

CONVERSATION AS PLANNED BEHAVIOR

Technical Note 203

December 1979

By: Jerry R. Hobbs
Computer Scientist
Artificial Intelligence Center

and

David A. Evans
Stanford University
Stanford, California

SRI Projects 7910 and 7500

Report Documentation Page

Form Approved
OMB No. 0704-0188

Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

1. REPORT DATE DEC 1979		2. REPORT TYPE		3. DATES COVERED 00-12-1979 to 00-12-1979	
4. TITLE AND SUBTITLE Conversation as Planned Behavior				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) SRI International, 333 Ravenswood Avenue, Menlo Park, CA, 94025				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES 52	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

CONVERSATION AS PLANNED BEHAVIOR

Jerry R. Hobbs
SRI International
Menlo Park, California

and

David Andreoff Evans
Stanford University
Stanford, California

1. Brief Overview

Perhaps the most promising working hypothesis for the study of conversation is that the participants can be viewed as using planning mechanisms much like those developed in artificial intelligence. In this paper, a framework for investigating conversation, which for convenience will be called the Planning Approach, is developed from this hypothesis. It suggests a style of analysis to apply to conversation, analysis in terms of the participants' goals, plans, and beliefs, and it indicates a consequent program of research to be pursued. These are developed in detail in Part 2.

Parts 3 and 4 are devoted to the microanalysis of an actual free-flowing conversation, as an illustration of the style of analysis. In the process, order is discovered in a conversation that on the surface seems quite incoherent. The microanalysis suggests some ways in which the planning mechanisms common in artificial intelligence will have to be extended to deal with conversation, and these are discussed in Part 5. In Part 6, certain methodological difficulties are examined. Part 7 addresses the problem that arises in this approach of what constitutes successful communication.

2. A Framework for the Investigation of Conversation

2.1. The Planning Mechanism

Research into problem-solving and planning has been one of the healthiest areas of artificial intelligence (Newell & Simon 1959, Fikes & Nilsson 1972, Newell & Simon 1972, Sussman 1975, Tate 1975, Sacerdoti 1974, 1977, Waldinger 1975). This work has dealt for the most part with single agents in simple microworlds performing only physical actions, such as the manipulation of a set of blocks on a table. Recently, however, there have been efforts to apply planning models to problems of discourse. These have taken three main tacks. First there is work on dialogs about plan-based activities. For example, Grosz (1977) and A. Robinson (1978) have studied dialogs between experts and apprentices repairing an appliance. The second main trend is in using planning models to determine the goals and plans of characters in a story. Among this work are Schank and Abelson (1977), Bruce and Newman (1978), Wilensky (1978), and Beaugrande (1980). Most relevant to the work described in this paper, however, is the third trend in planning and discourse -- the investigation of the planning that must go on in the production of utterances. Cohen (1978), Allen and Perrault (1978), and J. Moore (1978) have developed models for the planning of single speech acts. The goals of this paper are to go beyond the planning of single speech acts to the planning of longer stretches of conversation. In this it is related to the work of Levy (1979) describing how the goals of a speaker structure the explanation of some decisions just made.

Certain confusions often arise in discussions of the artificial intelligence approach to discourse because of the lexical ambiguity of "goal" and "plan". There are several intuitive senses of these words. The ones intended in this paper are as follows: A goal is a conceptualization of a specific state or class of states in the world and/or in himself that a person, consciously or unconsciously, strives to attain. A plan is some consciously or unconsciously constructed

conceptualization of one or more sequences of actions aimed at achieving a goal.

But in addition, "goal" and "plan" have become technical terms in artificial intelligence. The Planning Approach seeks to capitalize on this ambiguity by assuming some sort of correspondence between the intuitive and technical senses. But before we get into the correspondence, we need to define the technical terms. In this section a planning mechanism is defined, and in order to avoid some of the confusions, the definition will be given in somewhat greater detail than would ordinarily be required for readers familiar with the artificial intelligence literature. To establish the link between older planning research in artificial intelligence and the Planning Approach to discourse, two kinds of examples are given: examples typical in a blocks world, and examples that may prove useful in the domain of conversation. In Section 2.2, the nature of the correspondence between the intuitive and technical terminology is explored.

A planning mechanism consists of the following:

1. A formal language, such as predicate calculus, with a semantics that allows states in the world to be expressed in the language.

2. A goal, or set of goals. A goal is a logical formula in the formal language. Intuitively, it describes a condition the planning mechanism is to attempt to achieve, or a proposition it is to attempt to cause to be true. Examples of goals a planning mechanism might have are "on(BLOCKA,BLOCKB)" or "impressed-with(JOHN,ME)".

3. A set of actions, which can be described in the formal language. Some actions, though not necessarily all, are directly executable in the world by means of output devices. Typical actions might be to build a tower, or to impress the listener with one's intelligence. These would not ordinarily be directly executable actions, but would have to be decomposed, via the procedure of item 5 below, into more "primitive", directly executable actions. Typical directly executable actions might be to grasp a block, or to utter a sentence or effect a particular intonation contour.*

4. A set of "beliefs", about the world and about itself, expressed as axioms in the formal language. Especially important among the axioms are what may be called causal axioms, expressing facts about what causes or enables or tends to cause or enable what. For example, moving a block to the top of another block causes the first block to be on the second. A block having a clear top enables it to be grasped. In the domain of conversation, we would need axioms expressing, for example, the facts that humor generally causes the listener to have a favorable image of the speaker and that descriptions of mishaps are often humorous. We will frequently speak of such causal axioms as conversational strategies.

For convenience, we will also include under the heading of "causal axioms" those axioms that specify how one action "decomposes into" one or more other, more primitive actions. For example, moving block x to point y decomposes into grasping x, moving the arm to point y, and ungrasping x. The action of responding to the formula "How are you?" decomposes into the actions of saying "Fine" and the action of asking "How are you?" The more primitive actions may or may not be temporally ordered; in these two examples, they are. These axioms capture the notion of expressing actions at different levels of detail. It is possible for an action to have more than one decomposition.

Causal axioms play a key role in the planning process, as described in item 5. They provide the link between goals and actions.

5. A planning process, or a procedure for deriving a sequence of actions that will bring about the goal. For most of this paper, it will be sufficient to assume a fairly simple procedure, one that works from the top down employing "means-ends analysis". That is, given a goal G, the procedure searches through its causal axioms for axioms whose causal consequent matches the goal, i.e. axioms of the form "A cause G" where A is an action, or "S cause G" where S is a set of actions, perhaps with

* The overcareful reader will notice that the word "action" is used to refer to a type of event in the world, a token of that type, and a formal expression whose interpretation is that type. He should view these as examples of metonymy, not of imprecise thought.

constraints on temporal ordering. Where there is more than one such axiom, the procedure must (at some point) choose one of them. The problem of how that choice is made is addressed below.

For each action A that is chosen, the procedure then searches through the axioms for all axioms of the form "P enable A" to determine the preconditions for action A. P then becomes a subgoal to be satisfied in the same way as the original goal G.*

If an action A is not directly executable, it is decomposed into "more primitive" actions by means of axioms of the form "A decomposes-into S", where S is a set of actions, perhaps temporally ordered. Again, where several such axioms exist, a choice is made.

The procedure continues until it has derived a sequence of actions, all of which have all their preconditions currently true or satisfied by previous actions, and all of which are directly executable. **

The process is nondeterministic; there may be many ways of choosing an action or sequence of actions to satisfy a particular goal or subgoal. In this paper, we will mention some constraints on possible choices in Section 2.3, but we will not consider the problem of choosing among the various plausible options otherwise. Planning mechanisms in artificial intelligence generally use some heuristic evaluation function, but these tend to be highly ad hoc. There is a large body of

* By "state P enables action A" we mean "state P not holding causes action A not to occur". Expressing preconditions in this fashion causes our logic to be nonmonotonic (cf. McCarthy 1977, McDermott & Doyle 1978), in that adding a new axiom can invalidate a plan by adding a new precondition. Practically speaking, it is even worse, for our search for all axioms of that form may be limited by resources, so that a plan could be invalidated by axioms the planning mechanism already has. This seems a realistic reflection of the situation people find themselves in.

** In using causal axioms, we are moving away from the operators of Fikes & Nilsson (1972) to a formalism more like those used by Kowalski (1974) and Rieger & Grinberg (1977). No power is lost in this move since the formalisms are equivalent. Moreover, the causal axioms encode knowledge already required by the natural language interpretation and generation components that must be part of a total conversational system. Finally it seems easier to implement a more general control strategy with causal axioms.

research on multigoal, attribute-based evaluation and decision theory (Keeney & Raiffa 1976, Cochrane & Zeleny 1973), but it is not at all clear whether this could be folded into a model of the fine details of conversation. we will remain uncommitted on the choice functions used, in the belief that human choice is a mystery whose solution is not accessible to present-day cognitive science.

The sequence of actions produced in the planning process is sometimes called a plan. However, it will be more useful for us to use the term plan to refer to a tree-like structure that represents the derivation of this sequence of actions. The following example from the blocks domain should be adequate illustration: Suppose block C is on block A, and block B is standing by itself. The goal is to put block A on top of block B.




Figure 1.

