

Coherence Relation Assignment
K. Dahlgren
IBM/Los Angeles Scientific Center

Abstract. Three empirical studies of coherence in large corpora of commentary text are sketched, showing that cue phrases are infrequent, and that substantive coherence relations must be assigned in order to infer discourse structure. The notion of coherence is carefully defined in relation to the world, cognitive models of the world, and formal semantic representations of discourse. An efficient algorithm for assigning discourse coherence relations is described, which employs information from syntax, cue phrases, lexical items, formal semantics and naive semantics. The algorithm correctly assigns the coherence relations evident in an 8000 word corpus.

1.0 Introduction

Analyses of discourse structure in cognitive science have been dominated by two approaches. One examines syntactic and cue phrase information only, in order to avoid unsolved semantic puzzles [3, 13, 28, 32]. In this framework Grosz and Sidner have established that discourse is structured in a hierarchy, and that anaphora resolution is constrained by "focus spaces" or segments of discourse [13]. The other approach seeks to account for coherence---substantive relations between portions of a discourse [5, 11, 24, 17, 15, 27, 21, 31, 34, 36]. The coherence group divides into two camps, the top-downs [31], and the bottom-ups (all of the others). We are among the latter, that is, we aim to build discourse structure clause by clause, because a bottom-up approach will lead to more transportable, general results than will the (apparently) more tractable script scheme. This work is part of an ongoing project, NewSelector, for computational text understanding and precise text selection [5, 8]. The word meaning representation in NewSelector is based upon Naive Semantics (NS), a theory which identifies word meanings with commonsense theories of objects and events.

Studies of Coherence in Commentary. We have carried out three studies of coherence in an expanding corpus of *Wall Street Journal* (WSJ) commentary texts. Study 1 [5] of 8,000 words in six articles sought to examine what information is used in coherence relation assignment (CRA), and to determine whether syntactic markers and cue phrases were sufficient information for CRA. The coherence relation literature was reviewed [3, 10, 15, 17, 24], and 19 coherence relations (fully defined in [5]) which are relevant to the commentary genre were identified. In this paper we focus on just two of these, **cause** and **goal**. A coherence relation was assigned to each clause in the corpus by two judges. The syntactic and semantic properties of each clause were encoded. These properties included clause type, voice, mood, presence of negation, agentiveness of subject, type of subject and object, and aspectual class of the verb. The correlations between the coherence relation assigned to a clause, and syntactic/semantic properties of the clause form the basis of the algorithm described in Section 3. We found that the information used in CRA was: 1) syntax, 2) cue phrases, 3) lexical items, 4) tense, 5) aspect, 6) world knowledge.

Study 2 [6], of the same 8,000 words and 8,000 more, examined global coherence (or segmentation) to determine which factors influenced it. For each new S, the possibility of a new sister- or sub-segment arises. We found that change of coherence relation was the most reliable indicator of new segment and change of subject next most reliable. Other factors were

DAHLGREN

paragraph indentation, length of segment with the same coherence relation to some other, cue phrases, and event anaphors. Significantly, there was a segmenting cue phrase such as, "turning to..." in only 16% of the cases of a new sister segment. Clearly any computational system which looks only for direct cues will miss most of the structure. Substantive coherence relations, if they can be extracted, are powerful indicators of structure.

Study 3 [22] of the same 16,000 plus 4,500 more words of *WSJ* text examined personal pronouns, demonstratives and definite NP anaphora. We found that when we segment text as proposed in Study 2 [6], the resulting structure predicts constraints on anaphora resolution. This work supports [13] empirically and also shows that event anaphora has very different constraints from individual anaphora.

