

AD-A070 028

YALE UNIV NEW HAVEN CONN DEPT OF COMPUTER SCIENCE

F/6 5/7

THE CONCEPTUAL CONTENT OF CONVERSATION.(U)

MAY 79 R C SCHANK, W G LEHNERT

N00014-75-C-1111

UNCLASSIFIED

RR-160

NL

| OF |
AD
A070028



END
DATE
FILMED
7-79
DDC


MA070028

LEVEL II

Handwritten scribbles, possibly containing the numbers 6 and 5.



DDC
REPRODUCED
JUN 18 1979
RE C



DDC FILE COPY.

This document has been approved for public release and sale; its distribution is unlimited.

YALE UNIVERSITY
DEPARTMENT OF COMPUTER SCIENCE

79 06 15 003

DDC
RECEIVED
JUN 18 1979
C

The Conceptual Content of Conversation

by

Roger C. Schank and Wendy G. Lehnert

Research Report #160

1979

The research described here was done at the Yale Artificial Intelligence Project and is funded in part by the Advanced Research Projects Agency of the Department of Defense and monitored under the Office of Naval Research under contract N00014-75-C-1111.

79 00

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER #160	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) The Conceptual Content of Conversation		5. TYPE OF REPORT & PERIOD COVERED Technical Report
		6. PERFORMING ORG. REPORT NUMBER
7. AUTHOR(s) 10 Roger C./Schank Wendy G./Lehnert		8. CONTRACT OR GRANT NUMBER(s) 15 N00014-75-C-1111
9. PERFORMING ORGANIZATION NAME AND ADDRESS Yale University - Department of Computer Science 10 Hillhouse Avenue New Haven, Connecticut 06520		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS 12 32p.
11. CONTROLLING OFFICE NAME AND ADDRESS Advanced Research Projects Agency 1400 Wilson Boulevard Arlington, Virginia 22209		12. REPORT DATE 11 May 79
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office) Office of Naval Research Information Systems Program Arlington, Virginia 22217		13. NUMBER OF PAGES 26
		15. SECURITY CLASS. (of this report) Unclassified
16. DISTRIBUTION STATEMENT (of this Report) Distribution of this report is unlimited.		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE
9 Research rept.		
14 RR-160		
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)		
18. SUPPLEMENTARY NOTES		
19. KEY WORDS (Continue on reverse side if necessary and identify by block number) Natural Language Processing Models of Dialogue Knowledge Representation Inference Generation Discourse Analysis		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) X Conversations between humans constitute a challenging task orientation for theories of natural language processing. Problems of conceptual representation, predictive understanding techniques, and the recognition of conversational continuity are especially crucial for computational models of dialogue. Significant information exchanges often lie hidden beneath the surface communications of human conversation. This paper outlines a strategy for conversational analysis that utilizes multiple levels of		

information flow in order to characterize the conceptual content of a conversation. The role of predictive knowledge structures and contextual information is illustrated in an analysis of a sample conversation, along with an outline of conversational rules specific to particular levels of conversational analysis.

Accession For	
NTIS GRA&I	<input checked="" type="checkbox"/>
DDC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
By _____	
Distribution/	
Availability Codes	
Dist	Avail and/or special
A	

-- OFFICIAL DISTRIBUTION LIST --

Defense Documentation Center Cameron Station Alexandria, Virginia 22314	12 copies
Office of Naval Research Information Systems Program Code 437 Arlington, Virginia 22217	2 copies
Advanced Research Projects Agency Cybernetics Technology Office 1400 Wilson Boulevard Arlington, Virginia 22209	3 copies
Office of Naval Research Branch Office - Boston 495 Summer Street Boston, Massachusetts 02210	1 copy
Office of Naval Research Branch Office - Chicago 536 South Clark Street Chicago, Illinois 60615	1 copy
Office of Naval Research Branch Office - Pasadena 1030 East Green Street Pasadena, California 91106	1 copy
Mr. Steven Wong Administrative Contracting Officer New York Area Office 715 Broadway - 5th Floor New York, New York 10003	1 copy
Naval Research Laboratory Technical Information Division Code 2627 Washington, D.C. 20375	6 copies
Dr. A.L. Slafkosky Scientific Advisor Commandant of the Marine Corps Code RD-1 Washington, D.C. 20380	1 copy
Office of Naval Research Code 455 Arlington, Virginia 22217	1 copy

Office of Naval Research 1 copy
Code 458
Arlington, Virginia 22217

Naval Electronics Laboratory Center 1 copy
Advanced Software Technology Division
Code 5200
San Diego, California 92152

Mr. E.H. Gleissner 1 copy
Naval Ship Research and Development
Computation and Mathematics Department
Bethesda, Maryland 20084

