

DOCUMENT RESUME

ED 136 188

CS 003 190

AUTHOR Rubin, Ann D.; And Others
 TITLE A Process-Oriented Language for Describing Aspects of Reading Comprehension. Technical Report No. 13.
 INSTITUTION Bolt, Beranek and Newman, Inc., Cambridge, Mass.; Illinois Univ., Urbana. Center for the Study of Reading.
 SPONS AGENCY National Inst. of Education (DHEW), Washington, D.C.
 REPORT NO BBN-3427
 PUB DATE Oct 76
 CONTRACT 400-76-0116
 NOTE 41p.

EDRS PRICE MF-\$0.83 HC-\$2.06 Plus Postage.
 DESCRIPTORS *Artificial Intelligence; Computers; *Models; Prose; *Reading Comprehension; Reading Materials; *Reading Processes; *Reading Research; Reading Tests
 IDENTIFIERS *Center for the Study of Reading (Illinois)

ABSTRACT

Language comprehension is an immensely complex process involving the dynamic interaction of diverse sources of knowledge. In order to model this process, tools are needed which allow detailed specification of the process components. In this paper, the essential characteristics of a model of reading comprehension are discussed in the context of the development of a computer model of the processes involved. Specific examples of text are analyzed to illustrate some of the complexities. It is argued that such a model would be valuable both in the construction of tests and instructional materials and in the systematic study of reading. Issues in implementing the computer model are also considered.

(AA)

 * Documents acquired by ERIC include many informal unpublished *
 * materials not available from other sources. ERIC makes every effort *
 * to obtain the best copy available. Nevertheless, items of marginal *
 * reproducibility are often encountered and this affects the quality *
 * of the microfiche and hardcopy reproductions ERIC makes available *
 * via the ERIC Document Reproduction Service (EDRS). EDRS is not *
 * responsible for the quality of the original document. Reproductions *
 * supplied by EDRS are the best that can be made from the original. *

ED136188

CENTER FOR THE STUDY OF READING*

U S DEPARTMENT OF HEALTH,
EDUCATION & WELFARE
NATIONAL INSTITUTE OF
EDUCATION

THIS DOCUMENT HAS BEEN REPRO-
DUCED EXACTLY AS RECEIVED FROM
THE PERSON OR ORGANIZATION ORIGI-
NATING IT. POINTS OF VIEW OR OPINIONS
STATED DO NOT NECESSARILY REPRE-
SENT OFFICIAL NATIONAL INSTITUTE OF
EDUCATION POSITION OR POLICY

Technical Report No. 13

A PROCESS-ORIENTED LANGUAGE FOR DESCRIBING
ASPECTS OF READING COMPREHENSION

Ann D. Rubin, Bertram Bruce, and John S. Brown

Bolt Beranek and Newman Inc.

October 1976

University of Illinois
at Urbana-Champaign
1005 West Nevada Street
Urbana, Illinois 61801

BBN Report No. 3427

Bolt Beranek and Newman Inc.
50 Moulton Street
Cambridge, Massachusetts 02138

The research described herein was supported in part by the National Institute
of Education under Contract No. ~~MS-NIE-76~~ 400-76-0116.

*Designation as a Center pending approval.

TABLE OF CONTENTS

	<u>PAGE</u>
A. Overview	1
B. Essential Characteristics of a Model of Reading Comprehension	1
C. Why a Computer Model?	6
D. Potential Uses of a Language for Describing Reading Comprehension	8
E. A "Simple" Example	12
F. A "Garden Path" Analysis	21
G. Exploring Representation Issues	25
H. Implementation Issues	33
I. Conclusion	35
REFERENCES	36

A. Overview

Language comprehension is an immensely complex process involving the dynamic interaction of diverse sources of knowledge. In order to model this process we must have tools which allow detailed specification of the process components. Traditional approaches to the study of reading have met with only limited success. Our approach is based on the realization that a much richer variety of intellectual tools is required if we are to make significant progress in our understanding of the reading process. We propose to develop a language for describing aspects of reading comprehension which will facilitate construction of tests and instructional materials, and make possible a more systematic study of reading. The validity and usefulness of this language will be explored via the implementation of a computer model of aspects of comprehension for a particular text.

