

# ANALYZING THE STRUCTURE OF ARGUMENTATIVE DISCOURSE

Robin Cohen

Department of Computer Science  
University of Waterloo  
Waterloo, Ontario, Canada N2L 3G1

Consider a discourse situation where the speaker tries to convince the hearer of a particular point of view. The first task for the hearer is to understand what it is the speaker wants him to believe – to analyze the structure of the argument being presented, before judging credibility and eventually responding.

This paper describes a model for the analysis of arguments that includes:

- a theory of expected coherent structure which is used to limit analysis to the reconstruction of particular transmission forms;
- a theory of linguistic clues which assigns a functional interpretation to special words and phrases used by the speaker to indicate the structure of the argument;
- a theory of evidence relationships which includes the demand for pragmatic analysis to accommodate beliefs not currently held.

The implications of this particular design for dialogue analysis in general are thus:

- structure is an important feature to extract in a representation to control the processing;
- linguistic constructions can be assigned useful interpretations;
- pragmatic analysis is crucial in cases where the participants differ in beliefs.

## 1 THE PROBLEM AREA

Consider the task of designing an “argument understanding system”, a natural language understanding system (NLUS) where the input is restricted to arguments. Consider as well arguments constructed in a dialogue situation, where a speaker (S) tries to convince a hearer (H) of a particular point of view. The hearer patiently listens; hence, the input is “one-way communication”. The argument understanding system therefore plays the role of the hearer, and tries to analyze the structure of the argument being presented. This task is isolated as a necessary first step for a hearer, in order to be a successful participant in a conversation. In other words, the hearer must have some representation of what it is the speaker wants him to believe, before judging credibility and eventually responding.

Note that this language problem is relatively new and yet feasible. It is distinct from other NLU endeavors, such as story understanding, which appeal to a stereotype of content in order to reduce processing. In arguments, one is never sure what points the speaker will address; content can't be stereotyped. However, arguments have a defining characteristic – they are necessarily goal-oriented.

The speaker wants to convince the hearer of some overall point. Thus, there is an overall logical structure to the input and this fact may be used by a hearer to control analysis.

For our model, the representation for the structure of the argument is restricted to an indication of the claim and evidence relations between the propositions. The notion of evidence is discussed in more detail in section 4. A useful starting definition is: “A proposition E is evidence for a proposition C if there is some rule of inference such that E is premise to C's conclusion – in other words, there is some logical connection between E and C”.

In order to design an argument understanding system, what is then required is a computational model for the analysis of arguments. This in turn necessitates a theory of argument understanding, as a basis for the model. We suggest the following three components for the model:

- a theory of expected coherent structure, used to drive a restricted processing strategy. Analysis is kept to a computationally reasonable task, by limiting the input to be recognized to a characterization of expected coherent forms of transmission.

Copyright 1987 by the Association for Computational Linguistics. Permission to copy without fee all or part of this material is granted provided that the copies are not made for direct commercial advantage and the CL reference and this copyright notice are included on the first page. To copy otherwise, or to republish, requires a fee and/or specific permission.

0362-613X/87/010011-24\$03.00

- a **theory of linguistic clue interpretation**, including insight into both the occurrence of clue words and their possible function in overall discourse. Clue words are those words and phrases used by the speaker to directly indicate the structure of the argument to the hearer (e.g. connectives). Detecting clues can thus also serve to constrain the processing for the hearer. Moreover, it is important to have some facility for recognizing and interpreting clue words, to build a model that is robust enough to process a wide variety of input.
- a **theory of evidence relationships**. The most important observation is that pragmatic analysis is mandatory for an analysis model in order to recognize beliefs of the speaker, not currently held by the hearer. Evidence connections between propositions often appeal to unstated information not currently in the hearer's set of beliefs, but recognizable as an intended support relation on the part of the speaker.

## 2 RESTRICTED PROCESSING STRATEGY

Consider the following framework for the model. An argument is considered to be a set of propositions. The model is then designed to analyze the argument a proposition at a time, incrementally building a representation for the underlying structure. The representation developed is a tree of claim and evidence relations comprising the argument, where a claim node is *father* to its evidence *sons*. In order to assign an interpretation for a given proposition, one must thus simply assign it a place in the tree. In this way, one can tell to which propositions it serves as evidence and from which other propositions it receives support.

A key design decision is to separate the two main operations of determining for each proposition (i) *where* it fits with respect to the argument so far, and (ii) *how* it relates to some prior proposition. The question of how two propositions relate in this framework is one of verifying that an evidence relation holds between the propositions. This task is extracted and relegated to an **evidence oracle**, which, passed two propositions A and B, will respond "yes" or "no", as to whether A is evidence for B.

With the problem of evidence determination factored, the model must still cope with the question of where a proposition may fit. This is handled by characterizing possible coherent transmissions (ordering of propositions) on the part of the speaker and then limiting analysis to **reception algorithms** designed to recognize these coherent transmissions. In the section below we illustrate possible coherent strategies from the speaker, and present the associated reception algorithm required to recognize the input.

Note that the computational model for the analysis of arguments is designed with certain aims and limitations. In particular, the model is to provide for analysis of "spontaneous discourse", demanding construction of a representation for the argument as each new statement is

processed. The following restriction to the processing model is thus applied: the processor does not weigh all possible interpretations for a proposition; if the oracle sends back a yes answer to the question *Is P evidence for Q?*, then P is attached as a son to Q in the tree. In other words, the model simulates a hearer who does not have the luxury of "looking back" and re-interpreting previous statements. Moreover, the model aims to provide an interpretation for the current proposition, so that once an interpretation has been found (e.g. that P supplies evidence for Q) that proposition has been processed. The evidence relation is also taken to be transitive – i.e., if P is evidence for Q and Q is evidence for R, then P is evidence for R. (See Section 4 for more detail on the evidence relation).


### 2.1 PRE-ORDER

Pre-order transmissions are those where the speaker consistently presents a claim and then states evidence. The example below illustrates this form:

EX1:

- 1) Jones would make a good president
- 2) He has lots of experience
- 3) He's been on the city board 10 years
- 4) And he's very honest
- 5) He refused bribes while on the force

With the tree representation:




### 2.2 POST-ORDER

Another coherent strategy is post-order, where the speaker consistently presents evidence and then states the claim. Consider the comparable example below:

EX2:

- 1) Jones has been on the board 10 years
- 2) He has lots of experience
- 3) And he's refused bribes
- 4) So he's honest
- 5) He would really make a good president

With the representation:



In this example, the claim is the same as in EX1, that *Jones would make a good president*, but the evidence precedes the claim in each case – i.e., 1 is evidence for 2; 3 is evidence for 4; and 2 and 4 together are evidence for the final claim 5.