Captain Grace M. Hopper 1 copy
NAICOM/MIS Planning Board
Office of the Chief of Naval Operations
Washington, D.C. 20350

Mr. Kin B. Thompson 1 copy
Technical Director
Information Systems Division
OP-91T
Office of the Chief of Naval Operations
Washington, D.C. 20350

Advanced Research Project Agency 1 copy
Information Processing Techniques
1400 Wilson Boulevard
Arlington, Virginia 22209

Professor Omar Wing 1 copy
Columbia University in the City of New York
Department of Electrical Engineering and
Computer Science
New York, New York 10027

Office of Naval Research 1 copy
Assistant Chief for Technology
Code 200
Arlington, Virginia 22217

Captain Richard L. Martin, USN 1 copy
Commanding Officer
USS Francis Marion (LPA-249)
FPO New York 09501

Major J.P. Pennell 1 copy
Headquarters, Marine Corp.
(Attn: Code CCA-40)
Washington, D.C. 20380

The Conceptual Content of Conversation

by

Roger C. Schank

and

Wendy G. Lehnert

1. Introduction

Conversational interaction between people appears to be the most sophisticated of all natural language processing tasks. People draw on an unrestricted range of knowledge structures and rules specific to conversation in order to interpret what someone really means in a conversational setting. The process of communication is extremely complex, and it is possible to make many different responses to a sentence without breaking the conversational continuity. Sometimes the very act of not responding along one particular conversational line is often a statement in itself. A large portion of the conversational process takes place beneath the surface of the actual conversation. In trying to spell out the rules that people use in conversation, it is necessary to account for all this "hidden" communication; explanations that are limited to the surface interaction will not provide a sufficient perspective.

This work was supported in part by the Advanced Research Projects Agency of the Department of Defense and monitored under the Office of Naval Research under contract N00014-75-C-1111.

There are many levels of information flow when two people engage in conversation. Perhaps the best vehicle for examining what goes on under the surface of a conversation is to look at a conversation between emotionally related people who share a great deal of background information. We have chosen, therefore, to examine a hypothetical argument between a husband and wife:

A1: Why were you out so late last night?

B1: I went bowling with the boys.

A2: I thought you hated bowling.

B2: It's ok when I have some company.

A3: Aren't I company?

B3: It's not the same.

A4: Sure, because you can't pick up women at home.

B4: I don't pick up women at the bowling alley.

A5: Well, who says you go to the bowling alley?

B5: If I told you that's where I was, that's where I was.

A6: Then how come you smelled of perfume last night?

B6: What perfume? That was smoke.

A7: It sure was a funny kind of smoke.

B7: Well maybe it was.

A8: You'll get arrested if you do that in a bowling alley.

B8: We didn't do it in a bowling alley.

A9: Then where were you last night?

B9: All right. I was at Joe's house. We had a few beers and smoked some dope. I didn't want to tell you because I know you can't stand Joe.

A10: Liar! And hanging around with that creep. I want a divorce.

Having seen the entire exchange, it is possible to produce a paraphrase of the conversation which highlights its emotional and power-oriented aspects:

A1: Listen, this is important to me.

B1: ok.

A2: You're lying to me.
B2: No I'm not.

A3: I don't think you love me.
A3: That has nothing to do with it.

A4: You're cheating on me.
B4: I am not.

A5: I don't believe you.
B5: Enough of this!

A6: I have evidence.
B6: No you don't.

A7: I think I do.
B7: Think again.

A8: You're being stupid.
B8: No way.

A9: Aha! You trapped yourself.
B9: I give up.

A10: That does it. I've got you now.

While this paraphrase captures the emotional aspect of what occurred in the original conversation, it leaves out the substantive content. It is not a natural conversation which we would expect two people to have. What is he lying about? What is her evidence? In short, we have no idea what they're talking about beyond the fact that she is looking for an assurance of love or an admission of some sort.

The exchange on this level does have special continuity properties, however. Each of his statements seems to follow from hers in an acceptable manner. He seems to be responding to her. But many of her statements appear to come out of the blue. How does it make sense for her to say "You're lying to me," after he says "ok"? Why does she say "Aha! You've trapped yourself!" after he says "No way"? There is no continuity in these transitions. Our abstraction has

deleted some important aspects of the exchange which are necessary for its comprehension. At the same time, these discontinuities represent a strategic imbalance in the conversation. People on the defensive are usually in the position of responding to things; people on the offensive normally initiate things. In fact, if we consider the exchange in terms of offensive and defensive moves, it does appear that she is primarily an aggressive offensive player while he is largely defensive. It therefore makes sense that his statements address hers while hers often seem to ignore his.

