

AD-A063 063

YALE UNIV NEW HAVEN CONN DEPT OF COMPUTER SCIENCE

F/G 5/7

INFERENCE IN THE CONCEPTUAL DEPENDENCY PARADIGM: A PERSONAL HIS--ETC(U)

SEP 78 R C SCHANK

N00014-75-C-1111

UNCLASSIFIED

RR-141

NL

1 OF 1
AD
A063 063



END
DATE
FILMED
3 --79
DDC


(12)

LEVEL II

AD A063063

DDC FILE COPY

THIS DOCUMENT IS BEST QUALITY PRACTICABLE.
THE COPY FURNISHED TO DDC CONTAINED A
SIGNIFICANT NUMBER OF PAGES WHICH DO NOT
REPRODUCE LEGIBLY.




INFERENCE IN THE CONCEPTUAL DEPENDENCY PARADIGM:

A PERSONAL HISTORY

Roger C. Schank

Research Report #141

September 1978



YALE UNIVERSITY
DEPARTMENT OF COMPUTER SCIENCE

78 12 27 020

DISCLAIMER NOTICE

**THIS DOCUMENT IS BEST QUALITY
PRACTICABLE. THE COPY FURNISHED
TO DDC CONTAINED A SIGNIFICANT
NUMBER OF PAGES WHICH DO NOT
REPRODUCE LEGIBLY.**

INFERENCE IN THE CONCEPTUAL DEPENDENCY PARADIGM:

A PERSONAL HISTORY

by

**Roger C. Schank
Yale University
Department of Computer Science
New Haven, Connecticut 06520**

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.

DISTRIBUTION STATEMENT A

**Approved for public release
Distribution Unlimited**

14 RR- REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM	
1. REPORT NUMBER 141	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER 9 Research Rept.	4. TYPE OF REPORT & PERIOD COVERED Technical
6. TITLE (and Subtitle) Inference in the Conceptual Dependency Paradigm: A Personal History		6. PERFORMING ORG. REPORT NUMBER	
7. AUTHOR(s) 10 Roger C. Schank		8. CONTRACT OR GRANT NUMBER(s) 15 N00014-75-C-1111	
9. PERFORMING ORGANIZATION NAME AND ADDRESS Yale University - Dept. of Computer Science 10 Hillhouse Avenue New Haven, Connecticut 06520		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS	
11. CONTROLLING OFFICE NAME AND ADDRESS Advanced Research Projects Agency 1400 Wilson Boulevard Arlington, Virginia 22209		12. REPORT DATE 11 September 1978	
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office) Office of Naval Research Information Systems Program Arlington, Virginia 22217 1248p.		13. NUMBER OF PAGES 42	
15. SECURITY CLASS. (of this report) Unclassified		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE	
16. DISTRIBUTION STATEMENT (of this Report) Distribution of this report is unlimited. DISTRIBUTION STATEMENT A Approved for public release Distribution Unlimited			
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) conceptual dependency scripts plans goals natural language representation inference memory semantics artificial intelligence			
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) The history of inferences at the Yale Artificial Intelligence Project is discussed. A			

407051

JP

-- 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
Office of Naval Research Code 102IP Arlington, Virginia 22217	6 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

SESSION for	
3	White Section <input checked="" type="checkbox"/>
	Buff Section <input type="checkbox"/>
	UNFOUNDED <input type="checkbox"/>
JUSTIFICATION	
BY	
DISTRIBUTION/AVAILABILITY CODES	
Dist.	EXCL. OR SPECIAL
A	23 E4

Office of Naval Research Code 455 Arlington, Virginia 22217	1 copy
Office of Naval Research Code 458 Arlington, Virginia 22217	1 copy
Naval Electronics Laboratory Center Advanced Software Technology Division Code 5200 San Diego, California 92152	1 copy
Mr. E.H. Gleissner Naval Ship Research and Development Computation and Mathematics Department Bethesda, Maryland 20084	1 copy
Captain Grace M. Hopper NAICOM/MIS Planning Board Office of the Chief of Naval Operations Washington, D.C. 20350	1 copy
Mr. Kin B. Thompson Technical Director Information Systems Division OP-91T Office of the Chief of Naval Operations Washington, D.C. 20350	1 copy
Advanced Research Project Agency Information Processing Techniques 1400 Wilson Boulevard Arlington, Virginia 22209	1 copy
Professor Omar Wing Columbia University in the City of New York Department of Electrical Engineering and Computer Science New York, New York 10027	1 copy
Office of Naval Research Assistant Chief for Technology Code 200 Arlington, Virginia 22217	1 copy
Captain Richard L. Martin, USN Commanding Officer USS Francis Marion (LPA-249) FPO New York 09501	1 copy

Inference in the Conceptual Dependency Paradigm: A Personal History

PREFACE

During a summer workshop (funded by the Sloan Foundation) that we ran at Yale, I tried to present some of the views that we hold with respect to the problems of representation of meaning, the making of inferences, and the function of higher level descriptions of the structure of knowledge, to an audience primarily consisting of social and cognitive psychologists. Most of the participants in the workshop were interested in our ideas on this subject. However their background really had not prepared them to understand why we did what we did or how we came to do it. Consequently, I attempted to give them that background by retracing the steps in our research of the last ten years. I explained how we came to hold our current views on various subjects by showing what our initial assumptions were and how one position naturally led us to the next. Since most of the

participants in the workshop felt that these lectures provided the context that was necessary to help them to understand our current research, I felt that it would be of use to prepare a paper based on the lecture notes that I used in the workshop.

This paper therefore is necessarily sketchy. It is an attempt to present the outline of ten years of research and as such can only barely cover the issues. Furthermore, the paper is entirely biased towards the research within our own group. Much other work, some well known and other less known occurred before and during the work described here. To some extent this work affected our own, but by and large the work described here proceeded on its merry way untouched by very much from the outside. In this paper I shall attempt to show how and why we got to where we are today.