The associated reception algorithms for pre- and post-order are described in detail in Cohen (1981). Both can be shown to operate in linear time (linear in the number nodes of the tree) and so are quite efficient. The pre-order reception essentially finds an interpretation for the current proposition of an argument by searching for a father, up the right border of the tree. The post-order algorithm employs a stack; the current proposition tests to be father to the top of the stack; then all sons are popped off the stack, and the resulting tree pushed on the stack.

We trace the construction of the trees of EX1 and EX2 below, to provide details of the pre-order and post-order reception algorithms.

**EX1: PRE-ORDER**

Algorithm:

- The current proposition is NEW.
- The proposition immediately prior is LAST.
- 1) Try NEW as evidence for LAST.
- 2) If that fails, try NEW as evidence for each of LAST's ancestors, in turn, up to the root of the tree.
- 3) If the test in 1 succeeds, stop.

Consider as well a dummy root (D) for the tree for which all nodes are evidence (to place the first proposition).

candidates in order	test for NEW	evidence oracle response
		1,D
2,1,D	3 ev. for 2?	yes
3,2,1,D	4 ev. for 3?	no
	4 ev. for 2?	no
	4 ev. for 1?	yes
4,1,D	5 ev. for 4?	yes

**EX2: POST-ORDER**

Algorithm:

- Keep a stack of elements eligible to be evidence for a current proposition, with the latest one as TOP.
- To interpret the current proposition NEW:
  - 1) Test if TOP is evidence for NEW
  - 2) If yes, then pop TOP (TOP:= TOP - 1) and make it a son of NEW (build a tree under NEW); then repeat step 1 with NEW value of TOP.
  - 3) If no, make the tree with NEW as root the new TOP of stack (push NEW onto stack). In essence, all sons for a proposition are picked up at once.

stack elements	test for evidence	evidence oracle response
1	1 ev. for 2?	yes
tree:1 under 2	2 ev. for 3?	no
3, 1 under 2	3 ev. for 4?	yes
	2 ev. for 4?	no
3 under 4, 1 under 2	4 ev. for 5?	yes
	2 ev. for 5?	yes


**2.3 HYBRID**

As a first approximation to a general processing strategy, consider designing a reception algorithm to accept hybrid pre- and post-order arguments (i.e., any given sub-argument may be transmitted in pre- or post-order). An example of hybrid transmission is EX3 below.

**EX3:**

- 1) Jones would make a good president
- 2) He has lots of experience
- 3) He's been on the board 10 years
- 4) And he's refused bribes
- 5) So he's honest

With the representation:



Here, the first sub-argument is in pre-order, the second one in post-order, and the argument overall is still coherent.

The reception algorithm for accepting hybrid transmissions is basically a combination of techniques from the pre- and post-order receptions. Now, to process a current proposition both a father and possible sons must be searched for. But the search is still restricted – certain propositions get closed off as possible “relatives” to the current one (e.g., earlier brothers of an ancestor). Thus, the complexity of the algorithm is still reasonable; it can also be shown to be linear. Once more, see Cohen (1981) for further examples.

A full description of the hybrid algorithm is included below.

L is kept as a pointer to the lowest node of the tree still eligible to receive evidence. It is initially set to a dummy root node to which all other nodes succeed as evidence. Consider as well a labelling where the last proposition in the stream that succeeds as evidence for a node is stored as the rightmost son.

For each node NEW in the input stream do:

```

/* find father */
do while (NEW is not evidence for L
         and L is not dummy root)
    set L to father of L
end;
    
```

```

/* see if there are any sons to re-attach */
if (rightmost son of L is not evidence for NEW )
/* no sons to re-attach */
then do;
    attach NEW to L
    set L to NEW
end;
else
do;
/* some son will re-attach; attach all sons of L
which are evidence for NEW below NEW */
do while (rightmost son of L evidence for NEW)
    attach rightmost son of L to NEW
    remove rightmost son of L from L
end;
attach NEW to L
end;

```

Note that the hybrid algorithm, in searching for both sons and father to the current node, must contend with cases where a proposition is attached to a higher ancestor and must be re-attached to its immediate father. This occurs in a framework where the evidence relation is considered transitive. The kinds of orderings that involve this “re-location” are of the form: A, C, B – where C is evidence for B, B evidence for A (hence C also succeeds as evidence for A when tested).

For EX3:

candidates	test for evidence	evidence oracle response
D	1 ev. for D?	yes
1, D	2 ev. for 1?	yes
2,1,D	3 ev. for 2?	yes
3,2,1,D	4 ev. for 3?	no
	4 ev. for 2?	no
	4 ev. for 1?	yes
4,1,D	5 ev. for 4?	no
(3&2 “closed off”)		(4 is a case of 5, so can't assume 5 ev. for 4)
	5 ev. for 1?	yes
	4 ev. for 5?	yes
	(testing re-attach of sons)	

#### 2.4 SUMMARY OF PROCESSING STRATEGY

The processing strategy proposed for our model of argument analysis is designed to produce a selective interpretation. The particular restrictions to processing chosen for the model are, moreover, both

- useful for measuring the efficiency of the model, since expressed in a framework of an algorithm operating on a tree structure, where complexity of the algorithm can be studied and
- well-motivated, since based on a characterization of coherent input.

In fact, it is an expression of a theory of argument coherence that serves to produce a model that is both efficient and robust (can handle a wide variety of input).

Note that the particular restrictions suggested are drawn from analysis of a body of examples from rhetoric texts, together with some “naturally occurring” arguments (letters to the editor in newspapers). The aim is to characterize coherent transmission forms (ordering of propositions) that can be understood, without additional “clue words” to the underlying structure. We feel that the hybrid model is a good first approximation because examples of various forms of hybrid were encountered, and exceptions to the form all involved additional clue information. This leaves the study of clues and possible extended transmission forms to the next section. (In Cohen (1983) a few longer examples are run through the model to illustrate the forms it can accept and to motivate provision for the recognition of a wide variety of argument forms. Note that an actual implementation was not produced. There is now a scaled-down first implementation (Smedley 1986), which will be discussed further in section 6.)

### 3 LINGUISTIC CLUES

The second main component of the argument understanding model is a theory of linguistic clues. These are the words and phrases often used by the speaker to directly indicate the structure of the argument to the hearer. It is important to specify

- what kinds of clues exist
- the function of these clues in analysis – i.e., what interpretation can be assigned to a proposition that contains a clue word, and
- (a more difficult question) when clues are necessary to ensure comprehension of the argument structure by the hearer.