This paraphrased version of the original conversation focuses on three levels of conversational interaction: dominance, trust, and the interpersonal relationship of the participants. In our analysis of conceptual content we have isolated at least twelve parallel levels of communication, each of which characterizes a particular type of information exchange. In order to fully understand or participate in a conversation, it is necessary to "track" the conversation on each of these twelve levels. However, conversational continuity can be preserved by maintaining continuity on any one level. It follows that the "thread" of a conversation can be decomposed into as many as twelve "strands" (or dimensions), all simultaneously present and active. This being the case, it is easy to see how people often "miss the point," or misinterpret what they're told. A mistake on any one of these levels will result in a misunderstanding of some sort.

We can analyze any given conversation in terms of [1] the implicit information that is present underneath the surface of the conversation, or [2] the rules of conversation that are operating at

each of our twelve levels of interaction. We will discuss the above marital altercation in terms of these two viewpoints, but our presentation is not intended to suggest that any real division exists; in any process model for conversation these two aspects must be fully integrated.

2. Contextual Knowledge

When an anxious wife asks her husband about his whereabouts, she is ultimately concerned with something more than a literal account of his locational meanderings. She is really interested in something else. Without explicitly saying so, she wants to know if he is seeing other women, or bar hopping, or committing some other taboo that she considers to be a violation of his personal commitment to her. In reality, people are very idiosyncratic about their taboos and so the business of inference at this level is fairly risky. Some women are threatened by affairs, some are primarily concerned with compulsive gambling, while others may find certain recreational activities to be the most intolerable offense. The better we know the speaker, the more certainty we can place on our inferences because we have knowledge about that individual's personal situation, beliefs, and role-related expectations. Without personal knowledge of the speaker, we must fall back on stereotypes. A stereotype may or may not provide us with accurate assumptions for any given person; inferences based on stereotypic expectations can always be wrong. So while one can always argue about the validity of stereotypes, we are concerned here only with their function: stereotypes provide people with a set of assumptions that can be used when real knowledge of a particular

individual is lacking. Professors are smart. Cats are afraid of dogs. Poets are sensitive. Most stereotypes have evolved because they tend to hold true. Others have suspect origins. For example, a feminist might argue that negative stereotypes about women are a political device of the male establishment, designed to keep women in a position that is advantageous to men. But whatever their source or validity, all stereotypes are used the same way: to provide tentative assumptions in the absence of first-hand knowledge.

With this in mind, we will now consider the various types of knowledge that must be accessed in order to understand our husband/wife conversation. The first question,

A1: Why were you out so late last night?

•••
immediately requires a number of inferences in order to be properly interpreted. Because we know a woman is directing this question to her husband, we understand that she had an expectation about his arrival time which was violated. "So late" can only be interpreted in terms of her expectations. While we do not need to know exactly when he came home, or exactly when she thinks he should have come home, we accept her query as valid because we have thematic knowledge about the behavior of married couples. Married people expect to be with each other at night, and any deviation from this normative pattern deserves explanation. The same question would make far less sense in the context of a night clerk at a hotel confronting a registered guest.

Scriptal knowledge is referenced in his response:

B1: I went bowling with the boys.

A script-based inference at this point will turn out to be crucial later in the conversation: if he were bowling then he must have been at a bowling alley. When he subsequently claims that he wasn't (smoking dope) in a bowling alley, she demands to know where else he was. This slip-up on his part would not have signalled an inconsistency if script-based inferences about bowling had not been made.

Knowledge of goals and plans is needed to understand that

A4: Sure, because you can't pick up women at home.

is in fact an accusation. Had she countered with, "Sure, because you can't burn envelopes at home," it would be difficult to understand her rejoinder. (Why would anyone want to burn envelopes?) Picking up women on the other hand, is easily understood as a plan designed to satisfy the cyclic satisfaction goal of S-sex [Schank and Abelson 1977]. But the fact that picking up women makes sense is not sufficient for this to be an accusation. It is an accusation only because we know that the marriage role theme assumes a goal subsumption strategy for S-sex [Wilensky 1978] which furthermore precludes all other behavior directed towards the satisfaction of that goal. If he admits to this planful behavior, he is admitting more than a simple "Yes, I pick up women," or even "Yes, I pick up women for the purpose of sex." An admission of such behavior is more significantly a rejection of his marital agreement. With such a major

role theme at stake, we are not surprised to hear his denial. If we were in the context of a conversation between two men in which one accused the another of picking up women at a bowling alley, we would not expect such stubborn denials. On the contrary, thematic knowledge about macho behavior would turn any such accusation into a compliment. A partial overview of the thematic information needed to understand this conversation is presented below.