1966-1969

My initial research focussed on the representation of meaning as it would be used for the generation of natural language sentences. Since generation was the major problem in both linguistics and computational linguistics, this point of departure was not particular different from the established norms. The major difference was that my representations were intended to be psychologically correct (to the extent that that could be determined). This led me away from what I believed to be the many ad hoc entities that existed within transformational deep structures at the time.*1*

I thus began to think about the problem of representing

1 I started out in Artificial Intelligence as an undergraduate in mathematics (which included computer science) at Carnegie Tech (now Carnegie-Mellon) in the early sixties. I became interested in language, and, feeling that language would best be studied with those whose specialty it was, I went to study Linguistics. When I arrived at the University of Texas in 1966, Texas was a combination of two relatively unrelated paradigms. Transformational Generative grammar was beginning to become the dominant paradigm within the department of Linguistics itself. But, at the same time, an extremely large mechanical translation project that had flourished at Texas for some time, was present but drawing to a close, at least partially due to the ALPAC report which was issued that year. This meant that while there were a large number of computer oriented people in Austin whose interest was language, the power in the department itself, where I was a student, was held by the transformationalists.

As a computer type myself, I naturally assumed that the main issue in language study ought to be the representation of meaning. Not only was this not the view held by either of the paradigms present in Austin at the time, but it was a view whose opposite position was strongly held by all those around me. Syntax was the issue of the day in the Linguistics department itself, and at the Linguistics Research Center (LRC) where the MT work was being done, gigantic phrase structure grammars and rules that transferred structure from one language to another were what was being done.

My views came to me by way of Newell and Simon at Carnegie; (Actually this was quite indirect. I had never met either of them, but their work on GPS and more importantly their views on the nature of computer programs as theories was well known to most students at Carnegie), by Sheldon Klein who was my teacher at Carnegie and by Sydney Lamb who was Klein's teacher and whom I had met and interacted with at the Linguistics Institute held at the University of Michigan in the summer preceding my senior year at Carnegie.

My own view represented a sort of amalgamation of the views of all these people. I was encouraged in my effort by my advisor at Texas, Jacob Mey, and by my employer at Tracor (a company that allowed some of the LRC people to move there after the LRC funding dried up) Eugene Pendergraft.

meaning. In particular, as I was still interested in the computability of any representation that I came up with, I was especially concerned with the question of how a meaning representation could be of use in the generation of natural language sentences, and in the parsing of natural language sentences.

The first representation that I came up with looked a lot like English with arrows connecting it up. The arrows were gotten from dependency theory which had been written about by Hays and used quite a bit by Klein and to some extent Lamb. My contribution, as I saw it at that time, was to make the representation more conceptual. I reasoned that the dependency grammars being used at that time were too concerned with questions of whether the noun or the verb was really the head of the sentence and not enough concerned with the meaning of the sentence. Obviously, the main noun and main verb contributed equally to the conceptualization that underlied the sentence.*2* That is, both were necessary for the sentence to have meaning. I concentrated on issues such as these, coming up in the end with rules that described the make up of a conceptualization in terms of items more conceptual than nouns and verbs.

I then began to work on the problem of how my representation system would allow for random generation of English sentences (this being my computer-biased view of the work in Linguistics at the time and thus the field in which I saw myself doing battle)

2 That, of course, is the origin of the symbol \Leftrightarrow .

and how my system would work for the representation of languages other than English.*3*

To facilitate the former of these tasks, I had to invent what I called "a conceptual semantics" (Schank, 1968) which was basically a depository of world knowledge which prevented any random generator from generating sentences which meant nothing. The latter considerations of universality caused various English items such as prepositions to drop out of the representations. Even so, my representations bore a great deal of similarity to the surface properties of English. I was aware of this problem, but was far more concerned at the time with attempting to convince linguists that meaning considerations were important.*4*

At this point I was rather anxious to get a job, and largely through chance, found myself employed by Kenneth Colby at Stanford.*5* My job within his project was to create a parser that would allow his soon to be created*6* version of an automated psychiatrist to actually understand what the patient said. To do this, I attempted to reverse the rules that I had already written for generation out of Conceptual Dependency (CD).

3 My system being conceptual it ought to have been capable of handling this latter task. Furthermore, I could hardly have not worried about this issue since the emphasis in Linguistics was on what were then called "funny languages" (in my case I worked on Quiche and a little on Eskimo) and because MT was clearly in need of an interlingua that would facilitate translation.

4 I now know that there were some who were already convinced of that fact but they were not at Texas and were unknown to me. Fillmore's work was discussed but its properties were still far more syntactic than I had in mind.

Such a reversal of rules was established doctrine within computational linguistics and I assumed it would work.*7*

As an example of the kind of issue I was concerned about at that time consider the sentences:

I hit Fred on the nose

I hit Fred in the park

In order to parse these sentences correctly it is necessary to know where a person can be located. Here, "correctly" depended on what had to be represented in CD. There was a locative for entire conceptualizations and a "part of" relationship for objects, and either could be expressed in English with a locative prepositional phrase. To solve this problem I used the conceptual semantics I had invented for generation, and my rules that mapped from syntactic relationships to conceptual ones checked for acceptability each time a mapping was attempted. The

5 Colby had at this point recognized that his dreams of an automated psychiatrist depended on solving the natural language problem first. I of course agreed with this and found the idea of an automated psychiatrist fascinating. I had, until this time, never really thought about the higher level processes that were to operate on top of any meaning representation I came up with. I was interested, prior to my working for Colby, mostly in mechanical translation. Shortly after I arrived Colby and his project moved physically to the Stanford AI lab, thus cementing my involvement with AI.

6 It never actually existed in its planned newer form.

7 Of course, there were no semantic representations in computational linguistics at that time, those ideas about reversal had to do with syntactic rules. I saw no reason to not go directly to my new conceptual base, and anyway as I have said my conceptual base looked an awful lot like English in any case, so this did not seem to be that great a shift.

same thing was necessary for sentences such as:

I hit the boy with long hair

However, what had to be done to handle ambiguous sentence, was to add information in the conceptual semantics. The conceptual semantics for 'hit' consisted of information about the kind of objects to be found in various prepositional relationships. So, for 'hit' we had:

with - weapon object

by - no

on - part of <--PP

The final parse of this sentence put this additional information concerning the properties of "hair" in the actual representation. Although I did not use the term, this was in a sense the first class of inference to be made, and added to the meaning representation in my work. Since it had to be determined if the 'with' object for 'hit' was a weapon or a part of the object, that determination, once made, now became part of what had been understood and thus was part of what was meant. The can of worms that this adding of information not actually stated opened was tremendous. Additional information could be added to what had been said to form what had been meant. This was quite different from what had gone on in linguistics up until that time (and to some extent what still goes on in linguistics). I was saying that the meaning of a sentence was more than the sum of its parts. This heresy was not particularly appreciated when I

brought it up, although it probably had seemed obvious to those AI people that had looked at the problem.