This approach to the study of clue words is much more detailed than the initial suggestions of Hobbs (1976) on how to interpret a few connectives such as *and* in his framework. It is also distinct from the investigations of Reichman (1981) and (recently) Grosz and Sidner (1986). Grosz and Sidner acknowledge the existence of clues and discuss various discourse structures that can be formed in the presence of clues. Reichman also gives a longer list of clue words and the particular conversational moves they signal. But there is no systematic proposal for interpreting a clue word that may occur (classification), and there is little discussion of how to process some of these more complex discourse structures without clues (suggesting when clues are necessary). In this section we clarify more deeply some of the discussion of clues in Cohen (1983, 1984b).

#### 3.1 CLUES OF RE-DIRECTION

One type of clue is expressions that specifically re-direct the hearer to an earlier part of the argument. Consider:

**EX4:**

- 1) The city is a mess
  - 2) The parks are a mess
  - 3) The park benches are broken
  - 4) The grassy areas are parched
  - 5) Returning to city problems, the highways are bad too
- With the representation:



Here, the clue in proposition 5, *returning to city problems* signals additional support for proposition 1. In the absence of a clue, the specifications for the hybrid algorithm would have 5 test to be a son for 4, 2, and 1 (up the right border of the tree). So, adding clues should reduce the processing effort of the hearer. (In fact, if clues are consistently used by the speaker after long chains of evidence, the processing complexity of the reception algorithm can reduce from linear to real-time).


3.2 CONNECTIVES

Another popular type of clue word is the connective. We present a classification or taxonomy of connectives, and then associate a common interpretation rule within the claim and evidence framework for each category of the taxonomy. In this way, the interpretation of any proposition containing a connective can be determined on the basis of the clue word's classification. For example:

**EX5:**

- 1) This city is a disaster area
- 2) Houses have been demolished
- 3) Trees have been uprooted
- 4) As a result, we need national aid

With the representation:



Here, the connective *as a result* in proposition 4 belongs to a category known as inference, indicating that there should be some prior proposition that connects to 4 and serves as son to 4 – i.e., supplies evidence for “the result”. In this example, in fact, 1 acts as the son. The interpretation rules are necessarily default suggestions for translating the semantic relations between propositions into our claim and evidence classification. (See further discussion on the evidence relation (section 4) for possible exceptions).

The taxonomy and its associated interpretation rules are presented below.

Each category is a classification of connectives, made on the basis of semantic meaning of the connective – e.g., the parallel category would include all words that extended a list, including *next, then, first, secondly, thirdly*, etc. The set of classes was produced by considering the categories proposed in Quirk et al. (1972), and merging some classes that had similar semantics and suggested the same interpretation rule for claim and evidence. The inference category covers phrases that suggest one proposition can be inferred from another – e.g., *as a result, because of this*, etc. The detail category moves in the other direction, and includes connectives that specify further – e.g., *in particular, specifically*, etc. The summary category is used for phrases that conclude a list. Reformulation captures clue words that repeat an earlier idea – e.g., *in other words, once more*, etc. Finally, contrast covers the phrases that introduce comparisons, like *on the other hand* or *but*.

In conjunction with the semantic classes, evidence interpretation rules were then assigned as default interpretations. (See Section 4.3 for pragmatic “overrides” to the defaults). Note that some words may fall into more than one category, based on the meaning used – e.g., *then* meaning ‘next’ (in a list of actions) compared with *then* meaning ‘as a result’. Clue interpretation thus requires a classification process as well.


P is prior proposition; S is the proposition with the clue

CATEGORY	RELATION: P to S	EXAMPLE
parallel	brother	First, Second
inference	son	As a result
detail	father	In particular
summary	multiple sons	In sum
reformulation	son (& father)	In other words
contrast	brother or father	On the other hand

The taxonomy is described in detail in Cohen (1984b). We include discussion of one category here, as an example. The **detail** class is one case where connectives with a range of meanings were merged into one category. The title “detail” suggests that the connective will further specify some prior proposition. Included cases are: *for example, in particular, and as another instance*. The interpretation rule assigned to this category is that the proposition with the connective provides evidence for the earlier connecting statement. The motivation is that an accumulation of specific cases leads to a conclusion of a general statement (a form of reasoning used very often in naturally occurring arguments).

**EX6:**

- 1) The people in this town deserve a city-wide holiday
- 2) For example, Old Man Jones has worked non-stop since Christmas
- 3) And Mayor Flood is still recovering from all his efforts for the tornado relief
- 4) In short, all of us are tired



2 is son to 1 (detail class); 3 is also son to 1 (brother to 2) (parallel class); 4 is father to 2 and 3 (summary class):

3.3 THE FUNCTION OF CLUE WORDS


So far we have discussed two types of clue words that can occur in conjunction with arguments transmitted according to the specifications of the hybrid algorithm presented in section 2 (our characterization for coherent discourse). We point out that re-direction clues provide additional information concerning which of the eligible propositions is related to the current one, and that connective clues specify the kind of relation that must be found between the current proposition and one of the eligible priors.

In certain cases these restrictions will prevent some of the tests for evidence that would otherwise have been conducted, thus saving some processing effort. For example, consider the following:

EX7:

- 1) The city is in serious trouble
- 2) There are some fires going
- 3) Three separate blazes have broken out
- 4) In addition, a tornado is passing through

with the representation:



The clue in 4 prescribes an interpretation for 4 as brother to some prior proposition. This means that 4 must act as evidence for some different proposition. Thus, even though 3 is technically the first proposition to test out when interpreting a new proposition (NEW evidence for LAST is the test), in this particular case this option is not possible. Thus, one round of work for the evidence oracle is avoided. In fact, for this example, the test “4 evidence for 2” fails, and the final test of “4 evidence for 1” succeeds. (Note that the brother relation is tested by way of son relation to a (common) father. This is because the model only processes evidence relationships).

3.3.1 THE NECESSITY FOR CLUES

We now examine the use of clue words in conjunction with transmissions that violate the specifications of the hybrid base case. The hypothesis is that more complex transmissions can be accommodated by the argument analysis model provided there exist clues to assist the hearer in recognition. In these cases, the clues are there by necessity. Their function in the discourse is not to

merely add detail on the interpretation of the contained proposition, but to allow that proposition an interpretation that would otherwise be denied.


There is an advantage to adhering to the base case of acceptable argument structures and treating the use of clues with other transmission forms as exceptional. In the first place, this provides a framework for detecting one interpretation for an argument when another possible interpretation could be generated if further testing occurred. In other words, in this model a representation drawn using the rules of the hybrid reception will always be accepted as the intention of the speaker, unless clues specifically override possible tests.