B. Essential Characteristics of A Model of Reading Comprehension

Before discussing the uses and implications of a model of reading comprehension, we will discuss three characteristics of such models which we take to be essential. Briefly stated, such models should be multi-level, interactive, and hypothesis-based. Multi-level

implies that knowledge structures, which we call schemata, at several different levels are actively used in the reading process; traditionally-proposed levels include orthographic, phonological, lexical, syntactic and semantic. Clearly, higher-level knowledge sources such as inference rules (Rieger, 1975), social action theory (Bruce and Schmidt, 1974; Bruce, 1975a; Schmidt, 1975) and expectations about story structure (Rumelhart, 1975) are crucial components of the skilled reading process.

Interactive reflects our conviction that these varied knowledge sources interact in a heterarchical fashion; that is, although they may naturally form a knowledge hierarchy running from orthographic knowledge to expectations about story structure, communication is not limited to adjacent members of the hierarchy. The scenario proposed by some psychologists (Gough, 1972; LaBerge and Samuels 1974), which involves a visual input progressing linearly through the various knowledge levels to arrive finally at a "meaning", is not considered plausible. Instead, we will consider models which allow each knowledge source to put in its "two-cents' worth" at various points in the progression to comprehension of the text (Rumelhart, in press).

The coordination of this multitude of contributions requires a central structure which collects evidence for various interpretations of the text. We may generically

call such a structure a hypothesis and our models hypothesis-based models (Rubin, 1975). Two characteristics of hypotheses are important to mention here: (1) a hypothesis represents a possible interpretation which may later either be proven or disproven. (2) part of the structure of a hypothesis is the specification of those pieces of evidence which support or contradict it.

Several existing reading theories share significant properties with the general form described here. Goodman (1973) describes receptive language processes in general as hypothesis-based, defining them as "cycles of sampling, predicting, testing and confirming." He recognizes three levels of cues which readers use: graphemic, syntactic and semantic; these cue systems are used "simultaneously and interdependently." Productive reading is seen as requiring strategies which facilitate the selection of the most useful cues.

Smith (1973) also emphasizes the contribution of what he terms "nonvisual" information to reading. This nonvisual knowledge includes what people already know about reading, language and the world in general. He argues particularly that reading is not decoding to sound, but rather that semantic and other nonvisual processes intercede between visual processes and reading aloud.

Perfetti (1975) proposes at least three levels of sentence processing which obviously require corresponding levels of knowledge. He also focuses more explicitly on how the various component processes might interact, basing his overall conclusions on the fact that all the processes which occur during reading comprehension must share a "limited capacity processor."

Though our approach shares much with that of these and other investigators, there are also some differences in emphasis. We propose to be more explicit in the designation of different levels of knowledge sources, particularly in the area Goodman terms "semantic." We recognize at least the following types of knowledge: word semantics; knowledge of logical inference rules; discourse semantics; knowledge of social actions, their preconditions and outcomes; story schemata; understanding of various reading tasks; and strategic knowledge about how to use each of the above knowledge sources. In addition, we consider the explicit definition of the interaction between these knowledge components of the utmost importance and propose to investigate the possibility that some unskilled reading may be the result of not knowing how to use and interleave knowledge, rather than of a lack of knowledge itself!

A final emphasis of our theory-building will be to take the notion of hypothesis seriously, in particular the

notions that a hypothesis may be wrong and that at various points during the reading process it may be in a state of limbo, only partially specified, needing more evidence, or perhaps even uncertain because of conflicting evidence. Some researchers (e.g. Fodor, Bever, and Garrett, 1974) have tried to investigate the temporal course of reading comprehension with experiments such as phoneme monitoring; we intend to consider as well the possibility that as a consequence of some of the intermediate stages, the reader must "back up" and re-hypothesize about the meaning of a text. Goodman (1973) has noted that "proficient readers...are able to recover when they produce miscues which change the meaning in unacceptable ways." We will attempt to isolate these circumstances and define the methods skilled readers use to debug their hypotheses.

