A Model of Conversation Processing

Based on Micro Conversational Events

Massimo Poesio
University of Edinburgh
Centre for Cognitive Science
2 Buccleuch Place
Edinburgh EH8 9LW, Scotland, UK
poesio@cogsci.ed.ac.uk

Abstract

I present a theory of discourse interpretation based on the hypothesis that the common ground of a conversation contains a record not only of complete speech acts, but, more in general, of each action of uttering a contribution to the conversation: single words, word fragments, and fillers. I call the action of uttering a 'minimal' contribution a MICRO CONVERSATIONAL EVENT. This model can serve as the basis for accounts of reference resolution in spoken conversations, as well as the interaction between parsing, repair, and reference resolution.

The Problem

Traditional 'discourse models' such as (Webber, 1979; Grosz and Sidner, 1986; Kamp and Reyle, 1993) have been developed on the basis of facts about anaphoric reference and other forms of contextual dependency. The theories of anaphoric processing that include a detailed account of how the discourse model gets updated in response to a sentence / utterance, such as Discourse Representation Theory (DRT) (Kamp and Reyle, 1993) and related 'dynamic' theories, are all based on the assumption that interactions consist of complete, well-formed sentences. But in real conversations, complete and well-formed sentences are quite rare: conversations like in (1), from the TRAINS corpus,[1] include large numbers of repairs[2] and interruptions; one speaker's output typically alternates with the other speaker's acknowledgments.

(1)
```
9.1  M: so we should
9.2   : move the engine
9.3   : at Avon
9.4   : engine E
9.5   : to
10.1 S: engine E1
11.1 M: E1
12.1 S: okay
13.1 M: engine E1
13.2  : to Bath
13.3  : to /
13.4  : or
13.5  : we could actually move it
        to Dansville to pick up
        the boxcar there
14.1 S: okay
```

In (Poesio, 1994; Poesio and Traum, 1995) a version of Discourse Representation Theory was proposed, based on the idea from speech act theory and the AI work on intention recognition that the common ground includes a record of each speech act that occurred in a conversation. The theory is generalized here by relaxing the assumption that the input to discourse interpretation consists of complete sentences. This extension can serve as the basis for systems that have to engage in conversations such as (1). The theory is also meant to provide a foundation for theories of discourse interpretation in which repairs, grounding, and similar processes interact with processes such as reference resolution, scope disambiguation, and reference resolution; hypotheses about the form of such interaction are also discussed.

Discourse Models for Conversations

A Formal Model of Discourse

Discourse Representation Theory (DRT) (Kamp and Reyle, 1993) and related 'dynamic' theories such as Dynamic Predicate Logic) are the formal expression of intuitions also contained in discourse models such as Karttunen's or Webber's. The model of a discourse in DRT is a DISCOURSE REPRESENTATION STRUCTURE (DRS), a pair consisting of a set of DISCOURSE REFERENTS and a set of CONDITIONS (facts) that is typically represented in 'box' fashion, as in (3).

(2) Engine E2 is at Avon.

(3)

x w
engine E2(*x*) **Avon**(*w*) **at**(*x*,*w*)

The two advantages of the DRT model of discourse update are its clear semantics and the detailed account of how the discourse model gets updated by each utterance (the DRS CONSTRUCTION ALGORITHM). The contribution of a sentence to

[1]The TRAINS project at the University of Rochester (Allen et al., 1995) studies task-oriented conversations. The project involves both collecting a corpus of conversations between human beings involved in a task—the TRAINS domain is transportation of goods by train—and developing theories about the aspects of natural language interpretation and plan reasoning observed in these transcripts. An annotated subset of the conversations in the corpus is in (Gross, Allen, and Traum, 1993).

[2]For example, Hindle (1983) reports that approximately 10% of spontaneous utterances contain disfluencies including self-repair. Heeman and Allen (1994) report that the in TRAINS corpus, 25% of turns contain at least one repair.

the discourse model is computed in two steps: first, the previous discourse model is augmented with an 'uninterpreted condition' that provides information about the syntactic structure of the sentence; second, rewrite rules are applied to the DRS thus augmented. The rewrite rules add new conditions or new discourse referents to the DRS, or 'flesh out' the interpretation of the uninterpreted condition. The sentence *The boxcar is also at Avon and is hooked to it* results at first in the DRS in (5), with an 'uninterpreted condition'; at the end of the construction algorithm, the DRS in (6) is obtained. Note that the construction algorithm results in two new referents, for the boxcar and *it*; note also that the previously introduced discourse referent x is made available as antecedents of the pronoun u.[3]

(4) Engine E2 is at Avon. The boxcar is also at Avon and is hooked to it.