LOVE THEME
(X loves Y)

- 1 if X has free time then X will want to spend it with Y
- 2 if X goes to some social event, Y will ordinarily go along.
- 3 if X is out alone then X does not get involved with members of the opposite sex

HUSBAND ROLE THEME
(X husband of Y)

- 1 X goes to work
- 2 X comes home
- 3 X spends evening with Y
choices:
A) X&Y spend evening at home
B) X&Y go to MOVIES, RESTAURANT
BOWLING, DANCING etc
- 4 if X goes out at night X tells Y where and comes home on time
- 5 X tells the truth to Y

Also relevant here is the MARRIAGE theme which encompasses the HUSBAND theme as a subpart. In addition certain societal social rules are necessary. These are given below:

THE MARRIAGE THEME

X MARRIED TO Y

- 1 X loves Y
- 2 X marries Y
- 3 then either:
X&Y are unhappy
or X&Y are happy

SOME SOCIAL RULES

- 1 You may only love one person at a time
- 2 If you avoid doing something it's because you dislike it
- 3 People avoid people they dislike
- 4 People do what they like to do

IF X&Y are unhappy, then: |

X&Y fight |
 X may go out alone |
 X may find other woman |

this causes:

Y becomes unhappy |
 X&Y fight more |

The LOVE, HUSBAND, and MARRIAGE themes are accessed as soon as we understand that this is a conversation between a husband and wife. These themes effectively describe the context of the conversation and help us make appropriate inferences about the conversation. This contextual information is equally crucial for the first-hand participants of a conversation.

Anyone participating in a conversation must be thoroughly aware of the hidden implications lurking behind every statement. The import of a statement is often a function of the inferences we make on the basis of contextual information. Whenever a knowledge structures can be safely assumed, remarkably little surface communication is devoted to critical information which can be inferred from this common knowledge. People are not encouraged to state "the obvious." But there are many ways that implicit information can be obvious, and we must consider different levels of conversational communication in order to understand exactly what constitutes a valid conversational response.

3. Levels of Conversation

Our analysis of conversational content is motivated by the realm of reasonable responses people can make during a given conversation. Conversations can "turn" in a number of distinct directions, but some

transitions are far more natural than others. By characterizing reasonable responses in terms of conversational levels, we can arrive at a decomposition for the surface communication which captures the conceptual content of a conversational interaction.

TWELVE LEVELS OF CONVERSATION

1. Direct Q-A: This is the explicit surface level of communication where "literal" content is processed.
2. Knowledge State: Statements here fill in the gaps in the knowledge of the hearer or attempt to subtly elicit what the hearer really knows about a situation.
3. Dominance Games: At this level a speaker tries to get the upper hand in an adversary type conversation. Typical strategies here include putting the other speaker on the defensive, or blunting an attack with an accusation.
4. Emotions of A: Many statements can be interpreted as expressions of how the speaker is feeling at the moment. Emotional states are often explicitly described at the surface level of a conversation.
5. Emotions of B: This is the same as the one above it, except that it refers to the emotions of the other speaker.
6. Relationship of A and B: This is the level at which implicit or explicit statements are made about the relationship between the participants. This goes on continuously all the time (mostly implicitly) even in formal discussions.
7. Argument strategy: This level describes strategic exchanges when the conversation involves an argument about some disagreement.
8. Import: The true significance of a communication is often present only by inference. This level of analysis tracks the most significant aspect of conversational interactions.
9. Implicit Beliefs: Beliefs used as presuppositions for a statement are constantly being accepted (by not discussing them) or rejected (by making explicit corrections).
10. Points: The ultimate goal of a conversation can often be about something other than what it seems to be about. (Discussions between unattached men and women frequently have this flavor.) This level keeps track of any underlying goals and is responsible for interpreting information in terms of that goal orientation.

11. Topic Shifts: Statements such as 'Don't change the topic' are surface realizations of this level. Topics themselves are tracked and can be discussed.

12. Truth and trust: The truth, believability, or trust of the other participant are constantly monitored and sometimes surface.

The twelve levels of conversation each carries a set of rules which control responses at that level. We shall now consider what the rules are like for a few of the levels. In order to do so, will examine our example sentence A1: "Why were you out so late last night?" and we will consider the conceptual content of this question at various levels of conversational analysis. To begin we will look at the simplest level, level 1:

CATEGORY: level 1: Direct Q-A

INPUT ANALYSIS: why question.

HOW WE KNOW: rules for analyzing English are used.

RESPONSE REQUIREMENTS: Requires fact

WHERE FROM: the needed fact will be found in episodic event memory.

HOW WE KNOW WHAT TO SELECT: an assessment of A's knowledge state establishes what fact is missing.

HOW WE KNOW TO DO THAT: rules about answering why questions provide the algorithm.

In order to answer A1 at level 1, it must be analyzed as being a "Why Question". This is sufficient for beginning to create a response at level 1. This information is obtainable from the input English itself by processes specific to question answering [Lehnert 1978].