An important aspect of the above-described models which has practical implications for reading problems is the emphasis on structure-building. These structures or schemata are important for both the final representation of the meaning of the text and the intermediate hypotheses which are so crucial to attaining the final goal. Three classes of knowledge are necessary for building such structures. First of all, a reader must have sufficient information about the types of schemata which are possible at each level, how to recognize them and what implications they have for further processing. Second, there is a whole

body of knowledge which we might term strategic; it consists of information on how to use the structural knowledge, what priorities to use in evaluating hypotheses and what form the final "understood" structure should take. Third, there is knowledge about the purpose of reading the particular test, which can dramatically alter both the structural and strategic knowledge used.

C. Why a Computer Model?

The most important motivation for turning to the computer is the need for an appropriate language for expressing the theoretical constructs underlying the structure and use of schema theory and its interactions with lower-level knowledge sources. The comprehensiveness and utility of such a theory rests in part on how clearly one can specify these interactions so heterogeneous knowledge sources cooperate to produce "comprehension." How does one really define and represent the strategic knowledge controlling these interactions and verify that it has the desired effect?

We can talk loosely about these control structure issues in terms of passing messages back and forth between the various process levels as a way of controlling the interaction between high level hypothesis based processes and bottom-up data driven processes. An implemented

computer model, however, gives us much greater power for precise expression. It would provide us with a way of examining the consequences of modifying or deleting certain strategic rules. By hand (with paper and pencil) we could never keep track of the combinatorial interactions of all the processes involved. Processes can interact in subtle and unpredictable ways. A computer facility, however, provides an exhaustive system for carefully studying these interactions exhaustively.

We want to emphasize the importance of the influence of computational concepts and of actually implementing portions of our proposed model on the computer. To reiterate, a computer model is valuable for several reasons:

1. It forces us to be explicit in our design of representations, processes and strategies.
2. It provides a method for dealing with the complexity of the cognitive processes we are exploring; traditional pencil and paper methods fall short of this goal.
3. A computer model is an objective test of a theory; its operation is not influenced by the designer's prejudices and hopes.
4. It allows us to generate an exhaustive list of the possible processing paths in a given situation, rather than just the few which introspection discovers.
5. It provides at least some rudimentarily measurable quantities such as space and time requirements and number and type of inferences needed.

The BBN speech understanding system illustrates both the consequences of attempting to implement a complex

language-processing model and some of the techniques developed to deal with the problems encountered. One of the concepts developed in this context to deal with the interaction of low-level acoustic processes and higher-level syntactic and semantic ones has been that of "verification." The acoustic recognition procedures have a threshold by which they eliminate marginally-matched words in their preliminary processing of the input. If later the syntactic or semantic component proposes a word which the acoustic process did not discover in its initial scan, that word can be explicitly matched with less stringent requirements. We intend to take advantage of the insights already provided by work on the speech understanding system in our work on reading; such insights are indicative of the advantages of building and using computer models.

D. Potential Uses of a Language for Describing Reading Comprehension

A process-oriented language for describing reading comprehension has many potential uses in teaching and studying reading. Although it is not our goal to produce practical tools, we plan to test our model's feasibility by applying it to two real tasks: analyzing reading tests and scoring recall protocols.

The assessment of reading comprehension would be

greatly facilitated by a reading test which could determine whether or not a particular inferential skill had been mastered by a reader. By representing in the computer all of the relevant inference rules and world knowledge applicable to a small piece of a particular text, we could examine in detail all the possible applications and interactions of the rules which could lead to answering test questions. Each step of each solution path could be recorded. By examining the resulting solution space (i.e. the set of all solution paths) we could determine if all the answer paths used a particular mediating inference skill (such as rules about speech acts). These rules of inference will not be restricted to "logical" rules but will include such additional reasoning procedures which we know people use. To some extent this will be achieved by building on existing work on inference (e.g. Collins, Warnock, Aiello and Miller, 1975). We certainly cannot anticipate all conceivable ways a person might think in answering a given question. However, we claim such a computer model could be extremely useful for tracing out all the inferences -- both valid and hasty -- that could follow from the knowledge base of the model.