Thus, my point was that Chomsky was wrong in claiming that we should not be attempting to build a point by point model of a speaker-hearer. Such a model was precisely what I felt should be tackled. Linguists viewed this as performance and thus uninteresting. I took my case to psychologists and found them equally uninterested. Psychologists interested in language were mostly psycholinguists, and psycholinguists for the most part bought the assumptions of transformational grammar (although it seemed very odd to me that given the competence/performance distinction, psychologists should be on the side of competence).

1970

In 1970 we started to make our representations more conceptual.*8* Until this point our supposedly language-free representations had a great deal of language in them. We noticed a class of verbs (which we termed pseudo-state verbs), where the object of the sentence did not seem to be the same as the underlying deep object of the underlying action. In particular, our representations seemed to require us to put in a great deal more than was in the sentence in order to make conceptual sense. Thus for (1) below we had to make up something called 'create',

8 The 'we' consisted of Larry Tesler (a programmer who worked for Colby and was attempting to write our parser) and Sylvia Weber (a graduate student in the Computer Science department at Stanford.)

in (2) we had a particular sense of 'have' and in (3) all of a sudden we had 'truth' and 'saying' present. There did not seem to be any way to avoid this introduction of elements that weren't there initially if we were to represent the meaning of what had been said:

1) he wrote a book

by he
 he<==>create<--book<--
 ||
 writes

2) he desired Martha in the morning

he<==>want
 ^
 he<===>have<--Martha
 ^
 |
 morning

3) he doubted his wife

he<==>doubt
 ^
 x<===>true
 |
 \ /
 wife <==> say
 of ^
 ||
 he


Examining these representations, we began the search for some regularities in the representation that would give us a more canonical form. What we had until that point was so free form that we could create anything at any time. This did not seem very sensible. In particular, there was a problem of what sense of the various multiple sense verbs we had at any given time. We couldn't just continue writing 'have' the way we had done. There had to have been some underlying basic forms. We considered for

a while just writing subscripts on the verbs. So 'understand1' was equal to 'see3'. But which sense was more basic? And, more importantly, how many senses of a word would there turn out to be and what would their intersections be? In the case of partial overlap of senses there was a definite problem with the subscript method.


As a side issue at this time, we attempted to clean up the mess in which we had left our representation of prepositions. We had been using an arrow to mean any prepositional relationship, in the faith that higher level processes that used our representations (we really had no idea what they would be like the psychiatric model would have to deal with that problem) would figure out the true relationship that held between an action and its associated objects. We tried to think about what kinds of prepositional relationships there were. Location had long before been relegated to describing conceptualizations themselves, so it wasn't a candidate. 'Part of' relationships were used to describe objects, rather than the relationship between actions and objects so they weren't part of the problem either. By eliminating these two classes of prepositions (those that described entire events and those that solely described objects) we found that there were only three kinds of prepositional relationships: instrumental, directional, and recipient. These relationships then described the way an action could relate to an object in an event regardless of what preposition was being used. Since we were describing relationships and not prepositions, we realized that English could be considered to have a kind of null

preposition denoting objective relationships but that this did not indicate that this was any less of a relationship between action and object than the others. We knew that Fillmore had said similar things about syntactic relationships in English so we christened our relationships 'conceptual cases'. The differences between the two systems were a lot greater than their names suggested and in retrospect this was probably a poor choice of names (see Schank (1972) for a discussion of those differences).

This new system of cases immediately had ramifications throughout our entire system. Thus, for example we had previously represented 'I want money' as:



However, adding a recipient to this representation caused us to come up with the following representation:



That is, we knew that we had a Recipient here and it had to be 'I'. Similarly there had to be an Object because what else could 'money' be? It didn't seem like an actor. The actor was unknown but we knew he was the same person as the donor of the recipient case. Of course, the above diagram had a glaring hole. What was the action? Still this representation seemed to make a lot more sense than the one having money be an actor that did the action

'go'.*9* What was needed at this point was a name for our unknown action, and since it was obviously a kind of transfer of the money that was being done we called it 'trans'.

'Trans' helped us with other problems as well. It solved the partial overlap problem in the meaning of words such as 'give' and 'take' and 'buy' and 'sell'. Furthermore, it eliminated the need for elaborate transfer of meaning rules of the kind Katz (1967) had been proposing for mapping words like 'buy' into 'sell'. We began to wonder what other actions like 'trans' were around.

We began to look at other representations we had, using the conceptual case notions that we had invented. For example we had represented 'He heard me' as:

```

      I
    he<=>perceive<---ears
                        ^^ of
                        ||
                        he
                        & --->he
I <=>say <---X<--- |
                    ---<I

```

'Hear' obviously took a conceptual object of sound that got translated into a meaning by some process. 'Me' was not a sensible conceptual object for 'hear'. Similarly the conceptual form of 'hear' needed an instrument that had something to do with

9 I should point out that one of the basic maxims of CD was that there was an actor-action-object framework into which things should fit. This was part of the rules that I worked out in my thesis (Schank, 1969).

'ears'. Such considerations forced us to rearrange the sentence into a conceptual format that represented the fact that the meaning of an idea being transferred was the key action taking place.

We began at this point, to look more closely at the concept of an action. We attempted to classify the verbs we had been using according to the cases they took and properties of their objects. This left us with S(tate)-ACTs, P(hysical)-ACTs, E(motional)-ACTs, and so on (Schank et al., 1970). Using this classification for verbs, we could now predict the missing cases that were implicit and thus that had to be inferred. We continued to look for effective groupings that would facilitate inference. Thus, while we did not actually set out to discover primitives, the considerations that we had in representation issues forced us to come up with some workable classification of actions that fit within our framework.

Inference was not yet a major issue in this regard, but other problems forced us to focus on it. For example consider the sentence 'I fear bears' and our proposed representation of it at that time:

$$\begin{array}{c} I \leq \Rightarrow \text{fear} \\ \quad \wedge \\ \quad | \\ \text{bears} \leq \Rightarrow \text{harm} \leftarrow -I \end{array}$$

In the same paper where we were wrestling with the issue of representation of actions (Schank et al., 1970) we also introduced an idea we called "associative storage of concepts". In order to

adequately represent sentences of the above type, it was necessary to have available a conceptualization that could serve as the object of the verb 'fear'. (At this point we viewed such a verb as a kind of stative ACT. We later realized states were not ACTs but states of objects.) Obviously this conceptualization had to have in it both 'bears' and 'I' as part of the object of 'fear'. Here again we were faced with the question of what was the ACT? The answer we chose was an ACT called 'harm'. As we were not interested in primitives particularly this should not seem strange. The focus of our interest was: how were we going to find the concept 'harm' to add to our representation?