Suppose the planning mechanism has the following facts available:

```

move(x,y,z) cause on(x,z)
move(x,y,z) decomposes-into pick-up-from(x,y); put-on(x,z)
on(x,y) enable knock-off(x,y)
knock-off(x,y) cause cleartop(y)
on(x,z) enable pick-up-from(x,z)
cleartop(x) enable pick-up-from(x,y)
pick-up-from(x,y) cause cleartop(y)

```

("Move(x,y,z)" means the action of moving x from y to z. The ";" indicates temporal ordering. The rest should be self-explanatory.) Suppose furthermore that the actions "knock-off", "pick-up-from", and "put-on" are directly executable. The conditions "on(C,A)", "on(A,T)" where T is the table, "cleartop(C)", and "cleartop(B)" are all currently true in the world. Then Figure 2 shows the tree-like structure that

represents the plan that would be derived.* (Joined branches are conjunctive, unjoined branches are disjunctive, and the arrow represents temporal ordering.)




Figure 2.

At any moment during the execution of the plan, there is a directly executable action being executed or about to be executed. We will call the path from the top-level goal to that action the leading edge. This represents all the actions that are currently being performed and all the conditions the planner is currently attempting to achieve. The portion of the tree to the right of the leading edge represents the part of the plan that has not yet been carried out and can contain disjunctive branches, indicating that the planner has not yet chosen among the options. The portion to the left represents that part already carried out, and of course cannot contain disjunctive branches.

In the domain of conversation, the planning mechanism will begin with high-level conversational goals and use its causal axioms, including its conversational strategies, to generate a plan whose actions are utterances, gestures, and other conversational moves.

 Even this simple example is oversimplified. We have not worried about the hand being free to pick up a block nor about the problems that would arise if C were to be removed from A by placing it on B.

Typical conversational plans will require a number of steps to execute and may go awry at any point. Thus we must imagine the planning mechanism working in tandem with two other components -- a monitor and a debugger.

The monitor seeks to relate inputs from other participants in a conversation to the conversational plan, in order to extend the plan or judge its success. Research that has tried to develop ways of relating an utterance to a plan may be viewed as work on just such a component. Examples of this research will be found in Grosz (1977), A. Robinson (1978), Hobbs and J. Robinson (1978), Allen (1979), and Genesereth (1978). In the conversation analyzed in Part 4 of this paper, Y's reaction to X's moves (D5) and (D12) are interesting examples of monitoring.

In our planning, we are using causal axioms that are at best only plausible, and sometimes actions don't cause what they are expected to cause. If the monitor has learned new information that contradicts what was expected, a debugger must attempt to determine which of the causal axioms happened not to be true, to account for this by searching deeper into the knowledge base for factors not previously considered, and to call on the planner to generate a repair and a new plan.

The planning mechanism presented here is only an initial version. It is inadequate in several respects and will have to be complicated at least in the ways discussed in Part 5. There has been a significant amount of research on the problems of conflicting and interacting goals; these problems will not be addressed in this paper. Conflicting goals occur only once in the conversation analyzed in Part 4, and there it is assumed that the conflict is noticed and planned around.

2.2. The Working Hypothesis

Intuitively, the working hypothesis for the Planning Approach is that participants in a conversation may be viewed as using planning mechanisms very much like the one defined above for producing their

utterances and other conversational moves. This needs to be made more precise, and we will do so by giving three carefully constructed versions, H0, H1, and H2. The reader who lacks the taste for this sort of thing and promises not to nitpick can get the sense of these hypotheses from Figure 3.




Figure 3.

The three hypotheses make claims about correspondences between a person P on the one hand and a planning mechanism M on the other. They require certain ontological commitments, which will be specified. They all involve the notion of the interpretation of elements in the formalism of the planning mechanism, that is, the assertion of a correspondence between formal elements and things or processes in the mind or behavior of the speaker. The interpretation of a class of elements in the planning mechanism -- goals, actions, axioms, and so on -- is a mapping from that class to a set of things that are assumed to exist in the world. It will be convenient to use "move" to refer to one

of the person's utterances, gestures, or other conversational moves. References to the person's goals employ the intuitive sense of the word, to the planning mechanism's goals the technical sense.

The first hypothesis, H1, makes the ontological assumptions that there are such things as P's possible moves and his possible goals, and that it makes sense to speak of a move performed "in the service of" a goal. H1 then runs as follows: Given a goal G that P has, a move U performed in the service of G, and an interpretation of M's goals as possible goals of P and of M's actions as possible moves of P, then there is a goal G' of M whose interpretation is G, an action A of M's whose interpretation is U, and a plan derivable by M realizing G', in which A is the initial action.*

Less pedantically, this says that if a person makes a move in the service of some goal, then the planning mechanism can produce the same move from the same goal, by means of the planning process.

M's plan is then the formal derivation of the move. Because of the monitoring and debugging phases, M's plan can change during the course of a conversation. The sequences of plans for all the moves of all the participants in a conversation is the formal derivation of the conversation as a whole.

Intuitively, hypothesis H2 says that all the ways we have of talking about the planning mechanism are valid ways of talking about the speaker. The technical planning terminology is less metaphor and more an accurate description of reality.

The ontological assumptions of H2 are more radical. In addition to sets of possible goals and possible moves, we need to assume sets of P's possible beliefs and to agree that it makes sense to talk about P having

* There is a formal trick here. In many cases a speaker will be carrying out a plan in which move U is some action in the middle. We would then have to worry about the consistency of the previous actions in the plan with P's previous moves. To avoid this, we will assume that the previous portions of the plan are peeled away and their effects incorporated into the current state, and that we have a new plan about to be embarked upon.

a plan for realizing a goal, where the plan is effected by performing certain moves and consists of or is based on certain beliefs. We need furthermore to extend our interpretation of elements in the formalism to include the interpretation of M's axioms as P's beliefs and the interpretation of M's plans as P's plans. The latter must be stated in such a way that the structure of M's plan and P's plan are the same. The detailed definition of the interpretation of a plan will not be given here since it would require a repetition of the entire account of the planning process, stated as ontological assumptions about people, and a second repetition asserting the correspondences between a plan of M's and a plan of P's.

Assuming all of this, hypothesis H2 is as follows: Given a goal G of P's and a move U by P, and an interpretation of M's goals, actions, axioms and plans as possible goals, moves, beliefs and plans of P, then there is a goal G' of M whose interpretation is G, an action A of M whose interpretation is U, and a plan PL of P's for realizing G and a plan PL' derivable by M for realizing G', such that PL is the interpretation of PL', U is an initial move in PL, and A is an initial action in PL'.

H0 is a weaker hypothesis. The only ontological assumption required is that there is a set of P's possible moves. This should not be especially problematic for most readers. Then hypothesis H0 is as follows: Given a move U of P's and an interpretation of M's directly executable actions as possible moves by P, there exists a goal G' of M's, a directly executable action A of M's whose interpretation is U, and a plan derivable by M from G' in which A is the initial action.

In H0, the only correspondence assumed is between the moves performed by each. We do not assume a correspondence between M's goal and P's goal, nor indeed that such a thing as a person's goal exists.

Now some of the consequences: Between H0 and H1 falls the line that separates what might be called the "sociology of discourse" and what might be called the "psychology of discourse". With H0, we would have a purely formal description of observable behavior, in the same sense that

a simple grammar can be used to describe chess language. There would be no claims about psychological reality. Accepting H1 amounts to the decision that we are doing psychology, that we are investigating purposeful behavior. We are claiming that our plan is a possible mechanism for realizing a real goal in the speaker's mind. One who accepts H2 believes it is possible in principle to construct a "correct" blow-by-blow, computational account of what goes on in a speaker's mind.

Whether a researcher views himself as doing the sociology or the psychology of discourse seems to have consequences in what he looks for. For the sociologist of discourse, conversation can be studied as a "social object", in isolation from the cognitive processes of its participants, and abstract rules can be discovered that seem to characterize large classes of conversations. Typically, he tries to identify culturally defined discourse types and rules, that belong to a group without belonging to any particular member. By contrast, the psychologist of discourse makes conjectures about possible mental representations and processes implementing these discourse types and rules in individual speakers. For him, rules of turn-taking do not merely exist; people know and use them. He may even be interested in showing how the discourse types and rules belonging to the culture arise out of the typical goals, memory structures, and so on, of speaker/listeners.

An analogy with baseball may bring out the distinction more sharply. A sociologist of baseball is likely to be satisfied once he has discovered the official rules of the game, perhaps together with a few common strategies, such as "Never swing when the count is 3 and 0." An investigator taking the Planning Approach to baseball will attempt to build a mechanism that can play the game.*

* There is an intermediate position between H0 and H1 held by some sociolinguists, namely, that it is permissible to talk about some of a speaker's goals and beliefs, those of a distinctively linguistic or discursive character, but impermissible to "psychologize" about other goals.

In this paper we will be engaged in the psychology of conversation; we will assume H1. In doing a microanalysis of some fragment of conversation, the wisest strategy will be to assume H2 and aim for the most detailed correspondence possible between the formal derivation and what the person actually does, invoking evidence of the sort discussed in Part 6, where possible, to bolster our account. In defending the analysis, however, we will be committed only to H1.

