CONVERSATIONAL IMPLICATURES IN INDIRECT REPLIES

Nancy Green
Sandra Carberry
Department of Computer and Information Sciences
University of Delaware
Newark, Delaware 19716, USA

email: green@cis.udel.edu, carberry@cis.udel.edu

Abstract[1]

In this paper we present algorithms for the interpretation and generation of a kind of particularized conversational implicature occurring in certain indirect replies. Our algorithms make use of discourse expectations, discourse plans, and discourse relations. The algorithms calculate implicatures of discourse units of one or more sentences. Our approach has several advantages. First, by taking discourse relations into account, it can capture a variety of implicatures not handled before. Second, by treating implicatures of discourse units which may consist of more than one sentence, it avoids the limitations of a sentence-at-a-time approach. Third, by making use of properties of discourse which have been used in models of other discourse phenomena, our approach can be integrated with those models. Also, our model permits the same information to be used both in interpretation and generation.

1 Introduction

In this paper we present algorithms for the interpretation and generation of a certain kind of conversational implicature occurring in the following type of conversational exchange. One participant (Q) makes an illocutionary-level request[2] to be informed if p; the addressee (A), whose reply may consist of more than one sentence, conversationally implicates one of these replies: p, $\neg p$, that there is support for p, or that there is support for $\neg p$. For example, in (1), assuming Q's utterance has been interpreted as a request to be informed if A went shopping, and given certain mutual beliefs (e.g., that A's car breaking down would normally be sufficient to prevent A from going shopping, and

---

[1] We wish to thank Kathy McCoy for her comments on this paper.

[2] i.e., using Austin's (Austin, 1962) distinction between locutionary and illocutionary force, Q's utterance is intended to function as a request (although it need not have the grammatical form of a question)

that A's reply is coherent and cooperative), A's reply is intended to convey, in part, a 'no'.

(1) Q: Did you go shopping?
     A: a. My car's not running.
        b. The timing belt broke.

Such indirect replies satisfy the conditions proposed by Grice and others (Grice, 1975; Hirschberg, 1985; Sadock, 1978) for being classified as particularized conversational implicatures. First, A's reply does not entail (in virtue of its conventional meaning) that A did not go shopping. Second, the putative implicature can be cancelled; for example, it can be denied without the result sounding inconsistent, as can be seen by considering the addition of (2) to the end of A's reply in (1).

(2) A: So I took the bus to the mall.

Third, it is reinforceable; A's reply in (1) could have been preceded by an explicit "no" without destroying coherency or sounding redundant. Fourth, the putative implicature is nondetachable; the same reply would have been conveyed by an alternative realization of (1a) and (1b) (assuming that the alternative did not convey a Manner-based implicature). Fifth, Q and A must mutually believe that, given the assumption that A's reply is cooperative, and given certain shared background information, Q can and will infer that by A's reply, A meant 'no'. This paper presents algorithms for calculating such an inference from an indirect response and for generating an indirect response intended to carry such an inference.

2 Solution

2.1 Overview

Our algorithms are based upon three notions from discourse research: discourse expectations, discourse plans, and implicit relational propositions in discourse.

At certain points in a coherent conversation, the participants share certain expectations (Reichman, 1984; Carberry, 1990) about what kind of utterance is appropriate. In the type of exchange we are studying, at the point after Q's contribution, the participants share the beliefs that Q has requested to be informed if p and that the request was appropriate; hence, they share the discourse expectation that for A to be cooperative, he must now say as much as he can truthfully say in regard to the truth of p. (For convenience, we shall refer to this expectation as Answer-YNQ(p).)

A discourse plan operator[3] (Lambert & Carberry, 1991) is a representation of a normal or conventional way of accomplishing certain communicative goals. Alternatively, a discourse plan operator could be considered as a defeasible rule expressing the typical (intended) effect(s) of a sequence of illocutionary acts in a context in which certain applicability conditions hold. These discourse plan operators are mutually known by the conversational participants, and can be used by a speaker to construct a plan for achieving his communicative goals. We provide a set of discourse plan operators which can be used by A as part of a plan for fulfilling Answer-YNQ(p).

