

SRI International



WHY IS DISCOURSE COHERENT?

Technical Note 176

November 30, 1978

By: Jerry R. Hobbs, Computer Scientist
Artificial Intelligence Center

SRI Projects 5844, 7500, and 7910

The work reported herein was supported by the National Science Foundation under Grants No. MCS-76-22004 and MCS-78-07121 and by the Advance Research Projects Agency.

ABSTRACT

When people produce a discourse, what needs are they responding to when they make it coherent, and what form does this coherence take? In this paper, it is argued that coherence can be characterized in terms of a set of "coherence relations" between segments of a discourse. It is shown, from an abstract description of the discourse situation, that these relations correspond to the kinds of communicative work that needs to get done in discourse. In particular, four requirements for successful communication are isolated: the message itself must be conveyed; the message must be related to the goals of the discourse; what is new and unpredictable in the message must be related to what the listener already knows; and the speaker must guide the listener's inference processes toward the full intended meaning of the message. Corresponding to each requirement is a class of coherence relations that help the speaker satisfy the requirements. The coherence relations in each class are discussed and defined formally. Finally, a fragment of a conversation is analyzed in detail to illustrate the problems that face a speaker in trying to satisfy these requirements, and to demonstrate the role that the coherence relations play in the solution.

1. Introduction

Discourse tends to be coherent. Why should this be so? When people produce a discourse, what needs are they responding to when they make it coherent? How does a Speaker determine what he will say next when he wishes to continue a conversation coherently? When a Listener hears an utterance, how does he understand its relevance to the preceding discourse?

In this paper, I present a theory of coherence in which coherence is characterized in terms of a set of binary relations between a current utterance and the preceding discourse. The relations correspond to the kinds of work that needs to get done in the construction of a discourse. That is, for the Speaker, the question of what is needed next in the discourse is equivalent to the question of which relation should be used to continue it. For the Listener, the question of what the Speaker is trying to accomplish by continuing as he did is equivalent to the question of what coherence relation he used. We will look first at an abstract description of the discourse situation, in order to see what kinds of work needs to be accomplished in discourse and what classes of coherence relations this implies. Then we will examine a fragment of a seemingly unstructured conversation in minute detail, trying to discover what structure there is and how this structure aided the participants toward their discourse goals.

An example may serve to orient the reader to the framework I am assuming: When we see the text

(1) John can open Bill's safe. He knows the combination,

it strikes us intuitively as coherent. On closer inspection, we see that the second sentence is an elaboration of the first. Still closer inspection tells us that the two sentences describe the same condition from different perspectives, the first in terms of the state that can be effected, the second in terms of the means for effecting it.

In the theory of coherence presented here and in the computer system that implements it, I am trying to make explicit a mode of processing that could underlie these intuitions. The computer system must have access to a large knowledge base of facts about things. For instance, to deal with example (1) it must know that one can open a safe by dialing its combination, that having the ability to effect some state means knowing an action that will bring it about, and that dialing is an action. The system must be able to construct chains of inference. It must have definitions of the coherence relations in terms of the chains of inference that can be constructed from the propositional content of the sentences in the text. Finally it requires a procedure which seeks chains of inference to satisfy these definitions. For example, the coherence relation called Elaboration can be defined roughly as follows:

A segment of discourse S1 is an Elaboration of segment S0 if the same proposition P can be inferred from both S0 and S1, and one of the arguments of P is more fully specified in S1 than in S0.

In (1), S0 is the first sentence, S1 the second. The proposition P that can be inferred from both sentences is that John knows some action it is possible for him to perform that will cause the safe to be open. This action -- dialing the combination, rather than, say, blasting the safe with dynamite -- is more fully specified (implicitly, of course) in the second sentence than in the first.*

The claims of the theory of coherence can be summarized briefly as follows:

1) Coherence in discourse can be characterized by means of a small number of coherence relations which are definable in terms of the operations of an inference system. If an utterance strikes one intuitively as a coherent continuation of the discourse, then there is some coherence relation that holds between the utterance and some portion of the preceding discourse. If it strikes one as incoherent, no such relation exists.

*-----
* A fuller account of this example is given in Hobbs (1978).

2) Each of the coherence relations serves some communicative function.

This paper is an examination of the second claim. In Part 2, I try to clarify what is meant by coherence and what would count as an answer to the question the title asks. Part 3 is the heart of the paper. It is shaped like a pyramid: Section 3.1 is the top of the pyramid and contains my answer to the question, "What needs are the participants responding to when they make a discourse coherent?" The argument is structured as follows: There are four reasons for a Speaker to break a discourse into more than one clause. Corresponding to each of these reasons is a class of coherence relations connecting the discourse segments thus produced. The remainder of Part 3 forms the base of the pyramid. In it, I explicate each of the four classes of coherence relations in turn, thereby responding to the question, "What form does coherence take, given the answer of Section 3.1?"

Part 5 approaches the question of the title from another angle. It follows a Speaker through the production of eight utterances in a free-flowing conversation. At each point, it is shown how the needs described in Section 3.1 impinge upon the production process and how each need is satisfied by a coherence relation in the corresponding class. Part 4 describes the data that is analyzed in Part 5.

Other issues in the theory of coherence, including issues relating to the first claim, have been addressed in other papers. Hobbs (1976) gives a large list of the coherence relations, together with their formal definitions. Hobbs (1978) discusses some of the relations in more detail, with slightly modified definitions, and discusses their communicative functions. Hobbs (1976) and Hobbs (1977) present microanalyses of a Newsweek paragraph and an algorithm description, respectively.

The approach to coherence taken here is similar to that in work by Longacre (1976, 1977) and Grimes (1975) on relations between sentences, and in a more computational vein, by Phillips (1978a, 1978b), and work by Halliday and Hasan on conjunctive relations (1976).

2. Some Preliminary Remarks

2.1. Coherence, Cohesion, Relevance, and Understandability. I am attempting to turn "coherence" into a technical term, in much the same way as "grammaticality" is a technical term in generative grammar and "semantics" is a technical term in logic. It is necessary to make clear the intuitive sense of the word which is being explicated. There are two reasons for this. First, the reader will not be able to judge how successful the characterization of coherence is unless he has a clear notion of what sense of "coherence" is being characterized. Secondly, there is something a bit amiss in efforts to pre-empt quite general English terms and confine their meanings to special functions within a particular theory. Such efforts set the stage for future misunderstandings. It is important that the sense being characterized is strong enough to warrant the pre-emption.

I will call a discourse "coherent" if it exhibits structural relationships between its various segments, which depend on the propositional content of the segments. Contrasting "coherence" with three other terms -- "cohesion", "relevance", and "understandability" -- will clarify this.

In standard usage "cohesive" and "coherent" can be used interchangeably. Both mean "hanging together". However, "cohesion" has already been adopted as a technical term in discourse analysis. Halliday and Hasan (1976) use it to refer to the ways in which different sentences in a discourse can refer to the same entities. For example, the text

- (2) Jimmy Carter proposed a massive energy program. He used to cultivate peanuts.

is cohesive since both sentences refer to the same individual. But I would not like to call it coherent, except perhaps in the most special of contexts where we were able, for example, to construct some causal

connection between the two sentences or detect some significant contrast. The term "coherent" ought to be reserved for texts exhibiting more structure than example (2).

"Relevance", especially as Grice (1975) uses the term, is perhaps closest in meaning to "coherence" as I intend it. "Coherence" has been chosen, however, because "relevance" seems to be a property of a single utterance with respect to the preceding text, whereas it is convenient to speak not only of an utterance being "relevant" or a "coherent continuation", but also of an entire discourse being "coherent".

Finally, as with all words, "coherent" has been watered down, and has come to mean in some cases, simply "understandable". The sense of coherence that I am trying to explicate is a tighter sense than this. For instance, example (2) is certainly understandable, as is any text consisting of two semantically plausible sentences. But such texts should not be called coherent, for the individual sentences are understandable only in themselves. They contribute nothing to each other.

2.2. A Possible Misunderstanding. To some, the term "coherence" evokes images of high school composition classes. They may feel that coherence is an idealization that applies primarily to written texts and is realized mainly through the use of such superficial devices as conjunctions and syntactic similarity, but has nothing to do with the spoken language. As will be seen in the latter part of this paper, nothing could be further from the truth. The principles of coherence presented here operate in spoken discourse, and frequently involve quite deep semantic connections between superficially dissimilar sentences. In Part 5, for example, the sentences

You have to either draw or make things with the little
Q-tips. So she thinks she's going to win an eight
thousand dollar first prize.

are shown to be, among other things, an instance of the Parallel relation. Insofar as the coherence relations are identical to

prescriptive rhetorical principles, it is because the latter represent attempts to formalize what goes on in ordinary coherent conversation.