Despite the sociology - psychology distinction, whatever compelling explanations we come up with could still be of use to the sociologist of discourse if he reads them as purely formal descriptions of observable behavior. Moreover, the results from the sociology of discourse have a very important role in the program of research that is suggested by the Planning Approach and outlined in Section 2.3.

The planning metaphor, at the very least, provides an attractive vocabulary for describing conversation, for it seems to accord with the way we feel about our conscious moves and with what we are willing to attribute to our unconscious moves. The nondeterminism of the planning process allows room for our sense of free choice. Unlike more rigid formalisms, e.g. flowcharts, behavior outside the norm is not outside the system; rather it is a result of a less common option being chosen by the planner. Unlike the rule systems proposed by ethnomethodologists (e.g. Sacks, Jefferson and Schegloff 1974), the planning metaphor allows us to be explicit about the motives that lie behind the strategies we use. Among the various mechanistic metaphors of cognitive psychology, this one seems to detract the least from our humanity.

2.3. Style of Analysis and Program of Research

In Section 2.1, a planning mechanism was defined. It had five principal aspects:

1. Goals.
2. Actions.

3. Causal axioms, including conversational strategies, for connecting goals and actions.
4. Some unspecified means for choosing among the options presented by the causal axioms.
5. The planning mechanism itself.

In Section 2.2, we considered the possible interpretations of these aspects. In this section, we consider, in terms of the five aspects, the style of analysis the planning mechanism suggests and the program of research it indicates.

Briefly, the style of microanalysis is this: When we are confronted with a fragment of conversation to be analyzed, we make our best guesses, consistent with everything we know, about the participants' goals, the moves that occur in the conversation, the causal knowledge, including conversational strategies, the participants are using, influences on the choices they make, and the planning processes that seem to be taking place. If we can cast these into the formal language, we have a formal derivation, or explanation, of the conversation.

This is not a particularly radical recommendation. It is what we find in the best of sociolinguistic research (e.g. Labov & Fanshel 1977, Gumperz 1979). But whereas there it has the peripheral role of a mode of argumentation or a heuristic for discovery, in the Planning Approach it occupies a central role: it is the English gloss of the formal derivation of the conversation, toward which the entire investigation is aimed.

An individual microanalysis becomes more plausible if it is backed up by a substantial body of research, and here the five aspects appear once more. The areas of research that are required are on (1) the typical goals that participants have, (2) their actions or moves, (3) the most common conversational strategies, (4) constraints on the choices speakers make, and (5) the operation of the planning mechanism. This indicates a fivefold program of research. By good fortune, the first four are already thriving areas of research in the various fields

that study discourse, including linguistics, sociolinguistics, ethnomethodology, psychology, and natural language processing. The Planning Approach has therefore yielded a unified framework in which to view what has heretofore seemed a diverse collection of efforts.

The five areas are to:

1. Identify and classify the most common goals that participants in a conversation seek to satisfy. Halliday (1977) and Grosz (1979) have suggested a three-way classification. They identify ideational or domain goals, or goals external to the conversation, such as a task jointly engaged in (Grosz 1977, A. Robinson 1978, Hobbs & J. Robinson 1978, Allen & Perrault 1978), a plan jointly evolved (Linde & Goguen 1978), or an event jointly experienced; textual or discourse goals, including coherence goals, or the speaker's goals to structure the conversation in a way that will ease the listener's efforts in comprehension (Hobbs 1979) and goals to refer felicitously (Clark & Marshall 1978, Reichman 1978, Grosz & Hendrix 1979); and interpersonal or social goals, including the goal of "communing", or maintaining contact, and image goals, the speaker's desire to project or maintain a favorable image, or an image consistent with the role he has chosen to play (cf. Goffman 1974, chapter 14). In the microanalysis in this paper, image and coherence goals play the greatest role.

2. Identify the actions performed by speakers. This includes verbal actions such as use of a particular sentence structure or description or word, as well as non-verbal actions involving intonation (cf. Crystal 1969) and gesture (cf. Birdwhistell 1970, Argyle 1972). Some of these actions are examined in Part 5. It would in addition be useful to have some guidelines in identifying larger scale actions that span a number of turns.

3. Describe common conversational strategies. Many of these are unique to particular individuals, but others are common to large cultures or subcultures. Included are high-level strategies that may span a large number of utterances (cf. Goffman 1974, chapter 14, for a treatment of such strategies); mid-level strategies for, e.g.,

introducing a new topic, effecting transitions between topics, managing side sequences (Jefferson 1972), opening conversations and repairing the openings when they fail (Schiffrin 1977), passing up one's turn (Wiener & Goodenough 1977); as well as very local strategies for, e.g., indicating interest with eye gaze (Kendon 1967), using intonation contour to force a particular interpretation (Sag & Liberman 1975) or to indicate discourse structure (Bolinger 1972), using prosodic cues to indicate emotion (Gumperz 1979), suggesting an ironic outcome with the "Watch something happen" class of constructions (Fillmore 1979), or holding onto one's turn with a gesture or evaluating something negatively by one's choice of words, as we will see in Part 5. All of these strategies involve certain actions causing or tending to cause certain conversational goals to be satisfied, and ought to be expressible as causal axioms.

4. Identify and classify the most common modes of discourse, or "discourse types", viewed as constraints on the choices a speaker makes. A word of explanation: It is hopeless to try to account for why speakers make the choices they do. But their culture imposes certain constraints on the options they choose. Frequently these constraints are bundled together in the form of a discourse type. The effort to classify discourse types is therefore one way of investigating the constraints on a speaker's choices.

A great deal of work has already been done on classification by sociolinguists and others, who have investigated narratives (e.g. Labov & Waletzky 1967, Polanyi 1978), planning discourse (Linde & Goguen 1978), jokes (Sacks 1974), descriptions (Linde & Labov 1975, Chafe 1979), persuasion dialogs (Archbold 1976), disputes (Brenneis & Lein 1977), task-oriented dialogs (Grosz 1977), and helping dialogs (Mann, Moore & Levin 1977).

But a caution is in order here. It is possible, given a fairly rich collection of data, to make an arbitrary number of distinctions. There must be some constraints on the kinds of taxonomies we construct. One sometimes hears arguments that taxonomization must precede

formalization; an analogy advanced as an argument comes from biology: it would have been impossible for Darwin to conceive the theory of evolution if a taxonomy of the species had not first been constructed. But we can use the biology analogy against unconstrained taxonomizing. There are many principles of classification one can appeal to in classifying the species, for example, mode of locomotion. A taxonomy based on this principle could never have led to the theory of evolution. More important than classifying is identifying the most fruitful principle of classification.

The Planning Approach suggests just such a principle of classification: We distinguish a discourse type if a speaker knows he is employing that discourse type, and if that knowledge has a substantial effect on the planning process he engages in, i.e. influences the choices he makes among conversational strategies and the conversational moves he is likely to choose for realizing his goals. Jokes provide a good example; people know when they are telling a joke, and that constrains their next move quite narrowly.

5. Examine real data that will put pressure on the formalism. This can be illustrated by the example of the investigation of syntax. There have been a number of papers that propose new transformations, perhaps designed to handle a particular class of grammatical phenomena -- these were especially common in the early days of transformational grammar -- and there have been attempts to construct transformational grammars for entire languages. These efforts in syntax correspond to the first three efforts in our program of research. But the most influential papers in syntax have been the ones presenting examples that cause trouble for the current formalisms.

Similar examples need to be found for the Planning Approach, examples of fragments of conversation that are not easily handled by the planning mechanism we have defined. For this purpose, we have chosen a fragment of a dialog in which long-range plans are not easily discernible. The two participants are "just talking", in what seems at first to be a very random and incoherent manner. Most of the rest of

the paper will be devoted to a microanalysis of this fragment in terms of the goals and plans of the participants. In the process, we will see order emerge. We can begin to understand why what was said was said.

4. The Data to be Analyzed

The fragment of conversation to be analyzed 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 a manuscript that happens to be her dissertation and four large manila envelopes. She sits down and they begin the conversation shown 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. We 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, so in our analysis we do not have to worry about shared knowledge that we lack access to.

Non-verbal activity is bracketed. Brackets at the beginning of two successive lines indicate overlaps. Periods represent half-second intervals in which nothing is said.

- (D1) [Y displays dissertation.]
(D2) [Y displays four bulky envelopes.]
(D3) X: What's all this mail?
(D4) Y: My child is entering a Q-tips art contest.
(D5) [[X grins.]
[You see, you haaa-
(D6) Y: You don't have any children, obviously.
You must...
(D7) You have to either draw or make things with the little
Q-tips.
. .
(D8) So she thinks she's going to win an \$8000 first prize.
(D9) So I have to send in this trash for her.
. .
(D10) All these nice things made out of Q-tips.