Mann and Thompson (Mann & Thompson, 1983; Mann & Thompson, 1987) have described how the structure of a written text can be analyzed in terms of certain implicit relational propositions that may plausibly be attributed to the writer to preserve the assumption of textual coherency.[4] The role of discourse relations in our approach is motivated by the observation that direct replies may occur as part of a discourse unit conveying a relational proposition. For example, in (3), (b) is provided as the (most salient) obstacle to the action (going shopping) denied by (a);

(3) Q: Did you go shopping?
    A:a. No,
      b. my car's not running.

in (4), as an elaboration of the action (going shopping) conveyed by (a);

(4) Q: Did you go shopping?
    A:a. Yes,
      b. I bought some shoes.

and in (5), as a concession for failing to do the action (washing the dishes) denied by (a).

(5) Q: Did you wash the dishes?
    A:a. No,
      b. (but) I scraped them.

---

[3]in Pollack's terminology, a *recipe-for-action* (Pollack, 1988; Grosz & Sidner, 1988).

[4]Although they did not study dialogue, they suggested that it can be analyzed similarly. Also note that the relational predicates which we define are similar but not necessarily identical to theirs.

Note that given appropriate context, the (b) replies in (3) through (5) would be sufficient to conversationally implicate the corresponding direct replies. This, we claim, is by virtue of the recognition of the relational proposition that would be conveyed by use of the direct reply and the (b) sentences.

Our strategy, then, is to generate/interpret A's contribution using a set of discourse plan operators having the following properties: (1) if the applicability conditions hold, then executing the body would generate a sequence of utterances intended to implicitly convey a relational proposition $R(p,q)$; (2) the applicability conditions include the condition that $R(p,q)$ is plausible in the discourse context; (3) one of the goals is that Q believe that p, where p is the content of the direct reply; and (4) the step of the body which realizes the direct reply can be omitted under certain conditions. Thus, whenever the direct reply is omitted, it is nevertheless implicated as long as the intended relational proposition can be recognized. Note that property (2) requires a judgment that some relational proposition is plausible. Such judgments will be described using defeasible inference rules. The next section describes our discourse relation inference rules and discourse plan operators.

2.2 Discourse Plan Operators and Discourse Relation Inference Rules

A typical reason for the failure of an agent's attempt to achieve a domain goal is that the agent's domain plan encountered an obstacle. Thus, we give the rule in (6) for inferring a plausible discourse relation of Obstacle.[5]

```
(6)
If (i)  coherently-related(A,B), and
   (ii) A is a proposition that an agent
        failed to perform an action of
        act type T, and
   (iii) B is a proposition that
        a) a normal applicability condition
           of T did not hold, or
        b) a normal precondition of T
           failed, or
        c) a normal step of T failed, or
        d) the agent did not want to
           achieve a normal goal of T,
   then plausible(Obstacle(B,A)).
```

In (6) and in the rules to follow, 'coherently-related(A,B)' means that the propositions A and B are assumed to be coherently related in the discourse. The terminology in clause (iii) is that of the extended STRIPS planning formalism (Fikes

---

[5]For simplicity of exposition, (6) and the discourse relation inference rules to follow are stated in terms of the past; we plan to extend their coverage of times.

& Nilsson, 1971; Allen, 1979; Carberry, 1990; Litman & Allen, 1987).

Examples of A and B satisfying each of the conditions in (6.iii) are given in (7a) - (7d), respectively.

(7)   [A] I didn't go shopping.
  a. [B] The stores were closed.
  b. [B] My car wasn't running.
  c. [B] My car broke down on the way.
  d. [B] I didn't want to buy anything.

The discourse plan operator given in (8) describes a standard way of performing a denial (exemplified in (3)) that uses the discourse relation of Obstacle given in (6). In (8), as in (6), A is a proposition that an action of type T was *not* performed.

(8) Deny (with Obstacle)
  Applicability conditions:
    1) S BMB plausible(Obstacle(B,A))
  Body (unordered):
    1) (optional) S inform H that A
    2) Tell(S,H,B)
  Goals:
    1) H believe that A
    2) H believe that Obstacle(B,A)

In (8) (and in the discourse plan operators to follow) the formalism described above is used; 'S' and 'H' denote speaker and hearer, respectively; 'BMB' is the one-sided mutual belief[6] operator (Clark & Marshall, 1981); 'inform' denotes an illocutionary act of informing; 'believe' is Hintikka's (Hintikka, 1962) belief operator; 'Tell(S,H,B)' is a subgoal that can be achieved in a number of ways (to be discussed shortly), including just by S informing H that B; and steps of the body are not ordered. (Note that to use these operators for generation of direct replies, we must provide a method to determine a suitable ordering of the steps. Also, although it is sufficient for interpretation to specify that step 1 is optional, for generation, more information is required to decide whether it can or should be omitted; e.g., it should not be omitted if S believes that H might believe that some relation besides Obstacle is plausible in the context.[7] These are areas which we are currently investigating; for related research, see section 3.)