(5)

```
x w
----------------
engine E2(x)
Avon(w)
at(x,w)
[S [NP the boxcar]
[VP is also at Avon and is
hooked to it]]
```

(6)

```
x w y u
----------------
engine E2(x)
Avon(w)
at(x,w)

boxcar(y)
at(y,w)
hooked-to(y,u)
u = x
```

Conversation Representation Theory

Many recent theories of pronouns interpretation, definite description interpretation, and VP ellipsis resolution have been formalized in terms of a formal discourse model such as DRT, which makes precise claims about language and is easily implementable. But DRT lacks some crucial features required in a model of conversation. A model of the information used by processes such as reference resolution must include the pragmatic information used by such processes: for example, focusing information, and information about the subordination relations between utterances that result in DISCOURSE SEGMENTS (Grosz and Sidner, 1986). This information is not provided by the standard DRT model. Furthermore, the DRT model is based on the assumption that all utterances are assertions, whereas utterances in a real conversation can serve a much wider range of purposes.

The models of discourse developed in the AI literature on intention recognition include the needed pragmatic information. They are based on the assumption that utterances are actions in their own right, performed to achieve DISCOURSE GOALS whose hierarchical structure determines the structure of the discourse (Cohen and Perrault, 1979; Allen and Perrault, 1980; Grosz and Sidner, 1986). What is missing from these models is what DRT has: a precise specification of what is available for reference at a given point in the conversation, and a detailed account of the update process associated with a sentence. A combination of the two research traditions seems highly desirable.

Such a unified theory, CONVERSATION REPRESENTATION THEORY (CRT), was proposed in (Poesio, 1994) and further developed in (Poesio and Traum, 1995). The theory is based on the formulation of speech act theory adopted in Situation Semantics (Barwise and Perry, 1983), according to which the common ground consists of information about the DISCOURSE SITUATION (the set of actions performed by the participants to the conversation) rather than the DESCRIBED SITUATION (the situation which is the topic of the conversation); the participants in a conversations recover information about the latter from their information about the former. CRT preserves much of the formal machinery of DRT, including the update algorithm; but the DRS modeling the common ground as a whole (the ROOT DRS) is re-interpreted as consisting of information about the occurrence of speech acts or, as I will call them, CONVERSATIONAL EVENTS. Each utterance thus introduces two situations: a conversational event[4] and a described situation. The made-up mini-dialog in (7) is represented in CRT as in (8).

(7) A: Engine E2 is at Avon.
B: The boxcar is also at Avon and is hooked to *it*.

(8)

```
ce1 t1 ce2 t2 s
------------------------------------
                      x w
                      ------------
ce1: tell(A,B, s:     engine(x)      )
                      Avon(w)
                      at(x,w)
ce1 @ t1
                      y u
                      ------------
ce2: tell(B,A, s:     boxcar(y)      )
                      at(y,w)
                      hooked-to(y,u)
                      u = x
ce2 @ t2
```

The DRS in (8) represents a discourse situation in which two conversational events occurred, *ce1* and *ce2*, both of which are about the described situation *s*. The first two conditions assert that *ce1* is an event of A telling B that the described situation *s* includes an engine located at Avon, and that *ce1* took

[3]Space prevents a detailed discussion of the algorithm. See (Kamp and Reyle, 1993).

[4]In Situation Semantics, events are considered a special type of situation.

place at time $t1$. The third and fourth condition assert that $ce2$ is an event, taking place at time $t2$, of B telling A that boxcar y is also in Avon, and is hooked to u.

Anaphoric accessibility is preserved by modifying the interpretation of DRSs and by adding a new type of condition, indicated in (8) by expressions of the form $s{:}\phi$. A DRS is treated in CRT as a *situation type*, i.e., as denoting the set of situations which can be made to verify the conditions contained in the DRS once appropriate values for the discourse referents have been found. Conditions like s:K, then, assert that a situation s is of the type specified by the DRS K. More precisely: Expressions in CRT are assigned a value with respect to a situation, a variable assignment, and a set of CASES, one for each situation. An expression of the form s:K asserts that s is of the type specified by K, and furthermore, it shifts the parameters of evaluation so that the value of the discourse markers occurring in K is provided by the case associated with the value of s. This ensures that the DRS in the complement of $ce2$ is evaluated with respect to the same case that is used to evaluate the expressions in the complement of $ce1$, *e.g.*, that x is accessible from within the complement of $ce2$.