- (D11) And of course all the Q-tips will fall off.
 .
 .
 and... in the mail....
 .
- (D12) X: And it's all to be sent to Blair Nebraska, huh?
 (D13) Y: Yeah. This sounds really flaky though.
 (D14) I... I never heard of Blair Nebraska
 (D15) and you send it to a P.O. box.
 .
- (D16) So what happens too if I
 What happens if you have dishonest mailmen
 [X leans back in chair and crosses legs.]
 .
 and they see all these things going to an art contest, so
 they open it up and change it so that it's being sent
 from them? [Y leaning forward.]
 .
 .
 .
 .
 .
 .
- (D17) X: How would they change it?
 .
 Y: Well... Instead of...
 .
- (D18) Instead of the return address being my address they would
 put down their address, so they would win, you see.
 .
 .
- (D19) [Y picks up envelopes, revealing dissertation for the
 first time since (D1).]
 .
 .
 .
- (D20) Y: Not that my poor child is going to win.
 (D21) But anyway.
 (D22) X: I don't think anybody, except for a child, would want to
 enter a Q-tips art contest.
 (D23) [Both laugh. Y picks up dissertation and begins to leaf
 through it. Leans back. Shoulders relax.]
 (D24) Y: Well, maybe the postman has children.
 (D25) You never can tell.
 (D26) This is my dissertation. It's just been approved.

The gestures, eye gaze, and body positions accompanying the utterances were coded; some are discussed in Section 5.1. In addition,

Y was interviewed some time afterwards, when many of the problems discussed below had become apparent.

Like almost all transcripts of everyday conversation, this appears incoherent at first. This is especially acute since the conversation is of the "cocktail party" variety. The purpose of the conversation was just to talk. But it is precisely this that makes it good data. It provides an excellent minimal example of how conversation gets planned and of the structure that results, with little intrusion from the surrounding environment. Therefore, the structure we find here we would expect to find in any conversation. In fact, when we examine the conversation closely in terms of what X and Y are trying to accomplish, we discover quite an intricate structure.

5. Microanalysis of the Data

Since most of the action in this conversation is Y's, we concentrate on her plans. A complete treatment would require an analysis of the conversation from X's point of view as well.

We may assume Y has two principal goals -- to project a favorable image* and to cooperate with the experimental setup by maintaining the conversation.

The conversation can be divided into four episodes, each characterized by a different problem that faces Y. In the first, (D1)-(D2), Y attempts to introduce first her dissertation, then the mail, as a topic. In the second, (D3)-(D9), Y elaborates on the Q-tips. In (D10)-(D18), she tries to continue the conversation by fishing for a productive subtopic, finally hitting on the dishonest mailmen. In (D19)-(D26), due to the failure of this subtopic, she attempts to close the topic and again tries to introduce the dissertation.

Each of these episodes exhibits a high degree of internal coherence. Each provides a different example of the speaker's ability to manipulate the topic of conversation.

4.1. Attempting to Introduce Topics

Y's initial problem is to introduce a topic that will cast her in a favorable light. She has the material for it: Her dissertation has just been approved, and if they could talk about that, X would conclude she was at least intelligent enough to earn a Ph.D. degree.**

* Practitioners of certain modes of discourse about discourse are not licensed to speak of goals such as this, but that does not mean they do not exist and matter, nor that there is not evidence that can bear on what they are.

** It is more likely that Y wants to impress the unknown television audience. If she had known the sort of microanalysis the conversation would be subjected to, she would have known that was hopeless.

Here's where the Catch 22 comes in, however. If X believes Y is intelligent, then X will think favorably of Y. But if X believes Y has uttered something with the intention of causing X to believe Y is intelligent, the utterance will be interpreted as boasting, and will make X think unfavorably of Y. Thus for Y to introduce a topic that will lead to a positive image too directly is risky.


However, if she can get X to introduce the dissertation, she will have achieved the goal of talking about it without the side effect of boasting about it. To get X to introduce it, she can display it prominently, and at last, we have arrived at an action that is directly executable. Y waves the dissertation about a bit. X does not pick up on it, and the plan fails.

Another way to convey a favorable image is to project the image of a good mother, and talking about the good work of one's child is one way to do this. The problem as before is to introduce the topic, and the same hitch as before presents itself -- how to avoid boasting. The solution is the same as before. Y displays the envelopes, and the plan works as X asks, "What's all this mail?"

A broader look at the entire fragment of conversation seems to reveal a more complex goal structure here. During most of the fragment, Y goes through what has to be described as a mock checking sequence. She first picks up the envelopes, then she puts some of them in her lap, she checks the addresses, turns them over to check that they're sealed, then returns them to the table. We say "mock" because she has to have checked them before, and when she checks them now, she does so only in a very incomplete and uninvolved way. While she is checking the envelopes, she is also talking about them, and at certain points discussed below, her place in the checking sequence seems to partially generate the content of what she is saying. Then toward the end of the fragment she picks up the dissertation and begins checking that, but again in a very haphazard way. At the very end, she again begins talking about what she is checking -- now, the dissertation. This leads us to hypothesize that introducing the mail as a topic was in fact only

the first step in a new and more elaborate plan to introduce the dissertation.

Figure 4 illustrates the sequence of plans Y seems to have developed.



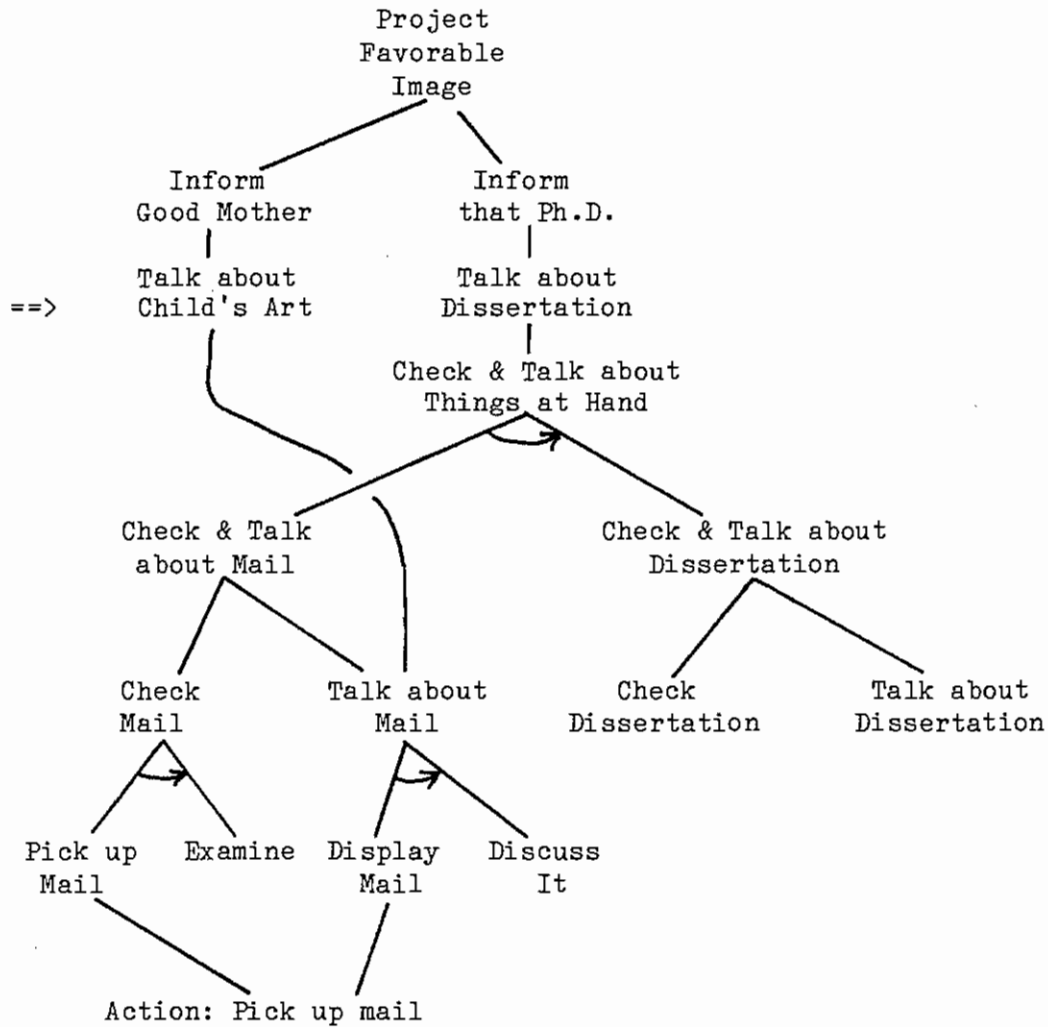


Figure 4

4.2. An Answer Perturbed

At first glance, Y's answer seems somewhat incoherent. But let's examine it more closely. To describe mail, one should describe its contents and destination,* so an answer might be

- (1) CONTENTS: The envelopes contain Q-tip designs.
- (2) DESTINATION: I'm sending them in to a contest.

* Or its source, depending on whether the stamp is postmarked or not.

In fact, (2) appears, almost as is, in (D9). But (1) is a bit unusual; the Q-tip designs require some explanation. Y must tell of the situation that gave rise to them -- the Q-tips art contest. Since this is also unusual, she has to elaborate on the nature of the contest and might, among other things, specify what the contestant must make or do (the entry) and something about the prize structure:

(3) My child is entering a Q-tips art contest.

ELABORATION:

(4) ENTRY: You have to draw or make things with Q-tips.

(5) PRIZE: There is an \$8000 first prize.