This opens the way to more objective scaling on a set of dimensions not normally used in test design, for example, measuring the amount of world knowledge required to answer a test question. One might attempt to measure this in terms

of the number of schemata invoked for a solution path and their degree of embedding. We could investigate the depth of inferencing required in terms of the shortest path, and this in turn could be used as the basis for a measure of the inferencing efficiency of particular solutions. We could measure some of the short term memory demands in terms of the amount of backtracking required or potentially required. In addition we might explore the possibility of devising more sophisticated measures of readability. Traditionally these have been based on more or less crude measures of sentence complexity, together with word frequency counts (e.g., Dale and Chall, 1948; Bormuth, 1967). A computer-based test analysis opens the door to much more varied and meaningful measures.

A computer model also has great promise for providing a partial solution to a long-standing problem in research in psychology and education, namely the problem of how to provide objective and reliable scores for free recall protocols. Currently, in many pivotal recall experiments, we must rely solely on the experimenter's good judgment in naming and classifying differences between the story and its recalled form. The partial solution we propose is to utilize a symbiotic person/machine system. The role of the computer model will be to specify a set of transformations between the original text and the recall protocol which maps one into the other (as far as possible). The role of the

human will be to determine what additional knowledge is required to complete the mapping. In other words, s/he will determine what knowledge is needed to account for idiosyncratic distortions, as well as those which are more widespread and predictable. Where s/he seems to figure prominently in recall p... can be entered into the system.

Using a computer model to help in scoring recall protocols is a good test of the model and may provide new insight into the analysis of recalls. A sophisticated scoring procedure must operate on a context larger than just isolated propositions of the text. For example, let us consider the simple proposition (in a text)

"Jane was watering the flowers."

which a subject recalls as:

"A little girl was watering her flowers".

If our scoring algorithm focussed exclusively on one proposition at a time (scoring proposition by proposition) then the first noun phrase might be scored as an over-generalization (Frederiksen, 1975). However, suppose somewhere later in the test there is the sentence:

"Her mother called to her to come in and pick up her dolls."

Then this later proposition interacts with the first (via an inference rule) yielding a highly plausible inference that Jane is, in fact, a little girl. It is precisely these

interactions that our process model can help account for and thereby make the task of the semi-automatic scoring of conceptual structures both objective and sensitive to various kinds of reading skills.

E. A "Simple" Example

In this section we will analyze an extended example, focussing on the types of inferences people make in answering test questions. The purpose of this and the next section is to illustrate the processes we expect our model to be able to explicate and, in particular, to emphasize the non-trivial nature of the reasoning necessary to understand even fairly simple stories. The piece of text we have chosen is representative of the sort of test item we hope to be able to handle: an apparently simple "story" and its related multiple choice questions. This example is taken from the Educational Testing Service (1960) Cooperative English Test of reading comprehension. In the context of the test, the story is followed by five multiple-choice questions to be answered on the basis of the passage. We will first discuss some of the knowledge and inferencing ability necessary to answer the questions, then consider more precise notions of how the information might be represented.

The inference mechanisms used in answering the test

questions are central to the issues we have just discussed; understanding these processes will help to provide techniques for the measurement of text difficulty and a method for specifying what each question is actually testing. In addition, we may expect some of the inferences pinpointed by the test questions to show up in recalls of the story, so our model will have to understand their derivation.

One major distinction we will see in the discussion of inferences below is that between linguistically-based and real-world-based (extra-linguistic) knowledge and inference. The former is language-specific knowledge which enables the reader to go from the printed words to his/her extra-linguistic knowledge. The latter is knowledge which the reader has primarily developed through experience, such as "when people yell, they are often angry."

Another point worth noting on a general level is the temporal nature of the comprehension process. Although the discussions of answering questions below do not explicitly deal with intermediate stages of reading the story, the order of sentences in a story obviously has an effect. For example, the reader needs to construct many partial hypotheses in the course of reading which cannot be completely specified until more of the story is read. Part of a reader's strategy may be to mark certain inferences as

"important to make as soon as enough information is present." A story which starts out, "Her father was a tyrant" should set up an expectation for the reader of resolving the reference to her. Such sequence-sensitive issues are noted in several places in the discussion below.