2.0 Coherence

Why Compute Coherence? The reasons for computing coherence are several. First, it is uncontroversial that text understanding requires segmentation [13,28]. These segments or focus spaces can only be found using coherence, as indicated in Study 2. Second, our Study 3 shows that by computing coherence, anaphora resolution can be significantly constrained. Otherwise, you either have to use brute force, which leads to a combinatorial explosion, and still indeterminate results, or try cues like indentation, which will be correct only 50% of the time, and cue phrases, which are present only 16% of the time. Third, more intelligent text understanding is made possible by coherence inferences. Text is telescopic, and the reader fills in the gaps. A computational system which models such inferences will be able to reflect much more accurately the human understanding of text. Fourth, when considering whether one text is relevant to another, the more naive inferencing computed, the more accurate the relevance reasoning will be. Finally, because of this last point, a computational system with coherence can answer many more queries accurately, as in (1)-(3).

- Why did John make a profit? --Because he invested. (1)
- Why did John invest? --In order to make a profit. (2)
- What did John invest in? --Typically, stocks or bonds. (3)

Problems with Coherence Theory. One of the reasons for the past emphasis on overt elements like structure and cue phrases, is that unsolved problems have plagued the coherence approach and made it a dubious notion. The first problem is one of *definition*. We have tried to rectify this below by giving a careful definition of what coherence is. Secondly, settling upon a set of coherence relations has been an elusive goal. Most studies have attempted to handle all genres with one big set [10, 15, 17, 24], but the set of relations varies with genre. The solution lies in assigning genre first, then computing coherence [27]. The third problem is informality. Much of the study of coherence has been descriptive [17,24] and has not attempted to provide a direct link between coherence theory and formal semantic theories. Recent developments in formal semantics provide a framework and ongoing research devoted to giving truth conditions for entire discourses as well as single sentences [20, 1, 35, 2]. It is now possible to integrate coherence structure with formal semantic representation (especially of temporal order and aspectual class which our Study 2 found are particularly significant indicators of coherence relations). The third problem is that no-one has proposed an algorithm for extracting coherence relations. This is because to do so, by all accounts, requires world knowledge, and

DAHLGREN

it has heretofore seemed impossible to encode world knowledge in a non-*ad hoc*, transportable and practicable way. Our theory of lexical representation, Naive Semantics, offers a solution.

What is Coherence? We consider the notion of coherence for all types of discourse, spoken and written. It is based upon the intuition that the discourse (4) seems to “hang together”, and that the discourse in (5) does not. Empirically, to the extent that discourses do not cohere, they are difficult to interpret and remember [16, 34, 11].

John invested heavily. He profited handsomely. (4)

John invested heavily. He ate pizza. (5)

The problems in coherence theory have been: 1) What is a coherent discourse? and 2) Which entities cohere: sentences, clauses, the propositions expressed, the real events denoted? In our view, a coherent discourse is one for which the hearer can build a cognitive representation such that the relations among events and individuals in the representation correspond with his understanding (theory) of the way actual world events and individuals relate.¹ Although the representation may contain a variety of types of sensory images, in general, the hearer’s “understanding” amounts to a naive (in the sense of [14]) theory of the causal and other structure of objects and events [11, 34].

Consider a formal semantic representation for the discourse (4), shown in (6) as a Discourse Representation Structure (DRS) after Kamp [20] and Asher [1]. It has the content that there was an individual John and two events, the first of John investing heavily, the second of John profiting handsomely. Notice that there are two inferences in the DRS, one of temporal order between the investing and the profiting ($r_1 < r_2$) [26], the other concerning anaphor resolution of “he” to John. The cognitive picture of the events constructed by the hearer (and presumably, intended by the speaker) is indicated in English on the right-hand side of (6). It includes all of the content of the DRS plus the inference that John’s goal in investing had been to make a profit. This goal inference is a coherence inference. The hearer brings the discourse into accord with his/her understanding or theory about investing. The reason for saying that “understanding” involves a theory (belief rather than knowledge) is that very often people, and cultures, are quite mistaken in such causal inferences. Nevertheless, they do use such structuring theories to manage the environment and to communicate via language. Because members of a subculture SHARE naive theories, the speaker can juxtapose just these two sentences, and know that the hearer will guess that John’s goal had been profit. In summary, speakers in a given genre make a discourse (and thereby their reporting of events) understandable by choosing to report events using certain verbs in a certain sequence. This choice in a well-structured discourse makes it maximally possible for the hearer to build a cognitive picture of these events which coheres. It will cohere to the extent that he/she can bring it into accord with her/his theories about the way the world works. The relationships in the naive theory of the world are causal, intentional, comparative, part-whole, etc.