Once this analysis is completed, the response requirements must be obtained. In this case we need a fact. Its availability must be obtained; here the fact should be in episodic memory. A choice algorithm must be given as well where choices are based on an assessment of the knowledge state of A.

To see better how this works we will examine some of the other levels in the same way. Level 3, dominance games, is responsible for tracking the power relationships in the conversation itself. That is, a conversation is sometimes a battlefield on which subtle tactics are being used. Conversations may address this level of conversation both implicitly and explicitly.

CATEGORY: level 3: Dominance Games

INPUT ANALYSIS: "You are going to be on the defensive now!"

HOW WE KNOW: 'Why questions' are candidates for level 3 analysis. If the reason being asked for addresses a role theme violation (or other potential inadequacy on the part of the hearer), then the above analysis applies.

WHY WE LOOKED THERE: A sufficient condition for a level 3 analysis is an inadequacy or failure to conform to expectations.

RESPONSE REQUIREMENTS: Level 3 responses are almost always under the surface. When one is put on the defensive, one has the choice of either returning the fire or accepting the defensive. The decision is based on the emotional relationship and power relationship of the participants. Once the choice has been made it affects the form and content of the response.

HOW WE KNOW WHAT TO SELECT: A consistency check on any presuppositions is needed to determine whether the attack is valid. Thus, we need to check facts as they exist in memory.

Our sample conversation can be viewed in many possible ways, one of which is as a kind of battle. Certain sentences overtly address the power gaming issue but nearly every sentence in an argument addresses this issue in some way. In this conversation A is attacking B. B is on the defensive initially, but B fires back in B2 by saying "It's (bowling) okay when I have company". Now this statement says many things, but on the level of dominance games, it is a covert attack on A's value as company. A responds explicitly to the attack but not to the dominance game. An explicit response to the dominance

game might have been, "Don't attack me just when I've got you!".

A retakes the offensive in A4 by averring that B is picking up women at the bowling alley, which is presumably an agreed upon bad thing. This puts B back on the defensive again in B4. This kind of jockeying for position in an argument is quite common. It is necessary, then, for a participant in such a conversation to be aware of what strategies are being pursued by his partner and how to respond to them. Understanding what argument strategies are being employed by a conversational partner and using such information in the creation of a response is an integral part of the conversational process.

Another conversational level is level 4, which can be construed to be a conversation about the emotional state of one of the participants in the conversation:

CATEGORY: level 4: emotions of A
INPUT ANALYSIS: "You don't love me anymore"
HOW WE KNOW: Violations of themes are always checked. When some are violated they imply that their preconditions are no longer valid. Here, HUSBAND rule 3 applies and we can infer the above analysis from its violation.
RESPONSE REQUIREMENTS: Assure, sympathize, ignore, etc.
WHERE FROM: What to do is determined by the hearer's emotional responses.
HOW WE KNOW WHAT TO SELECT: It is necessary to check the truth of implicit assertions such as this. The appropriateness of an explicit response on this level is dependent on the current emotional relationship and the desires of the speaker in this argument.

Here again, although the conversation may never explicitly be about how A is feeling, it may be so implicitly. In fact, in this conversation, as in many conversations between emotionally related people, the conversation might be construed to be almost exclusively

about that issue. This is of importance to the people in the conversation because they may be expected to respond on level 4 at any given time.

In our sample conversation A3 ("Aren't I company?") can be considered to be a statement at level 4 meaning 'You don't love me, do you?' B's response to that, ("It's not the same"), can be understood as explicitly ignoring the issue, which in a sense, is a qualified "yes you're right." This analysis can be seen as part of the reason for the escalation of matters by A in A4. At this point she goes directly on the attack, stating her suspicions overtly. This is consistent with B's previous refusal to discuss the level 4 aspect of the conversation. That is, B has rejected A by not reassuring her here. To see how level 4 can actually appear on the surface, we need only change response B3 to "I still love you." Such a response would make no sense at any other level of conversational analysis or with any other rules of continuity. Yet it is nevertheless an appropriate response here because level 4 has been implicitly brought up by A and can therefore be explicitly addressed by B if he so chooses. This can only be done by tracking this level continuously throughout the conversation.

People can also talk about their relationship in a conversation. Here again, this conversational level is usually implicit but it is also frequently brought to the surface. The rules here are as follows:

CATEGORY: level 6: relationship of A and B
INPUT ANALYSIS: "Our relationship is in trouble"
HOW WE KNOW: Rule 1 of the marriage theme has been violated.

WHY WE LOOKED THERE: In talking to someone, we evaluate the implications of what is said with respect to our relationship with that person. The input sentence violates one of the rules for marriage. A violation of an interpersonal thematic rule implies that a relationship may be shifting.

RESPONSE REQUIREMENTS: Confirm, deny, or ignore.