2.3. Why Binary Relations? The assumption that the relations are all binary requires some justification, for it seems clear when we adopt the Olympian perspective of a linguist examining a text or transcript, that there are n-ary relations. A general proposition can be exemplified by a number of instances. I am viewing the problem from a procedural point of view however, from the point of view of the Speaker producing a sequence of utterances and from the point of view of a Listener trying to comprehend an utterance (cf. Levy 1978). At a given moment in the discourse, the Speaker is producing what we will call the "current utterance"; the Listener is trying to understand it. In both of their minds there is some memory of the preceding discourse; it is structured somehow, and for the time being I will assume this structure can be represented as a tree whose terminal nodes are semantic representations of the clauses* of the preceding discourse. The problem for the Listener in understanding the relevance of the current utterance to the ongoing discourse is to understand with which portion of the preceding discourse the current utterance is linked and what the nature of that linkage is. For the Speaker the problem of continuing in a relevant fashion is one of choosing the right portion of the preceding discourse to expand on and of choosing an appropriate way to expand on it. In either case, we are dealing with a binary relation between the current utterance and some portion of previous discourse. This is true even of an n-ary relation, for then we may view it as a binary relation between the current utterance and the sequence of previous arguments of the n-ary relation.

This picture will have to be modified somewhat in two ways. First of all, there are cases in which we reserve judgment on the relevance of an utterance for a short time, with the expectation that the utterance will be subordinate somehow to an utterance yet to come which itself will be relevant to preceding discourse. Consider a simple example: If

* Perhaps elliptical.

we are asked a question, the most immediate relevant response is a direct answer. However, we may respond by starting to tell a story, giving the setting, the main actors and so forth. If we do this, the questioner will generally assume that all of this is building up to a story whose point will somehow answer the question (Labov & Fanshel 1977).

The second modification is simply that all discourse is not coherent. In written texts and in certain formal speaking situations it is common for a single structure to cover a large portion of the discourse. But in casual conversational situations, we typically find islands of coherence of varying sizes with more or less successful attempts to bridge between them.

2.4. Other Views of Coherence. The view of coherence presented here may be compared and contrasted with three other views of coherence. The first is that a text is coherent if it continues to talk about the same entities, that is, if it is cohesive. This sometimes takes the form that the coherent stretches of discourse are the stretches characterized by the same "topic". Where "topic" is identified as "the entity which the discourse is about", this is hardly better than equating coherence with cohesion. For example, we could say Jimmy Carter is the topic of (2). It is certainly true that if a text is coherent, it is cohesive, but elsewhere (Hobbs 1978) I have argued that the causal relationship goes the other way -- discourse is cohesive because of the nature of coherence.

A more sophisticated notion of "topic", proposed by Keenan and Schieffelin (1976), is that the topic is the proposition about which some claim is being made or elicited. I have no quarrel with this position, and in fact what I present may be interpreted as an exploration of the fine structure of topic, so defined.

A second and rather substantial view of coherence is that discourse is coherent if the utterances can be seen as actions in some overall plan designed to achieve some goal. The problem of recognizing coherence is then a matter of deducing the Speaker's plan and finding

the place where the current utterance fits into it. This is stated in a strong form by Allen and Perrault (1978) and Cohen (1978). In A. Robinson (1978), comprehension in task-oriented dialogues is guided by the plan for performing a complex task. This approach is applied to indirect answers to questions by Hobbs and J. Robinson (1978). Labov and Fanshel (1977) are presumably subscribing to this view when they say that coherence is at the level of actions rather than utterances.

Where both participants in a conversation know the plan that is being worked through, and both know that the other has the plan as well, and all the utterances can be related to this plan, we would certainly want to call the conversation coherent. Similarly, if the Listener is able to deduce the Speaker's plan and responds appropriately, the response is coherent. Whatever the precise nature of the inference operations, I would expect them to include this plan deduction capability.

But the plan deduction theory is too strong. Participants in a conversation frequently have access only to the most superficial goals of the other participants, and yet are able to proceed coherently. To take a simple example, if a stranger walks up to me on the street and asks me the time, it seems more reasonable for me to tell him the time than to ask, "Why do you want to know?" Yet the latter could be viewed as a reasonable step in plan deduction, since we are trying to get at his ultimate goals. A weaker theory of coherence than the plan deduction theory is required to explain all those conversations which seem coherent, even though we can see on close inspection that the participants are talking past each other.

A third view of coherence has been suggested by Chafe (1978): The structure of discourse reflects somehow the structure of memory, either the pre-existing structure of the contents of memory, or the structure of the way we are reminded of things by other things.* I find this idea attractive, and would not object if the reader saw this paper as a detailed exploration, along somewhat different lines than Chafe has pursued, of how discourse reflects the structure of memory.

* In a footnote in Hobbs (1978) I made a similar suggestion.

I have one reservation about this idea however. There is a certain amount of utter randomness in the way we are reminded of things. Occasionally we encounter someone who clearly talks about things as he is reminded of them. Insofar as we are not able to follow the links he is following, perhaps because they are personal and not accessible to us as listeners, we will judge him scatterbrained. We may be more or less forgiving depending on the formality of the occasion and the interest of what he says. But we would not call his conversation coherent. So the very most we would want to say is that the coherence relations are conventionalized ways of being reminded of things. They are those ways of travelling through our mental maps that we can reasonably expect a listener to follow.

2.5. Procedural vs. Empirical Adequacy. Finally, I should make some remarks about the aims of this research. Those who study language find themselves on the horns of dilemma. We might label the horns "empirical adequacy" and "procedural adequacy". On the one hand, we would like our models to be faithful representations of the way people actually process language, the knowledge they actually use, and the goals they are actually driven by; on the other hand, we want procedures that solve the discourse problems we are faced with. In most cases, we can't have both. If we adhere to empirical adequacy and do not go beyond what our data entitles us to, we will not solve the problems, for we simply do not have access to all that we need; we will be condemned to sterile theories. If we try to construct procedures that work, we are on shaky empirical grounds; we will be accused of being "ad hoc". The problem of pronoun resolution is an example. We can write fairly good procedures for resolving pronoun references, but they depend on a great deal of highly detailed world knowledge, which we have no empirical justification for assuming the listener possesses or uses. If we use only data we can be sure of, we cannot construct the procedures.

The usual way out of this dilemma for workers in artificial intelligence is to aim for procedural adequacy and to settle for empirical plausibility. That is what I will do. In Part 5, I present

an analysis of how a particular conversation could have been produced. The ultimate test of this analysis is not whether I can provide conclusive evidence that it in fact was was produced in the way I describe. Rather, it is whether the analysis is plausible and involves procedures that are actually computable and solve the problems they are designed to solve.

3. The Discourse Situation and Coherence Relations

3.1. The Discourse Situation. At the typical moment in a discourse, there is a Speaker and at least one Listener. The Speaker is speaking because he perceives some difference between what he and his Listener know, believe, imagine, or desire (the message), and because he believes that reducing this difference will serve some purpose (the goal). The Listener applies an active inference process to understand the message and to link what is new in it to what he already knows. It is thus part of the Speaker's job to provide the necessary linkage and to try to manipulate the Listener's inference process to lead him to the correct interpretation. This description of the discourse situation will enable us to categorize the coherence relations according to their communicative function.

The most fundamental question in an investigation of coherence is "Why are any discourses longer than one clause?" The obvious answer is that the participants have more to say than conveniently fits into a single clause. But what needs to get done in discourse that would lead us to say that a series of utterances constitutes a single discourse rather than a sequence of contiguous discourses? The above description of the discourse situation points up four kinds of work that the participants need to accomplish in discourse, or four reasons a discourse may exceed a clause in length. Corresponding to each of the four reasons for adding an utterance is a class of relations that link utterances.

1) The message itself may include too much for a single clause, and yet may describe a coherent set of states or events in the "world",* so that we would not want to call it several discourses.** This leads to

* Here I intend "world" to be interpreted broadly enough to include anyone's counterexamples to the notion that discourse is about things in the world. For example, I would include cultural and mathematical facts, as well as facts about the ongoing discourse if the discourse takes itself as the subject matter.