We claim that coherence belongs in a cognitive inference module, not in syntax or semantics [6]. Temporal order and anaphora resolution belong in the compositional semantics because they are explicitly and linguistically marked. Coherence is a gradient phenomenon in that the better-structured the discourse, the more readily and reliably a hearer will make coherence inferences. Similarly, the more knowledgeable and tuned in the hearer, the more accurately

Discourse Representation Structure

```

u1,e1,e2,now,r1,r2
-----
John(u1)
e1 invest(u1)
heavily(e1)
e2 profit(u1)
handsomely(e2)
r1 < now
e1 ⊂ r1
r1 < r2
e2 ⊂ r2
    
```

Cognitive Representation (6)

John first invested heavily
and then profited
handsomely.

John's goal in investing had
been to make a profit.

he/she will recover the speaker's intended cognitive representation. Thus CRA requires cognitive reasoning which goes beyond what the discourse says directly. In this paper we are offering two innovations: an account of the relationship between formal semantics and coherence, based upon [2], and an empirically-constructed algorithm for CRA.

In comparison with other work, ours is similar to van Dijk and Kintsch [34] in that they define coherence by whether sentences in a discourse describe related facts in some possible world, and they assume that large amounts of world knowledge are employed in building a cognitive model of a discourse. We differ in defining coherence as relating discourse events, rather than as relating sentences. Furthermore, we clarify the question of truth conditions as opposed to naive (or heuristic) inference regarding discourse interpretation. And we provide an algorithm. We draw upon Hobbs [16] and Mann and Thompson [24] for coherence relations. However, we define them as relating discourse events in a cognitive event model, rather than as relating utterances, clauses, or spans of discourse, as they do. For them, coherence is essentially a property of presentation style, of the speaker's intended effect on the hearer. In contrast, for us, coherence is essentially a property of mental models [19] which finds its origin in belief about relationships among real events. Our approach appeals to cognitive strategies and beliefs people use all of the time, whether thinking verbally or not. Yet another view defines coherence in terms of the speaker's goals [13,15]. We agree with Polanyi [27] that this aspect of coherence belongs in a level of theory above that of cognitive models of interpretation of discourse.

What are Coherence Relations? In particular, what are the relations **cause** and **goal**? A coherence predication $cause(e1,e2)$ is a speaker/hearer theory about the causal connections between events. Starting with the basics, what is an event (denotationally)? One standard view, with many problems, is that an event is a spatially and temporally located occurrence in the real world. A more sophisticated view of events is as concrete particulars individuated by their causes and consequences [9]. This would explain the existence of real events in the world. Now consider a real event, say a car falling over a cliff. This can be broken into two causally related events as in (7).

e1. The car came to the edge. e2. The car fell over the edge. (7)

DAHLGREN

The car came too close to the edge, a critical point was reached in which the car was no longer balanced on the edge, and it fell over. We take cause as defined in [9], where cause is a two-place predicate relating events such that if the same event (with all of the relevant situation included) should reoccur, the caused event would reoccur. On this analysis we can assume the existence of causation in the world.

But in terms of the cognitive phenomenon of events and their structure in discourse, we must explain human interpretation of the actual world, not the world itself. Even the direct observation of some real event involves observer interpretation, minimally, of the idea that each frame or pixel he sees is part of the same real event and not a series of different events. The observer views that a certain thing is happening, e.g., that car is falling off of that cliff, is a theory---often a conscious verbal theory. "Oh, that car is falling off of that cliff". When the observer thinks, "Oh, an accident", we have even more theorizing and interpretation as actor goals (or lack of them) are inferred. So the predication of events is epistemological, heuristic, and retractable.