Next, consider that a speaker may wish to inform the hearer of an aspect of the plan by which she accomplished a goal, if she believes that H may not be aware of that aspect. Thus, we give the rule in (9) for inferring a plausible discourse relation of Elaboration.

----

[6] 'S BMB *p*' is to be read as 'S believes that it is mutually believed between S and H that *p*'.

[7] A related question, which has been studied by others (Joshi, Webber & Weischedel, 1984a; Joshi, Webber & Weischedel, 1984b), is in what situations is a speaker required to supply step 2 to avoid misleading the hearer?

(9)
If (i) coherently-related(A,B), and
  (ii) A is a proposition that an agent
     performed some action of act type
     T, and
  (iii) B is a proposition that describes
      information believed to be new to
      H about
    a) the satisfaction of a normal
       applicability condition of T such
       that its satisfaction is not
       believed likely by H, or
    b) the satisfaction of a normal
       precondition of T such that its
       satisfaction is not believed
       likely by H, or
    c) the success of a normal step of T,
       or
    d) the achievement of a normal goal
       of T,
then plausible(Elaboration(B,A)).

Examples of A and B satisfying each of the conditions in (9.iii) are given in (10a) - (10d), respectively.

(10)   [A] I went shopping today.
  a. [B] I found a store that was open.
  b. [B] I got my car fixed yesterday.
  c. [B] I went to Macy's.
  d. [B] I got running shoes.

The discourse plan operator given in (11) describes a standard way of performing an affirmation (exemplified in (4)) that uses the discourse relation of Elaboration.

(11) Affirm (with Elaboration)
  Applicability conditions:
    1) S BMB plausible(Elaboration(B,A))
  Body (unordered):
    1) (optional) S inform H that A
    2) Tell(S,H,B)
  Goals:
    1) H believe that A
    2) H believe that Elaboration(B,A)

Finally, note that a speaker may concede a failure to achieve a certain goal while seeking credit for the partial success of a plan to achieve that goal. For example, the [B] utterances in (10) can be used following (12) (or alone, in the right context) to concede failure.

(12) [A] I didn't go shopping today, but

Thus, the rule we give in (13) for inferring a plausible discourse relation of Concession is similar (but not identical) to (9).

66

```
(13)
If (i) coherently-related(A,B), and
   (ii) A is a proposition that an agent
        failed to do an action of act
        type T, and
  (iii) B is a proposition that describes
        a) the satisfaction of a normal
           applicability condition of T, or
        b) the satisfaction of a normal
           precondition of T, or
        c) the success of a normal step of T,
           or
        d) the achievement of a normal goal
           of T, and
   (iv) the achievement of the plan's
        component in B may bring credit
        to the agent,
then plausible(Concession(B,A)).
```

A discourse plan operator, Deny (with Concession), can be given to describe another standard way of performing a denial (exemplified in (5)). This operator is similar to the one given in (8), except with Concession in the place of Obstacle.

An interesting implication of the discourse plan operators for Affirm (with Elaboration) and Deny (with Concession) is that, in cases where the speaker chooses not to perform the optional step (i.e., chooses to omit the direct reply), it requires that the intended discourse relation be inferred in order to correctly interpret the indirect reply, since either an affirmation or denial could be realized with the same utterance. (Although (9) and (13) contain some features that differentiate Elaboration and Concession, other factors, such as intonation, will be considered in future research.)

The next two discourse relations (described in (14) and (16)) may be part of plan operators for conveying a 'yes' similar to Affirm (with Elaboration).

```
(14)
If (i) coherently-related(A,B), and
   (ii) A is a proposition that an agent
        performed an action X, and
  (iii) B is a proposition that normally
        implies that the agent has a goal
        G, and
   (iv) X is a type of action occurring
        as a normal part of a plan to
        achieve G,
then plausible(
        Motivate-Volitional-Action(B,A)).
```

(15) shows the use of Motivate-Volitional-Action (MVA) in an indirect (affirmative) reply.

```
(15) Q: Did you close the window?
     A: I was cold.
```

```
(16)
If (i) coherently-related(A,B), and
   (ii) A is a proposition that an event E
        occurred, and
  (iii) B is a proposition that an event F
        occurred, and
   (iv) it is not believed that F followed
        E, and
    (v) F-type events normally cause
        E-type events,
then plausible(Cause-Non-Volitional(B,A)).
```