The update algorithm of DRT is (minimally) modified as follows. The first step in the DRT construction algorithm is replaced by the application of a CONVERSATIONAL EVENT GENERATION RULES that update the existing model of the discourse situation by adding two new discourse referents (a conversational event and a time) and two conditions, one asserting the occurrence of the appropriate locutionary act (one of **tell**, **ask**, and **instruct**), the other recording the time at which it occurred. The process of DRS construction can then proceed much as in (Kamp and Reyle, 1993), except that the whole process of rewriting takes place within the DRS that is the complement of the locutionary act. DRSs like (8) are the final result of this process.

The model in (7), although a very simplified representation of the conversation in (8), also includes the pragmatic information required to perform reference resolution and speech acts-based reasoning, such as inferring subordination relations. The subordination relations between conversational events can be inferred by reasoning about the agents' intentions, all information that, as well, could be included in the model of the conversation (although for simplicity hasn't been represented here). Although the descriptions of the conversational events in (8) are very abstract, they can serve as the basis for reasoning processes such as those discussed in (Allen and Perrault, 1980; Traum and Hinkelman, 1992); in fact, various algorithms of this sort have been developed in connection with TRAINS.

The model of interpretation underlying CRT involves the generation of multiple hypotheses in parallel, some (or all) of which are then discarded. The model, based on a 'prioritized' version of default logic similar to Brewka's (1991), and appears to be roughly compatible with the findings of the literature on lexical disambiguation, syntactic disambiguation, and reference resolution (Seidenberg et al., 1982; Gibson, 1991) and with cognitively motivated computational models of sentence such as Jurafsky's SAL (Jurafsky, 1992) and COMPERE (Eiselt, Mahesh, and Holbrook, 1993). The model does not include a commitment to the way the hypotheses are pruned.

Micro Conversational Events

Basic Idea

The central idea of this paper is that a model of the common ground such as the one in CRT, in which the common ground is a representation of the discourse situation, generalizes naturally into a model in which the apparently complex effects of utterances such as those in (1) can be explained rather simply. If we accept that the common ground includes a record of each utterance, it is but a small step to assume that the common ground contains a record of the events of uttering units smaller than sentences, such as single words, 'fill-phrases' like *hmm*, or word fragments. In other words, it becomes easier to think that each minimal contribution to a conversation[5] constitute a 'micro-speech act' —or, as I will call them, a MICRO CONVERSATIONAL EVENT—whose occurrence leads to a 'micro-update' of the common ground.[6]

For example, let us assume that speaker A utters the word-form *the*. I propose that as a result the common ground gets updated with the following information:

(9)

$\ldots ce_1\ \mu ce_1\ \mu t_1$
...
μce_1: **utters**$(A, \langle s, \text{"the"}\rangle)$
μce_1 @ μt_1
sub$(\mu ce_1, ce_1)$
...

This update is the result of a MICRO-CONVERSATIONAL EVENT GENERATION RULE that replaces the conversational event generation rules described in the previous section. The information added to the common ground by such a rule includes the fact that a micro-conversational event μce_1 occurred, uttered by A, and with argument the word-form *the*. The fact that μce_1 is subordinate to a conversational event ce_1 also becomes part of the common ground. The expression $\langle s, \text{"}the\text{"}\rangle$ denotes the value of the second element of the pair, interpreted with respect to the case associated with s.

The update just presented may result in further updates as other processes are 'triggered' by the occurrence of the micro-

[5] I will assume here for simplicity reasons that the 'minimal incremental units' are words. Other hypotheses made in the literature are that phonemes or prosodic phrases serve as minimal incremental unit.

[6] The 'micro-utterances' hypothesis has been made in work on Situation Theory such as (Gawron and Peters, 1990), although not in relation with work on processing spoken conversations. In that literature the hypothesis is motivated on theoretical grounds, rather than to explain phenomena such as those discussed in this paper, and its details haven't been worked out yet.

conversational event. If the utterance fragment is recognized as a lexeme, at least lexical interpretation and disambiguation take place. The formulation of an hypothesis about lexical interpretation can be expressed as a rewriting operation that replaces the word form with its lexical interpretation. (More than one such hypothesis will be generated in the case of lexical disambiguation.) In this example, the assertion that A uttered the word form *the* will be replaced by the assertion that A uttered the determiner **the**, as shown in (10). Expressions such as $\langle s, [_{\text{Det}}\ \textbf{the}] \rangle$ are logical expressions, whose 'underspecified' semantics is described in (Poesio, 1994).