How shall we analyze observer construction of causal inferences? Given that there are cases where one event causes another (such as, the car comes too close to the edge, and falls over the cliff), in the cognitive interpretation of events (and of texts reporting events), people actively theorize and hypothesize about the causal and intentional structure of events as they unfold [11]. People are tentative about such inferences, but in order to function, they must guess. Such guesses are naive theories about the causal and intentional structure of events. Furthermore, the inferred causal structure is an important correlate of memory for the events [11]. In example (7), such a naive hypothesis would be formed when the observer uses naive physics to figure out that (e1) and (e2) are causally related. The point is that in understanding, interpreting and labelling observations of real world events, members of a Western culture (and probably of any culture) infer that the critical point was reached, and that the coming close to the edge (in the end) caused the falling over the edge. The same holds for discourse and text understanding. As they read, people infer causal structure and intentional structure (the goals of agents) [11]. In other words, cause, goal and enablement are salient relations hypothesized by readers about text events. Revising the definition, then, a coherence relation is a naive theory of the relation between events introduced into a discourse. It is a binary predicate whose arguments are discourse individuals, discourse events or states or sets thereof (see Asher [2]). A coherence theory (predication) arises from naive theories about the causal and other structure of the world.²

Naive Semantics. Turning to the third problem with coherence theory, we briefly describe our solution, a way of representing commonsense knowledge. Naive Semantics (NS) [5] is a theory of word sense meaning representation in which associated with each sense of each content word (noun, verb, adjective) is a naive theory of the sort of object, action or property named by the word. NS rejects a theory in which word meanings are broken up into atomic primitives which directly play a truth conditional role in sentence and discourse interpretation. Instead, these are supposed to contribute non-monotonically to the meaning representation, and further, they involve many unfired inferences. These rich naive theories are generalizations. Though not "true", they are and must be close enough to true, enough of the time, for

DAHLGREN

people to refer correctly to real objects and events, and to communicate using language. NS representations are rich and open-ended. Anything at all can be there. The number of feature values could be as large as the number of words in English [33]. Word meanings are seen as the names of concepts, and concepts are mental representations which may take many forms - visual, motor, tactile or verbal. Obviously, at the present stage of computation, we are limited to the verbal aspects. NS representations of noun concepts come from the results of psycholinguistic studies of object concepts in the prototype theory [29, 4]. In principle, such representations could be quite extended. In practice, we use the first 1.5 minutes of subjects' freelisting of properties. These are typed (e.g., color(red)). An example is *banker*, shown below in English translation.

Typically, a banker is a well-dressed educated male who works in an office in a bank. He is trained in mathematics. He is dedicated, civic-minded, and has high status. He functions as a financier who lends money. Inherently a banker is a person in authority whose function is to engage in business and handle money.

For verbs, we use the approach of Graesser [12], where it was shown that people conceive of actions in terms of their implications, such as cause, goal, result, location, manner and so on. We also classify verbs aspectually after Vendler [23]. The verb entry is based upon typical and inherent implications. An example is *invest*:

Typically, investing is done with capital in the form of money or other asset. One invests in stock, commodities, and real estate. Later, one may sell it or use it as collateral. Inherently, a sentient invests with the goal of making a profit.

NS provides a means of representing the world knowledge attached to English words in a general, non-*ad hoc* way, resulting in transportable representations. From the engineering point of view, these representations, though painstaking, are in fact feasible, and they go a long way toward providing the information necessary for syntactic disambiguation [7], word sense disambiguation [5], relevance reasoning and coherence [6]. Fortunately, it is not necessary to encode all of commonsense knowledge in order to achieve significant and useful results in text understanding. With our independently derived representations, we succeeded with all of the CRA's we need NS for in the Study I corpus. We found and confirmed that discourse cue phrases, syntax, compositional semantics (tense and aspect) and NS (or conceptual knowledge) all contribute to discourse coherence.