The story we will use as the basis of our discussion is the following:

"Alice!" called a voice.

The effect on the reader and her listener, both of whom were sitting on the floor, was instantaneous. Each started and sat rigidly intent for a moment; then, as the sound of approaching footsteps was heard, one girl hastily slipped a little volume under the coverlet of the bed, while the other sprang to her feet and in a hurried, flustered way pretended to be getting something out of a tall wardrobe.

Before the one who hid the book had time to rise, a woman of fifty entered the room and, after a glance, cried, "Alice! How often have I told you ~~not~~ to sit on the floor?"

"Very often, Mommy," said Alice, rising meekly, meantime casting a quick glance at the bed to see how far its smoothness had been disturbed.

"And still you continue such unbecoming behavior."

"Oh, Mommy, but it is so nice!" cried the girl. "Didn't you like to sit on the floor when you were fifteen?"

The first question is:

1. Alice's companion was
 - A a girl
 - B her brother
 - C the family dog
 - D a doll

The information necessary to answer this question is essentially contained in the fragment... "one girl hastily slipped a little volume under the coverlet of the bed, while the other sprang to her feet..." Using basically linguistic knowledge about gender and the implications of "the other", we can infer that two girls are involved in the action. However, we only discover that one of them is, indeed, Alice when the "woman of fifty" reprimands her by name and that discovery is contingent on understanding direct address, another linguistic ~~feature~~. Note that the very first sentence of the story ~~sets~~ the expectation that someone in the story is ~~named~~ Alice and that part of the comprehension process will ~~involve~~ discovering who it is.

The second question ~~is~~ more complex in its involvement of real-world knowledge:

2. When Alice heard the approaching footsteps, she probably was:

- E angry
- F alarmed
- G puzzled
- H amused

Several pieces of evidence go into the inference that Alice was most probably alarmed. At one level, we may look at various words used to ~~describe~~ Alice, that she "started

and sat rigidly intent" certainly suggests alarm. But this is not sufficient in itself, and comprehension requires setting up a hypothesis designating this description and supporting evidence. This hypothesis might be confirmed or refuted by further sentences in the story. Alice's later being "meek" gives support to the alarmed hypothesis, but only if we have some structure which relates the two. The real reason that we believe Alice is alarmed is that we know she is feeling guilty and is afraid her mother will discover the book hidden under the covers. Many parts of the story contribute to the "guilt" hypothesis: besides the above-mentioned phrases, the fact that one girl hid the book while the other girl pretended to be occupied with the wardrobe, is a link to the reader's non-linguistic knowledge of such situations. It is the cumulative effect of such details that supports the "Alice was feeling guilty" hypothesis.

The third question is:

3. We may infer that Alice is:
- A. stupid and resentful
 - B. very much in love
 - C. fifteen years of age
 - D. a spoiled child.

The phrasing of this question alerts us to the fact that inference will be important. In fact, deciding that Alice is fifteen is risky at best and in no way "provable". We decide she is fifteen because we know of a strategy: "if you're being blamed for something, attempt to elicit the

sympathy of the blaming authority by getting them to admit they've done the same thing." In order to infer that this strategy is being applied here, we must first realize that Alice is being blamed for sitting on the floor, a conclusion which follows fairly directly from the mother's first question and Alice's meek response. Then we must note that, in speaking to her mother, Alice has added a piece of information to the description of her action which (under this hypothetical persuasion strategy) indicates she is herself fifteen. It is worthwhile noting that almost all of these conclusions are based on the reader's understanding of the implications of social actions and speech acts. For example, although Alice's final remark is syntactically a question, its real purpose is to persuade, not to gain information. Neither is Alice's mother's "How often have I told you not to sit on the floor?" really a question. The inference of guilt is based on our knowledge of the social conventions surrounding the speech acts as well as our knowledge of mother/child relationships.