(17) shows the use of Cause-Non-Volitional (CNV) in an indirect (affirmative) reply.

```
(17) Q: Did you wake up very early?
     A: The neighbor's dog was barking.
```

The discourse relation described in (18) may be part of a plan operator similar to Deny (with Obstacle) for conveying a 'no'.

```
(18)
If (i) coherently-related(A,B), and
   (ii) A is a proposition that an event E
        did not occur, and
  (iii) B is a proposition that an action F
        was performed, and
   (iv) F-type actions are normally
        performed as a way of preventing
        E-type events,
then plausible(Prevent(B,A)).
```

(19) shows the use of Prevent in an indirect denial.

```
(19) Q: Did you catch the flu?
     A: I got a flu shot.
```

The discourse relation described in (20) can be part of a plan operator similar to the others described above except that one of the speaker's goals is, rather than affirming or denying p, to provide support for the belief that p.

```
(20)
If (i) coherently-related(A,B), and
   (ii) B is a proposition that describes
        a typical result of the situation
        described in proposition A,
then plausible(Evidence(B,A)).
```

Assuming an appropriate context, (21) is an example of use of this relation to convey support, i.e., to convey that it is likely that someone is home.

(21) Q: Is anyone home?
A: The upstairs lights are on.

A similar rule could be defined for a relation used to convey support against a belief.

2.3 Implicatures of Discourse Units

Consider the similar dialogues in (22) and (23).

(22) Q: Did you go shopping?
A:a. I had to take the bus.
b. (because) My car's not running.
c. (You see,) The timing belt broke.
(23) Q: Did you go shopping?
A:a. My car's not running.
b. The timing belt broke.
c. (So) I had to take the bus.

First, note that although the order of the sentences realizing A's reply varies in (22) and (23), A's overall discourse purpose in both is to convey a 'yes'. Second, note that it is necessary to have a rule so that if A's reply consists solely of (22a) (=23c), an implicated 'yes' is derived; and if it consists solely of (22b) (=23a), an implicated 'no'.

In existing sentence-at-a-time models of calculating implicatures (Gazdar, 1979; Hirschberg, 1985), processing (22a) would result in an implicated 'yes' being added to the context, which would successfully block the addition of an implicated 'no' on processing (22b). However, processing (23a) would result in a putatively implicated 'no' being added to the context (incorrectly attributing a fleeting intention of A to convey a 'no'); then, on processing (23c) the conflicting but intended 'yes' would be blocked by context, giving an incorrect result. Thus, a sentence-at-a-time model must predict when (23c) should override (23a). Also, in that model, processing (23) requires "extra effort", a nonmonotonic revision of belief not needed to handle (22); yet (23) seems more like (22) than a case in which a speaker actually changes her mind.

In our model, since implicatures correspond to goals of inferred or constructed *hierarchical* plans, we avoid this problem. (22A) and (23A) both correspond to step 2 of Affirm (with Elaboration), Tell(S,H,B); several different discourse plan operators can be used to construct a plan for this Tell action. For example, one operator for Tell(S,H,B) is given below in (24); the operator represents that in telling H that B, where B describes an agent's volitional action, a speaker may provide motivation for the agent's action.

(24) Tell(S,H,p)
Applicability Conditions:
1) S BMB plausible(
Motivate-Volitional-Action(q,p))
Body (unordered):
1) Tell(S,H,q)
2) S inform H that p
Goals:
1) H believe that p
2) H believe that
Motivate-Volitional-Action(q,p)

(We are currently investigating, in generation, when to use an operator such as (24). For example, a speaker might want to use (24) in case he thinks that the hearer might doubt the truth of B unless he knows of the motivation.) Thus, (22a)/(23c) corresponds to step 2 of (24); (22b) - (22c), as well as (23a) - (23b), correspond to step 1. Another operator for Tell(S,H,p) could represent that in telling H that p, a speaker may provide the cause of an event; i.e., the operator would be like (24) but with Cause-Non-Volitional as the discourse relation. This operator could be used to decompose (22b) - (22c)/(23a) - (23b). The structure proposed for (22A)/(23A) is illustrated in Figure 1.[8] Linear precedence in the tree does not necessarily represent narrative order; one way of ordering the two nodes directly dominated by Tell(MVA) gives (22A), another gives (23A). (Narrative order in the generation of indirect replies is an area we are currently investigating also; for related research, see section 3.)