3.0 Coherence Relation Assignment Algorithm

The CRA algorithm was developed by examining the information-bearingness of each of the factors found in Study I relative to each other for each coherence relation. Those with a high informational load were included as factors to be considered during processing. In addition, we considered the most efficient ordering of tests for the factors. Information-bearingness results follow. Cue phrases such as *in order to* for **goal** and *because* for **cause** [3] are decisive for CRA where present, but are only present in 9% of local coherence and 16% of global. Similarly, certain specific lexical items such as the verbs *contrast* or *oppose* for the **contrast** relation are highly indicative, but rare in the data. As for syntax, most constructions are merely suggestive of coherence relations. For example, a main clause is more likely to

<p>Table I. Discourse Coherence Algorithm</p> <p>Source and Target cohere under Relation if</p> <p>syntactic tests return Relation, or</p> <p>connectives indicate Relation, or</p> <p>Relation = comment if comment tests succeed, or</p> <p>Relation = import if import tests succeed, or</p> <p>causal tests return Relation and</p> <p>not both Source and Target are stative and</p> <p>source is temporally before target</p> <p>or</p> <p>Relation = situation if situation activity tests succeed, or</p> <p>Relation = sequence if sequence tests succeed and</p> <p>Source and Target are telic and</p> <p>source is temporally before target.</p>
--

introduce an argument of certain coherence relations. However, such tendencies are not helpful in building an algorithm. A small number of syntactic structures, on the other hand, are decisive. These are the comparative for **contrast**, a generic sentence for a **generalization**, verb ellipsis for a **parallel** or **contrast**, relative clause, participial and appositive for **description**.

Turning to formal semantics, tense alone is not informative for CRA, because clauses in the simple past tense introduce events which can bear any coherence relation to other discourse events. Clauses in the simple present introduce events which bear all but one (**reported event**) of the coherence relations to other discourse events. However, temporal order, that is, whether or not the events or states introduced in two clauses overlapped in time, is informative. **Cause**, **goal**, **elaboration** and **comment** require temporal precedence, while **parallel**, **contrast**, **generalization**, **description**, and others can relate fully overlapping events or states. Thus lack of temporal order can be used to exclude the possibility of certain coherence relations.

Aspect is more decisive than temporal order in CRA. Aspect refers to the temporal perspective, the continuity and completion, of a clause. One of the two clauses introducing discourse events must be telic for certain causal coherence relations to hold between the events. For other relations to hold, clauses must be clause activity or clause stative. In NS, the aspect of a verb is listed as part of its lexical entry. But context affects aspect, so the aspect of the entire clause must be computed [25], taking into account factors such as progressive verb marking and quantified or unspecified subject or object. An algorithm for clause aspect assignment is under development.

A final factor which influences discourse coherence inferences is commonsense knowledge, which is required when a pair of clauses provide no or insufficient cue phrases, syntactic properties, temporal order or aspect information for CRA. An example is (4), which has two simple past tense clauses, both clause-telic, both main clauses with no discourse cues. These properties are consistent with **sequence**, **cause**, **goal**, **enablement**, **elaboration**, **import** or **comment**. Here NS can be used for CRA. The NS representation of the verb *invest* is powerful enough to drive the inference that $\text{goal}(e_1, e_2)$. In the corpus, independently derived NS representations are sufficient for CRA in all of the cases where it is needed.

DAHLGREN

Coherence Relation Assignment Algorithm. The local algorithm, shown in Table I, considers each clause (Source) in relation to the others in a segment one at a time (Target). Another (global) algorithm builds the segment tree [6]. The information the local algorithm uses are syntactic properties of the source clause, connectives in either the source or target, the temporal order of the events in the source and target, NS information associated with the verbs, semantic information such as types of adverbials, mood, and agentiveness in the clauses. The algorithm was hand tested in the original corpus with 97.5% accuracy. A test on an additional corpus of 8,000 words is in progress.

Applying the Algorithm. Finally, we step through the algorithm on a complex sample text, which is a paraphrase of one of the articles in our corpus.