To explain, consider the following example:

EX8:

- 1) The park benches are rotting
- 2) The parks are a mess
- 3) The highways are run down
- 4) (And another problem with the parks is that) the grass is dying
- 5) This city is in sad shape

Without the clue phrase in 4, re-directing to proposition 2, to add more evidence out of turn, a coherent representation could be built just the same, as below:




If the speaker intends 4 to add detail to the parks problem, he cannot expect the hearer to make this connection without more information, simply because a more effortless interpretation of 4 is possible (and the speaker should realize that it is this representation that the hearer will draw).

Another important reason for separating the base case of acceptable transmissions is to allow for input that can be characterized as somehow a coherently generated plan of the speaker. Recall that the proposition analyzer will continuously call on an evidence oracle to determine if some proposition A acts as evidence for some other proposition B. Suppose there were no restrictions in the testing of evidence relations. Then, tests for evidence that would be interpreted as positive would return this response, regardless of when asked. In other words, a totally random display of propositions would result in the same representation for the argument as the reception of a coherently ordered presentation. Consider the following example:

EX9:

- 1) Yogi's been a shrewd manager
- 2) He hired a hitter who now bats .400
- 3) He traded in a pitcher who is now 0 and 12
- 4) But he got involved in drug deals to the players
- 5) And that's inexcusable for a manager
- 6) He really needs to be axed

with the representation:



(where 5 and 1 are contrasting evidence to 6).

**EX9B:** The argument as above, presented in the order::  
2,5,3,6,4,1

- 1) Yogi hired a .400 hitter
- 2) (And) that's inexcusable for a manager
- 3) He traded in a pitcher who is now 0 and 12
- 4) Yogi really needs to be axed
- 5) He got involved in drug deals to the players
- 6) He is a shrewd manager

This argument now appears incoherent. One reason is that there is contrast overall that must be clearly separated. In any case, the ordering does not conform to the specifications of the hybrid and as such is a candidate for an unacceptable transmission.

This example suggests that the use of clue words, though helpful to signal exceptional transmissions, should still be studied as a systematic process of the speaker to assist the hearer in comprehension. In EX9B, could any number of clue words be added to still make this recognizable? Consider the following attempted repair to EX9B:

**EX9C:**

- 1) Yogi hired a .400 hitter
- 2) But he's done some things inexcusable for a manager
- 3) Although he also did other smart things like trading a pitcher who is now 0 and 12
- 4) No, Yogi really has to be axed
- 5) He got involved in drug deals to the players
- 6) Though he still is a shrewd manager

The question is whether an argument of this form would still be judged coherent. Our preference will be to specify particular types of exceptional transmissions that may be judged coherent. To this end, we have tried to isolate a few specific cases where clues can be used in conjunction with coherent orderings beyond the specifications of the hybrid algorithm. These are highlighted in the next subsection below.

One more point about the last example above is that it emphasizes our concern that the analysis of arguments be done efficiently. In this case, we want to avoid making tests for evidence relations that could not exist between certain propositions as part of a coherent input. In other words, we don't want the model to simply test out all possible combinations of evidence relations (even though this would only be  $n*n$  vs.  $k*n$  number of operations), because an input that is not coherent would then be accepted.

We would also not want to derive computationally a representation for an argument that involved more computational effort, if a simpler interpretation of the

same argument could be derived. (Again, the speaker is to assume that the hearer will not be searching unnecessarily). This is illustrated in EX8 above.

### 3.3.2 CASES OF NECESSITY


Three kinds of acceptable extended transmission strategies are studied in Cohen (1983): **parallel evidence**, **multiple evidence** (a proposition acting as evidence for several claims, in restricted conditions) and **mixed-mode sub-arguments** (with evidence both preceding and following a claim). We present an example of the parallel construction below. See Cohen (1984b, 1983) for further examples.

**EX10:**

- 1) The city has problems
- 2) The parks are a mess
- 3) The highways are a mess
- 4) The buildings are a mess
- 5) Here's some evidence for the fact that the parks are a problem: the benches are broken
- 6) As for the highways, they're full of potholes

...

With the representation:



Here, the argument breaks the rules of hybrid transmission by adding evidence for an earlier brother (proposition 2); however, this shift is signalled with a phrase of intention in proposition 5, and the hearer may then expect a parallel expansion of additional support for each of the earlier brothers, in turn. (Note that without the clue, the argument structure derived would simply record all of 2 to 6 as evidence for 1, in the same vein as EX8).

This example illustrates another interesting feature of clues – the variety of possible surface forms that can signal the same evidence relation between propositions. In EX10, the clue in 5 could also have been *The problem with the parks is...* or *I will now explain why the parks are such a problem – ....* A range of explicitness is thus possible. In cases other than this parallel construction, in fact, a signal to a relation between propositions may be advocated by the simple use of an anaphor. For example:

**EX11:**

- 1) The mayor hasn't done much for this city
- 2) He doesn't seem to want to do much
- 3) That man is a complete loser

Here, the phrase *that man* signals a link to *the mayor*. It is difficult to decide whether the phrase qualifies as a clue word. The problem is determining a "bottom line" – i.e., "can't every sentence be seen to have some clue, from semantics, to its interpretation within the argument?". For now, we do not consider the cases of anaphora as above.

### 3.4 SUMMARY AND FUTURE DIRECTIONS OF CLUE THEORY

Our theory of clue interpretation so far has outlined the following principles:

- Clues may occur with expected coherent transmissions or to signal exceptional cases.
- Connective clues can be assigned common interpretation rules according to the semantics of the clue.
- To distinguish helpful versus necessary clues, the preference will always be to recognize the hybrid transmission; if clue rules or semantics force an exceptional reading, only certain exceptional structures should be accepted.
- In all cases, a reading that can be derived with less computational effort will always be taken as the intended reading.
- The cases where clues are necessary to force a certain interpretation provide insight into the function of clues; their use in conjunction with acceptable transmissions suggests a function of additional processing reductions.

In order to recognize clues and incorporate their interpretation into a larger model that derives argument structure, we propose a separate clue interpretation module, interacting with the basic reception algorithm and the calls to the evidence oracle. Exactly how this module would function is left as future work. We do have a few initial insights, to suggest that clue interpretation is not only quite useful (argued previously in this section) but feasible.

For connectives, the clue can be recognized from a classification. Then, determining whether a related proposition available from the list of eligibles is in fact related can be tested, according to the restrictions of the interpretation rule (as suggested in EX7). Further, the required semantic relation to the prior proposition can be passed as additional information to the oracle. The problem is that this oracle must do some kind of search for connections between facts and axioms of a knowledge base. How this semantic analysis is achieved depends on the underlying representation, but additional semantic constraints should help to restrict operations.