WHERE FROM: To respond here, we must check to see how we feel about the relationship and decide if it is worth pursuing that line of discussion directly. Indirect reassurance can be made by arguing that a theme violation has not actually occurred. This is what B does here. A's analysis of B's response should perceive a level 6 denial: ("No, our relationship is not in trouble, you do not know the facts").

The relationship between the individuals in a conversation is liable to become an explicit topic of conversation in exactly the same way that the emotions of A were. For example, in A10 we see an explicit statement about the relationship of A and B ("I want a divorce"). In a sense this entire conversation is about the relationship of A and B. We can imagine that the rules for deciding what to say for both A and B used information about what they perceived the other to be saying about their relationship throughout the conversation.

In addition to rules about A and B themselves, an important part of the implicit conversation taking place are the thematic rules (discussed in section 2) themselves. When these rules are invoked by the speaker, the hearer either (1) implicitly acknowledges their validity by responding on other levels or (2) contradicts them by an explicit rejection.

CATEGORY: level 9: beliefs

INPUT ANALYSIS: "Husbands are supposed to be home with their wives at night"

"Husbands should tell wives what they do"

HOW WE KNOW: Checking the husband role theme we find a violation of rules 3 and 4. Statements that implicitly refer to a role theme violation covertly assert one's belief in that role theme.

WHY WE LOOKED THERE: Beliefs about husbands' behavior are relevant when a wife is talking to a husband. The husband role theme thus constitutes a background against which inputs are checked.

RESPONSE REQUIREMENTS: To ignore this level is to accept covert beliefs. A belief must be countered if it is not shared.

WHERE FROM: To do this, one must consult one's own beliefs.

HOW WE KNOW WHAT TO SELECT: Contradictory beliefs are selected to be explicitly output if they exist. Thus if B does not share A's belief he can say something like "Being married doesn't mean you are in jail you know."

At this level of analysis, many implicit beliefs are being discussed in the conversation. A belief-oriented interpretation of the beginning of this conversation would look like the following:

A1: Husbands should tell wives what they do.
B1: That's true.

A2: People don't do things they don't enjoy.
B2: True, but good company helps make dull things interesting.

A3: Yes, but a wife should be a good companion.
B3: Friends are more interesting companions than wives.

This aspect of conversation focuses on the implicit agreements or explicit disagreements concerning knowledge-based presuppositions inherent in a statement. Another level of analysis concentrates on underlying accusations or assertions about credibility. The level of trust between two people is a function of credibility, and implicit acceptances or rejections of information can reveal whether a speaker is trusting or suspicious. In this conversation trust is a key (but implicit) point for A1:

CATEGORY: level 12: truth and trust

INPUT ANALYSIS: "Why don't you tell me where you go? It's causing me to not trust you."

HOW WE KNOW: After finding the violations of rules 3 and 4 in the husband theme, trust is inferred to be low. B can respond to the relationship problem or belief problem (above) or to the issue of trust.

RESPONSE REQUIREMENTS: Giving required information (if it is safe) helps foster trust. The trust issue can also be confronted head on. Such a confrontation will have an affect on the dominance game. Thus "Why don't you trust me?" is legitimate here and also puts A on the defensive.

A trust-oriented interpretation of the beginning of this conversation would procede as follows:

- A1: Why don't you tell me where you go?
It's causing me to not trust you.
B1: I'll tell you. There is no reason not to trust me.
A2: No reason? You are lying right now.
B2: No I'm not.

Using these rules that we have outlined, we can now examine some of the other sentences in this conversation with an eye towards their conceptual content on various levels of conversational analysis:

B1: I went bowling with the boys.

B1 responds to A1 primarily at the direct Q-A level. By responding to the question in a straightforward manner, B1 also signals an acceptance of role theme assumptions underlying A1 on the level of beliefs. The other five questions proposed by A1 are not addressed by this response.

A2: I thought you hated bowling.

Here are some possible interpretations of A2 using the thematic and social rules mentioned above, together with the conversational levels they address (L=love, S=social, M=marriage, H=husband):

- 1: You wouldn't go bowling if you loved me. L2,S2,S3: levels 4,6,9
2: You must be trying to avoid me. S3: levels 4,6,9
3: I still don't understand why you won't spend time with me. H3,L2: levels 4,6,9

4: People don't do things they
hate- I don't believe you. S4: levels 9,12

5: What were you really doing? H5: levels 3,12

The actual response of B to A2 addresses some of these levels
explicitly (E) and some implicitly (I):

B2: It's ok when I have some company.