(10) μce_1: $\boxed{\textbf{utters}(A, \langle \text{s}, [_{\text{Det}}\ \textbf{the}] \rangle)}$

The same process is repeated after each utterance fragment. Assume this next utterance fragment is again a lexical item, the word *engine*. After lexical disambiguation, the common ground contains the additional information:

(11)

$\ldots \mu ce_2\ \mu t_2$
...
μce_2: $\boxed{\textbf{utters}(A, \langle \text{s}, [_{\text{N}}\ \text{engine}] \rangle)}$
μce_2 @ μt_2
$\textbf{sub}(\mu ce_2, ce_1)$
current-CE = ce_1

At the end of an utterance unit, a 'macro' conversational event (i.e., a speech act in the sense traditionally used in the AI literature, such as 'inform', 'request', etc.) may then be obtained out of the micro conversational events by means of reasoning processes such as those proposed in the AI literature, except that in the present model such inference processes do not require the utterance of a complete sentence.

Incremental Parsing

In order to discuss more in detail how this model can be applied to account for the complex interactions between parsing, repair, and reference resolution displayed in conversations such as (1), several assumptions about how parsing and repair work need to be discussed; I will do this in the next section. In order to provide an account of the other-initiated repair in 10.1, however, the model of update just sketched only has to be supplemented with the assumption that at least some discourse interpretation processes can work off micro-conversational events without waiting for an 'utterance completion' signal.

After $\mu ce2$ has been added to the common ground, a parser working in an incremental fashion[7] and operating on the subset of micro-conversational events that have been successfully processed by lexical disambiguation[8] may determine that a Noun Phrase just occurred.[9] This information may be added to the common ground in the form of a condition to the effect that a micro-conversational event μce_3 of uttering an NP with constituents μce_1 and μce_2 occurred. After this update, the common ground contains not only the fact that a determiner and a noun were uttered, but also the fact that they combined to form a noun phrase.

(12)

$\ldots\ldots \mu ce_3\ \mu t_3$
...
μce_3: $\boxed{\textbf{utters}(A, \langle \text{s}, [_{\text{NP}}\ \mu \text{ce}_1\ \mu \text{ce}_2] \rangle)}$
μce_3 @ μt_3
$\textbf{sub}(\mu ce_3, ce_1)$
...

The parser may not have enough evidence to combine the lexically interpretable micro-ces into micro-ces describing larger phrases; in this case, no larger constituent is added to the common ground.

Incremental reference resolution

The other-initiated repair of utterance 10.1 in (1) may be taken as evidence that definite description interpretation takes place incrementally, as well. At that point in the conversation M has clearly decided that *the engine at Avon* is not *engine E*. This suggests that after processing utterances 9.2 and 9.3 M has enough information to start processing the definite NP, hence, that the reference resolution process can be triggered by the occurrence of a micro-conversational event of uttering a definite NP.

Except for the fact that they are initiated by a micro-conversational event rather than a complete conversational event, the processes concerned with finding the referent of a definite description and with updating the common ground if that referent is identified proceed much as discussed in (Poesio, 1994). The discourse referents associated with the situation *s* may serve as the antecedent of the definite.

Applications

Parsing with Filler Phrases and Fragments

In addition to providing an explanation of how referential expressions can be interpreted before a sentence is completed, the theory just discussed suggests a simple model of the processing that takes place when 'parsable' input such as lexical items alternates with 'unparsable' input such as fill phrases or word fragments. Two such cases are illustrated by the examples in (13) and (14). (All of these examples are from the TRAINS corpus.)

(13)
```
145.1   M: no we  /
145.2    : we took the sh/
145.3    : uh dadadada
```

[7] Evidence that the parser works in this way is discussed, for example, in (Crain and Steedman, 1985; Jurafsky, 1992).

[8] See below.

[9] It is immaterial for the present purposes whether this decision involved pragmatical information or was arrived at on the basis of syntactic information only, in a completely modular fashion. I.e., the present model does not commit us either to a 'modular' or to an 'integrated' view of sentence processing.

```
        145.4   : no still /
        145.5   : still take the shortest route
(14)    39.4    : so then send the
        39.5    : send the /
                [2sec]
        39.6    : send the orange car .. back
```

In (13), the utterance of the initial parts of a noun phrase is followed by 'filling phrases' like *uhhh*. In (14), the speaker utters *send* and *the* three times each before continuing its utterance.