** Phillips (1977b) makes a similar point when he says that "a description of events in the world that are perceived as having some unity should constitute a coherent discourse."

the question of what gives coherence to a set of events in the world. To get a handle on this problem, let's consider a few examples. If someone walks by in the hallway, and then I get up to sharpen my pencil, we would not in general call this a coherent set of events. On the other hand, if my pencil breaks and I get up to sharpen it, it is a coherent set of events, for they are causally related. But causality is too strong. If someone walks by in the hallway and I buttonhole him to ask him a question, these are coherent events, but we can't say my asking was caused by his walking by. Surely I had other motives. We can't even say his walking by enabled the asking, for I could have sought him out. But we can say the first event "set up" the situation for the second, in the sense that a particular intrinsic feature of the second event -- in this case, the time and location -- were as they were because of the change effected in the first event. The second event was occasioned by the first. The second event could have happened anyway, but it happened the way it did in fact happen because of the first event.* A related type of coherence between events is when the second event effects a change out of a state which is an intrinsic part of the first event. An example is when after standing in the hallway answering my question, the person I buttonholed walks on. The second event -- walking on -- effects a change of location, where the initial location was an intrinsic part of the first event -- standing and answering.

These weak causal, or strong temporal, relations between events will form the basis of the first class of coherence relations between segments of discourse, what I will call the Strong Temporal Relations.

2) The discourse derives from a goal or set of goals the Participants have, and it may help realize these goals to a greater or lesser degree. The discourse is therefore continually subject to evaluation by the participants, to judgments as to its effectiveness in carrying out its goals. Thus, in the second class of coherence relations is Evaluation.

* Balzer, Goldman and Wile (1977) refer to this as a "producer-consumer" relation. The first event produces a state that is "consumed" by the second.

3) The Speaker needs to provide the Listener with a linkage between the message and what the Listener can be expected to know already. Since there are severe limits on the amount of descriptive and explanatory material that can be included in one clause, this fact gives rise to what will be called Linkage Relations.

4) Comprehension is not simply a matter of the Speaker depositing a proposition in the Listener's head. It involves an active inference process, in which, among other things, the Listener must infer the specific from the general or the general from the specific, in order to zero in on the Speaker's full intended meaning. By choosing and ordering his utterances in a particular fashion, the Speaker can exercise some control over this inference process by supplying or modifying the appropriate framework for their interpretation. The class of relations this fact gives rise to will be called Expansion Relations, because where they obtain, the current segment of discourse can be seen as expanding on a previous segment. They all indicate ways a Speaker can move or fail to move between specific and general statements.

Sections 3.2 - 3.5 fill out the details in the picture sketched in this section. They explicate a taxonomy of coherence relations that is illustrated in Figure 1.

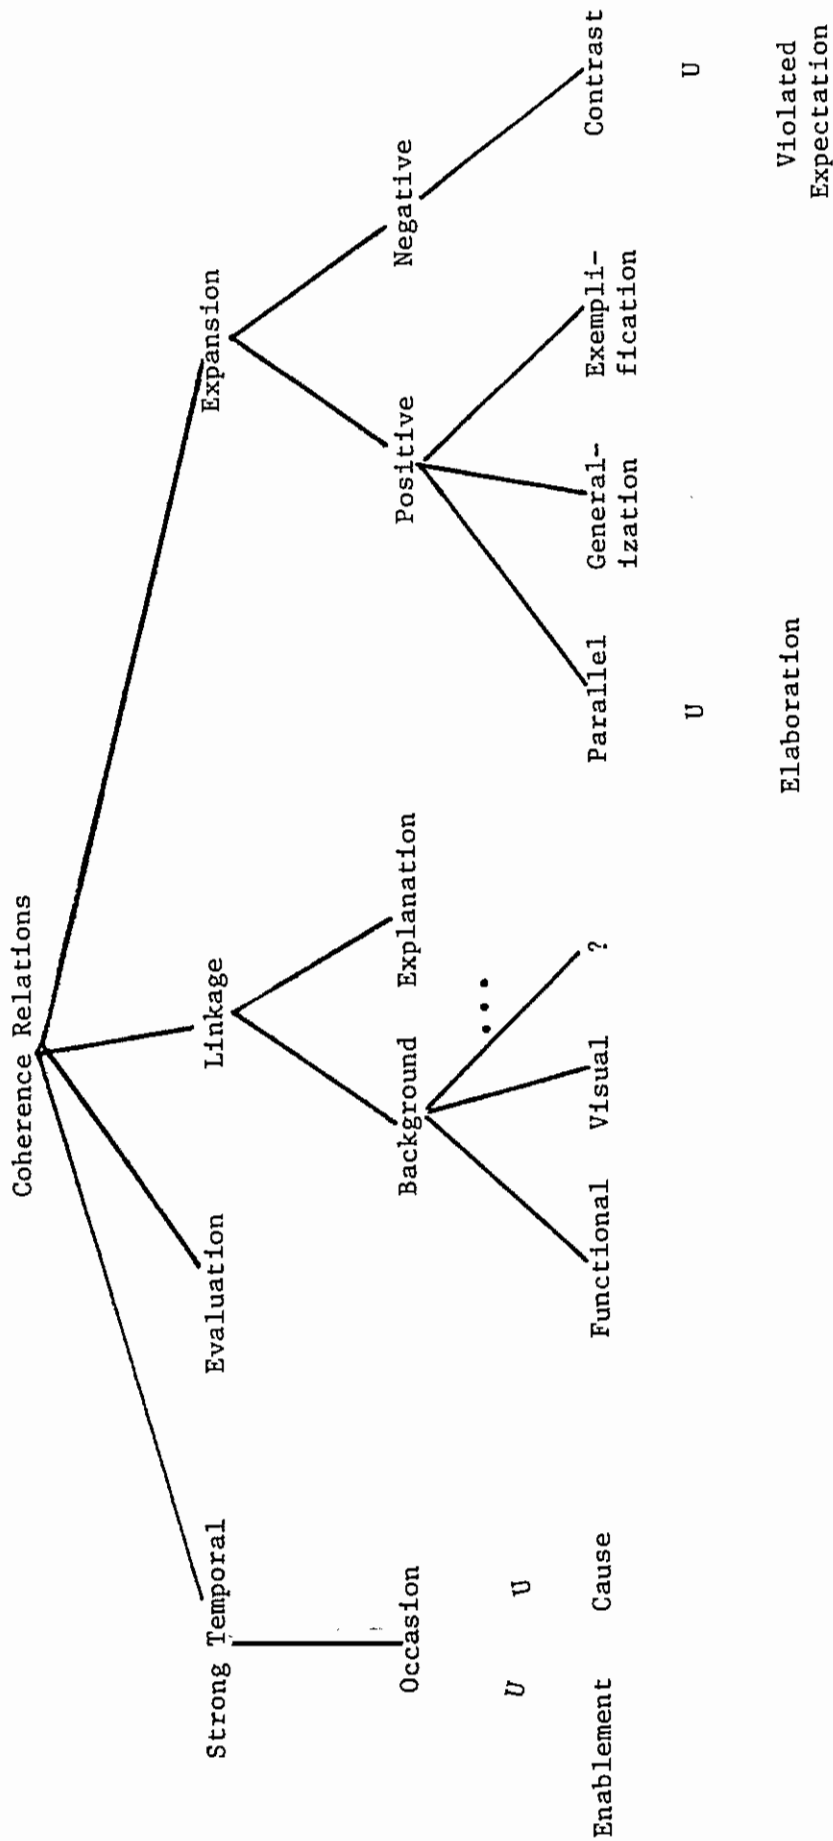


Figure 1

3.2. Strong Temporal Relations. A message relating a coherent sequence of events has to be chopped into separate clauses. The weakest relations that can obtain between them are what will be called the Occasion relations.* Intuitively, the event asserted in the first clause occasions the event asserted in the second. More precisely, we can define the relation as follows:

An Occasion relation holds between discourse segments S0 and S1 if

- a. S0 asserts a change whose final state is implicit in S1, or
- b. Implicit in S0 is a state which is the initial state of a change asserted by S1.

Some examples are

- (3) Decrease N by 1. If it is zero, reset it to MAX.
- (4) Walk out the door of the building. Turn left.
- (5) He noticed the broken connection in the control mechanism, and took it to his workshop to fix.