Note that Deny (with Obstacle) can not be used to generate/interpret (22A) or (23A) since its body can not be expanded to account for (22a)/(23c).[9] Thus, the correct implicatures can be derived without attributing spurious intentions to A, and without requiring cancellation of spurious implicatures.

---

[8] To use the terminology of (Moore & Paris, 1989; Moore & Paris, 1988), the labelled arcs represent *satellites*, and the unlabelled arcs *nucleii*. However, note that in their model, a nucleus can not be optional. This differs from our approach, in that we have shown that direct replies are optional in contexts such as those described by plan operators such as Affirm (with Elaboration).

[9] Determining this requires that the end of the relevant discourse unit be marked/recognized by cue phrases, intonation, or shift of focus; we plan to investigate this problem.

Affirm (with Elaboration)

I went shopping

(Elaboration)

Tell (MVA)

(Motivate-Volitional-Action)

Tell (CNV)

I had to take the bus

(Cause-Non-Volitional)
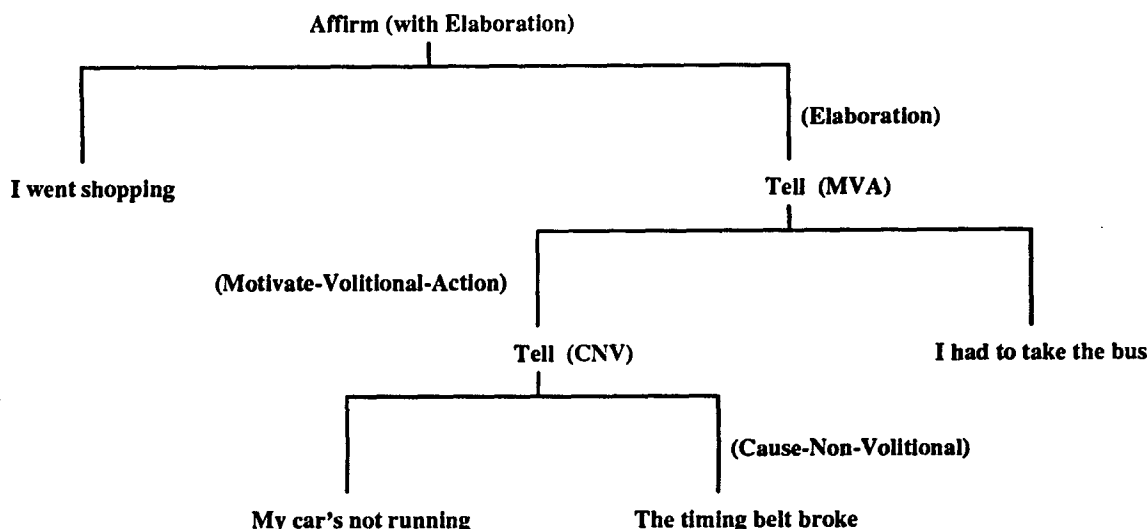
My car's not running

The timing belt broke

Figure 1. A Sample Discourse Structure

2.4 Algorithms

Generation and interpretation algorithms are given in (25) and (26), respectively. They presuppose that the plausible discourse relation is available.[10] The generation algorithm assumes as given an illocutionary-level representation of A's communicative goals.[11]

(25) Generation of indirect reply:
1. Select discourse plan operator: Select from the Answer-YNQ(p) plan operators all those for which
   a) the applicability conditions hold, and
   b) the goals include S's goals.
2. If more than one operator was selected in step 1, then choose one. Also, determine step ordering and whether it is necessary to include optional steps. (We are currently investigating how these choices are determined.)
3. Construct a plan from the chosen operator and execute it.

---

[10] We plan to implement an inference mechanism for the discourse relation inference rules.

[11] Note that A's goals depend, in part, on the illocutionary-level representation of Q's request. We assume that an analysis, such as provided in (Perrault & Allen, 1980), is available.

(26) Interpretation of indirect reply:
1. Infer discourse plan: Select from the Answer-YNQ(p) plan operators all those for which
   a) the second step of the body matches S's contribution, and
   b) the applicability conditions hold, and
   c) it is mutually believed that the goals are consistent with S's goals.
2. If more than one operator was selected in step 1, then choose one. (We are currently investigating what factors are involved in this choice. Of course, the utterance may be ambiguous.)
3. Ascribe to S the goal(s) of the chosen plan operator.