Although one could try to complicate the grammar in order to account for such phenomena, the model of discourse update just presented allows for a simpler explanation of the processing that goes on in these cases. According to this account, the parser is not the first module to operate on the input, and therefore the other modules do not get only what the parser 'lets through', as it were. Instead, everything that is uttered becomes part of the conversational record, and the parser only works on the subset of this input that it is actually competent to deal with. The micro-conversational events that are not recognized as lexical items are either 'skipped' or dealt with by other processes, as we shall see below. The following interaction between the parser and the rest of sentence processing can be envisaged on the basis of the micro-conversational events model:

1. Every new micro-conversational event activates the lexicon, that attempts to classify the new input.

2. Whenever the input is recognized as a proper lexical item, the parser is activated. The parser used in TRAINS, discussed in (Allen et al., 1995), makes use of a chart to generate all interpretations in parallel and relies on statistical information to choose the best one. All interpretations of lower plausibility are discarded. The only difference I am assuming here is that the parser is invoked on the input after every word or after every phrase boundary, and therefore this process of parallel hyothesis generation and statistical choice is repeated every time, much as in Jurafsky's SAL architecture (Jurafsky, 1992).

3. The parser updates the discourse model by adding a new micro-conversational event as discussed above whenever it has obtained a single hypothesis about a portion of the utterance. When more than one hypothesis are statistically equivalent, no micro conversational event is added.

Thus, in the case of (13), the fill-phrase *sh uh dadadada* is simply skipped over by the parser since it cannot be assigned a lexical entry. (The more complex interaction in (14) is discussed below.) This architecture results both in a simpler theory and in computational advantages (in that complicating the grammar always results in increased ambiguities).

Parsing and Repairs

The micro-conversational events model also allows for a simple theory of the interaction between parsing and a 'local' model of repair such as the algorithm in (Heeman and Allen, 1994), in which speech repairs are detected and corrected online and using only local clues. The assumptions about the structure of repairs adopted in the model are shown in (15) (from (Heeman and Allen, 1994), p. 297).

```
(15)    go| to|  oran-| um| go| to| Corning
        m1| m2|  x|  int| et| m1| m2|
```

These models work by identifying the so-called INTERRUPTION POINT (identified in (15) by a -), skipping over EDITING TERMS such as um, identifying the REPARANDUM, i.e., the parts of the input to be replaced, and replacing the reparandum with a REPAIR PATTERN. It is typically (although not always) the case that the repair pattern contains one 'matching element' for each element in the reparandum, as shown in (15).

What is important for my purposes is, first of all, and most obviously, that the process of repair affects the parser, in the sense that in the syntactic structure eventually assigned to an utterance the repaired terms have been replaced with the terms of the repair pattern. In order for this to happen, it is necessary for there to be an exchange of information between the parser and the repair module. Furthermore, this exchange cannot be mono-directional (e.g., the repair module always working off the output of the repair module), because entire phrases can be replaced, not just lexical items, as in the following examples, from Heeman and Allen's paper:

(16) After the orange juice is at- the oranges are at the OJ factory

(17) How would that- how long would that ...

In (16), the noun phrase *the oranges* replaces the NP *the orange juice*; in (17), the wh-phrase *how long* replaces the wh-phrase *how*. Both of these examples can be reconduced to the simplest form of repair (one-to-one replacement) if the repair module is affected by the output of the parser. In order for the interaction between parser and repair module to be bidirectional, a common level of representation is needed. The micro-conversational events level is the required level of representation.

Repairs and Reference

Perhaps the most convincing argument in favor of the hypothesis that repairs take place at the micro-conversational event level is the fact that the processes concerned with reference interpretation and those concerned with repair interact. The fragment in (1) contains two examples of such interaction. One, the case of other-initiated repair in 10.1, has been discussed above.

A second example of interaction between repair and reference, and furter evidence that the common ground is updated before a sentence is completely interpreted, is the fresh start in (1) (utterances 13.4 and 13.5). S replaces the proposal introduced in 9.1-13.2 with a new one, but in doing so he[10] as-

[10]The manager in this conversation was played by a male student

sumes that *the engine at Avon, engine E1* is part of the common ground. If the repair process were to take place before discourse referents are established and reference resolution is performed, the referent would be removed, and we would end up with a pronoun without antecedent.