The answer was through associative storage of concepts. What we meant by this was that there had to be some connection between fear and bears that would allow us to infer 'harm' as the missing ACT. Quillian (1966) had used an idea of a linked network of concepts that could be searched from two paths in order to find their shortest intersection. This idea had been used for disambiguation, but it now seemed that it could be extended for use here as well.

However, that seemed like a lot of work for so little. When we looked at other examples of the phenomenon we were trying to account for, an easier solution presented itself. For example, the sentence 'I like books' clearly needed something about 'I read books' inside the conceptualization that represented its meaning. It was obvious that this could be done simply by listing 'books' in the dictionary as a 'READ object'. All that

would then be required was an empty slot requiring an ACT. CD gave us that when a stative ACT was recognized, so all we had to do at that point was to look in the dictionary for an ACT associated with the object we had available. This did not solve the problem when the object was not the source of the inference. A functional object like a 'book' could well be listed as a 'READ object', but what were we to do when 'bears' or 'Nixon' was the object of a stative ACT? Since these objects were not functional in the same way, it seemed that the missing ACT would have to be supplied as a part of the meaning of the word 'fear'. Here again, we had, without quite intending to, decomposed the meaning of a word (fear) into more basic elements (fear plus expected harm). The reason this had happened was again attributable to the requirements we had put on CD with respect to what slots there were in a conceptualization and how they were to be filled. so, we were left at this point with a representation like:

```

I fear Nixon
  I<=>fear
    ^
    |
  Nixon<=>do
    ^^^
    |||
something <=>harm<---I

```

Thus, at this point we were now freely adding to our representation concepts that were not present in the English sentence in the first place, and perhaps more importantly, concepts that were only probably part of the meaning. These were the first explicit inferences that we had.

1971

In 1971, we began to focus on the problem of the inference of intentions. We got into this problem because of a peculiar use of language that we happened to come across, that we realized it was crucial for an reasonable understanding system to handle. The example was:

Q: Do you want a piece of chocolate?

A: I just had an ice cream cone.

Clearly, it is necessary to understand the answer here as meaning 'no'. In attempting to figure out how to do this, we realized that it was necessary to fill out the structure of the conceptualizations underlying both sentences so that a match could be made from the answer to the question. To do this required inferences, ones that were different from the "fill in the ACT" ones we had been working on. Thus we needed a structure like:

```

      want
      ~~~
      |||
trans  |||
      ~~~
      |||
      eat
      ~~~
      |||
satisfied

```

To get this structure we had to postulate that when a trans was present, the object being transed might enable an actor to perform the usual functional ACT done to this object.

Furthermore we had to examine the result of this action, because whatever state it caused was the key for the pattern match. That is, a paraphrase of this question might be 'Do you want me to trans you an object which is edible so you can eat it so that will make you feel some feeling (full, happy etc.)? The answer would then be 'I already have that feeling because I just did an action (here 'had' has to be inferred to be 'eat') that resulted in that feeling. To do all this required a new set of resultative and enabling inferences, and began to cause us to focus on the question of what kinds of inferences there were and where they came from.

One of the first issues however was the potential use of such inferences however. Since we were primarily concerned with parsing at this stage, we focussed initially on the issue of what expectations there were in processing that came from places other than the CD or syntactic expectations themselves.*10*

We looked at an example of a conversation where a person in a fit of anger at his wife, asks for a knife from a friend and when he is refused it says:

I think I ought to...

The question we asked was: what different kinds of things do you

10 Expectations were the key idea behind how our parser was supposed to work. The parser Tesler designed (Schank and Tesler, 1969) was intended to 'laugh' at the sentence 'I saw the Grand Canyon flying to New York' because it would have had its expectations violated. Riesbeck (1975) of course later did this in a more serious way.

expect at this point? We isolated these (Schank, 1971):

1. sentential - a verb coming
 2. conceptual - a entire conceptualization is coming
 3. context - "ought to have fish"
excluded by fighting context
something violent expected
 4. conversational - inference of reason person is talking
why tell someone about your future
violence unless you want them to stop
 5. memory - what kind of person is John?
should we take his anger seriously?
 6. culture - what happens in situations of this sort
- memory structure inferences used

it?

?

These questions started us looking seriously at what else was going on in understanding besides parsing. Clearly we needed a memory full of facts about the world to do any sensible understanding. At this point our focus began to change. The issues of representation and parsing still existed of course, but memory and inference were obviously at least as crucial.

Around this time, I met Abelson, who was working on beliefs as was Colby. I began to see that beliefs had a great deal to do with the processing of language. My group*11* began to attack this problem in a number of ways. Hemphill (1975) began to work on identifying how parsing was influenced by beliefs implicitly referred to in a text. I, as usual, concentrated my efforts on representation. In particular, it was necessary, in order to

11 By this time, a number of students had begun working with me. Goldman, Rieger, Riesbeck, Hemphill and Weber finished degrees with me at Stanford and were active in weekly meetings held during this time to discuss these issues. A number of other students who didn't finish contributed as well.

handle the above example, to postulate a set of beliefs that could account for our expectations about an actor's behavior. To understand that John was not likely to want to now sit down and be friendly in the above example we needed to know that when you're angry you don't like to be with the people you are angry with. This was represented as:

```

one <==>do
  1 ^^^
    |||
one <==>angry
  2 ^^^
    |||
one <==>want
  2 ^
    |
    one <==>interact
      1
      ^
      one
      2

```

Beliefs of this sort were useful for predicting the future actions of an actor. Adding beliefs to the representation changed the idea of inference as just added information that would help in the parsing of a sentence. It suggested that we had to concentrate on problems having to do with the representation of information in memory, and with the overall integration of incoming data with a given memory model. It thus became clear that natural language processing was a bit of a misnomer for our enterprise. What we were doing was not essentially different from what Colby or Abelson were doing. That is, we had to deal with the problem of belief systems in general. But, added to that, was the problem of representation

of meaning, knowledge, and memory in general.

1972

The integration of all these problems caused us to deal with sentences whose meaning was a product of the combination of all these issues. For example, "He acts like Harry" means different things if Harry is a cat, a child or an aged man. What is the correct representation for the meaning of such a sentence? Clearly it cannot be determined in any way apart from the memory structures its meaning relies on. Similarly, the sentence "He is doglike in his devotion" means nothing if there is no belief about the devotion or lack of it of dogs available in memory.