Y begins this orderly answer. She says (3) and then begins her elaboration (4). But she is interrupted, in a way that changes the rest of her answer significantly.

While just beginning (4) 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. Y must therefore justify her involvement if she is going to maintain a favorable image. She does so by saying

(D6) You don't have any children, obviously.

The implicit line of reasoning is -- if X had children, then X would understand Y's situation, and hence not evaluate negatively. So (D6) is an accusation of ignorance as a defense against the negative evaluation. This move is examined more closely in section 6.1.

The next utterance (D7) is unaffected by the interruption, since it was entirely planned out before. There are several indications of this in her gestures. The rest of her answer, however, does seem affected in subtle ways.

The next utterance

(D8) So she thinks she's going to win an \$8000 first prize,

is quite problematic. It does convey the information in (5). But (5) is not really an essential part of the background information for the answer to X's question, for it does not explain anything that is out of the ordinary.

One possible explanation for Y's saying (D8) is that 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, (D8) functions as a further retort to X's negative evaluation.

The next utterance, "So I have to send in this trash for her," completes the answer. But it also defends against X's evaluation. We will examine how in Section 6.1.*

4.3. Searching for Something to Say

It is now X's turn to talk, but he doesn't, so Y must continue in a way that coheres with what has just been said. Her first attempt involves an inappropriate elaboration (D10), uttered in a forceless, offhand manner. But it is also coherent to say "what happens next." (This has been called the Occasion coherence relation (Hobbs 1978)). Sending in the Q-tip designs provides the occasion for them to fall off. Hence, (D11) continues coherently.

She has now tapped into a productive topic, so she thinks -- possible mishaps to the designs on their way to contest headquarters. At this point X interposes with the remark that it is all going to Blair, Nebraska.

It is good to pause here to look at what X has been doing all this time, for the conversation has quite a different structure from his

* This segment of the conversation is examined in much greater detail in Hobbs (1978).

perspective. At the beginning, he leaned forward and asked his question, "What's all this mail?" Mail is characterized by its contents and its destination. He reacted a bit to Y's description of the contents. Then he examined the envelopes on the table, and noted that it was all going to Blair, Nebraska. It is likely that this is no more than a follow-up to his question about the nature of the mail.

But Y, rather than deducing the place of this remark in X's conversational plan, such as it is, incorporates it into her own. She takes Blair Nebraska to be an example of sending the packages out into the unknown, and states the generalization of which Blair Nebraska is one example, namely, that the situation is flaky. Then she gives a further example, that the destination is a post office box. At this point, she pops up to the general topic of mishaps, of which the Q-tips falling off and the strange destination are two examples, and gives her third example, which she apparently believes will turn out to be a productive subtopic -- the dishonest mailmen.

Figure 5 illustrates this development.




Figure 5.

The source of the dishonest mailmen scenario, the fragment's closest approach to literature, gets at the heart of the creative use of language, and remains a mystery. In the interview we attempted to get some insight into this process, and got instead a further burst of such creativity:

Int: Do you have any idea what made you think of this as a next thing to talk about?

- (1) Y: It just followed naturally from the discussion before. Those are the kinds of things I think about at night. I mean. . . If [X] didn't say anything I could y'know. . . You could continue on and start talking about the problems of the bureaucracy of the post office and to their uniforms and whether or not they should carry mace and problems of attack dogs. I mean you could go on forever.

One conjecture we could make is that the mock checking sequence she seems to be going through prompts her to consider all the things that could go wrong. Just before she proposed the dishonest mailmen scenario, she was checking the address and return address on one of the envelopes. What could go wrong with the return address is that someone could change it.

In (D16) Y confronts X, demanding a response with her direct "what if" question. There is a pause of 3 1/2 seconds. This is very long for a conversation like this, and it has a humorous effect on most viewers of the videotape. X has not been interested in the whole topic -- in Y's words during the interview "Q-tips aren't a big grabber in his life". As soon as she began to talk about the dishonest mailmen, he leaned back in his chair, threw his arm over the back of the chair, and crossed his legs, in a kind of defensive withdrawal. Then in response to Y's question, X does one of the worst things he could do with this topic -- he takes it seriously, and consequently dismisses it as a possibility.

Y responds by trying to construct a serious means by which the disaster could happen, and finally it becomes apparent that the topic has failed on all counts -- it is not productive of further conversation and it is not making her look good. She decides to cut bait and introduce a new topic, her original choice, the dissertation.

4.5. Escaping from a Failed Topic

She now faces the final topic-manipulation problem that we will examine -- how to escape the current topic. When asked a question, she must answer if she is to cohere. If what is said to her reflects unfavorably on her, then she should retort. Finally, it is incoherent to suddenly switch topics -- or insofar as it is coherent, it is an admission of the failure of the previous topic.

Y is thus faced with three subgoals in pursuit of maintaining the conversation in a way that will make her look good -- she must salvage the current topic by arguing for the scenario's plausibility, close the current topic, and introduce the dissertation as a new topic. These three goals interweave in her next sequence of utterances and actions. She now displays the thesis for the first time since (D1), by removing the envelopes from on top of it. She has already defended plausibility in (D18), so she is free to close the topic. One way to do this is to deny the relevance of the topic to practical affairs, which she does with "Not that my poor child is going to win." At this point however, X won't let go. He responds to the whole idea with "I don't think anybody, except for a child, would want to enter a Q-tips art contest." This challenge puts Y back in the position of having to retort and then close again before introducing the new topic. She retorts with "Maybe the postman has children", thereby denying the force of his argument, and then says "You never can tell," indicating that it is beyond their means at present to settle the question. She has thereby closed the topic again. In (D23) she has already picked up the dissertation and started to leaf through it, making the introduction of it as a topic less of a break with ongoing events. Then she says (D26), "This is my dissertation", thus succeeding in her original goal.

Y utters (D26) with a relatively flat intonation and whispered delivery, conveying a strong sense of "triumph". This makes sense within our top-down exposition of the context of the utterances in terms of Y's long-term plans. It is interesting to note, however, that a bottom-up analysis that confined itself strictly to linguistic and discursive goals could not explain this sense of triumph. (D26)

apparently follows abruptly on (D24) and (D25), "Well, maybe the postman has children. You never can tell." Yet it does nothing to expand or comment on what has preceded. This might serve to establish (D26) as a topic-initiating or topic-shifting utterance, but does not account for its "goal-achieved" delivery.* Indeed, nowhere in the text itself can we find justification for or foreshadowing of that sense of (D26).

As soon as we consider the non-verbal sequences that accompany the utterances, we find evidence of a larger pattern in which (D26) occupies a natural place. At the very beginning of the fragment, the dissertation is displayed prominently before being placed at the bottom of the stack of envelopes, which then become the object of the mock checking activity. Uncovering (discovering) the dissertation again at (D19), after the envelopes have been checked, occurs not accidentally immediately before her first attempt to shift topics (D20 and D21). This leads us to postulate that the same higher-level goals which serve to initiate the non-verbal activity at the beginning of the fragment are satisfied at its end when the dissertation is finally introduced explicitly into the conversation. This marks the achievement of the original goal and closes that portion of the conversation that we torment with our microanalysis.

* Even if it were argued that the mere shifting of topics away from something that had become unfruitful and awkward constitutes sufficient grounds for "triumph", it is necessary to posit goals of a higher-level than, for example, a discourse level "shift topic". If this were not the case, we would expect all topic shifts to have this sense, or, alternatively, we would have to consider topic-initiating or topic-shifting utterances with a "triumph" sense to be in free variation with those without such a sense.

5. Computational Mechanisms

5.1. Multiple Acts in Single Utterances

In contrast to robot planning, where a single goal is realized by a sequence of actions, in conversation a single utterance frequently effects multiple goals. This is because a single utterance is not a single act but a composite of many acts, each of which can realize separate goals. In this section, two illustrations of this are given. In the first, gestures, eye gaze, and body position are used for realizing the speaker's goals. In the second, we see that the various lexical choices that go into making a sentence provide loci at which diverse goals can operate.

The first example is (D6)

(D6) You don't have any children, obviously.

and its accompanying gestures. We can assume that Y has three goals while uttering (D6). She has to PROTECT herself from the negative evaluation. She wants to take the offensive and RETORT, and she wants to HOLD the floor for the continuation of the answer she has already begun. These three goals are realized in a variety of ways in a very complex sequence of gestures. (See Figure 6 below.)

At the beginning of utterance (D4), both participants are looking down at the envelopes. Halfway through (D4), Y looks up at X. At the 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 hunching her shoulders and laughing (PROTECT), and the segment that is analyzed here begins.

During utterances (D3) and (D4), Y's body is at a moderate angle, leaning slightly forward but not too far. She rocks a bit when she laughs, and as she begins utterance (D6), she leans forward slightly.

She remains at that angle until resuming her answer in (D7), at which point she reassumes her former position. It is as if her body position is bracketing the side sequence (D6), and her forward angle seems to accord with her aggressive retort (RETORT).