3 Comparison to Past Research

Most previous work in computational or formal linguistics on particularized conversational implicature (Green, 1990; Horacek, 1991; Joshi, Webber & Weischedel, 1984a; Joshi, Webber & Weischedel, 1984b; Reiter, 1990; Wainer & Maida, 1991) has treated other kinds of implicature than we consider here. Hirschberg (Hirschberg, 1985) provided licensing rules making use of mutual beliefs about salient partial orderings of entities in

the discourse context to calculate the scalar implicatures of an utterance. Our model is similar to Hirschberg's in that both rely on the representation of aspects of context to generate implicatures, and our discourse plan operators are roughly analogous in function to her licensing rules. However, her model makes no use of discourse relations. Therefore, it does not handle several kinds of indirect replies which we treat. For example, although A in (27) could be analyzed as scalar implicating a 'no' in some contexts, Hirschberg's model could not account for the use of A in other contexts as an elaboration (of how A managed to read chapter 1) intended to convey a 'yes'.[12]

(27) Q: Did you read the first chapter?
     A: I took it to the beach with me.

Furthermore, Hirschberg provided no computational method for determining the salient partially ordered set in a context. Also, in her model, implicatures are calculated one sentence at a time, which has the potential problems described above.

Lascarides, Asher, and Oberlander (Lascarides & Asher, 1991; Lascarides & Oberlander, 1992) described the interpretation and generation of temporal implicatures. Although that type of implicature (being Manner-based) is somewhat different from what we are studying, we have adopted their technique of providing defeasible inference rules for inferring discourse relations.

In philosophy, Thomason (Thomason, 1990) suggested that discourse expectations play a role in some implicatures. McCafferty (McCafferty, 1987) argued that interpreting certain implicated replies requires domain plan reconstruction. However, he did not provide a computational method for interpreting implicatures. Also, his proposed technique can not handle many types of indirect replies. For example, it can not account for the implicated negative replies in (1) and (5), since their interpretation involves reconstructing domain plans that were *not* executed successfully; it can not account for the implicated affirmative reply in (17), in which no reasoning about domain plans is involved; and it can not account for implicated replies conveying support for or against a belief, as in (21). Lastly, his approach cannot handle implicatures conveyed by discourse units containing more than one sentence.

Finally, note that our approach of including rhetorical goals in discourse plans is modelled on the work of Hovy (Hovy, 1988) and Moore and Paris (Moore & Paris, 1989; Moore & Paris, 1988), who used rhetorical plans to generate coherent text.

---

12 The two intended interpretations are marked by different intonations.

4 Conclusions

We have provided algorithms for the interpretation/generation of a type of reply involving a highly context-dependent conversational implicature. Our algorithms make use of discourse expectations, discourse plans, and discourse relations. The algorithms calculate implicatures of discourse units of one or more sentences. Our approach has several advantages. First, by taking discourse relations into account, it can capture a variety of implicatures not handled before. Second, by treating implicatures of discourse units which may consist of more than one sentence, it avoids the limitations of a sentence-at-a-time approach. Third, by making use of properties of discourse which have been used in models of other discourse phenomena, our approach can be integrated with those models. Also, our model permits the same information to be used both in interpretation and in generation.

Our current and anticipated research includes: refining and implementing our algorithms (including developing an inference mechanism for the discourse relation rules); extending our model to other types of implicatures; and investigating the integration of our model into general interpretation and generation frameworks.

References

Allen, James F. (1979). *A Plan-Based Approach to Speech Act Recognition*. PhD thesis, University of Toronto, Toronto, Ontario, Canada.

Austin, J. L. (1962). *How To Do Things With Words*. Cambridge, Massachusetts: Harvard University Press.

Carberry, Sandra (1990). *Plan Recognition in Natural Language Dialogue*. Cambridge, Massachusetts: MIT Press.

Clark, H. & Marshall, C. (1981). Definite reference and mutual knowledge. In A. K. Joshi, B. Webber, & I. Sag (Eds.), *Elements of discourse understanding*. Cambridge: Cambridge University Press.

Fikes, R. E. & Nilsson, N. J. (1971). Strips: A new approach to the application of theorem proving to problem solving. *Artificial Intelligence, 2*, 189–208.