Given that we understand, at least sketchily, how we might conclude that Alice is fifteen, we are still faced with an important problem in understanding how we can answer this question. The problem is one of control structure: how do we choose this particular reasoning path out of all the possible ones to follow? In this case, reasoning backward from the question is clearly important. Good test-takers

read over the possible answers to multiple-choice questions and use them to guide their detailed thinking. In this case, in considering answer C the reader's attention can be directed to the final paragraph where there is a reference to age, and reasoning continues from there. To understand the distinction between inferences made while reading the story and those made in response to questions, one might describe Alice just after reading the story compared with a description given after answering the questions. Mention of Alice's age would be much more common in the second description; although the information necessary to infer her age is present in the story itself, the actual inference is probably not made (or not remembered) unless explicitly asked for.

There is more evidence of question-directed inference in the fourth question:

4. When she heard her name called Alice was evidently
- E reading to herself
 - F reading aloud
 - G lying in bed
 - H making her bed

We know fairly directly that a "reading aloud" is taking place from the phrase "the reader and her listener." (This is not really a trivial inference and working it out in detail might make a good first goal for a representation.) By following the chain of references through the next several sentences, we can infer that it was Alice who hid the book. However, we have no reason to

believe that Alice was reading rather than listening; the fact that she hid the book is suggestive, but not confirming. A "process of elimination" strategy is necessary to answer the question. In this case, the other three possible answers are easy to rule out and we conclude that it was Alice. One implication of this example is that a child may do better on a reading test because s/he uses certain strategies which might be termed test-taking skills. These strategies are examples of reading with a goal, and they must be considered part of the knowledge necessary to perform well on such reading tests. The existence of such question-based inference strategies also points out a weakness in determining the difficulty of a text in vacuo, i.e., outside of a task definition. It is easier to check whether or not a given fact is consistent with a story than it is to answer a more general question.

Finally, the fifth question:

5. Alice was worried about the appearance of the bed because

- A she had neglected to make it up
- B her companion had been sitting on it
- C her companion was hiding under it
- D she was afraid her mother might find the book

Answering this question is closely related to answering questions 2 and 3; it requires a global understanding of the story and the interaction between Alice and her mother. Even understanding that Alice was worried about the bed's appearance requires being able to interpret the story in

terms of guilt, wrong-doing and anger. It is fairly easy to infer that Alice hid the book under the coverlet; the final move to comprehending the relationship of this action to her mother requires some real-world knowledge about why people hide things. In a little more detail, the inferential process might proceed as follows:

Fact from story: Alice hid the book under the coverlet.

Real-world knowledge: People hide things so that other people won't find them.

Hypothesis: Alice hid the book so her mother wouldn't find it.

Evidence: Alice hid the book when she heard her mother approaching. (From the beginning, Alice knew who it was, although we didn't). When her mother was in the room, Alice was worried about the bed.

Real-world knowledge: Hiding something means you worry about the other person finding it when they're around.

Conclusion: Alice was afraid her mother might find the book.

The process of inference in texts, even short, fairly simple ones, is clearly complex. None of the questions in this story can be answered without a significant corpus of facts about social situations, human emotion and motivation. Just as essential is knowledge about reference, focus of sentences, and the implications of direct and indirect address. A preliminary exploration of representation in Section G below illustrates both the complexity of the necessary inferences and our preliminary approach to handling them.

F. A "Garden Path" Analysis

To illustrate the use of and need for a detailed process model of text comprehension we will now examine an example of a subject "comprehending" the Alice story. An adult was read the story, asked the questions, and then asked to summarize the episode. The example shows how a single overlooked fact leads to catastrophe in terms of the answers to the multiple choice questions. This observation alone is surprising, but it also nicely illustrates the far reaching consequences that a single piece of data can have in a hypothesis-driven scheme of reading.

The subject answered two out of the five questions "correctly" for a "comprehension" score of 40%. Examining the hypotheses this subject reported in her summary, we found that she had carefully and properly articulated a "garden path" hypothesis (that is, one which is plausible except for some easily-overlooked piece of refuting evidence.)