These texts do not quite involve enablement. In (3) you could have checked for zero before decreasing N; in (4) you could have turned left before leaving the building; in (5) he could have taken the control mechanism to his workshop before noticing the broken connection. But all of these changes would have altered the outcome of events in serious ways.

Stronger versions of this relation are Enablement and Cause. The knowledge available to a language user will include axioms encoding our knowledge of preconditions, enablements, causes and effects. We will say that two segments S0 and S1 are linked by Enablement if we can infer that the state or event asserted in S0 enables the state or event asserted in S1. An example is

* In previous papers, I have referred to this relation as "Overlapping Temporal Succession". I have abandoned this term because of possible confusions with Longacre's (1977) weaker relation "Temporal Overlap". I have also referred to the relation previously as "Then", but this leads to confusion with the sort of weak temporal sequence illustrated by examples (6) and (7) below.

John got off the subway, and walked up to the street.

John could not have walked up to the street without first getting off the subway.

Segment S0 is a Cause of segment S1 if a causal chain can be found from the state or event asserted in S0 to that asserted in S1. Two examples are

John cheated on the final. I flunked him.

My child is entering an art contest. So I have to send this in for her.

There are certain texts which are coherent, and describe events that are temporally ordered, but which I would not like to call coherent by virtue of the relations in this class, for they contain not a hint of causality or enablement. Some examples:

- (6) John brushed his teeth and combed his hair.
- (7) Set stack A empty and set link variable P to T.

Here the prior action in no way effects the outcome of the subsequent action. The order of the events could have been reversed without any change in ultimate outcome.

When the context of the discourse demands that we decide upon the temporal ordering of the events in cases like these, we follow the usual convention that clause order reflects event order. But otherwise we view them as examples of the Parallel relation described below. That is, they are coherent to the extent that they exhibit the Parallel relation, not because the second follows the first in time. Under this treatment, (6) is coherent since both clauses describe John's morning activities, and (7) is coherent since both clauses initialize a data structure. Example (8) below is not coherent, because it does not exhibit the Parallel relation, even though the clauses describe events that could plausibly follow each other in time:

- (8) My grandmother brushed her teeth. Jimmy Carter met with Congressional leaders.

3.3. Evaluations. An Evaluation is a segment of discourse from which we can infer a relation between some other segment of discourse and a goal of the discourse.* For example, when a narrative is introduced by

The most remarkable thing happened,

it asserts a relation between the narrative to come and the goal of saying something interesting. When someone apologizes for a story that fell flat with

It was funny at the time,

he simultaneously admits that the story failed to satisfy the goal and gives his reason for having believed it would satisfy it.

Evaluations frequently come from the Listener, and are thus interactional moves as well as ways a Speaker himself can continue coherently. They can let the Speaker know whether or not what he is saying is interesting, too unusual to be appropriate, and so forth. They can take the form of an utterance, a laugh at something intended to be funny, a laugh at something not intended to be funny, a gesture, or as will be seen in Parts 4 and 5, a grin.

3.4. The Linkage Relations. The third class of relations are those that arise out of the need to link what the Speaker says that is new and remarkable with what is known or believed to be known to the Listener. That is, sometimes a discourse contains more than one clause because the required descriptive material overflows the restrictive confines of a single clause. There are two aspects of this -- Background and Explanation.

* This of course is not a precise definition, for "a goal of the discourse" is not defined. For this, we need a formalization of the picture I drew in Section 3.1 of the discourse situation, that includes a model of the discourse participants as agents who can plan utterances to achieve goals. Work along this line is in progress.

Some examples of the Background relation are as follows:

- (9) There once was a cobbler in a small Bavarian village.
He worked day and night to keep his wife and seven
children fed.
- (10) Turn left at the Prudential Building. It's the only tall
building around.

In (9), the first sentence provides background material for the second. In (10) the second sentence provides further descriptive material on the Prudential Building, and thus can be considered as Background for the first sentence. An utterance which provides background material for some other segment does so by describing one or more of the entities involved in that segment. In narratives, a common use of the Background relation is to present the setting and the characters. In algorithm descriptions, a common use is to describe the relevant data structures. In general, background material may precede or follow the segment that it is background for, but it should not follow by much.

Defining the Background relation is somewhat problematic, however, for we cannot allow free modification. Background material cannot be just any further information, or else coherence is reduced to a matter of cohesion. For instance, in example (2) we could view the second sentence as background for the first or the first as background for the second. Restrictions are needed.

There are two reasons I know of for allowing information to be counted as coherent background material. The first is that the information will be functional in the subsequent discourse. For example, consider

A man was standing on a ladder, picking pears. He was in
a foul humor.*

We expect the man's foul humor to be significant somehow in the rest of the story. Of course, one can't in general distinguish this case before the end of the discourse.

* This example and the next are modifications of examples discussed in Chafe (1978).

The second reason is that the information contributes to the visual image that is being created. In

A man was standing on a ladder, picking pears. He wore a red bandana.

we would not be surprised if the red bandana played no further role in the story. It is a functionally irrelevant detail which nonetheless adds a vivid quality and verisimilitude to the account.

The second coherence relation in the Linkage class is Explanation. This occurs when an event, state or object present in the message is somehow unusual or inexplicable. To connect it with what the Listener already knows, the Speaker provides a causal chain from some normal situation to what is unusual. Consider the following example:

(11) He was in a foul humor. He hadn't slept well that night.
His electric blanket hadn't worked.

For most people, being in a foul humor is not the usual. The situation could use some explanation. Not sleeping well is a plausible cause, but itself could use an explanation. The electric blanket not working is a reasonable explanation, and the situation of a manufactured object not working is common enough that our curiosity about causes is satisfied. If we wished further explanation, we could make reasonable guesses ourselves.

Example (11) raises the interesting issue of how far back our explanations should go. There is no clearly defined notion of what is "sufficiently normal". For example, each of the sentences of (11) could itself have served as an ultimate normal cause:

He was in a foul humor. He hadn't slept well that night.

He yelled at his secretary. He was in a foul humor.

Furthermore, we could have carried out the causal chain farther.

He was in a foul humor. He hadn't slept well that night.
His electric blanket hadn't worked. The control mechanism was broken.

It is not clear in each case how far the chain needs to be carried. It is certain only that it needs to be carried at least to something which is normal or known to be known.

The reader may think it strange that I have included causality in two different categories, both as a variety of Strong Temporal relation, and as Explanation in the class of Linkage relations. But it seems to me that it plays two very different roles. As a Strong Temporal relation, it carries a story or the description of a process forward in time. It would be bizarre to tell a story backwards, however, so that backward causal relations require a special accounting, such as that given here.

3.5. The Expansion Relations. The relations within this class all indicate ways the Speaker can move between specific and general statements. They can be placed in an array as follows:

	Specific <u>-> Specific</u>	Specific <u>-> General</u>	General <u>-> Specific</u>
Positive:	Parallel	Generalization	Exemplification
Negative:	Contrast	--	--

There are also a limiting case of Parallel -- Elaboration -- and a limiting case of Contrast -- Violated Expectation. We will consider each of these relations in turn.

1) A move from a specific instance of some class to another specific instance of the same class. This is the Parallel relation. It can be defined roughly as follows:

Segments S0 and S1 are in Parallel if we can infer some predication p(a) from S0 and p(b) from S1, where a and b are both members of some independently definable superclass.

An example is

Set stack A to empty and set link variable P to T.

From the first clause we can infer "Initialize stack A", and from the second "Initialize link variable P", where A and P are both members of the independently definable set of data structures.

One frequent function of the Parallel relation is to suggest or support the generalization which the two segments are specific instances of. The relation often links segments which together function as an Exemplification or Elaboration* of a more general statement.

2) A move from a specific instance of some class to the general statement for the class as a whole. This is Generalization. It is defined roughly as follows:

S1 is a Generalization of S0 if $p(a)$ can be inferred from S0 and $p(A)$ can be inferred from S1, where a is a member or subset of A.

An example is

Dewey was certain he would win. All the Republicans were overconfident.

3) A move from a general statement to a specific instance. This is Exemplification. It is defined roughly as follows:

S1 is an Exemplification of S0 if $p(A)$ can be inferred from S0 and $p(a)$ from S1, where a is a member or subset of A.

An example is

Republicans were confident of their chances. Dewey was certain of victory.

The predication "believes will win election" is made first of the set of Republicans, then of the member of the set of Republicans, Dewey.

Exemplification frequently functions to substantiate or clarify a general statement.

* Discussed below.