For re-direction clues, a processor would first have to recognize the appropriate phrase used. Some standard list (e.g., *returning to, on the topic of*) may be specified as a start. Then, the phrase should suggest some particular semantic content to the prior proposition (e.g., *returning to the parks* mentions *parks* as central). Now it is the form of representation of the propositions which may influence what is acceptable on a list of eligibles. If this semantic processing can be done very broadly, some calls to the oracle may be avoided, and this would be an improvement, assuming the oracle's operations encompass a more extensive search.

## 4 EVIDENCE DETERMINATION

The third main component of the argument analysis model is a theory of evidence, to govern the verification of evidence relations between propositions.

An initial definition for evidence offered in Section 1 is: "A proposition P is evidence for a proposition Q if there is some logical connection from P to Q – i.e., some rule of inference such that P is premise to Q's conclusion". The main problem in establishing evidence relations is that not all the premises are stated. For example, one common rule of inference used in arguments is Modus Ponens, of the form:  $P \rightarrow Q, P$  therefore  $Q$ . The way this rule is most often used, the speaker will simply state P and Q and leave out the major premise " $P \rightarrow Q$ ", expecting the hearer to be able to fill in the unstated connection to recognize the evidence relation from P to Q. Omitting certain premises is referred to as Modus Brevis and studied in Sadock (1977).

We list below the rule of inference **frames** included for our model. Each rule has a slot for major premise, minor premise, and conclusion, to be filled by stated or unstated propositions, in recognizing an evidence relationship.

#### RULE OF INFERENCE FRAMES:

	Major	Minor	Conclusion
Modus Ponens	$P \rightarrow Q$	P	Q
Modus Tollens	$P \rightarrow Q$	$\sim Q$	$\sim P$
Modus Tollendo Ponens	$P \vee \sim Q$	Q	P
Modus Ponendo Tollens	$P \vee Q$	Q	$\sim P$

The form most often used is Modus Ponens. When the major premise is missing, this is the rule of inference to consider as the intended link from P to Q.

#### EX12:

- 1) The Jays had a fantastic team this year
- 2) All their players had averages over .250

fill:

- 3) If a team has all players with averages over .250, then that team is fantastic

The common form for arguments, then, is one where the hearer must supply missing statements in order to establish the connections for the representation of the argument.

There are several possible Modus Brevis forms for each frame above. The possible missing parts are classified below for the case of Modus Ponens.

#### MODUS BREVIS FORMS (MODUS PONENS):

	Given	Conclusion
Normal	$P \rightarrow Q, P$	Q
Missing Minor	$P \rightarrow Q$	Q
Missing Major	P	Q (most popular form)
Only Major	$P \rightarrow Q$	(assume rest)
Only Minor	P	(assume rest)

#### 4.1. OVERVIEW OF ORACLE'S PROCESSING

It is important to demonstrate that the part of processing relegated to the evidence oracle within the overall model is not insurmountable, to defend the model as useful. In this section we examine more closely the operation of the



oracle, opening up the “black box” just enough to suggest how it operates.

In general, the oracle is asked to test two propositions to be in an evidence relationship, responding “yes” or “no” to a question of the form: “is P evidence for Q?”. We sketch the operation of the oracle for the example below:

**EX13:**

- 1) Joey is a dangerous
- 2) He is a shark

Assuming some resolution of anaphora, etc., a crude representation of the example in terms of predicates and arguments is:

P: Shark(j); Q: Dangerous(j) (j: Joey).

A Modus Ponens template would be of the form: S(j), for-all (x) ( S(x)  $\rightarrow$  D(x) ), therefore D(j) (S:is-shark, D:is-dangerous). (Note that we are not addressing certain questions here such as the significance that *is a shark* is definitional, while *is dangerous* is really assertional).

Recall that the argument analysis model seeks to recognize intended argument structures. So, in this example the hearer can at least recognize that “P is evidence for Q” would follow through if *All sharks are dangerous* were believed (i.e., for-all(x) (S(x)  $\rightarrow$  D(x) )).

It is thus proposed that the oracle:

1. identify missing premises (the Modus Brevis form of the argument being presented);
2. verify plausibility of these missing premises (that they could be intended by the speaker to be believed by the hearer); and
3. conclude that an evidence relation holds if the missing premises are plausible.

For step 2, we try to specify more precisely in the remainder of this section the kind of tests the hearer can apply. A summary is provided below:

- a) Identify the missing premise within a knowledge base of shared knowledge.
- b) Identify a “relaxed version” of the missing premise within own private knowledge.
- c) Identify the missing premise within a model of the speaker’s beliefs.
- d) Judge the beliefs of a hypothetical third party (which could be simplified if the bottom line is “believe it, unless there’s reason to strongly doubt, from within one’s own beliefs”).

We have found in simulations of the model on a variety of examples that most of the tests for evidence can be answered through (a) and (b). We hypothesize that, given a specification of a knowledge base, the search for connections between propositions can be controlled. This hypothesis would be best verified with an implementation of the oracle and extensive analysis of examples, and could be the focus of the next stage of our implementation (beyond Smedley (1986)). (The two large examples dissected in Cohen (1983) do have this property.)

In addition to the problem that the major premise may be unstated is the problem that this premise should really be tempered by the beliefs of the speaker. In other words, the missing major premise that the hearer must fill in is really of the form: H believes that S wants H to believe (P  $\rightarrow$  Q). In other words, this premise is not necessarily one of the hearer’s beliefs. It is important to emphasize the importance of pragmatic processing in establishing evidence relations. The tree of claim and evidence relations built as a representation for the argument is really an indication of the plan of the speaker, in the sense that each evidence relation recorded is the hearer’s conception of a support connection intended by the speaker.

Note that it is difficult to specify how a plan of a speaker is determined during analysis. What we are advocating is to interpret the intention of each proposition of the argument, the other propositions for which it provides evidence. The result of processing the entire discourse is not a complete plan of the speaker, in the sense that each of the “steps” could be executed and the top level goal (convince the hearer of some overall point) would then follow. It is more an indication of the motivation behind each utterance towards the ultimate goal of convincing the hearer. The difficulties in plan inference for discourse are discussed in more detail in Grosz and Sidner (1986), and are in fact a topic of our current concern (see discussion of future work in Section 6).

There is in fact a whole spectrum of problems the hearer must face in recognizing evidence relationships between propositions. The four main tests for the hearer can be described as:

- use logic,
  - relax the logic,
  - stereotype the speaker,
  - judge plausibility (reason as a “hypothetical person”).
- We illustrate these possible operations with examples below.