There was only one linguistically-based mistake: she failed to connect "one girl ..., while the other ... " with the idea of two girls. Therefore in her recall, Alice both hid the book and went to the wardrobe. Like most readers, the subject felt obliged to account for why the book was secret; she assumed that it had to be a diary. The sequencing of hypotheses along the way to comprehension can

sometimes drastically alter the final understanding of the text. This subject paid more attention to Alice and her motives in understanding why she hid the book than do most readers; usually readers think the mother would consider reading the book to be sufficient cause for blame. Also, she reported getting the idea that the voice belonged to Alice's mother because it called (and didn't yell or cry,) and on reading tests "Mothers always call, children always yell." Most subjects would have to wait until Alice addresses her as "Mommy."

Then came the first question. One of the answers has to be right, and who would you read your secret diary to? A doll is safest. Little girls do read to their dolls, and a fantasy world is the safest place for secrets. Since the subject didn't identify "the reader and her listener" with "one girl ..., while the other", the usual path to answering this question was blocked. Therefore she was obliged to rely on a longer chain of more tenuous question-time inferences.

The second question was answered conventionally; as detailed in the last section, Alice hurried to hide the book, so she must have been alarmed.

The third question, beginning "We may infer that," suggested to the subject that further inferences were called for. Having already concluded that Alice was fifteen years

old, she regarded that conclusion as explicitly stated, not inferred. Here again, the supposition that Alice was reading her secret diary figures prominently in the audit trail of steps to the conclusion. Alice could most plausibly be "very much in love" because that would be recorded in her diary, and a girl of fifteen would especially not want her mother to know that.

The fourth question was answered reasonably given the episodic structure set up to answer the first question. This structure says that when her name was called, Alice was reading to "her listener," the doll. The subject chose to describe it as "reading to herself" rather than "reading aloud" because the doll was only being read to in Alice's imagination. "Alice was evidently reading to herself."

The fifth question, like the second, tests the reader's understanding of Alice's fear of discovery. The subject displayed no misunderstanding here.

So a deeper analysis of reading done by the subject revealed much better reading skills than were measured by the five questions. Just one omission crept in when she missed "one girl ..., while the other," possibly because the clause in the ellipsis requires so much processing, possibly because, as she later said, the phrase "the reader and her listener" implied to her that one was capable of talking, while the other was not. The rest of her "troubles" were

all the result of a behavior that actually is part of skilled comprehension, the amalgamation of explicit and implicit information in the narrative.

The multiple-choice design of the test also contributes unnecessarily to the confusion since one of the four sentence completions must be correct, and that sentence is bound to have presuppositions which will get integrated into the reader's overall story interpretation.

Thus, a "wrong" answer for question 1 strengthened the diary hypothesis, which was therefore trusted again in question 3. Her answer to question 4 was based on her answer to question 1. Indeed, from the subject's point of view all of the questions were based on understanding Alice's diary: its audience, its import, its content, and its secrecy. Yet, far from failing to understand the story, the subject demonstrated great skill (if perhaps a little haste) in jumping to conclusions. She "deserved" to have missed only the first question which tested whether the reference to the two girls had been established.

We believe that only by carefully representing the linguistically- and conceptually-based knowledge used in reading to the depth described can we faithfully perceive what skills are involved in reading, where they are absent, and even eventually how and in what order they may be taught. This is a detailed scientific undertaking which

requires the use of a computer to marshal all the relevant information at once. It is one thing to build a speculative blackboard model of the information used in comprehending a single story; it is quite another to design a process with the clarity of attention to find its way through the space of possible reasoning steps to an actual scenario of text comprehension. As we saw with the above example, it is not the end result, but how you get there that counts.

G. Exploring Representation Issues

The development of an improved language for describing comprehension requires major inputs from a variety of sources. Some of the effort must be directed towards gathering and analyzing previous work on representations of knowledge, as in Bruce (1975b). Some must go into informal recall and question-answering experiments of the kind discussed in the previous section, followed later by more rigorous tests. Much of the work is purely of the "pencil and paper" variety, wherein notions of representation, control structure and so on, are examined for adequacy and consistency. This type of work is exemplified in Rubin (1975), Bruce (1972), Nash-Webber and Bruce (1976), and Bobrow and Brown (1976). Finally, much of our work will be done in the context of computer modeling. Later in this section we illustrate the general form of our techniques by means of a tentative (and limited) analysis of one line of

the Alice story.