Levine, (e_1) charged with SEC violations last May, (e_2) was convicted (e_3) and sentenced here yesterday. Levine (e_4) had engaged in extensive insider trading. He (s_5) was greedy and (s_6) wanted more money. Levine's light sentence (s_7) reflects (e_8) an attempt by the court to (e_9) reward cooperation in such cases. The judge (e_{10}) said that Levine's (e_{11}) cooperation (e_{12}) had influenced him in his favor. Critics (e_{13}) argued that light sentences (e_{14}) will result in more violations.

The coherence relations in this text that we explain are **reported event**(e_2), **sequence**(e_3, e_2), **situation activity**(e_1, e_2), **situation-activity**(e_4, e_2), **cause**(s_5, e_4), **goal**(s_6, e_4), **import**(s_7, e_2), **comment**(e_{10}, e_2), and **comment**(e_{13}, e_2). We use the notation C_i to denote the clause which introduces an event e_i or a state s_i . The first clause to be considered is C_2 in relation to the participial clause C_1 . Referring to the algorithm in Table I, we see that the main clause C_2 will designate a **reported event** because C_2 will fail the syntactic tests (it is not a relative clause, appositive, nor any of the syntactic structures the algorithm looks for). There is no connective, no verb of saying for the comment test, and no modal, conditional, interrogative or import verb for the import test. When the algorithm tries in the causal tests to prove that C_2 expresses a **reported event**, it will succeed. Next, the algorithm considers C_1 in relation to C_2 . Here tests succeed on the source clause C_1 . Since C_1 is a participial it must be either **description** or **situation**. As C_1 contains a time adverbial, it is designated **situation**. Note that the time adverbial in the main clause C_2 did not result in the same assignment. Next the algorithm considers C_4 in relation to C_2 . **Reported events** (C_2) are tried first as targets in commentary, because the commentary genre revolves around them. Considering C_4 in relation to C_2 , syntactic tests, connectives, comment tests and import tests all fail. There is only an indirect relation between breaking the law and being convicted, so causal tests fail. Now the algorithm tries **situation_activity**, and succeeds because C_4 is in the perfect. Next the algorithm tries C_5 in relation to C_2 . All tests up through import fail. Now the algorithm tries causal tests and finds in the NS representations that *greed* can cause people to break the law, so it assigns **cause**(s_5, e_2). Similarly, it finds that a typical goal of breaking the law is making money, so it assigns **goal**(s_2, e_4). Turning to C_7 in relation to C_2 , the syntactic, connective and comment tests fail. The import test succeeds because (s_7) overlaps (e_2) in time, and *reflect* is an import verb. Finally, the two comment clauses C_{10} and C_{13} are discovered because they fail the syntactic and connective tests, and they contain non-performative verbs of saying.

DAHLGREN

Notes

1. We would claim that this is true even in the interpretation of metaphors and fiction.
2. This point of view has evolved in discussions with N. Asher, C. Lord, J.P. McDowell, B. Partee and E.P. Stabler, Jr. and the Symposium on Discourse Coherence and Segmentation, University of Texas, 1989.

References.

