wilbert P. Spooren, and Leo G. Noordman. 1992. Toward a taxonomy of coherence relations. *Discourse Processes* 15:1–35.
- Schegloff, Emanuel A. 1972. Sequencing in conversational openings. In J. J. Gumperz and Dell H. Hymes, editors, *Directions in Sociolinguistics*. Holt, Rinehart and Winston, New York, pages 346–80.
- Schegloff, Emanuel A. 1979. Identification and recognition in telephone conversation openings. In G. Psathas, editor, *Everyday Language: Studies in Ethnomethodology*. Irvington, New York, pages 23–78.
- Siegel, Sidney. 1956. *Nonparametric Statistics for the Behavioral Sciences*. McGraw-Hill, New York.
- SRI Tapes. 1992. Transcripts of audiotape conversations. Prepared by Jacqueline Kowto under the Direction of Patti Price at SRI International, Menlo Park, CA.
- Stenström, Anna-Brita. 1984. Questions and responses in English conversation. In Claes Schaar and Jan Svartvik, editors, *Lund Studies in English* 68. CWK Gleerup, Malmö, Sweden.
- Thomason, Richmond H. 1990. Accommodation, meaning, and implicature: Interdisciplinary foundations for pragmatics. In Philip R. Cohen, Jerry Morgan, and Martha Pollack, editors, *Intentions in Communication*. MIT Press, Cambridge, MA, pages 325–363.
- Walker, Marilyn A. 1993. *Informational Redundancy and Resource Bounds in Dialogue*. Ph.D. thesis, University of Pennsylvania.
- Young, R. Michael, Johanna D. Moore, and Martha E. Pollack. 1994. Towards a principled representation of discourse plans. In *Proceedings of the Sixteenth Annual Meeting of the Cognitive Science Society*, pages 946–951.