Her eye gaze also accords with the aggressive reply. The usual behavior for eye gaze was catalogued by Kendon (1967) and is quite apparent in our record of the Q-tips conversation. Typically, a speaker will look down for the first part of an utterance, as though planning it out. During the last part, a speaker will generally look up at the listener, as though monitoring its effect. Y, however, after looking down during her laughter, looks up at X simultaneous with the beginning of utterance (D6). It seems reasonable to attribute this marked behavior to the goal RETORT. Then just before the end of the utterance she looks down at the mail again. This could be to HOLD her turn.

While saying (D7)

(D7) You have to either draw or make things with the little
Q-tips,

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.

This sequence of gestures interacts with the interruption (D6) in a curious way. 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 during (D6): She first pushes her hair back with her left hand (PROTECT), but at the same time her right hand remains in position, index finger pointing down at a slight angle. Then she pushes her hair back with her right hand (PROTECT), while her left hand reassumes 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 (D7), which she has already planned and begun (HOLD).


Then we come to the false start, "You must-". This looks like a simple performance error until we notice the following: It took the left and right hands exactly the same amount of time to push the hair back. Y completed the utterance (D6) slightly before the right hand completed its half of the gesture. Thus the right hand was not yet free for the gestural accompaniment for utterance (D7). Yet Y did not want to risk losing the floor through a momentary silence. It seems reasonable to conjecture that the false start was generated by the goal HOLD to make the timing come out right.

Finally the immediate repair to Y's plan is accomplished, and Y is ready to continue with her interrupted answer. Realizing a goal we might call CONTINUE, Y moves her hands so that the index fingers are pointing directly at each other, returns the body to its middle position, and initiates utterance (D7).

The image that suggests itself is of the modalities -- eye gaze, hands, body position, and so on -- as separate conveyor belts passing a single station, and of the goals as agents at this station. The goals have various material at their disposal for actualizing themselves. That is, there are causal axioms of the form "A cause G", where G is a goal and A is an executable action in one of the modalities. A goal can load its material onto the belt when there is a match between what the goal can use and what is appropriate for a given modality, or conveyor belt, and when no stronger goal has already taken charge of the modality by filling it with its own material.

Figure 6 illustrates this process. The horizontal lines represent the different modalities. The numbers on the TIME line are seconds since the beginning of the segment. The four relevant goals are written across the top of the figure. An arrow from a goal to the beginning of an action on some modality indicates that that action was placed there by that goal. For instance, the arrow from the goal HOLD to the action "on mail" on the EYES track means that Y directed her gaze to the mail to HOLD onto her turn.

Figure 4



Utterance (D9)

(D9) So I have to send in this trash for her,

also effects multiple goals and suggests a similar mechanism. (D9) simultaneously answers a question, explains the motivation for an action, disavows the same action, and is humorous.

There seem to be at least three goals operating at this point. First, if Y is to cohere, she must complete her ANSWER of the question (D3) by conveying the information in (2). Moreover, the goal of defending against the negative evaluation remains, leading to two subgoals: She wants to show that her involvement results from some inexorable external circumstances (call this MOTIVATE), and to DISTANCE herself from the events by indicating that they are not a serious concern of hers.

Realizing all these goals in a single utterance suggests a variation on the above mechanism: View the utterance as a conveyor belt with slots for each of its elements. The goals compete to fill these slots with lexical or syntactic material that will aid their own realization. Filling a slot, as before, is a matter of finding a match between what the sentence requires and the resources a goal has available. (A more pedestrian description would speak of looping through the slots, and for each slot, looping through the goals, and so on.)

Figure 7 illustrates this process:

DECISIONS:	UNMARKED:	ACTUAL:	SOURCE:
Conjunction	--	so	MOTIVATE
Subject	I	I	ANSWER
Aspect	Prog	have to	MOTIVATE
Verb	send	send	ANSWER
Destination	in	in	ANSWER
Object	these	this trash	DISTANCE
Beneficiary	--	for her	MOTIVATE

Figure 7.

The goal ANSWER has determined the unmarked propositional content of the utterance at what Thompson (1977) has called the strategic level. At Thompson's tactical level, certain unmarked lexical choices are displaced with material provided by other goals. Thus, MOTIVATE supplies the conjunction "so" to indicate that Y's sending in the envelopes is due to her child's high expectations. MOTIVATE also replaces the present progressive tense with "have to"; Y's obligation excuses her for carrying around an armful of Q-tips. "For her" in the beneficiary slot indicates the circumstances behind the obligation.

Since those things one takes seriously, one necessarily values, one way to create DISTANCE is to evaluate the Q-tip designs negatively. Noun choice is a rich resource for such evaluations. The word "trash" means material with no value, and fits the bill perfectly.

An utterance needs to be viewed not as a single action, but as a bundle of actions, happening simultaneously or in quick succession. We need to determine the principle actions, or choice points, that go into the making of the typical sentence. The following seems a plausible version of a catalog of actions: For the "verbal" aspects of the utterance, the speaker chooses one or more propositions to convey. He chooses the properties to convey for referential purposes. He may decide to plow some of the information to be asserted into the presuppositional structure of the sentence. He may choose a "non-neutral" grammatical structure. He must choose lexical items for each of the slots in the message thus constructed. For the non-verbal aspects of the utterance, he must choose an intonation contour; this itself may involve several decisions. Finally, each of the modalities -- eye gaze, body position, gesture -- must be filled in some appropriate way, even if only with a neutral option.

Each of these aspects of an utterance are resources the speaker can utilize to realize his goals.

5.2. Bidirectional Planning for a Next Utterance

Notice something about the segment from (D4) to (D16). First the daughter makes something to put into the envelopes (D4, D7), then Y sends them in (D9), they are in the mail (D11), they arrive in Blair Nebraska (D14), and are put into a P.O. box (D15) by mailmen (D16). What we have is the most mundane story imaginable of mail going to its destination. But Y has infused humor into this framework at every point, transforming dull raw material into the stuff of a good conversation. This is very suggestive about how utterances get planned.

A bidirectional search for a plan works not only top down from the goal to the actions, but also bottom up from whatever moves are currently possible to the goal. In the blocks world, this would be a bad idea, for there are simply too few constraints on the next move. But in conversation, the tendency, once a schema is tapped, to follow the natural flow from one event to the next may constrain the possible next moves enough to make a bidirectional search feasible.

For example, we can imagine the initial stages of planning utterance (D11), "And of course all the Q-tips will fall off ... in the mail," going as follows: Top-down: Since the beginning, Y has had the goal of being humorous as a way of projecting a favorable image. Bottom-up: The next step in the mail scenario after the sending is that the envelopes are "in the mail" for a while. The design's fragility was a concern in packing and is called to mind because of Y's mock checking sequence, so associated with this step in the scenario is the possibility that the Q-tips might pop off. This is recognized as a mishap, and possible mishaps are a source of humor. Thus we have the link between the top-down and bottom-up searches for a plan, and the next step in the schema, as modified, will serve her purposes.

Once tapped into, the "mishaps" strategy gives a productive way of transforming further steps in the mundane scenario into conversational material. Bottom-up, we ask what happens next; top-down, we ask how this can be made interesting.*

* Hayes-Roth et al (1979) examine the problem of planning errands and show the need not just for bidirectional planning, but for what they call opportunistic planning. In addition to top-down and bottom-up components, there are also "middle-out" components.

6. Methodological Difficulties

The problem with the foregoing microanalysis is obvious. Plausible as it is, it is still speculative. How much represents processing the participants were doing and how much this author's invention? How can we minimize such speculation?

There are a number of kinds of evidence we can use to bolster our accounts of conversational planning, and several have been used in this instance. None is perfect. Among them:

(1) Interviewing the participants afterwards: One obtains mixed results with this. The participants are typically quite certain and probably reliable in identifying the referents of referential phrases (Mann et al 1975), quite unreliable in offering insights into the minor bumbblings and ineffective moves that microanalysis tends to turn up. Perhaps the most interesting phenomenon in interviews is what might be called the "Doppelganger effect". As Y sits and watches the videotape, X on the videotape tells a joke, Y on the videotape laughs, and Y watching the videotape laughs in exactly the same way, at exactly the same point. In other variations, Y expanded on topics in the interview that had been cut short in the original conversation. Her excursus on the perils of being a mailman (1) is an example of this.*

(2) Use oneself as a participant: Instead of interviewing, one simply introspects. Labov & Fanshel (1977, chapter 11) mention difficulties with this.

(3) Choosing situations in which there is an "authority" on the participants' mental states: In Labov & Fanshel's analysis of a therapy session, the therapist is the expert on the patient's motivations, while

* The argument that an interview merely produces more data, no different in kind from the original conversation, is not valid, because we take different perspectives on the two. In the original, we are concerned with gestures, false starts, hesitations, intonation, repairs, and so on. We are not in the least interested in the truth value of the utterances. In the interview, it is only the truth value that concerns us.

psychoanalytic theory provides the explanations of the therapist's behavior. In the task-oriented dialogs examined by Grosz (1977), there are well-defined goals and a highly structured knowledge base, and it is reasonable to assume the task model we would construct is the same as the one the participants have internalized; this at least gives us authoritative access to the participants' domain goals and plans. Concerning the mental states of those recorded in the Watergate transcripts, investigated by Linde & Goguen (1978), volumes have been written and can be used as corroborating evidence.

d. Using videotape: This is extremely difficult to transcribe, and the additional information is harder to pin down than the linguistic data. We know much less about the significance of gestures and eye gaze than about the meanings of sentences. The additional information frequently disambiguates and clarifies, however. It is a common experience for one not to be able to make any sense at all out of a transcript, and to have it make perfect sense when watching the videotape.