In attempting to catalog the kinds of beliefs that were being used in natural language expressions, we came across two (Schank, 1974) that used inference quite frequently. These were:

Pay back in kind.
Pay back not in kind.

Such beliefs were used in sentences such as:

John threw a hammer at Bill vengefully.

Mercifully the king only banished the knight
for killing his favorite horse.

This kind of analysis meant that we could predict other actions from what we determined was a reasonable inference. That is we could make inferences from inferences. So, 'vengefully' told us that Bill may have done something to John, and that John really wanted to hurt Bill and so on. Such inferences were an important part of the understanding process.

The cataloging of such random facts was not within our view of how to attack a problem, however. Instead we returned to attempting to make more rigorous the CD representations we were using so that we could better establish what was within the domain of a system like CD and what was outside of it. To do this, we considered the nature of the ACTs we had been using. At that point we had been using 'trans' and a hodgepodge of others that suited us. To remedy this situation we looked at the mental verbs which we had, to this point, virtually ignored.*12*

The significance of the primitive ACTs for us was that we could now be sure that we had a given agreed-upon representation for most of the sentences we were dealing with. This made our system usable by the large group of students who were beginning to concern themselves with programming systems that could communicate with each other. Further, we now knew what was in the bounds of the theory and what was not. We knew that to do the kind of work we were interested in, a canonical form was necessary. We were not so concerned with the ultimate correctness of that system as we were with its usability. No other canonical form existed, and transformational deep structure representations, which were the major well known alternative, neither adequately represented meaning nor were in any sense canonical. The most important part of the primitives of CD for

12 We formed a special group to consider this problem, consisting of Goldman, Rieger, Riesbeck and myself. Eventually, each of these students came up with one mental ACT and defended it to the others. This left us with three mental ACTs, CONC, MBUILD and MTRANS invented by the above people respectively. This gave us a total of sixteen primitive ACTs.

us then were that they facilitated our getting on to the more interesting problems at hand. They did this because they gave us a language in which to describe those problems.

1973

The most important problem was inference. The first paper on the complete set of primitives we had (Schank, 1973) made that clear in its title: 'The Fourteen Primitive Acts and Their Inferences'. The single most important fact about the primitive ACTs was that they helped to organize the inference problem. No primitive ACT meant anything in the system at all, other than the conceptualizations that might come to exist as inferences from it. Primitive ACTs served to organize the inference process, thus giving us a starting point from which to attack the problem.

We began to concern ourselves therefore with two principle kinds of inference, results from ACTs and enablements for ACTs. Then, having exhausted the inferences derivable from the ACTs themselves, we began to attempt to categorize the kinds of inferences that needed to be made in general. In Schank and Rieger (1974) we delimited the following kinds of inference:

1. Linguistic Inference (done before parsing is over)

buy - infers "money" as object
hit - infers "hand" as object

2. ACT Inference

whenever an actor and object were present in a CD
an action had to be inferred. Thus for "I like books"
"read" is inferred.

3. TRANS-ENABLE

For sentences such as "John wants a book" it is necessary to infer an ATRANS which then enables an ACT to take place. That ACT can be determined by ACT inference.

4. Result Inference

For any given ACT its results can be easily determined. "John went to N.Y." implies he got there. "Mary gave Bill a book" implies Bill has it and Mary no longer does. These come from PTRANS and ATRANS respectively.

5. Object Affect Inference

Inferences come from the interaction of objects and ACTs. In "John hit Mary with a rock" we infer that Mary is damaged and the rock is not. In "John ate an egg" we infer that the egg has been transformed. These inferences come from the ACT.

6. Belief Pattern Inference

When we see "John hit Mary" we infer that Mary must have angered him by doing something. This inference is gotten from a belief. The belief is accessed by matching a pattern containing "intentional damage" which is written in CD. This inference does not come from ACTs therefore but from states.

7. Instrumental Inference

We can infer instruments for ACTs that we have found. Thus INGEST implies PTRANS as instrument. PTRANS implies MOVE or PROPEL as instrument and so on.

8. Property Inference

Did Nixon run for President in 1863? This can best be determined by examining properties such as whether Nixon is alive or if it is election year before doing any exhaustive memory search. The inference of those underlying propositions means that preconditions for actions must be known. That is simple when the actions are simple ACTs. For concepts such as an election it is harder.

9. Sequential Inference

Results of the combination of two sentences can bring up new information. Thus "All redheads are Obnoxious" followed by "Mary has red hair", makes the latter statement more contentive. Such additional content depends on knowledge about Mary's personality. The

correct inference requires a memory search.

10. Causality

"John hit Mary and she died" or "John hit Mary" followed by "John died" both imply causality. Such causality information can only be determined by examining the resultative properties of the ACT.

11. Backward Inference

We can often determine a new fact by pondering how an old fact came to be. Thus "John knows where Mary

is" implies that he saw her or someone told him. This comes from the enablements for an action.

12. Intention

It is important to know why people do what they do. Thus intentions and motivations had to be inferred. We really did not know how to do that.

Rieger, Riesbeck, Goldman and I began to design a computer implementation of these ideas in 1972 which resulted in the MARGIE system (Schank et al., 1975).^{*13*} During the implementation of these ideas our views on parsing, generation and inference were altered by the task of attempting to specify precise algorithms for these processes. In particular for our discussion of inference, Rieger created a new classification of inferences based on his experiences with MARGIE (Rieger, 1974). These were:

- | | |
|------------------|--|
| 1. Specification | - unmentioned particulars are inferred |
| 2. Causative | - reasons for an action |
| 3. Resultative | - results of an action |
| 4. Motivational | - motivations for an action |

^{*13*} Actually this is not quite accurate. The MARGIE system was never intended to work as a whole system, and it was not designed that way. Jerry Feldman suggested that we put it all together. Prior to his suggestion we just had three unrelated student projects meant to cover three areas of interest.