Discussion

The model of update sketched in the previous sections is a simple extension to standard DRT, yet the result is a model of discourse update that is appropriate for real conversations. It can serve as the basis for theories about the interface between repair, incremental parsing, and reference, as sketched in the previous section.

Ideas such as providing a level at which different modules can interact are familiar from the work on 'blackboard architectures' (Hayes-Roth, 1985). Some aspects of the model, such as the way syntactic and lexical information are integrated, make it compatible with cognitive models of sentence processing such as Jurafsky's. One way to see the contribution of the present work is to spell out the details of a conversation processing model consistent with such an architecture, especially as far as the semantic details of reference are concerned.

The model has been so far tested only by means of hand simulations over the TRAINS corpus. Implementations of the modules discussed here (parsing, repair, reference, and DRT update) already exist, however, and work on an integration is under way.

References

Allen, J. F. and C. Perrault. 1980. Analyzing intention in utterances. *Artificial Intelligence*, 15(3):143–178.

Allen, J. F., L. K. Schubert, G. Ferguson, P. Heeman, C. H. Hwang, T. Kato, M. Light, N. Martin, B. Miller, M. Poesio, and D. R. Traum. 1995. The TRAINS project: a case study in building a conversational planning agent. *Journal of Experimental and Theoretical AI*, 7:7–48.

Barwise, J. and J. Perry. 1983. *Situations and Attitudes*. The MIT Press.

Brewka, G. 1991. *Nonmonotonic Reasoning: Logical Foundations of Commonsense*. , edited by. Cambridge University Press.

Cohen, P. R. and C. R. Perrault. 1979. Elements of a plan based theory of speech acts. *Cognitive Science*, 3(3):177–212.

Crain, S. and M. Steedman. 1985. On not being led up the garden path: the use of context by the psychological syntax processor. In D. R. Dowty, L. Karttunen, and A. M. Zwicky, editors, *Natural Language Parsing: Psychological, Computational and Theoretical perspectives*. Cambridge University Press, New York, pages 320–358.

at the University of Rochester.

Eiselt, K. P., K. Mahesh, and J. K. Holbrook. 1993. Having yoour cake and eating it too: Autonomy and interaction in a model of sentence processing. In *Proc. 11th National Conference on Artificial Intelligence*, pages 380–385, Washington, DC.

Gawron, J. M. and S. Peters. 1990. *Anaphora and Quantification in Situation Semantics*. Lecture Notes, volume 19. CSLI.

Gibson, E. 1991. *A Computational Theory of human linguistic processing: memory limitations and processing breakdown*. Ph.D. thesis, Carnegie Mellon University, Pittsburgh.

Gross, D., J. Allen, and D. Traum. 1993. The TRAINS 91 dialogues. TRAINS Technical Note 92-1, Computer Science Dept. University of Rochester, June.

Grosz, B. J. and C. L. Sidner. 1986. Attention, intention, and the structure of discourse. *Computational Linguistics*, 12(3):175–204.

Hayes-Roth, B. 1985. A blackboard architecture for control. *Artificial Intelligence*, 26:251–321.

Heeman, P. A. and J. F. Allen. 1994. Detecting and correcting speech repairs. In *Proc. of the Association for Computational Linguistics*.

Hindle, Donald. 1983. User manual for FIDDITCH, a deterministic parser. Technical Memorandum 7590-142, Naval Research Laboratories.

Jurafsky, D. 1992. *An On-Line Computational Model of Human Sentence Interpretation: A Theory of the Representation and Use of Linguistic Knowledge*. Ph.D. thesis, University of California at Berkeley, Department of Computer Science.

Kamp, H. and U. Reyle. 1993. *From Discourse to Logic*. Dordrecht: D. Reidel.

Poesio, M. 1994. *Discourse Interpretation and the Scope of Operators*. Ph.D. thesis, University of Rochester, Department of Computer Science, Rochester, NY.

Poesio, M. and D. Traum. 1995. A multi-purpose model of conversational context. In *Proc. of the IJCAI Workshop on Context and Knowledge Representation*, Montreal, August.

Seidenberg, M. S., M. K. Tanenhaus, J. Leiman, and M. Bienkowski. 1982. Automatic access of the meanings of ambiguous words in context: some limitations of knowledge-based processing. *Cognitive Psychology*, 14:489–537.

Traum, D. R. and E. A. Hinkelman. 1992. Conversation acts in task-oriented spoken dialogue. *Computational Intelligence*, 8(3). Special Issue on Non-literal Language.

Webber, B. L. 1979. *A Formal Approach to Discourse Anaphora*. New York: Garland.