#### 4.2 LOGIC AND RELAXED LOGIC

In example EX14, all the premises of the Modus Ponens argument are present. The hearer should realize that 1 and 2 are evidence for the claim in 3. Then EX15 illustrates the more typical case of missing major premise. If the hearer fills in *All machine candidates win*, the connection from 1 to 2 can be seen. The problem is that the speaker probably believes something more along the lines of: *Most machine candidates win*. And yet, one couldn’t record a Modus Ponens relation in the argument with a quantifier such as *most* instead of *for all*. Thus, the hearer must use some relaxation to the rules of logic to recognize the evidence relation. The detection of evidence through “relaxed logic” can be accomplished by having the hearer judge the unstated connection as a plausible generalization, based on a few known cases. For example, if the hearer tries to recognize an evidence relation from 1 to 2 in EX16 below, by filling in *All*

*sharks are dangerous*, and the hearer doesn't believe this "axiom" but knows of a few sharks that are dangerous, he may reason that the missing major premise is reasonable.

**EX14:**

- 1) Aristotle is a man
- 2) All men are mortal
- 3) So Aristotle is mortal

**EX15:**

- 1) Bilandic will win
- 2) He's the machine candidate

**EX16:**

- 1) Joey is a shark
- 2) So, he is dangerous

The idea of recognizing an intended connection from some other conversant based on one's own beliefs is not necessarily simple to implement. There has been some recent work by Pollack (1986) that suggests a more concrete foundation for this operation. Pollack discusses the problem of inferring a questioner Q's plan from his discourse. A first process has the responder R ascribing to Q a belief about some connection ("conditional generation relation") that she herself believes true. Occasionally, R will need to recognize a connection that is not one of her beliefs. Then Pollack suggests there is a rule where "R ascribes to Q a belief about a relation between act-types that is a slight variation of one she herself has". In particular, one slight difference possible has Q believing a stronger conditional generation relation, which is missing one of the required conditions.

This related research is quoted here, not to argue that this problem is solved, but to acknowledge that it is important to specify this "relaxed" connection between one's own beliefs and those attributed to another. Some groundwork is being laid with more formal descriptions of plans such as Pollack's.

#### 4.3 DIFFERENCE IN BELIEFS

The other type of problem faced by a hearer in recognizing evidence relations arises because of difference in beliefs between speaker and hearer. As mentioned, the hearer is actually discerning intended relations on the part of the speaker, and must often reason outside his personal framework of knowledge. Again, an issue is raised of how to discern intentions from discourse. Some work has been done at the level of one utterance (e.g., Allen 1979). We are mostly concerned with advocating the inclusion of reasoning beyond one's own beliefs, without a full theory of how to infer another person's beliefs. Instead, we advocate a simplified framework, discerning evidence relations and allowing a connection to be drawn as intended if it is plausible to the hearer. For future work, we are studying how to specify this process more precisely. (See also Grosz and Sidner 1986).

The point is that the hearer is still able to recognize connections he does not believe. In EX17, the hearer should be able to understand an evidence relation from 1 to 2, upon filling in a missing premise of the form "If a person stands for apple pie and Mom then he is great". If the hearer does not believe this statement himself, he may still consider it to be a reasonable belief of the speaker; having stereotyped knowledge of the speaker's views may thus be useful.

**EX17:**

- 1) Reagan is great
- 2) He stands for apple pie and Mom

Finally, if the hearer is testing a possible evidence relation between two propositions, does not believe the missing premise, and has no prior knowledge of the speaker, the best option available is to try to judge the plausibility of the unstated information. Essentially, the hearer must postulate new facts (which he may not be sure he wants to also believe) and consider relationships between these facts as plausible or not. It is in this sense that he adopts a "hypothetical person's" beliefs. Note that it is often the case that one will accept new facts unless something from one's own beliefs suggests a contradiction. In this sense, the judgement of plausibility does relate back to the hearer's own beliefs.

An example with an implausible missing connection is EX18 below. If the hearer tests 2 as possible evidence for 1, a major premise of the form "All sharks like tap dancing" would establish the relation. But the hearer would not regard this as a plausible belief of the speaker, and so would fail to recognize an evidence relation between the two propositions.

**EX18:**

- 1) Joey likes to tap dance
- 2) He is a shark

The problem of judging plausibility is difficult. To make this process more computationally tractable, one suggestion is to incorporate into the model some tracking of mutual belief between speaker and hearer. (See Cohen (1978) for further discussion on the use of mutual belief in natural language processing.) Then, certain tests for evidence relations can be blocked in the oracle, based on mutual belief.

For instance, rules can be postulated regarding the use of claims and evidence that are mutually believed. Two sample rules are:

- (i) "If P is mutually believed, it can't be used as claim".
- (ii) "If  $\sim$ P is mutually believed, P can't be used as evidence".

In addition, some of the default interpretation rules associated with the taxonomy of linguistic clues can be overruled by pragmatic considerations. The idea is to possibly override the default semantic interpretations of evidence relations otherwise specified, by testing whether

propositions are already mutually believed. Note that the idea of a “pragmatic override” is also employed in Gazdar (1979) for the problem of determining presuppositions. The importance of pragmatic processing for argument analysis is once more emphasized, as it is a critical component to the difficult procedure of judging plausibility.

While considering mutual belief will help to eliminate some potentially difficult tests for the oracle, the model would require a more detailed specification of the maintenance and use of mutual belief. This is left as a topic for future work. Some current ideas are explored in more detail in Cohen (1985).

#### 4.4 SUMMARY OF EVIDENCE THEORY

The “theory” of evidence relationships, developed to specify the operation of the evidence oracle component of the argument analysis model, really presents insight into the problems relevant to evidence relations, rather than offering solutions. Still, the fundamental question of how connections between propositions can be verified has not been studied to any extent by other researchers. It is extremely worthwhile to acknowledge that it is not sufficient to indicate what relations do occur, without also suggesting how these relations could be established during analysis.

In addition, we have provided some insight into the more general question of how to accommodate a possible difference in beliefs between conversants in a natural language dialogue processing system. We also suggest that a less sophisticated oracle can be constructed that merely searches known facts and axioms, possibly including relaxations, to handle a large amount of naturally occurring arguments.

## 5 RELATED WORK

### 5.1 ARGUMENT UNDERSTANDING

Other researchers in natural language have studied arguments, in particular. However, the focus of the research in each case has been different. Birnbaum and the group at Yale (Birnbaum et al. 1980, McGuire et al. 1981) study two-way communication, developing an argument graph to display the points raised by both conversants. This graph is then used to determine the best moves on the part of an adversary, to challenge the position of the other conversant. Thus, the question of what responses to generate is investigated. On the other hand, there is little insight into how a hearer can detect the points being raised by the speaker, to construct this argument graph. Our focus has thus been on this preliminary problem to argument understanding.