Two examples will illustrate this. In Section 4.3, X's utterance (D12) was analyzed as a follow-up question on the nature of the mail. Some people have argued, however, that he could have been feeding Y material for her "strange mishaps" topic. The plausibility of the latter fades when one views the tape and sees the coherence of gesture and body position between X's initial question and (D12), and the break between that and the grouping of gestures and body position that begins when he discovers and rejects Y's development of that topic.

The second example is from Turner (1976), cited by Wootton (1975). In the first moments of a conversation between a therapist and a new patient, the following exchange occurs:

T6: What do you do?

P6: I'm a nurse, but my husband won't let me work.

T7: How old are you?

Turner takes T7 to be a comment on P6, criticizing the patient for not taking responsibility herself. Wootton suggests the more mundane interpretation that T7 is intended to elicit essential background information. This ambiguity could well be resolved by intonation, gesture, and body position. For example, stress on "old" would favor the mundane interpretation, stress on "are" the comment interpretation.

(5) Looking for distributional regularities: This is possible for certain relatively simple phenomena, such as eye gaze (Kendon 1967). But for more abstract rules, such as "To distance yourself from something, evaluate it negatively," it is so difficult to recognize the goal and the action themselves that there seems little hope for large-scale studies of their cooccurrence, especially since many rules are specific to particular microcultures.

The best any of these methods can do is to eliminate some interpretations. They can never reveal the truth. From a theoretical point of view, however, we have an escape:

A theory of conversation would concern itself with utterances that are appropriate in particular contexts to particular conversational goals. Since this data is not exhaustively presented, the theory would have to make predictions to verify that it covers the data, so we need to be precise about what we can expect our theory to predict. We cannot expect predictions of the utterances, given only the context and the speaker's goals, any more than we can expect a theory of syntax to predict utterances, given only the speaker's intent to speak grammatically. A goal can be realized in many ways, and the mystery of human choice intervenes. The most we can hope for is to predict the set of possible utterances. But this set exists as data only in the form of its characteristic function,* the appropriateness judgments of a competent observer. It is this that the theory should predict, just as a theory of syntax predicts grammaticality judgments.

* The characteristic function of a set S is a function f such that $f(x)=1$ if x is in S, 0 otherwise.

The best observer is someone with the greatest possible access to the context of utterance and the speaker's conversational goals. But this is just the speaker herself. In studying real conversation, we may assume utterances to be appropriate in context unless there is strong evidence to the contrary. We assume inappropriateness only with the greatest reluctance. The fact that a competent speaker uttered a sentence and did not retract it is generally the best appropriateness judgment we have. This assumption gives us a very large collection at least of positive judgments.

Deciding to predict appropriateness judgments makes our job easier. To predict an utterance, we would have to show why a derivation of the utterance from the conversational goals was chosen over derivations of all other possible utterances. To predict the appropriateness judgment, we need only show that some derivation exists.

In brief, we will never be able to say what went on in the actual production of the conversation, only what could and couldn't have gone on. But in this, our situation is no different from the rest of cognitive science. The best we can do is to know all we can and to tell a story that contradicts nothing we know.

8. Do People Talk to Each Other?

It is by now a truism that that comprehension involves deducing the speaker's intention. But this notion has received insufficient analysis. It is coherent as it stands in a framework which views an utterance as having a single literal meaning and a single intention or speaker's meaning, which may or may not be the same. But in a framework that replaces a single intention with many goals, at many different levels, in a highly structured, ongoing, changing plan, it becomes problematic. There are two difficulties that arise immediately, one representing a sophistication in people that the standard view fails to capture, and one a lack of sophistication in people it fails to excuse.

The first problem is -- at what level must the listener understand the speaker's plan. Consider the extremes: It is certainly the case that a listener must discover that a speaker's goal in asking "What time is it?" is to find out the time. On the other hand, it is not necessary to discover that a person's goal in telling you a story is to make you feel positive toward him, and in fact, to respond too directly to this global goal with, for example, "I like you" would be an abrupt move. The speaker has a whole range of goals, some of which it is necessary to respond to and some of which it is inappropriate to respond to, and we need to develop a finer sense of which is which.

A plausible beginning of an answer is that it is appropriate to respond to goals you are intended to recognize, and the speaker will provide you with adequate signals to discern these (cf. Cohen 1978). Thus, when A asks B "Do you have a watch?" A intends for B to understand that A's goal is to learn the time, but A does not intend for B to understand that A's goal is to learn whether he will be late to a concert. But this is still a bit too simple. It is appropriate to reply to a lie with an accusation that it is a lie, even though the liar did not intend his goal of deceiving to be discerned.

The second problem is that very frequently in quite normal conversations, the participants are too involved in their own goals to address each other's goals. In fact, the more one investigates

conversation, the more it seems that people talk past each other, for precisely that reason. We will give three examples.*

The first is the Q-tips conversation. For X, the fragment of conversation we have investigated breaks into two episodes, evident from body position as well as from content. In the first, X is leaning forward trying to determine the nature of the mail. In the second, he leans back and rejects the topic of the dishonest mailmen, first by his silence, then by his overly literal questions. This structure does not mesh well with the structure of Y's side of the conversation, and the mismatch shows up in two examples cited above: What for him is a mere confirmation of the mail's destination (D12) is for her an example of strange mishaps, and when she tries to escape from the dishonest mailmen topic, he pulls her back in (D22). As the conversation continues, the mismatch continues. While speaking of the dissertation, Y wants to discuss its substance, X wants to know if he is cited.

The second example comes from a dialog, collected by Grosz (1977), conducted over terminals between an expert and an apprentice engaged in repairing an air compressor. In one stretch several minutes long, the apprentice thinks the conversation is about the trouble she is having loosening a bolt. But in fact the expert is trying to get her to stop using the pliers because it will strip the bolt. The apprentice wasn't aware of this discrepancy until she was interviewed some time later.

The final example comes from a dialog between a radio talk show host H and a woman W who calls in to tell about her worst blind date.** For the first half of the dialog, H is trying to turn everything W says into a joke, while W is trying to get on with her story. It is clear that H doesn't expect his typical caller to have a good story and feels he must entertain his audience at the caller's expense. Suddenly, when W says she stole her date's car, H takes interest, and makes her repeat herself twice. From then on, H is trying to extract more good material

* Gumperz (1979) also gives a striking example of this.

** We are indebted to Bill Mann for making this transcript available to us.

from W. But she has already told her story and now only elaborates it with mundane details.

Given our account of the mechanisms of interaction, none of this should be surprising. A participant in a conversation is viewed as a planning mechanism whose behavior is occasionally altered because of input produced by other such planning mechanisms. It is true that his plans may involve the goals, plans, and beliefs of the other participants, and in fact it may be one of his most urgent goals to aid the others toward their goals. But all of this is seen from inside the black box, and when the details of processing are focussed upon, it is such a close-up view that the other seems almost to disappear. His nature and his goals are imperfectly understood and become relevant only by becoming part of or interfering with the speaker's own plan.

This view is similar to the "toolmaker metaphor for communication" suggested by Reddy (1979) as an antidote to the standard "conduit metaphor". In the toolmaker metaphor, each participant is viewed as living in his own kind of world. The messages he gets from others are only the sparsest blueprints of objects designed for the sender's world. The receiver must reconstruct from the blueprints an object that will be useful in his own world. In the toolmaker metaphor, failure to communicate is not aberrant; it is the norm.

Nevertheless, communication is the ideal toward which conversation aims, and there are no doubt essential properties of conversation that arise out of the nature of that experience. The more elaborate view of a speaker's goals and plans presented here enables us to address in a more detailed way what it is to communicate. While fragments such as the one analyzed in this paper, coming from the beginning of a conversation in which topic and status are being negotiated, provide a good challenge for the Planning Approach, they are not sufficient for investigating the nature of communication. For this, we need to study examples of conversations in which communication in fact succeeds, rare though they be.

Acknowledgments

We have profitted from discussions with Armar Archbold, Joe Goguen, Barbara Grosz, Will Leben, Charlotte Linde, Bob Moore, Livia Polanyi, Ann Robinson, Jane Robinson, Earl Sacerdoti, Richard Waldinger, and the other members of the TINLUNCH discussion group and the Bay Area discussion group on discourse, and from the Zeitgeist at SRI. The work was supported in part by the National Science Foundation under Grants No. MCS-76-22004 and MCS-78-07121 and by the Defense Advanced Research Projects Agency under Contract No. N00039-79-C-0118 with the Naval Electronic Systems Command.