We have gone from specific to specific, specific to general, and general to specific. We next need to consider the same moves negated.

4) A move from a specific instance of some class to the negation of the property for some other specific instance of the same class. This is Contrast. One variety can be defined roughly as follows:

Segments S_0 and S_1 are in Contrast if $p(a)$ can be inferred from S_0 and $\text{not-}p(b)$ can be inferred from S_1 , where a and b are members of some independently definable superset.

An example is

You are not likely to hit the bull's eye, but you are more likely to hit the bull's eye than any other equal area.

From the first clause S_0 , we can infer that whatever probability q counts as likely is greater than the probability p of hitting the bull's eye. From the second clause S_1 , we can infer that the typical probability r of hitting those other equal areas is not greater than the probability p of hitting the bull's eye. The predicates "greater than p " and "not greater than p " are true of q and r , respectively, and q and r are members of the independently definable set of probabilities.

One of the functions of Contrast is to inhibit generalizations the Listener may be tempted to draw. It is a means of fine tuning one's message.

The other two cases we might expect are logical impossibilities. They are a move from a specific instance to the negation of the general statement, and a move from a general statement to the negation of a specific instance. Where we do encounter texts which seem to fit this pattern, we need to reinterpret the text into the simple Contrast. For example, in

Dogs are stupid, but Lassie is smart,

we must reinterpret the first clause to be something like

Dogs other than Lassie are stupid,

in which case we would have the Contrast pattern.*

Besides giving the Speaker control over the generalizations the Listener draws, it is possible that this whole structure of relations has something to do with the structure of memory. That is, a Speaker applies these moves because they are the ways in which he can move most easily through his mental map, and because, as a consequence, they will minimize the processing required of the Listener.

There are two limiting cases of coherence relations that are worthy of separate mention. In the Parallel relation, if the entities a and b are identical, we have the coherence relation called Elaboration. Both segments imply pretty much the same thing. However, since the match is on an inferred proposition $p(a)$, the chains of inference which lead to this proposition could contain quite different information. The definition and an example, given in Part 1, are repeated here for the reader's convenience.

* Longacre (1977), however, lists this as the separate pattern of Exception.

A segment of discourse S1 is an Elaboration of segment S0 if the same proposition P can be inferred from both S0 and S1, and one of the arguments of P is more fully specified in S1 than in S0.

(1) John can open Bill's safe. He knows the combination,

One common function of Elaboration is clarification. Another function is to convey the same message from two different perspectives. Text (1) illustrates the most typical case: The first clause presents the message from a global perspective, the second contains details of a procedural nature. Although this relation is called Elaboration, I mean also to include under this heading such trivial moves as pure repetitions, repairs, tag questions, and the like.

One of the prototypical moves in conversation is the Answer to a Question. This can be viewed as a kind of Elaboration, in that an Answer elaborates on the propositional content of the Question in a way specified by the Question. The inferences required to recognize an Answer as appropriate are the same as those required to recognize an Elaboration.

The other limiting case occurs in the Contrast relation when the entities a and b are identical. That is,

Segment S0 implies p(a) and segment S1 implies not-p(a).

This of course is a contradiction, and we must back away from one of the conclusions. This is the Violated Expectation relation. We accept the second inference and reject the first. For example, in

This paper is weak but interesting,

from the first clause we would normally infer that it is unpublishable, from the second we can normally infer that it is not unpublishable. It is the second conclusion the Speaker intends to stand behind.

Another variant on Contrast occurs as a conversational move. Participant A makes a statement that implies $p(a)$. Participant B responds with a statement that implies $\text{not-}p(a)$. This is a Disagreement, or, as I call it in Part 5, a Retort.

The taxonomy of coherence relations, developed here and illustrated in Figure 1, is not necessarily meant to be complete, but if any further classes of coherence relations are proposed, either they must be located within this description of the discourse situation or the description itself must be elaborated.

4. The Data to be Explained

In Part 5, a fragment of conversation is analyzed. The microproblems that face one of the participants are pointed out, and it is shown how the coherence relations come to her aid. In this section the fragment is described.

It comes from the beginning of a videotaped conversation between a man X and a woman Y.* The man enters the room first and sits down. Several minutes later the woman enters carrying four large manila envelopes in her arms. She sits down and after an exchange of pleasantries, they begin the conversation shown in Figure 2. (The utterances are labelled D1 - D9, since they are referred to frequently below.)

Both people are very much aware of the TV camera on the other side of the room and of the microphones on the table in front of them. They appear rather nervous as a result, although Y disclaims any nervousness. It is likely that both are concerned about projecting a favorable image, or at least not projecting an unfavorable one, and Y at least evinces concern about maintaining the conversation. I do not think this setting makes the data less natural, for such concerns are hardly unusual in conversational encounters. The two have met each other only briefly before, and this is their first lengthy conversation.

*-----
* I am indebted to Will Leben and Dave Evans for making this videotape available to me.

- (D1) X: What's all this mail?
- (D2) Y: My child is entering a Q-tips art contest.**
You see, you haaa-
[Laughs]
- (D4) You don't have any children, obviously.
You must
- (D5) You have to either draw or make things with the little Q-tips.
- (D6) So she thinks she's going to win an eight thousand dollar
first prize.
- (D7) So I have to send in this trash for her.
[Pause]
- (D8) All these nice things made out of Q-tips.
- (D9) And of course all the Q-tips will fall off in the mail.

Figure 2.

Examination of the videotape reveals much that is not shown in the transcript. From the transcript, it looks like X initiates the topic of the mail. In fact, the videotape shows that Y spends the few seconds before utterance (D1) moving the envelopes about, and in the second before (D1), lifts them high enough to hide half of X's face from the camera. Thus, while X seems to initiate the topic by asking the question, the tape shows that Y could well have induced him to ask it by means of her gestures.

At the beginning of utterance (D2), both participants are looking down at the envelopes. Halfway through (D2), Y looks up at X. At the

** For the international audience, a Q-tip is a small stick with a cotton wad at each end, used for cleaning out one's ears, manufactured by Johnson and Johnson, Inc.

beginning of "You see", X looks up at Y. He is clearly suppressing a smile, and in the next second he breaks into a grin. X responds immediately by laughing.

X's grin can be interpreted as an evaluation of what Y has said, with at least some negative connotations. Y responds with embarrassment, first laughing while looking down and shifting a bit in her seat, then while saying (D4) pushing her hair back with each of her hands in turn. As discussed in Part 5, X's grin has a strong effect on Y's subsequent utterances, and therefore must be counted as a move in the conversation. (Call it (D3).)

During utterances (D1) and (D2), Y's body is at a moderate angle, leaning slightly forward but not too far. As she begins utterance (D4), she leans forward slightly, and remains at that angle until resuming her answer in (D5), when she resumes her former position. It is as if her body position is bracketing the side sequence (D4).

While saying (D5) she goes through a fascinating sequence of gestures. On the word "either" her two hands are in front of her, with the two index fingers pointing at each other, as though to pose the two alternatives. On the word "draw", she draws a circle in the air with her left index finger. On the word "make", both her index fingers are pointing downward toward the envelopes, and on "Q-tips" she grasps the sides of the envelopes, which she then shuffles for the remainder of the fragment shown in Figure 2.

It is interesting to see how this sequence of gestures interacts with the interruption of (D4). As she says "You haaa-" her hands are moving toward each other with the index fingers pointing at a slightly downward angle, as though preparing for the gesture associated with "either". Then the following happens while she is pushing her hair back during (D4). She first pushes her hair back with her left hand, but at the same time her right hand remains in position, index finger pointing down at a slight angle. Then as she pushes her hair back with her right hand, her left hand resumes that position, index finger pointing downward. It is as if the hand not pushing her hair back, is holding

the floor for the next utterance (D5), which she has already planned and begun.

This visual information will aid our analysis in Part 5.

5. The Data Explained

5.1. The Logical Answer. X's question is "What's all this mail?" Since there are unusual elements in Y's answer, she must provide a certain amount of background material. It is instructive to see what a direct, fully informative, one-sentence answer would look like:

- Envelopes
- (11) containing sculptures
 - (12) made out of Q-tips by my child
 - (13) for an art contest she is entering,
 - (14) (sponsored by the manufacturers of Q-tips)
 - (15) in which one must draw or make things with Q-tips
 - (16) (in hopes of winning an \$8000 first prize),
 - (17) which I am sending to the Q-tips company for her.