5. Enablement - enablements for an action
6. Function - when an object is mentioned its potential use is determined
7. Enable Prediction - what ACT will be enabled by ACT or STATE
8. Missing Enablement - why someone can't do an ACT
9. Intervention - reasons for actions that prevent harm
10. Action Prediction - predict ACT from object
11. Knowledge Propagation - determining who knows what
12. Normative inferences - determining normal states of the world
13. State Duration - how long an ACT or STATE goes on
14. Features - who can be expected to do what
15. Situation - events are imbedded in larger events
16. Utterance Intent - why people do what they do and say what they say

In Schank (1975) we attempted to further codify the kinds of inferences that were available for a given ACT. For example, we listed these rules for the ACT PROPEL:

```

X<-->PROPEL<--Y<--|
                    |----<Z
                    |
                    |----<W

```

1. TRANS is implied if (Obj is not fixed)
2. Object in directive case (Z) is negatively affected if PHYSCONT is present and sizes are right.
3. If Z is human then X may have been angry at Z
4. If Y is rigid and brittle and nonfixed, and speed of instrumental ACT is great, then Y will be NEGPHYSST

In general we noted that there were two major kinds of inferences:

Forward ---> what consequences from an ACT?

Backward ---> why an ACT and what enabled an ACT.

At this point we began to take seriously the problem of codifying the kinds of causal relations that there were.*14* This work was crucial to the inference problem since, we had come to believe that the major inferences were (forward) consequences and (backward) reasons. Thus the primary task of the inference process was the filling in of causal chains. We identified four kinds of causal links, RESULT, REASON, INITIATE and ENABLE. RESULT and ENABLE were the forward and backward causal rules for physical ACTs, and REASON and INITIATE were the forward and backward links for mental ACTs. We also added the rule that ACTs could only result in states and only states could enable ACTs. This had the consequence of making our causal chains and thus our CD representations both very precise and very cumbersome. The precision was of course important for any canonical form, but the cumbersomeness was obviously a problem that needed to be dealt with.

As an example of the kinds of representations we were creating by doing this, consider the representation that we now had for a sentence such as "John's cold improved because I gave

14 I had wanted to do this for some time, but really had not had the opportunity. In 1973 I took a year off from Stanford and went to Lugano, Switzerland where a new Institute was starting up, taking Riesbeck and Goldman with me. We had little to do there but think and I was able to finish the book on MARGIE and start working on issues that I had not had time for before.

him an apple."

```

              o      D |----->John
I <==>PTRANS<---apple <---|
  ^^^                      |-----<
  ||| r
  |----->LOC (John)
apple <===|
  ^^^
  ||| E      o      D |----->John
John <==>INGEST <--- apple <---|
  POT                      |-----<
  ^^^
  ||| I      o |--->I want to eat the apple
John <==>MBUILD <---|
  ^^^                      |---<
  ||| R      D |----->John
John <==>INGEST <---apple<---|
  ^^^                      |-----
  ||| r
John <==>|----->HEALTH (XTY)
  P |-----<HEALTH (X)

```

One of the advantages of all this detail aside from those already mentioned is that it provided a facility for tying together sentences in a text. Thus, a paragraph will frequently consist of a series of conceptualizations that can be related by their implicit causal connections.

1974

We began, therefore, to work on the problem of representing text. This was, after all, the major issue all along. We were not particularly interested in isolated sentences out of context. Such sentences were probably the root of many of the solutions and the problems with those solutions found by transformationalists and computational linguists. People do not understand sentences in a null context. Why then did our theories try and deal with out of context sentences? The answer

was obviously that this was thought to a simplification that would facilitate research. But the problem was really significantly changed by this supposed simplification. Certainly parsing sentences in context is a more reasonable problem with respect to word sense disambiguation than is parsing out of context.

We had never dealt with texts of more than one sentence before because we just did not know how to represent them. Now, with the idea of causal chains, we could tie together texts in terms of their causal relations. Such a tying together, when attempted on real texts (Schank, 1975) helped to explain certain memory results (particularly those of Bartlett, 1932). Now we had a theory that said that a crucial piece of information had many causal connections and an irrelevant piece of information no causal consequences.

The work on causal connectedness gave us a theory that was helpful in explaining problems of forgetting and remembering, and helped tie together text. However, it could not explain how to tie together texts whose parts were not relatable by chains of results and enablements. Something else was needed for those situations.

The something else was obvious once we thought about it. The answer was scripts.*15* That is, scripts are really just prepackaged sequences of causal chains. Some causal chains are used so often that we do not spell out enough of their details for an understander to make the connections directly. Scripts

are a kind of key to connecting events together that do not connect by their superficial features but rather by the remembrance of their having been connected before. The prototypical script we chose was to describe what goes on in a restaurant. In a restaurant we cannot infer from entering a restaurant the causal connection to either ordering or paying. Because speakers assume you know that they do not bother to mention it. There is a causal chain there, but inferring it bit by bit is impossible, so scripts are necessary.

1975

We set about testing our assumptions about how scripts would facilitate the processing of connected text by building SAM (Script Applier Mechanism).^{*16*} SAM became a kind of inference maker because what it was doing was filling out the specific implicit events in a causal chain describing a static situation. Cullingford (1977) described the inferences that SAM made as:

^{*15*} The word script was originally used by Abelson for something different than, but related to, the current notion of a script. The concept of a script was invented very shortly after I arrived at Yale from Lugano. Abelson and I were discussing issues of mutual interest as soon as I arrived. At one time or another, Rieger, David Levy, and Allan Collins were also present. Minsky's notion of a frame was at that time known by us, and it had some influence in the finalization of the notion.

^{*16*} As soon as I arrived at Yale, I assembled a group of students interested in AI. In January 1975 they began to work with me. We thought for a while in seminars about what scripts were like, and in early May began to put together SAM. The initial version worked about six weeks later just in time to be shown at TINLAP I. The group that put SAM together was Richard Cullingford, Richard Proudfoot, Walter Stutzman, Wendy Lehnert, Gerry De Jong, and Chris Riesbeck.

1. Reference specification
2. Causal Chain connection
as part of Script Applier

Other kinds of inferences were:

- 1 - Immediate results: in order to match PTRANS
to the PATTERN - IS LOC

inferences about result were needed "Welcomed at
Peking Airport" generates PTRANS which can then
match \$VIPVISIT

- 2 - Mental ACT

if X says Y then Y-- thus, pronounced dead

implies that a person is dead.

service good implies customer knows it.

Thus activates \$TIP

- 3 - Location

a location described by one method is the
same as that described by another

"customer sat down"

"waiter went to the table"

need to infer that the seat is near table

Enclosure: John went to the hospital.

He was treated in the Emergency Room.

need to infer that Emergency Room is in hospital

- 4 - Movement

"John picked up a magazine. Then he went into
the living room."