REFERENCES

- Allen, J. 1979. A plan-based approach to speech act recognition. Technical Report No. 131/79, Dept. of Computer Science, University of Toronto.
- Allen, J. & C. Perrault 1978. Participating in dialogues: Understanding via plan deduction. Proceedings, Second National Conference, Canadian Society for Computational Studies of Intelligence, Toronto.
- Archbold, A. 1976. A study of some argument-forms in a persuasion dialog. University of Southern California Information Sciences Institute. September 1976.
- Argyle, M. 1972. Non-verbal communication in human social interaction. In R. Hinde (Ed.), Non-Verbal Communication, 243-269. Cambridge, England: Cambridge University Press.
- Beaugrande, R. de. 1980. The pragmatics of discourse planning. To appear in Journal of Pragmatics, Vol. 3. 1980.
- Birdwhistell, R. 1970. Kinesics and Context. Philadelphia: University of Pennsylvania Press.
- Bolinger, D. 1972. Relative height. In D. Bolinger (Ed), Intonation. Middlesex: Penguins.
- Brenneis, D. and L. Lein 1977. "You Fruithead": A sociolinguistic approach to children's dispute settlement. In S. Ervin-Tripp and C. Mitchell-Kernan (Eds), Child Discourse, 49-65. New York: Academic Press.
- Bruce, B. and D. Newman. 1978. Interacting plans. Cognitive Science, Vol. 2, 195-233.
- Chafe, W. 1979. The flow of thought and the flow of language. In T. Givon (Ed.), Syntax and semantics, Vol. 12. New York: Academic Press.
- Clark, H. & C. Marshall 1978. Definite reference and mutual knowledge. Paper presented at the Sloan Workshop on Computational Aspects of Linguistic Structure and Discourse Setting. University of Pennsylvania, May 1978.

- Cochrane, J. and M. Zeleny (Eds) 1973. Multiple Criteria Decision Making. Columbia SC: University of South Carolina 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.
- Crystal, D. 1969. Prosodic Systems and Intonation in English. Cambridge, England: Cambridge University Press.
- Fikes, R. & N. Nilsson. 1971. STRIPS: A new approach to the application of theorem proving to problem solving. Artificial Intelligence, Vol. 2, 189-208.
- Fillmore, C. 1979. Innocence: A second idealization for linguistics. Proceedings, Fifth Annual Meeting, Berkeley Linguistics Society, 63-76. Berkeley, California.
- Genesereth, M. 1978. Automated consultation for complex computer systems. Ph. D. Thesis, Harvard University, Cambridge, Mass.
- Goffman, E. 1974. Frame analysis. New York: Harper.
- Grosz, B. 1977. The representation and use of focus in dialogue understanding. Stanford Research Institute Technical Note 151, Stanford Research Institute, Menlo Park, CA. July 1977.
- Grosz, B. 1979. Utterance and objective: Issues in natural language communication. SRI International Technical Note 188, SRI International, Menlo Park, CA. June 1979.
- Grosz, B. and G. Hendrix 1979. A computational perspective on indefinite reference. SRI International Technical Note 181, SRI International, Menlo Park, CA.
- Gumperz, J. 1979. The sociolinguistic basis of speech act theory. In J. Boyd and S. Ferrara (Eds), Speech Act Ten Years After. Milan: Versus.
- Halliday, M. 1977. Structure and function in language. In T. Givon (Ed.), Syntax and semantics, Vol. 12. New York: Academic Press.
- Hayes-Roth, B., F. Hayes-Roth, S. Rosenschein and S. Cammarata 1979. Modeling Planning as an incremental, opportunistic process. Proceedings, International Joint Conference on Artificial Intelligence, 375-383. Tokyo, Japan. August 1979.
- Hobbs, J. 1979. Why is discourse coherent? To appear in F. Neubauer (Ed.) Coherence in natural language texts.

- Hobbs, J. & J. Robinson 1978. Why Ask? SRI Technical Note 169, SRI International, Menlo Park, California. October 1978.
- Jefferson, G. 1972. Side sequences. In D.N. Sudnow (Ed.) Studies in social interaction, 294-338. New York: Free Press.
- Keeney, R. and H. Raiffa 1976. Decisions with Multiple Objectives: Preferences and Value Tradeoffs. New York: Wiley.
- Kendon, A. 1967. Some functions of gaze-direction in social interaction. Acta Psychologica, Vol. 26, 22-63.
- Kowalski, R. 1974. Logic for problem solving. Memo No 75, Imperial College, University of London. March 1974.
- Labov, W. and D. Fanshel. 1977. Therapeutic discourse. New York: Academic Press.
- Labov, W. and J. Waletzky 1967. Narrative analysis: Oral versions of personal experience. In J. Helms (Ed), Essays on the Verbal and Visual Arts, 12-44. Seattle: University of Washington Press.
- Levy, D. 1979. Communicative goals and strategies: Between discourse and syntax. In T. Givon (Ed.), Syntax and semantics, Vol. 12. New York: Academic Press.
- Linde, C. & J. Goguen 1978. Structure of planning discourse. Journal of Social and Biological Structures, Vol. 1, 219-251.
- Linde, C. & W. Labov 1975. Spatial networks as a site for the study of language and thought. Language, vol. 51, no. 4, 924-939.
- Mann, W., J. Moore and J. Levin 1977. A comprehension model for human dialogue. Proceedings, International Joint Conference on Artificial Intelligence, 77-87, Cambridge, Mass. August 1977.
- Mann, W., J. Moore, J. Levin and J. Carlisle 1975. Observation methods for human dialogue. University of Southern California Information Sciences Institute Research Report 75-33. June 1975.
- McDermott, D. and J. Doyle 1978. Non-monotonic logic I, A.I. Memo 486, Artificial Intelligence Laboratory, Massachusetts Institute of Technology. August 1978.
- McCarthy, J. 1977. Epistemological problems of artificial intelligence. Proceedings, International Joint Conference on Artificial Intelligence, 1038-1044, Cambridge, Mass. August 1977.
- Moore, J. 1978. Modelling communication as a problem solving activity. Unpublished manuscript.

- Newell, A., J. Shaw and H. Simon 1959. Report on a general problem-solving program. Proceedings of the International Conference on Information Processing. Paris: UNESCO House.
- Newell, A. and H. Simon 1972. Human Problem Solving. Englewood Cliffs, NJ: Prentice-Hall.
- Polanyi, L. 1978. The American story: Cultural constraints on the meaning and structure of stories in conversation. Ph. D. Thesis, Dept. of English Language and Literature, University of Michigan, Ann Arbor, Mich.
- Reddy, M. 1979. The conduit metaphor -- A case of frame conflict in our language about language. In A. Ortony (Ed), Metaphor and Thought, 284-324. Cambridge, England: Cambridge University Press.
- Reichman, R. 1978. Conversational coherency. TR-17-78. Aiken Computation Laboratory, Harvard University, Cambridge, Mass.
- Rieger, C. & M. Grinberg. 1977. The declarative representation and procedural simulation of causality in physical mechanisms. Proceedings, International Joint Conference on Artificial Intelligence, 250-256, Cambridge, Mass. 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.
- Robinson, J. 1978. Purposeful questions and pointed answers. Presented at 7th International Conference on Computational Linguistics, Bergen, Norway, August 1978.
- Sacerdoti, E. 1974. Planning in a hierarchy of abstraction spaces. Artificial Intelligence, vol. 5, no. 2, 115-135.
- Sacerdoti, E. 1977. A structure for plans and behavior. New York: Elsevier.
- Sacks, H. 1974. An analysis of the course of a joke's telling in conversation. In R. Bauman & J. Sherzer (Eds.), Explorations in the Ethnography of Speaking, 337-353. London: Cambridge University Press.
- Sacks, H., G. Jefferson and E. Schegloff 1974. A simplest systematics for the organization of turn-taking for conversation. Language, vol. 50, no. 2, 696-735.
- Sag, I. and M. Liberman 1975. The intonational disambiguation of indirect speech acts. Papers from the Eleventh Regional Meeting, Chicago Linguistic Society, 487-497.

- Schank, R. and R. Abelson 1977. Scripts, plans, goals, and understanding. Hillsdale N.J.: Laurence Erlbaum Associates.
- Schiffrin, D. 1977. Opening encounters. American Sociological Review, Vol. 42, No. 5, 679-691.
- Sussman, G. 1975, A computer model of skill acquisition. New York: Elsevier.
- Tate, A. 1975. Interacting goals and their use. Proceedings, International Joint Conference on Artificial Intelligence, 215-218, Tblisi, Georgia, USSR. September 1975.
- Thompson, H. 1977. Strategy and tactice: A model for language production. Papers from the Thirteenth Regional Meeting, Chicago Linguistic Society, 651-668.
- Turner, R. 1976. Utterance positioning as an interactional resource. In R. Wilson (Ed), Ethnomethodology, Labelling Theory, and Deviant Behaviour. London: Routledge & Kegan Paul.
- Waldinger, R. 1975. Achieving several goals simultaneously. SRI Artificial Intelligence Center Tech Note 107, Stanford Research Institute, Menlo Park, CA. July 1975.
- Weiner, S. and D. Goodenough. 1977. A move toward a psychology of conversation. In R. Freedle (Ed.), Discourse production and comprehension, 212-225. Norwood NJ: Ablex.
- Wilensky, R. 1978. Understanding goal-based stories. Yale University Dept. of Computer Science Research Report 140. September 1978.
- Wootton, A. 1975. Dilemmas of discourse. London: George Allen & Unwin Ltd.