To describe mail, one must say what the contents are and who the recipient is to be, since what mail is is precisely contents of a package or envelope sent to some recipient. This accounts for (11)-(12) and (17).

But further description is required, and this raises the question of where the description should stop. The answer given in Section 3.4 is that the description must explain all that is non-normative in terms of the normal, for a link must be provided from what the Speaker says to what the Listener already knows. When this is done, the Listener will be able to supply further expansions on his own if he wishes, for this is part of what normality entails. To see this, consider two other possible situations and their appropriate answers. If it were an ordinary art contest and the daughter's entry were sketches, Y could have said only, "Sketches my child is sending to an art contest." Art contests are normal and sketches are normal entries. On the other hand, if it were an ordinary art contest and the daughter's entry were Q-tips sculptures, Y would have been obliged to supply some motivation for the daughter's choice of that medium.

In the present case, lines (14) and (15) provide sufficient explanation of the medium. Line (14) does not appear in the actual response, but it can be inferred from the phrase "Q-tips art contest."

A description of a contest would have to include at least three pieces of information, determined by what a contest is:

1. What the contestants must make or do in order to enter (the entry),
2. the sponsors who judge the entries, and
3. the prizes awarded by the sponsors, which the contestants enter the contest with the goal of winning.

Thus, (14) - (16) count as a description of the contest. But while (16) is part of the description, it is not really a necessary part of Y's answer. It is probably not out of the ordinary, and does not help to explain the non-normative nature of the mail. I will speculate briefly below about why this information is included in the actual answer.

5.2. The Actual Response Begins. The actual response contains all the information in (11) - (17), but (11) - (17) could not be said in casual conversation. It has to be chopped and rearranged. Let us imagine several ways this could happen. For convenience, we will label the essential parts of the answer:

Answer (A): The envelopes contain Q-tips sculptures.

Background 1 (B1): The Q-tips sculptures were made by my child as her entry for a Q-tips art contest.

Background 2 (B2): In the Q-tips art contest, one must draw or make things out of Q-tips.

Since background information for a statement can either precede or follow the statement, there are four possible orders:

A + B1 + B2:

These envelopes contain Q-tips sculptures. My child made them for a Q-tips art contest. In this contest, you have to draw or make things out of Q-tips.

A + B2 + B1:

These envelopes contain Q-tips sculptures. There's a Q-tips art contest in which you have to draw or make things out of Q-tips, and my child is entering it.

B2 + B1 + A:

There's a Q-tips art contest in which one must draw or make things out of Q-tips. My child is entering it, and these envelopes contain her entry.

B1 + B2 + A:

My child is entering a Q-tips art contest. One must draw or make things out of Q-tips. These envelopes contain her entry.

Y chose the last of these structures. To try to decide why would be pure speculation, but several relevant factors can be mentioned.

It is difficult to lexicalize the objects that above are somewhat unhappily called "Q-tips sculptures". Saying B2 before A eliminates this problem, favoring the last two structures.

Secondly, in order to say B2, one has to introduce the Q-tips art contest in an existential construction or an adverbial. It is just as easy to introduce it by saying the child is entering the contest. This favors the first and last structures in which B1 precedes B2.

Third, by placing the answer A last, Y is more likely to hold the floor until all the information is imparted. This favors the last two structures.

The final factor is a matter of conversational style. In addition to having the immediate goal of answering the question, it is likely that Y has a further goal of carrying on a good, or even witty, conversation. With the envelopes, she certainly has the material to do so. Her animation during her answer is a clear indication that she is intending to do more than merely answer the question. She is trying to tell a good story. A common device that adds a touch of suspense to a story is to state the non-normal before explaining it. From this consideration, the first structure could have given the answer the most punch. For example, the exchange could have gone as follows:

X: What's all this mail?

Y: Q-tips. My child's entering a Q-tips art contest....

The third structure would be the flattest. This factor combined with the lexicalization problem would dictate the fourth structure.

Y begins her answer with this structure:

(D2) My child is entering a Q-tips art contest.
You see, you haaa-

When she completes B2 (in utterance D5), she gives it a form that makes it more than just Background Material. It becomes an Elaboration.

(D5) You have to either draw or make things with the little Q-tips.

Utterance (D5) asserts that one takes on the obligation of doing what the contest requires, and that is equivalent to entering the contest. (D5) does not merely modify an element in the previous clause as the Background relation would require; it restates the assertion of that clause in a more fleshed out form, thereby fulfilling the Elaboration relation.

Before completing (D5) however, she is interrupted, in a way that changes the rest of her answer significantly.

5.3. The Interruption and its Effects. While just beginning her Elaboration on the nature of the contest, she looks up, the smile that X has been trying to suppress breaks into a grin, and they both laugh. His reaction to the notion of a Q-tips art contest is a negative evaluation of sorts. Although it would be impossible to say what the "propositional content of the smile" would mean, it is certainly true that the smile indicates an negative relationship between what Y has just said and her goal of maintaining a favorable image, thus fulfilling the Evaluation coherence relation of Section 3.3. To see this, imagine Y had said in perfect earnestness that she herself was entering a Q-tips art contest. For X to grin in the same way would have been rude, and Y would have had to choose between taking offense or justifying her involvement. As it is, she still must justify her involvement, but she has the additional option of distancing herself from the situation. She does both by saying

(D4) You don't have any children, obviously.

This is a conclusion, and X can be expected to recover the line of reasoning that led to the conclusion. It would go something like the following:

(18) Obligations to children often cause one to do things one wouldn't ordinarily do.

If you had children, you would know this, and would thus understand my involvement with a Q-tips art contest, and thus would not evaluate negatively.

You have evaluated negatively.

Therefore, you don't have any children, obviously.

Premise (18) serves both to motivate Y's involvement and distance herself from it. Utterance (D4) thus serves as a Retort to the negative evaluation (D3).

The next utterance (D5) seems to be unaffected by the interruption. The reason for this is very probably that it was entirely planned out before. There are several indications of this. She had already begun the utterance before the interruption, and when she resumed it, she did so with exactly the same words. As noted in Part 4, she also resumed the same body position, and she completed the sequence of gestures that she had already begun and had held on to during the interruption. The rest of her answer, however, does seem affected by the negative evaluation.

The interruption could have had two kinds of effects on the subsequent discourse. First, it could have prompted Y to provide further justification of and further distancing from the situation. Second, it could have disrupted the orderly production of her answer in subtle ways that are perhaps best described in terms of two different models of the flow of discourse.

Chafe (1978) has distinguished between "hierarchical" and "flow" models of discourse production. In the hierarchical model, a discourse

is viewed as a tree-structured object in which an utterance's place in the tree is the prime determinant of the form of the utterance. In Sections 5.1 and 5.2 it was a hierarchical model of our fragment that was derived. In the flow model, one takes more seriously the fact that a discourse takes place over time. Here the chief influence on the form of an utterance is seen to be the content and effectiveness of the previous few utterances. The hierarchical model explains well-planned-out discourse, which tends to be dominated by the Expansion relations; the flow model better characterizes rambling discourse, which is frequently dominated by Strong Temporal relations. Of course, the two models are not incompatible, for where a discourse seems hierarchical, we could view the flow model as describing how a Speaker walks through his hierarchy, and where flow factors seem to dominate, we could say we have a very flat hierarchy.

The remainder of our fragment provides a nice illustration of the interaction between the two perspectives. We might speculate that the interruption perturbed the orderly hierarchical progression of the answer just enough to allow flow factors to enter the picture.

The next utterance

(D6) So she thinks she's going to win an eight thousand dollar first prize,

is quite problematic. On the one hand, hoping to win the first prize is an essential component of entering a contest, along with making whatever entry the contest requires. In that sense, (D6) stands in a Parallel relation to (D5) and both together function as an Elaboration of (D2).

But (D6) is not an essential part of the background information for the answer to X's question, for it does not explain anything that is non-normative. This leads us to ask why Y said (D6). There are two possible explanations. For one thing, not only is hoping to win the contest an integral part of entering, but the continuance of this hope is a causal consequence of the entering. There is a natural flow from a mention of entering to a mention of the hope. This would explain the conjunction "so".

Furthermore, the daughter's high expectations provide a very strong motivation for Y to take the trouble to mail the entries. One does not like to shatter one's child's dreams. For this reason, (D6) functions as a further retort to X's negative evaluation.