Archbold (Archbold 1976, Archbold and Hobbs 1980) is most concerned with evaluative arguments, those with strong underlying ideologies. For example, Lenin’s speeches are appropriate sample input. Thus, the difficult question of recognizing differing opinions is a focus to Archbold’s investigations. In addition, he studies

text rather than discourse, allowing for a deeper review (re-reading) of the input in order to derive an analysis.

Weiner (1979) describes a representation for arguments that is also a tree structure, with a variety of links possible. His main concern is to characterize types of argument structures, for use in the generation of explanations. There is thus little attention on the problems encountered in deriving argument structures during analysis.

Weiner’s (1979) model for the structure of explanation bears some resemblance to the representation described for arguments here. Weiner claims that natural explanation can be regarded as a series of transformations of an underlying tree structure that represents the abstract form of the argument being developed. The ways in which support can be supplied are listed more extensively, including examples, alternatives, etc. How the tree can be built up relies on tracking a node that is “in focus”. The fact that determining the relations between statements may make use of clue words is mentioned briefly as well. Basically, some of the features we advocate appear in this research. But we are trying to provide more insight into operational questions such as:

- How do you determine the (best) relation between propositions?
- How is the focus set? and
- When are clue words likely to occur?

By contrast, Weiner concentrates on how to generate explanations using his precisely specified characterization.

Reichman (1981) is concerned with a larger problem of producing a model of discourse (not just arguments), but her approach should handle arguments as well. The core of the model is an ATN grammar for parsing and generation, coupled with a representation of “context spaces” containing conversational moves. The conversational moves provide a classification of larger components of discourse (not just single propositions). For example, there is an extensive study of a “challenge” operation. Since Reichman’s aims are broader than ours, the lower level issues we address of verifying evidence and studying the necessity for clue words are not considered for the model. Moreover, there is an intentional lack of concern with pragmatic processing, another crucial feature our model. Instead, Reichman presents a model for the analysis of a variety of two-person interactions.

In sum, our efforts in argument understanding are worthwhile because we focus on the necessary first step in argument analysis – determining the intended structure, or “what the argument is about”. We study the possible structural relations between propositions, and investigate the difficult issue of verifying evidence relationships. The importance of pragmatic analysis to recognize whole classes of arguments that involve differing beliefs is stressed in our work. And finally, the use and interpretation of clue words is addressed. It is worth

noting that the differences in our existing studies can possibly be exploited by pooling efforts and suggesting a powerful general model for argument analysis.

### 5.2 REFERENCE RESOLUTION AND FOCUS

Some of the work on using focus for reference resolution contains similarities to the model presented here for analyzing argument structure. Sidner (1979) maintains a focus stack of possible items in focus and an alternate focus list to support shifts of focus. The candidates for resolving references are thus restricted and ordered. The point is that these restrictions are drawn from a characterization of coherent discourse, the same approach taken for our control of processing. Grosz (1977) presents a model of focus spaces which may be used for several purposes, including the resolution of definite noun phrase resolution. The spaces are organized into a hierarchy, thus similar to our tree representation for argument structure. Both active and open spaces are tracked, similar to our tracking candidates eligible to be relatives to the current proposition.

Because of similarities in the representations and techniques for controlling search for interpretations, it is worth investigating as future work the precise relationship among coherence, reference resolution and focus determination for dialogues (some of this is being done (Grosz and Sidner 1986)).

### 5.3 PSYCHOLOGICAL RESEARCH

Although our model is not designed according to psychological studies of discourse comprehension, there are some interesting parallels with existing psychological research. Labov and Fanshel (1977) investigate therapeutic discourse, dialogue between a psychologist and his patient. The research describes several properties of the arguments advanced by the patients including: the use of poor logic, the tendency to omit premises in arguments, a variety of transmission forms (claims before and after), and the existence of statements about the structure of the argument (clues). Since our characterization of input provides for all these forms, it strengthens our case for having a robust model.

Geva (1981) investigates the usefulness of flowcharting text structure to assist students in comprehending the underlying structure. The top-down influence of building and using a representation is mentioned as important. The fact that many texts do not follow a strict linear ordering of connections between statements again confirms our concern with a variety of possible coherent transmissions.

In brief, discovering psychological experiments that agree with the constraints of our model serve to defend its design. Further, some suggestions we make about computational measures of discourse processing may serve to inspire new experiments into the nature of human processing. So, the relationship with psychology should be a two-way exchange, and suggests future work.

## 6 USEFULNESS OF THE THEORY

The computational model for the analysis of arguments, as described in the previous sections, is built on a theory of argument understanding; it, in turn, can be used as the basis for an implementation of an argument understanding system. One suggested real-life application area is a complaint bureau for department stores. Future work could include a full implementation of the model, and fine-tuning the design by selecting a particular application area for arguments.

Although there is no complete implementation of the model to date, an overview of a possible design is presented here, to indicate how the various components of the model could come together into one integrated "system". (Note that an initial implementation of the model does exist now, written in Prolog (Smedley 1986). But this program merely tests the various reception algorithms described in section 2. The evidence oracle is replaced by a "query the user" facility. Nonetheless, the groundwork is in place for a future implementation that tests the other components of the model).

In Figure 1, there are three main modules: the Proposition Analyzer, Clue Interpreter, and Evidence Oracle. The Proposition Analyzer takes as input the argument itself and produces a representation of its underlying structure. For each proposition of the argument the Proposition Analyzer attempts to assign it a location in the representation tree, indicating to which other proposition it relates (provides evidence for or receives evidence from). The Proposition Analyzer may call on the Clue Interpreter in the presence of clues, to assist in the interpretation. In addition, once an eligible relative to the current proposition is selected, the decision of whether an evidence relation exists is made by the Evidence Oracle, which is passed the two propositions and responds with a yes or no answer. The Evidence Oracle has available, a knowledge base of shared facts and, if possible a model of the speaker. Moreover, if certain beliefs of the speaker can be extracted during the tests for unstated premises, the model of the speaker may even be updated by the Evidence Oracle, to aid in the processing of later propositions.

In the absence of an implementation, the model can still be defended as a useful prescription of analysis of arguments. This is accomplished in Cohen (1983) by hand simulations of a variety of examples, to demonstrate robustness, and analysis of the complexity of the processing algorithms, to demonstrate efficiency.