1. Asher, N. 1987. A Typology for Attitude Verbs and their Anaphoric Properties. *Ling. and Phil.* 10:125-198.
2. Asher, N. 1989. *Abstract Objects and Anaphora in Semantics*. Manuscript.
3. Cohen, R. 1984. A Computational Theory of the Function of Clue Words in Argument Understanding. *COLING* 251-258.
4. Dahlgren, K. 1985. The Cognitive Structure of Social Categories. *Cognitive Science* 9:379-398.
5. Dahlgren, K. 1988. *Naive Semantics for Natural Language Understanding*. Boston: Kluwer Academic Press.
6. Dahlgren, K. 1989. Formal Properties of Discourse Segmentation. In preparation.
7. Dahlgren, K. and J. McDowell. 1986b. Using Commonsense Knowledge to Disambiguate Prepositional Phrase Modifiers. *Proc. AAAI 86*.
8. Dahlgren, K., J.P. McDowell, and E.P. Stabler, Jr. 1989. Knowledge Representation for Commonsense Reasoning with Text. Forthcoming in *Computational Linguistics*.
9. Davidson, D. 1967. Causal Relations. *J. Phil.* 64:692-703.
10. Fox, B. 1984. *Discourse Structure and Anaphora in Written and Conversational English*. UCLA Dissertation.
11. Graesser, A. 1981. *Prose Comprehension Beyond the Word*. New York: Springer-Verlag.
12. Graesser, A. and L. Clark. 1985b. The Generation of Knowledge-Based Inferences during Narrative Comprehension. In G. Rickheit and H. Strohner, eds., *Inferences in Text Processing*, Amsterdam: North-Holland.
13. Grosz, B. and C. Sidner. 1986. Attention, Intentions and the Structure of Discourse. A Review. *Computational Linguistics* 7:85-98. 12:175-204.
14. Hayes, P. J. 1985. The Second Naive Physics Manifesto. In J. R. Hobbs and R. C. Moore, eds., *Formal Theories of the Commonsense World*, Norwood, NJ: Ablex.
15. Hirst, G. 1981. Discourse-Oriented Anaphora Resolution: A Review. *Computational Linguistics* 7:85-98.
16. Hobbs, J.A. 1979. Why is Discourse Coherent? SRI Tech. Note #176.
17. Hobbs, J. 1985. On the Coherence and Structure of Discourse. CSLI Report # CSLI-85-37.
18. Hopper, P. and S. Thompson. 1980. Transitivity in Grammar and Discourse. *Language*, Vol. 56, pp. 251-299.
19. Johnson-Laird, P.N. 1983. *Mental Models*. Harvard U. Press.
20. Kamp, H. 1981. A Theory of Truth and Semantic Representation. In J. Groenendijk, Th. Janssen, and M. Stokhof, eds., *Formal Methods in the Study of Language*, Amsterdam: Mathematisch Centrum, 277-322.
21. Lockman, A. and A.D. Klappholz. 1980. Toward a Procedural Model of Contextual Reference Resolution. *Discourse Processes* 3:25-71.
22. Lord, C. and K. Dahlgren. 1989. Tracking Participants and Events in Newspaper Articles. In preparation.
23. McDowell J. and K. Dahlgren. 1987. Commonsense Reasoning with Verbs. *Proc. IJCAI*.
24. Mann, W. and S. Thompson. 1987. Rhetorical Structure Theory: A Theory of Text Organization. ISI Reprint Series: ISI-RS-87-190.
25. Moens, M. and M. Steedman. 1987. Temporal Ontology in Natural Language. *Proc. ACL* 1-7.
26. Partee, B. 1984. Nominal and Temporal Anaphora. *Ling. and Phil.* 7:243-286.
27. Polanyi, L. 1988. A Formal Model of the Structure of Discourse. *J. Pragmatics* 12:601-638.
28. Reichman, R. 1985. *Getting Computers to Talk Like You and Me*. Cambridge, MA: MIT Press.
29. E. Rosch and B. B. Lloyd, eds., *Cognition and Categorization*. New York: Erlbaum.
30. Sanford, A.J. and S.C. Garrod. 1981. *Understanding Written Language*. New York: Wiley and Sons.
31. Schank, R. C. and R. P. Abelson. 1977. *Scripts, Plans, Goals and Understanding*. Hillsdale, NJ: Erlbaum.
32. Schiffren, D. 1987. *Discourse Markers*. Cambridge U Press.
33. Schubert, L. K., R.G. Goebel, and N.J. Cercone. 1979. The Structure and Organization of a Semantic Net for Comprehension and Inference. In N.V. Findler, ed., *Associative Networks*, New York: Academic Press.
34. Van Dijk, T. and W. Kintsch. 1983. *Strategies of Discourse Comprehension*. New York: Academic Press.
35. Wada, H. and N. Asher. 1986. BUILDERS: An Implementation of DR Theory and LFG. *Proc. COLING* 540-545.
36. Wilks, Y. 1975. Preference Semantics. In E. Keenan, ed., *Formal Semantics of Natural Language*. Cambridge U. Press.