Thus, there are three possible functions of this utterance. The phrase "an eight thousand dollar first prize" elaborates on the nature of the contest. "She thinks she's going to win" provides motivation for the situation. "So" makes the causal connection between the entering and the thinking explicit. It is quite reasonable to suppose that the utterance in fact served all three functions.*

The next utterance completes the answer:

(D7) So I have to send in this trash for her.

But here also there is an interesting contrast in what is accomplished by the information content of the sentence and the form it takes. The information that needs to be conveyed to answer the question is

These envelopes contain her entry for the contest.

This can be inferred from (D7) together with the clear reference of "this trash" to the envelopes, the bulky appearance of the envelopes, and the fact that (D7) comes after (D2) and (D5) as an answer to (D1). The logical answer is thus complete.

But the form of the sentence may do two other things as well. First the "so" could be a signal of the causal relation between entering the contest and sending the envelopes, so that (D7) functions as a reasonable continuation in the flow of the story of an event.

An alternative is that "so" indicates the causal connection between the daughter's high expectations and Y's having to send the envelopes. Insofar as this is true, (D7) functions as a further justification of

* Labov & Fanshel (1977) close their 233-page microanalysis of a fragment of a therapeutic session with the heartening information that the patient eventually recovered. I regret to report that Y's daughter did not win the \$8000 first prize.

Y's involvement in this situation, and thus as an expansion of her defense against the negative evaluation.

This interpretation is bolstered by Y's use of "have to send" rather than "am sending". This choice further distances her from responsibility for the negatively evaluated situation. She is not in complete control of her actions.

It is interesting to note that Y further distances herself from the situation by the choice of certain deprecating expressions -- "the little Q-tips", "this trash", and in the next two utterances, "all these nice things" and "of course" to introduce a disaster. These all come after the interruption.

5.4. Stretching It Out. Utterance (D7) completes the answer and closes off all the open issues in the conversation so far. This leaves the two people with the problem of continuing the conversation in some fashion. It would normally be X's turn to speak, but he shows no indication of contributing anything himself. So after a pause, Y continues with

(D8) All these nice things made out of Q-tips.

This is an expansion on "this trash" of the previous utterance, but it clearly adds no new information. It is said in a muted off-hand manner, as Y sits back, places two of the envelopes in her lap and begins to inspect them. The utterance clearly functions as a filler. It is not so much a coherent continuation, as it is a cohesive one. Reference is made to a previously mentioned entity, but nothing is said about the entity that relates to what has gone before.

Finally in (D9) she says

(D9) And of course all the Q-tips will fall off.

This initiates a new topic -- the coming trials of the Q-tips sculptures -- which will occupy the next 21 utterances. Utterances (D7) and (D9) stand in an Occasion relation to each other, in that sending the

envelopes sets up the situation for them to fall off in the mail. It thus continues the temporal and causal chain that was a superficial organizing principle for utterances (D2) - (D7). Nevertheless, it does not continue the answer to (D1), and thus is merely a way of exploiting a coherence relation to enter into a new topic, that is, into a new, more highly structured island of coherence.

5.5. Representing the Structure. Figure 3 represents the structure of the discourse as we have analyzed it. The binary relations are represented in one of two ways, illustrated in Figures 4a and 4b.




Fig. 4a




Fig. 4b

S0 and S1 are two segments of discourse linked by relation R, and A and B are labels for the two parts of the relation. The style of Figure 4a is used when a hierarchical organizing principle seems dominant, i.e. when it seems part of a well-planned out structure. The style of Figure 4b is used when "flow" factors seem prevalent, i.e. when the form or content of S1 seems a spur of the moment response to the content or effect of S0.

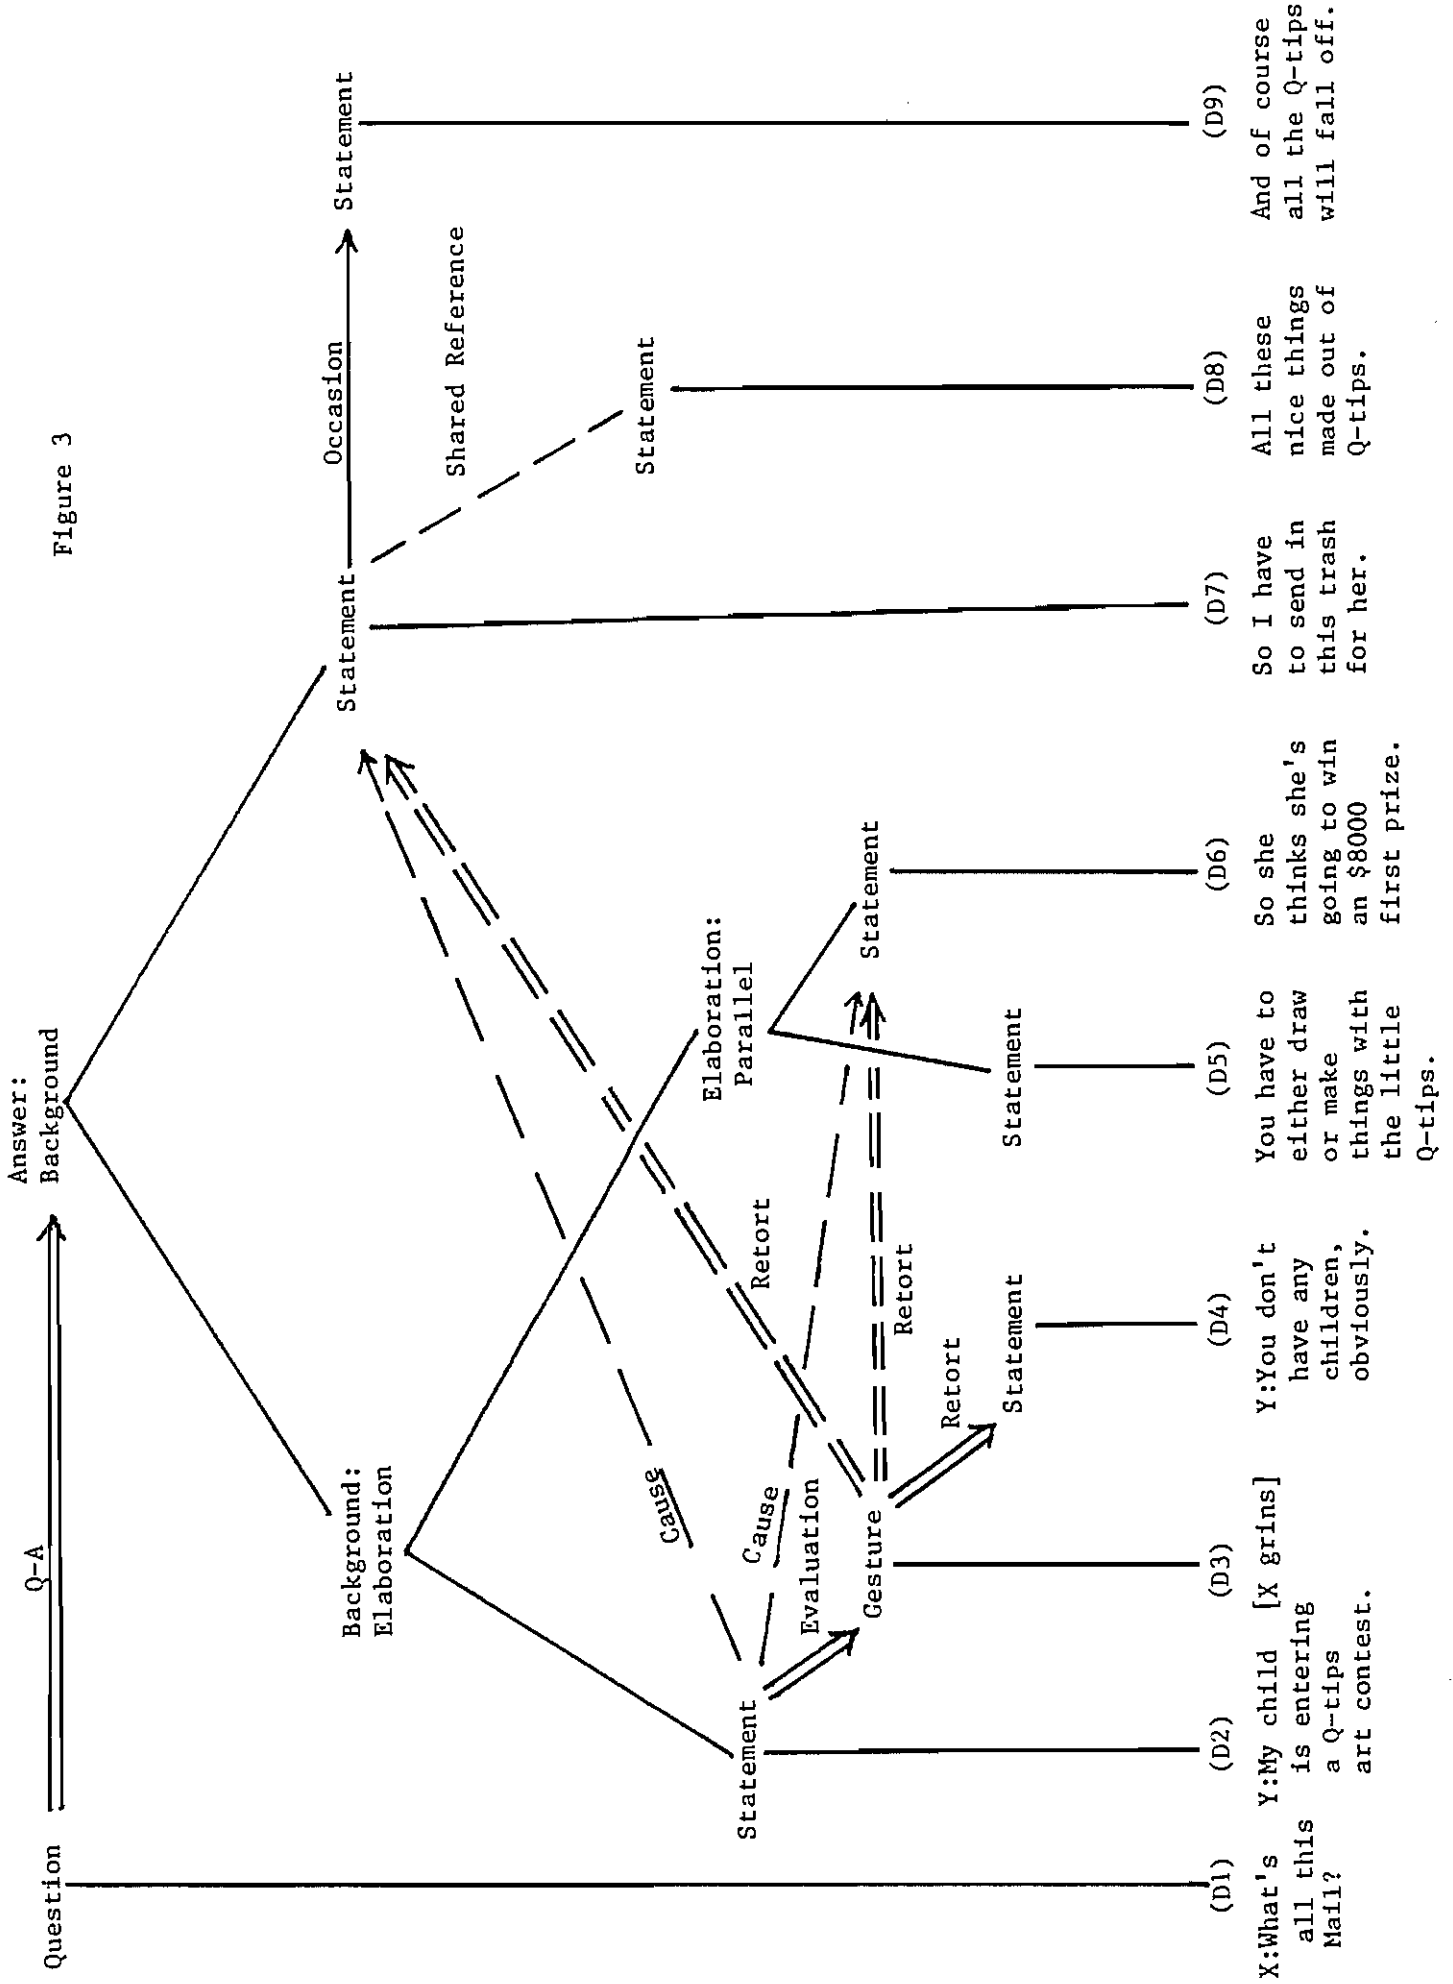
When an element of a higher relation is itself a complex structure, its node is labelled first with the role it plays in the higher relation, followed by a colon, followed by the name of the complex structure. Thus, the node labelled