Another argument for the usefulness of the argument analysis model is that the theories developed for the model may be applied to the solution of other language understanding problems. As a result, the study of arguments may be viewed as a worthwhile exercise in the study of language. Some examples of the wider applicability of the model are:

- It has been shown that extracting the underlying structure of discourse is useful to study the complexity of




Figure 1. System design.

analysis. In the model, the separation of where and how propositions relate has provided a means of monitoring the number of calls to an inference engine, apart from the more difficult measurement of the process of actually filling in missing inferences. Hopefully, some characterization of structures for language problems other than arguments would be extremely beneficial.

- It has been shown that certain linguistic constructions serve a function in facilitating the analysis process for the hearer. Developing common interpretation rules for various “linguistic clues” should continue for several language understanding tasks.
- Some insight has been offered into how communication can proceed despite differing beliefs of speaker and hearer. The ideas outlined for recognizing beliefs similar to one’s own, for judging plausible generalizations, should extend to other problems where reasoning beyond one’s current beliefs is required.

In addition, very few researchers seem concerned with truly “low-level” operations, determining not just “what’s a good representation for discourse” but also how this representation can be derived, the specification of some algorithm for processing. It is in this domain that we feel our research is making a contribution.

For future work, we are currently developing a model for discourse analysis in general, based on the principles of this model’s design. A hypothesis worth investigating from the existing model is that the resulting representation serves to outline both the linguistic structure of the discourse and the intentional structure (the speaker’s intentions behind utterances). The idea is that determining evidence relations in an argumentative discourse may best be described as uncovering the intended uses of utterances (e.g., speaker utters P in order to get hearer to believe Q), hence reflecting the plan of the speaker. But this main function of deriving intentional structure must be performed in conjunction with testing “logical” connections between propositions, and recognizing clues, thus isolating linguistic structure (or grouping into segments). We are interested in specifying a processing model for discourse understanding that operates at the level of individual utterances, in the manner of the argument model, to gain insight into how to derive linguistic and intentional structure simultaneously. This research is

of significance to the current work of Grosz and Sidner (1986).

#### ACKNOWLEDGMENTS

This work was supported in part by the Natural Sciences and Engineering Research Council of Canada. I am grateful to the anonymous referees for their valuable comments and to Ray Perrault for his initial supervision of this research.

#### REFERENCES

- Archbold, A. 1976 A Study of Some Argument Forms in a Persuasion Dialogue. ISI internal document, Marina del Rey, California.
- Archbold, A. and Hobbs, J. 1980 Notes on “The Analysis of Evaluative Argumentation in Text”: A Report on On-Going Work. SRI unpublished draft, Menlo Park, California.
- Birnbaum, L.; Flowers, M.; and McGuire, R. 1980 Towards an AI Model of Argumentation. *Proceedings of American Association for Artificial Intelligence (AAAI) Conference*: 313-315.
- Cohen, P. 1978 On Knowing What to Say: Planning Speech Acts. Technical Report No. 118, Computer Science Department, University of Toronto, Toronto, Ontario, Canada.
- Cohen, R. 1980 Understanding Arguments. *Proceedings of Canadian Society for Computational Studies of Intelligence (CSCSI) Conference*: 272-279.
- Cohen, R. 1981 Investigation of Processing Strategies for the Structural Analysis of Arguments. *Proceedings of the Association for Computational Linguistics (ACL) Conference*: 71-75.
- Cohen, R. 1983 A Computational Model for the Analysis of Arguments. Ph.D. thesis, Technical Report No. 151, Computer Systems Research Group, University of Toronto, Toronto, Ontario, Canada.
- Cohen, R. 1984a A Theory of Discourse Coherence for Argument Understanding. *Proceedings of Canadian Society for Computational Studies of Intelligence (CSCSI) Conference*: 6-10.
- Cohen, R. 1984b A Computational Theory of the Function of Clue Words in Argument Understanding. *Proceedings of 10th International Conference on Computational Linguistics (COLING 84)*: 251-258.
- Cohen, R. 1985 The Need for Pragmatics in Natural Language Understanding. *Proceedings of CSCSI-sponsored Theoretical Advances in Natural Language Understanding Conference*.
- Gazdar, G. 1979 *Pragmatics: Implicature, Presupposition and Logical Form*. Academic Press, New York, New York.
- Geva, E. 1981 Flowcharting Expository Texts and Reading Comprehension. Paper presented at Annual Meeting of American Educational Research Association.
- Grosz, B. 1977 The Representation and Use of Focus in Dialogue Understanding. SRI Technical Note No. 151, Menlo Park, California.
- Grosz, B. and Sidner, C. 1986 The Structures of Discourse Structure. Report No. 6097. Bolt, Beranek and Newman (BBN) Cambridge.

- Massachusetts. (also Report No. CSLI-85-39; Attention, Intentions, and the Structure of Discourse, *Computational Linguistics* 12(3): 175-204.
- Hobbs, J. 1976 A Computational Approach to Discourse Analysis. Research Report No. 76-2. City University of New York Department of Computer, New York, New York.
- Labov, W. and Fanshel, D. 1977 *Therapeutic Discourse*. Academic Press, New York, New York.
- McCarty, L. and Sridharan, N.S. 1981 A Computational Theory of Legal Argumentation. Research Report LRP-TR-13, Rutgers University Lab for Computer Science, New Brunswick, New Jersey.
- McGuire, R., Birnbaum, L. and Flowers, M. 1981 Opportunistic Processing in Arguments. *Proceedings of the International Joint Conference on Artificial Intelligence. bf one;(IJCAI) Conference*: 58-60.
- Pollack, M. 1986 A Model of Plan Inference that Distinguishes Between the Beliefs of Actors and Observers. *Proceedings of the Association for Computational Linguistics (ACL) Conference*: 207-214.
- Quirk, R. et. al. 1972 *A Grammar of Contemporary English*. Longmans Co., London, England.
- Reichman, R. 1981 Plain Speaking: A Theory and Grammar of Spontaneous Discourse. Report No. 4681, Bolt, Beranek and Newman (BBN), Cambridge, Massachusetts..
- Sadock, J. 1977 Modus Brevis: The Truncated Argument. *Papers from the 13th Regional Meeting, Chicago Linguistics Society*: 545-554.
- Sidner, C. 1979 Towards a Computational Theory of Definite Anaphora Comprehension in English Discourse. AI Lab Report TR-537, Massachusetts Institute of Technology (MIT), Cambridge, Massachusetts.
- Smedley, T. 1986 An Implementation of a Computational Model for the Analysis of Arguments: An Introduction to the First Attempt. Research Report CS-86-26, University of Waterloo Department of Computer Science, Waterloo, Ontario, Canada.
- Weiner, J. 1979 BLAH, A System which Explains Its Reasoning. Technical Report TR 79-14, University of New Hampshire Computer Science Department, Durham, New Hampshire.