

Charles L. Hedrick
 Coordinated Science Laboratory
 University of Illinois at Urbana-Champaign
 Urbana, Illinois 61801

This paper discusses the need in a natural language understanding system for a model of the speaker and of the conversation process itself. Most current programs use models of the domain of discourse to supply the knowledge necessary to understand what is being talked about. (See, e.g., Schank[1973] or Charmak[1972 J.]) For example, Schank's system might contain a restaurant "script", listing all of the actions one normally does at a restaurant, from entering and being seated to paying the check and leaving. Such systems appear to be based on the idea that natural language texts are mostly assertions about the domain of discourse, and can be mapped into additions to a data base involving this domain. However, examination of real English texts suggests that much of it has other purposes.

Consider the following example, taken more or less at random from the 27 May 1977 issue of the Christian Science Monitor: "As the first heat waves roll in and Good Humor trucks jingle for customers, a most sticky question is being thrashed out in the back halls of government: What is ice cream?" As a first approximation, we may say that this is mostly an exercise in wit. One pales at the thought of an A.I. program churning away on it, bringing out the association between ice cream sales and hot weather, and retrieving the various facts needed to understand Good Humor's marketing strategy. This passage is intended somehow to impress the reader, but not to be translated directly into entries in his "world model" in the same way as more prosaic parts of the story.

Another kind of problem shows up in the following dialog, given by Schank[1973, p. 189]:

[John meets his friend Fred on the street.
 Fred is holding a knife. ...]

John: I could use a knife right now.
 [agitated tone]

Fred: What's the matter?

John: Damn Mary, always on my back. She'll be sorry.

Fred: I don't think a knife will help you.

The actual subject matter of the conversation is never explicitly mentioned. It is John's rhetorical threat to use a knife on Mary. Nevertheless that threat is so clearly implicit in the dialog that the latter may be considered to contain a sort of conceptual anaphoric reference

Current address: Computer Science Dept., Rutgers, New Brunswick, N.J. 08903

to it. This sort of undertone is, of course, far more prevalent in conversation than in newspaper articles. Indeed there seems to be a tendency in conversation to mention any sensitive topic only by indirect or ambiguous language.

Let us consider what inferences would be needed to process the dialog above. Superficially, "I could use a knife right now" is a simple fact. However such a statement is usually intended as a request. In this case, it is not a serious request, but rather an expression of John's agitation. Presumably the tone of voice is what allows Fred to realize this. With this understanding, "What's the matter?" is not really asking what John wants the knife for, which might be the interpretation in other contexts. Instead it refers to whatever has upset John, and is really just a sympathetic encouragement to go on. Thus "Damn Mary, always on my back" is not the non sequitur it might appear to be. It is what John intended to say originally, if he got a sympathetic response. "She'll be sorry" is the most difficult part. It must first be recognized as a threat. Then it must be connected with the knife, which makes the content of the threat clear. Finally, in the context of this discussion, the threat must be recognized as only rhetorical.

We propose that a more complete analysis be done of the ways that readers and hearers use utterances. We suspect that in a typical passage, much of the text serves to direct attention and otherwise control the course of the interaction between the speaker and listener. Actually, existing programs have mechanisms that are relevant to the sort of processing we have in mind. These programs "explain" observed actions by fitting them into known patterns and inferring motives for the actors. (See, for example, Schmidt and Sridharan[1977].) However such a system must be applied to explaining what the speaker is doing, as well as what is going on in the domain he is talking about. It must also have a variety of ways to respond to an utterance, other than treating it as data to be added to memory.

References

Eugene Charniak, "Toward a Model of Children's Story Comprehension", AI TR-266, Artificial Intelligence Laboratory, MIT, Cambridge, Mass., December, 1972.

Roger C. Schank, "Identification of Conceptualizations Underlying Natural Language", in Roger C. Schank and Kenneth Mark Colby, eds., Computer models of Thought and Language, w. H. Freeman and Co., San Francisco, pp. 187-247, 1973-

Charles F. Schmidt and N.S. Sridharan, "The Plan Recognition Problem: A Hypothesize and Revise Paradigm," working paper CBM-TR-77, Computer Science Dept., Rutgers, New Brunswick, N.J., March, 1977.