1: I went with other people. 2E (updates knowledge state)

2: I didn't go with you you'll notice. L2,H3
7I (takes offensive here)
9E (modifies her belief about his belief)
4I (implicitly says that her emotions do not
matter)
6I (yes the relationship is in trouble)

3: I had fun without you.
9E (wives are not the only company for husbands)
6I (you do not satisfy my companionship needs)

A's response ("Aren't I company?") addresses levels 4,6 and 9. A3 can
be seen as a response to the implicit statements made above. The fact
that multiple interpretations seem possible indicates that a facility
for subjective understanding [Carbonell 1979] must be at work. As
third party observers, we are far more conscious of the subjective
factors involved than we would be if we were first-party participants.

4. Conversational Topics

The notion of "topic" has been conspicuously absent from our
discussion thus far. But in fact, the strategy of multi-leveled
analysis carries some strong consequences for the notion of topics
[Schank 1977]. To see how, let us return to our conversation. The
first sentence is:

(A1) Why were you out so late last night?

If we ask ourselves what the topic of this sentence is, we find that

we have a difficult time defining it. Is the topic here B's lateness, or his whereabouts last night, or the reasons for both of those or what? A simpler question to ask is what the expected response of A1 is. That is, what does A expect to hear back? This is the major determinant in B's response, so from a process model point of view it is a much more germane question than that of determining the topic of A1. In a sense then, since A1 is a very specific question, its topic is exactly equal to its expected response. As we have said, the key element in creating a response is the attempt to ascertain what the speaker had in mind for a response. When the speaker indicates directly what response he had in mind, then the rules for responding can work from there. For the moment, we will say that the topic of a sentence is an expected response.

Topic of A1: = REASON for C1 where C1 is the event

C1: B IS OUT LATE LAST NIGHT

The problem here is with the idea of REASON. Clearly there are a great many reasons for a given event. The range of possible responses to A1 depends on the particular concept or concepts in C1 for which REASONS are being requested. In fact, appropriate responses depend crucially on the interpretation of the underlying meaning of A1 as determined by the hearer.

For example, any of these seven possible responses to A1 would be all right in the appropriate context:

1. I was out bowling.
2. Listen, I was home early the last seven nights in a row!
3. Come on, I still love you.

4. I just didn't have a chance to call you.
5. It is not necessary to spend every night with your wife.
6. I didn't know that you cared to go with me.
7. It was for business reasons.

If, as we have assumed, these seven sentences can be appropriate responses in the right context, then it is incumbent on a theory of dialogue to explain this phenomenon. Notice that the notion of topic as expected response will not work here. The expected response may well be the reasons for C1 (B's lateness), but this will not explain why response 3 could ever be uttered and taken as relevant to the conversation. The reason that it is relevant has to do with the context of the conversation. To make this more explicit, we will define context for our purposes as the extant themes in any situation. (A theme is defined in Schank and Abelson, 1977 as a kind of goal generator for interpersonal relationships, societal relationships, and certain character traits.)

The themes that are relevant here are the HUSBAND role theme and the LOVE theme, both of which help to establish the prior context of the situation. Basically the HUSBAND role theme which is used in A1 implicitly deals with the expectations that a wife can reasonably make on a husband in our society. The LOVE theme deals with what lovers expect from each other. The relevant pieces of these themes for the purposes of this conversation were outlined above.

The use of this thematic information, plus some standard rules of inference, can help to explain why the above possible responses can actually make sense. They make sense because the initial question (A1) has been reconstructed by the hearer (B) into a new question or statement that he has decided to respond to. The role of question

reconstruction has been investigated on a lower level of conceptual analysis [Lehnert 1978], but the reconstructions we are proposing here involve a far more radical transformation of meaning. We will now attempt to outline this process for each of these seven instances.

The original question was:

A1: Why were you out so late last night?

The first response above is actually a response to a reconstructed question of the form:

1: Where were you last night?

This question is obtained by straightforward inference procedures of the following sort: A1 is examined and answered internally. In order to express an answer of the sort "Because the bowling lasted a long time" it is necessary for B to assess the knowledge of A to find if she knew that "bowling" was where he had been. When the answer to this is discovered to be negative, the question is internally reconstructed to be "where were you last night" using the conversational rule that all logically prior questions must be answered first.

2: Why weren't you with me last night?

This question is obtained by first doing the reconstruction necessary for (1) and then checking the HUSBAND role theme to see if there is any relevant rule there that might be the source of the reconstructed question. Rule 3 from the husband role theme is found (Husbands are

supposed to spend the evening with their wives). This rule transforms A1 into (2).

3: You don't love me anymore.

Two rules that match the reconstructed (1) are also found in the LOVE theme. LOVE rules 1 and 2 both speak to the issue of what someone who loves someone is supposed to do in the evening. Since C1 violates this rule the underlying LOVE theme itself is open to question. (Notice that this does not happen with role themes such as husband. We would not expect that his actual status as husband can be questioned. But LOVE is an interpersonal theme, and since interpersonal themes can change more easily, they can be called into question, thus producing (3)).