Background:
Elaboration

indicates that the segment of discourse it subsumes functions as Background to some other statement in a higher Background relation, and itself has the structure of a statement followed by an elaboration of the statement.

Double lines are used to indicate relations between moves by different speakers. Dashed lines indicate a more tenuous relationship between utterances.

Figure 3



6. Summary

We have isolated some of the needs people are responding to when they make a discourse coherent. They are describing what they perceive as a coherent sequence of states or events in the world. They want to keep firmly before them the goals they are trying to achieve through the discourse. The Speaker seeks to link what is new and unpredictable in what he says with what he can expect the Listener to know already. He seeks to manipulate the Listener's processes of comprehension in order to focus him precisely on the detailed picture that is being painted with the broad brush of everyday language.

These needs are real, and the coherence relations, or coherence moves, are real resources at the Speaker's command that satisfy these needs. They are not applied just once as a coherent segment of discourse is planned out. Rather, they operate continuously as the interaction progresses. They are sometimes realized jointly, sometimes they conflict with each other or deflect each other, and sometimes they are exploited even where they are not dictated. Taken together, they turn a discourse into the intricately structured object that it is.

Acknowledgments

I am indebted to Wally Chafe, Barbara Grosz, Ann Robinson, and Jane Robinson for useful discussions leading to ideas expressed in this paper, and I am especially indebted to Dave Evans for his many insights during the long hours we spent analyzing the videotaped conversation. I am of course solely responsible for any errors.

REFERENCES

- Allen, J. & C. Perrault 1978. Participating in dialogues: Understanding via plan deduction. Proceedings, Second National Conference, Canadian Society for Computational Studies of Intelligence, Toronto.
- Balzer, R., N. Goldman, & D. Wile 1977. Informality in program specifications. Proceedings, International Joint Conference on Artificial Intelligence, 389-397, Cambridge, Mass. August 1977.
- Chafe, W. 1978. The flow of thought and the flow of language. To appear in T. Givon (Ed.), Syntax and semantics, Vol. 12. New York: Academic Press.
- Cohen, P. 1978. On knowing what to say: planning speech acts. Technical Report No. 118, Department of Computer Science, University of Toronto. January 1978.
- Grice, H. 1975. Logic and conversation. In P. Cole and J. Morgan (Eds), Syntax and semantics, vol. 3, 41-58. New York: Academic Press.
- Grimes, J. 1975. The thread of discourse. The Hague: Mouton.
- Halliday, M., & R. Hasan 1976. Cohesion in English. London: Longman.
- Hobbs, J. 1976. A computational approach to discourse analysis. Research Report 76-2, Department of Computer Sciences, City College, City University of New York. December 1976.
- Hobbs, J. 1977. From "well-written" algorithm descriptions into code. Research Report 77-1, Department of Computer Sciences, City College, City University of New York. July 1977.
- Hobbs, J. 1978. Coherence and coreference. SRI Technical Note 168, SRI International, Menlo Park, California. August 1978.
- Hobbs, J. & J. Robinson 1978. Why Ask? SRI Technical Note 169, SRI International, Menlo Park, California. October 1978.
- Labov, W. and D. Fanshel. 1977. Therapeutic discourse. New York: Academic Press.
- Levy, D. 1978. Communicative goals and strategies: Between discourse and syntax. To appear in T. Givon (Ed.), Syntax and semantics, Vol. 12. New York: Academic Press.

Longacre, R. 1976. An anatomy of speech notions. Ghent: The Peter de Ridder Press.

Longacre, R. 1977. The paragraph as a grammatical unit. Symposium on Discourse, UCLA. November 1977.

Phillips, B. 1977a. Discourse connectives. KSL-11. Department of Information Engineering, University of Illinois at Chicago Circle. March 1977.

Phillips, B. 1977b. A calculus of cohesion. Paper presented at the Fourth LACUS Forum, Montreal, Canada. August 1977.

Robinson, A. 1978. Investigating the process of natural-language communication: A status report. SRI Technical Note 165. SRI International, Menlo Park, California. July 1978.

Figure 3