To answer "Where is the magazine?" requires
knowing that small objects move with you.
(imagine "table" substituted for newspaper)

We began to wonder about where scripts came from. In
thinking about this we came up with the idea that plans gave rise
to scripts and that goals gave rise to plans. Coincidentally
Meehan was developing a story generator that served as a vehicle
for developing our ideas about plans and goals.*17*

Meehan's TALESPIIN made inferences in the course of telling a story, in order to keep track of the world model. Meehan (1976) found that Rieger's 16 classes were of little use in this task because the most interesting question was what is affected when a fact enters memory? Meehan had to keep updating his world model every time he generated a new fact. This meant making inferences about the consequences of every fact that he generated and using those consequences to affect the continuation of the story. Meehan found the following kinds of inferences useful for his task:

1. CONSEQUENCES

ATRANS ----> POSS
 FIXED ----> (let go -->fall)
 INGEST ----> object gone
 person satisfied

2. REACTIONS

What people will feel about things and people
 ATRANS what you want---->dislike
 THREATEN ----> person will do ACTION if afraid
 HUNGER ----> actor wants to fix it
 LOC
 (underwater) ----> get out
 flattery ----> like and trust

1976-1978

The last three years have found us developing the system of plans, goals, themes and scripts for use in understanding systems. This work produced many working systems (Carbonell, 1978 Wilensky, 1978, DeJong, 1978) and has greatly broadened our ideas

17 Meehan had been working on a story generator of a different sort when I arrived at Yale. The plans and goals work that we were doing forced him to reconsider his problem and approach it from a different angle.

about inference. We now believe the following:

There are a great many possible levels of description. Each of these levels is characterized by its own system of primitives and conceptual relationships. Inferences occur at each of these levels. Thus, for every set of primitives there are a set of inferences that apply to them. These levels have been described in Schank and Abelson (1977) and will not be dealt with in any detail here. We currently use the following levels and inferences on those levels:

1 - Micro CD

All events in a story can be connected at a level where every event is connected to the events that follow from it, and to the states enable it subsequent events. This produces a very detailed causal chains made up of the events and states that were actually mentioned in the text as well as those that had to be inferred in order to complete the chain. Thus, the Causal Chain made by the low level expression of facts is one part of understanding. Thus, in order to read a magazine, you must: ATRANS it; OPEN it; ATTEND to it; and MTRANS from it. When any one of these events is discerned the others must be inferred.

2-Macro-CD

Another type of causal chain exists at the Macro-CD level. There, events connect to other states and events in the same way as they did at the micro-CD level but the level of description is

different. Thus, going to Boston enables eating in a Boston restaurant at the macro-CD level. But, at the micro-CD level, the locations would have to be further specified. Actually going to Boston results in being in Boston which enables beginning to look for and go to a restaurant. This latter level of description can regress in infinite detail where, for example, walking is enabled by putting one foot in front of the other. Thus, the level of detail of inferences is extremely important and is dependent on the purposes the understander has in mind.

Thus, there are two levels of causal chains that apply in the magazine situation.

ATRANS

^^^

||| r

POSS

^^^

||| E

ATTEND

^^^

||| rE

MTRANS

This is MACRO-CD.

MICRO-CD is concerned with opening the magazine, holding it turning the pages etc. Each of those ACTs also uses causal chains but at a much more detailed level. Neither one of these levels of description is more correct than the other.

For causal chaining then, the needed inference types are:

What Enables

What Results

What are Reasons

What Initiates

These apply at both at the macro level and at the micro level.

3. Filling in Missing Information

For every object and person we hear about we are always tracking where they are; the state they are in; what they know and believe; and how they feel. All these inferences are possibly appropriate at any given time. Thus, other kinds of inference types that are necessary are:

Locational specification

Object specifications

Emotional specifications

Belief Specifications

4. Scripts

Scripts are an important part of the understanding process. Thus, the inferring of the presence of scripts and of the unstated parts of scripts is an important part of the understanding process. The following kinds of inference are significant:

filling in missing causal chains in a script

inferring what script is being used

inferring what unstated script was used instrumentally

Thus, when we hear that 'John robbed the liquor store,' it is appropriate to ask how he got there, how he got in, where he got his weapon, and so on. Such inquiries are a part of the inference process since it is only by knowing what we don't know that we can seek to infer it.

One of the main problems with reference to inferences about scripts is the question of why is a script being pursued. This leads to the problem of inferring plans.

5. Plans

For any given event, it is often important to know the motivations and intentions of the the actors in that event. This means knowing the plans being pursued by an actor. Thus it is necessary to make the following kinds of inferences:

Inferring the planbox being used

Why was a particular planbox chosen?

Inferring facts about an actor given

his choice of plans & planboxes

Inferring other plans an actor is likely

to pursue to get his goal

Predictive inferences about future planbox choices

What goal is he operating under?

This last inference leads to another class of information that spawns new inferences.

6. Goals

Detecting the presence of a goal causes the following goal based inferences to be made:

Why was this goal chosen?

What is in conflict with it?

Can it be subsumed?

Given this goal, what other goals can we infer?

Under what circumstances will it be abandoned?

Actually these inference types represent only the tip of the numerous kinds of goal based inferences that have been isolated by Wilensky (1978) and Carbonell (1978).

Here again since goals are dominated by themes, detecting what theme is present and making the appropriate inferences is necessary.

7. Themes

The theme based inferences include finding out:

What goals will be generated next?

What themes are likely to coexist with the given one?

Are there any conflicts in themes?

How might theme conflicts that are detected be resolved?

Where did this theme come from?

OUR PRESENT ANALYSIS OF INFERENCE TYPES

The inference types we have used are rather similar whether we are referring to scripts, goals, plans, themes, or whatever. Inside Conceptual Dependency structures, or knowledge structures or our recently invented triangular structures (Schank and Carbonell, 1978), or probably any reasonable representation of knowledge, the following general inference rules apply:

1. SPECIFICATION: Given a piece of an event, what else can be specified about the rest of the pieces?
2. MOTIVATION: Why did an event happen? Why this event and not another? What did the actor believe he was doing?
3. ENABLEMENT: What was necessary for the event to occur?
4. RESULTS: What are the results or effects of this event?
5. STRUCTURE : What higher level structure does this fit in?
6. OTHER EVENTS: What other events are known to cooccur with this event? What could not have happened if this event happened?