Gazdar, G. (1979). *Pragmatics: Implicature, Presupposition, and Logical Form*. New York: Academic Press.

Green, Nancy L. (1990). Normal state implicature. In *Proceedings of the 28th Annual Meeting*, Pittsburgh. Association for Computational Linguistics.

Grice, H. Paul (1975). Logic and conversation. In Cole, P. & Morgan, J. L. (Eds.), *Syntax and*

Semantics III: Speech Acts, (pp. 41–58)., New York. Academic Press.

Grosz, Barbara & Sidner, Candace (1988). Plans for discourse. In P. Cohen, J. Morgan, & M. Pollack (Eds.), *Intentions in Communication*. MIT Press.

Hintikka, J. (1962). *Knowledge and Belief*. Ithaca: Cornell University Press.

Hirschberg, Julia B. (1985). *A Theory of Scalar Implicature*. PhD thesis, University of Pennsylvania.

Horacek, Helmut (1991). Exploiting conversational implicature for generating concise explanations. In *Proceedings*. European Association for Computational Linguistics.

Hovy, Eduard H. (1988). Planning coherent multisentential text. In *Proceedings of the 26th Annual Meeting*, (pp. 163–169). Association for Computational Linguistics.

Joshi, Aravind, Webber, Bonnie, & Weischedel, Ralph (1984a). Living up to expectations: Computing expert responses. In *Proceedings of the Fourth National Conference on Artificial Intelligence*, (pp. 169–175)., Austin, Texas.

Joshi, Aravind, Webber, Bonnie, & Weischedel, Ralph (1984b). Preventing false inferences. In *Proceedings of Coling84*, (pp. 134–138)., Stanford University, California. Association for Computational Linguistics.

Lambert, Lynn & Carberry, Sandra (1991). A tripartite plan-based model of dialogue. In *Proceedings of the 29th Annual Meeting*, (pp. 47–54). Association for Computational Linguistics.

Lascarides, Alex & Asher, Nicholas (1991). Discourse relations and defeasible knowledge. In *Proceedings of the 29th Annual Meeting*, (pp. 55–62). Association for Computational Linguistics.

Lascarides, Alex & Oberlander, Jon (1992). Temporal coherence and defeasible knowledge. *Theoretical Linguistics, 18*.

Litman, Diane & Allen, James (1987). A plan recognition model for subdialogues in conversation. *Cognitive Science, 11*, 163–200.

Mann, William C. & Thompson, Sandra A. (1983). Relational propositions in discourse. Technical Report ISI/RR-83-115, ISI/USC.

Mann, William C. & Thompson, Sandra A. (1987). 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*. PhD thesis, University of Pittsburgh.

Moore, Johanna D. & Paris, Cecile (1989). Planning text for advisory dialogues. In *Proceedings of the 27th Annual Meeting*, University of British Columbia, Vancouver. Association of Computational Linguistics.

Moore, Johanna D. & Paris, Cecile L. (1988). Constructing coherent text using rhetorical relations. In *Proc. 10th Annual Conference*. Cognitive Science Society.

Perrault, Raymond & Allen, James (1980). A plan-based analysis of indirect speech acts. *American Journal of Computational Linguistics, 6*(3-4), 167–182.

Pollack, Martha (1988). Plans as complex mental attitudes. In P. Cohen, J. Morgan, & M. Pollack (Eds.), *Intentions in Communication*. MIT Press.

Reichman, Rachel (1984). Extended person-machine interface. *Artificial Intelligence, 22*, 157–218.

Reiter, Ehud (1990). The computational complexity of avoiding conversational implicatures. In *Proceedings of the 28th Annual Meeting*, (pp. 97–104)., Pittsburgh. Association for Computational Linguistics.

Sadock, Jerrold M. (1978). On testing for conversational implicature. In Cole, P. & Morgan, J. L. (Eds.), *Syntax and Semantics*, (pp. 281–297)., N.Y. Academic Press.

Thomason, Richmond H. (1990). Accommodation, meaning, and implicature: Interdisciplinary foundations for pragmatics. In P. Cohen, J. Morgan, & M. Pollack (Eds.), *Intentions in Communication*. Cambridge, Massachusetts: MIT Press.

Wainer, Jacques & Maida, Anthony (1991). Good and bad news in formalizing generalized implicatures. In *Proceedings of the Sixteenth Annual Meeting of the Berkeley Linguistics Society*, (pp. 66–71)., Berkeley, California.