4: Why don't you tell me what you do?

Rule 4 from the HUSBAND theme is used here. This works the same way that rule 3 was used for (2) above.

5: Our relationship is in trouble.

The MARRIAGE theme has a kind of script for good and bad marriages. The "unhappy" path in the MARRIAGE theme outlines some characteristics of a bad marriage. When one of these is matched (B going out alone) it can be the source of a question. Thus the reconstruction process comes up with (5).

6: Don't you like spending time with me?

Sometimes themes can combine to produce a question. To reconstruct

(6) both rule 1 from the LOVE theme and rule 3 from the HUSBAND theme are used.

7: Husbands are supposed to be home with their wives at night.

There are many ways to actually reconstruct a question using the same rules. (7) is an example of HUSBAND role theme rule 3 applied slightly differently. The question has been reconstructed into a statement that is a reiteration of the thematic rule itself. In this case an answer can be a discussion of the validity of the rule or of conditions on the rule that allow for exceptions. In a sense then, what we would have here is a discussion of the rule itself. A says that husbands should be home, and B says 'well, not every night'.

What we can see from all this is that a conversational topic can span a number of conceptual information types. There is a sense in which a question is being responded to on one level, while it is being explicitly ignored on another. Another way of saying this is that there are a variety of different conversations going on at once. Given this perspective, conversational topics must exist at many levels concurrently. A decomposition of topics into multiple levels of analysis would account for the great flexibility conversations display in terms of what are usually thought to be topic shifts. If a topic is identified in terms of conversational levels, the problem of analyzing many transitions which appear to be topic shifts, will reduce to a problem of choosing one conceptual level over another when formulating responses.

5. Conclusions

A multi-faceted approach to conversational analysis is necessary in order to build a program that is capable of sensible conversational response. Since people do not restrict themselves to a single level of conceptual information, it follows that multi-level analysis will be crucial for conversational computer programs as well. The twelve levels of conversation that we have presented may not be the most appropriate decomposition strategy; we will be able to determine a correct set of them only by attempting to actually construct a program that analyzes sentences and comes up with appropriate responses. Clearly such a program will need a data base of world knowledge, episodic knowledge, goals and emotional relationships in order to converse effectively. We are currently working on specifying exactly how it will all work.

By way of conclusion, our final example is intended to illustrate the complexity of this problem by examining A10. At this point in the conversation, A could explicitly respond to B on nearly any conversational level. The last statement of B was:

B9: All right. I was at Joe's house. We had a few beers and smoked some dope. I didn't want to tell you because I know you can't stand Joe.

This can be responded to on each level as follows:

1. Direct Q-A: How can you stand Joe?
2. Knowledge State: Why didn't you tell me that in the first place?
I really like Joe, I just never told you about my change of heart about him.
3. Dominance Games: Well I've got news for you. I've been seeing Joe when you actually do go bowling.
4. Emotions of A: I can't love a person who lies to me like that.
5. Emotions of B: You couldn't love me and be friends with him.
6. Relationship of A and B: I want a divorce.
7. Argument strategy: You mean you created this giant argument out of

just that little thing?

8. Import: That's what you were worried about? I don't care about that.
9. Implicit Beliefs: You were smoking dope? That's immoral. I won't stand for it.
10. Points: That's my point exactly. You are a liar.
11. Topic Shifts: (The topic has not been shifted so this is also inappropriate.)
12. Truth and trust: How can I trust you after you lied to me like that?

In addition to what we have listed above, it is also possible to combine many of the levels in one response. (Indeed it was quite difficult to avoid doing that in what was written above.)

To effectively model conversation then, it is necessary to find the levels at which we people operate, the rules people use to relate inputs to these levels, the methods of response generation appropriate at each of these levels, and the rules people use to select from alternative responses. When we have understood how to do all this, we will have made a start at analyzing the process of conversation and will then be ready to attempt the construction of an automated conversationalist.

References

Carbonell, J. G. (1979). Subjective understanding: Computer models of belief systems. Research Report #150. (PhD thesis). Department of Computer Science, Yale University, New Haven, Conn.

Lehnert, W. G. [1978]. The Process of Question Answering. Lawrence Erlbaum Associates, Hillsdale, New Jersey.

Schank, R. C. [1977]. Rules and topics in conversation. Cognitive Science. Vol. 2, No. 1.

Schank, R.C. and Abelson, R.P. [1977]. Scripts, Plans, Goals and Understanding: An Inquiry Into Human Knowledge Structures, Lawrence Erlbaum Associates, Hillsdale, New Jersey.

Wilensky, R. [1978]. Understanding goal-based stories. (PhD thesis) Technical Report #140. Department of Computer Science, Yale University, New Haven, Conn.