These six inference types then are what we have. Scripts, plans and so on fit in as events in the above description. Thus, we can ask for SPECIFICATION, MOTIVATION, ENABLEMENT, RESULTS,

STRUCTURE and OTHER EVENTS for a script, a plan, a goal, or a theme or probably any other higher level structure we are likely to invent.

Inference then, is the fitting in of new information into a context that explains it and predicts other facts that follow from it. Since these explanations can occur at many levels inference is a very complex problem and one we expect to continue working on in an attempt to find out how people understand and how computers could understand.

Overview

I have attempted here to demonstrate how our ideas evolved and why they are where they are today. Since this theory evolution is ongoing it should be clear that the conclusions we have reached about inferences here are probably also just stopping points in the evolution of a theory. Nevertheless there are some things we can conclude from all this. In particular, there are patterns from which we can get a glimpse of the future.

As I have stated, this work started out as a linguistic theory, albeit one with a computer-based bias. Linguists have explicitly rejected it as a possible linguistic theory (see for example Dresher and Hornstein, 1976). In one sense they are right. The phenomena we have become interested in over the years are not particularly phenomena of language per se. Rather, they are phenomena having to do with the processing of language in general and the issue of the representation of knowledge in

particular. Thus, as we have moved away from linguistics over the years, we have become more involved with psychology.

At the same time as this work was going on, the field of Artificial Intelligence has been evolving too. When I first arrived at the Stanford AI lab, the major issues in AI were theorem proving, game playing and vision. Natural language was not considered to be a serious part of AI until Winograd (1971) presented the AI community with SHRDLU. This work contributed substantially to the evolution of AI. The major concern of AI would now seem to be the issue of the representation of knowledge, which of course makes the work in natural language processing quite central.

In the future I expect many of the relevant fields will begin to become less separate. AI must come to terms with the fact that it is concerned with many issues that are also of interest to philosophers. I hope that the cooperation will be of more use than was the head butting that has gone on between AI people and linguists. (Although this too has changed as the more liberal forces in linguistics have become both stronger and more interested in AI.) Also, the interaction between psychologists and AI people should continue to flourish. The work of Rower et al (1978) and Smith (1976) has already served to bolster the relationship between our group and cognitive psychology.

And what will happen to our theories in the future? I can only say that many of our ideas on parsing, the separation of inference from other processes, generation, and memory are

rapidly changing. We will, of course, continue to use the same methodology of the free form speculation approach to theory building, modified by our experiences testing out these theories on the computer.

References

- Bartlett, R. (1932). Remembering: A Study in Experimental and Social Psychology, Cambridge University Press, London.
- Bower, G.H., Black, J.B. and Turner, T.J. (1978). Scripts in Text Comprhension and Memory. To apper in COGNITIVE PSYCHOLOGY, pp. 78-139.
- Carbonell, J.G. (1978). Politics: Automated Ideological Reasoning. COGNITIVE SCIENCE, (2), pp. 27-51.
- Cullingford, R.E. (1977). Script Application: Computer Understanding of Newspaper Stories. Ph.D. Thesis, Yale University.
- DeJong, G.F. (1978). Automated Partial Understanding of Connected Natural Language Text. Ph.D. Thesis (forthcomming), Yale University.
- Dresher, B.E. and Hornstein, N. (1976). On Some Supposed Contribution of Artificial Intelligence to the Scientific Study of Language. COGNITION (4), pp. 321-398.
- Hemphill, L. (1975). The relationship of language and belief: With special emphasis on English 'for' contructions. Ph.D. Thesis, Department of Linguistics, Stanford University.
- Katz, J. (1967). Recent issues in semantic theory. Foundations of Language, Volume 3.
- Meehan, J.R. (1976). The Metanovel: Writing Stories by Computer. Ph.D. Thesis, Yale University.
- Rieger, C.J. (1975). Conceptual Memory. In R.C. Schank, ed. Conceptual Information Processing, North Holland Publishing Corporation, Amsterdam.
- Riesbeck, C.K. (1975). Conceptual Analysis. In R.C. Schank, ed. Conceptual Information Processing, North Holland Publishing Corporation, Amsterdam.
- Quillian, M.R. (1968). Semantic Memory. In M. Minsky, ed. Semantic Information Processing, M.I.T. Press, Cambridge.
- Schank, R.C. (1968). Outline of a conceptual semantics for generation of coherent discourse. TRACOR-68-472-U, Tracor, Inc., Austin, Texas. Also printed in Mathematical Biosciences 5, 1969.
- Schank, R.C. (1969). Conceptual dependency as a framework for linguistic analysis. LINGUISTICS 49.

Schank, R.C. (1971). Finding the conceptual content and intention in an utterance in natural language conversation. Proceedings of the Second International Joint Conference on Artificial Intelligence, London.

Schank, R.C. (1972). "Semantics" in conceptual analysis. LINGUA, 30.

Schank, R.C. (1973). The fourteen primitive actions and their inferences. AIM-183, Department of Computer Science, Stanford University.

Schank, R.C. (1974). Adverbs and belief. LINGUA, 33(1), pp. 45-67.

Schank, R.C. (1975). Conceptual Information Processing. North-Holland Publishing Corporation, Amsterdam.

Schank, R.C. and Tesler, L. (1969). Inference and a Conceptual Parser for Natural Language. Proceedings of the International Joint Conference on Artificial Intelligence. Edited by Walker and Norton, Washington, D.C.

Schank, R.C., Tesler, L., and Weber, S. (1970). SPINOZA: Conceptual case-based natural language analysis. Stanford AI MEMO 109, Department of Computer Science, Stanford University.

Schank, R.C. and Rieger, C.J. (1974). Inference and the computer understanding of natural language. Artificial Intelligence 5, pp. 373-412.

Schank, R.C. and Abelson, R.P. (1975). Scripts, plans and knowledge. Proceedings of the Fourth International Joint Conference on Artificial Intelligence, Tbilisi, USSR.

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

Schank, R.C. and Carbonell, J.C. (1978). Re: The Gettysburg Address: Representing Social and Political Acts. Research Report No. 127. Department of Computer Science, Yale University.

Smith, E.E., Adams, N. and Schorr, D. (1976). Fact Retrieval and the Paradox of the Expert. Stanford University.

Wilensky, R. (1978). Understanding Goal-Based Stories. Ph.D. Thesis, Yale University.

Winograd, T. (1971). Procedures as a representation for data in a computer program for understanding natural

language. TR-84, MIT Project MAC, Massachusetts Institute
of Technology, Cambridge.