How can we characterize the diverse knowledge needed for reading so that it can be used by a computer program? How do we make the knowledge explicit so that the resulting model tells us something about reading comprehension? Can the knowledge representation structures be made flexible enough to accommodate varying theories about reading so that they can be compared? Answering these and related questions will be a major focus for our work.

Previous and ongoing work at BBN which deals with various areas connected with language provides us with a powerful set of technical tools. This work includes reliable and established software for handling semantic networks (used extensively in the SCHOLAR system (Carbonell and Collins, 1974; Collins, et al. 1975) and the SOPHIE system (Brown and Burton, 1975)) and for building augmented transition network parsers, as well as techniques for using and building procedural representations. In addition the BBN speech understanding project (see Nash-Webber and Bruce, 1976) has some 50 person-years of experience in dealing with interacting processes. Tools and experience of this kind mean that the design and implementation of our model will not require us to start from the very beginning.

In order to show in a more concrete, albeit simplified manner, what such a model might look like, we will use the

notion of boxes which contain information and point to other boxes. In fact many of these boxes can be regarded as schemata, but they also represent high level control processes, temporary storage locations, etc.

We need to represent in boxes all of the orthographic, phonemic, syntactic, semantic and pragmatic information which might be retained and used by a reader of a text. We also need to represent a substantial amount of knowledge not given by the text, e.g., schemata about people, places, and things, knowledge of speech acts and social actions; knowledge about the context and purpose of the reading task, and so on. Given this knowledge representation we can then attempt to analyze the text, the answering of questions on the text, and recall protocols of the text.

For example, consider the first word in the Alice story:

"Alice!"

A possible box representation for this word (actually, a manifestation of the word, "Alice", which is itself distinct from the concept, <Alice>) is shown in Figure 1. Note that this box becomes meaningful only when we show the boxes it points to. For example, Word1 is a manifestation of "Alice", as shown in Figure 2. The positional significance of "Alice" is indicated by the FirstWordOf pointer. One indication that such information is retained comes from

WORD1	
MANIFESTATIONOF	ALICE
QUOTED?	T
PRECEDINGWORD	NIL
FOLLOWINGWORD	NIL
FOLLOWING PUNCTUATION	[]
FIRST LINE	LINE1
FIRST PARAGRAPH	PARAGRAPH1

Figure 1

ALICE	
INSTANCEOF	WORD
MANIFESTATIONS	WORD1
LETTERS	A, L, I, C, E
PHONEMES	[/'A L əs/]
MORPHEMES	[ALICE]

Figure 2

PERSON1	
INSTANCEOF	PERSON
FIRSTNAME	ALICE

Figure 3

informal recall ~~and~~ have done on the Alice story. In every case the ~~word~~ remembered correctly as the first word.

Now, seeing "Alice" at the beginning of a text, the reader is likely to ~~infer~~ that there is a person, whose (first) name is "Alice". Whether this person is being called to, shouted ~~at~~ just named is not clear without reading more. Still ~~we~~ build the structure shown in Figure 3. Finally, ~~the~~ reader produces structures for the syntax of "Alice!" and ~~for~~ the utterance itself (as opposed to the words making up ~~the~~ utterance).

It should be clear at this point that for a single manifestation of a word there is a lot of information to organize and remember. One thing that helps is that these boxes are highly interconnected, forming a network-like structures as shown in Figure 4. The box labeled "Tree1" is simply the top box for a whole set of boxes representing pertinent syntactic information (e.g. parse trees).

A complete representation of even the first sentence of the Alice story would not be appropriate here. Instead let us assume that the details at the orthographic, phonemic and syntactic levels are given and focus on the conceptual representation, remembering, however, that the interactions across levels may be crucial to comprehension. For the first sentence we might get the conceptual representation

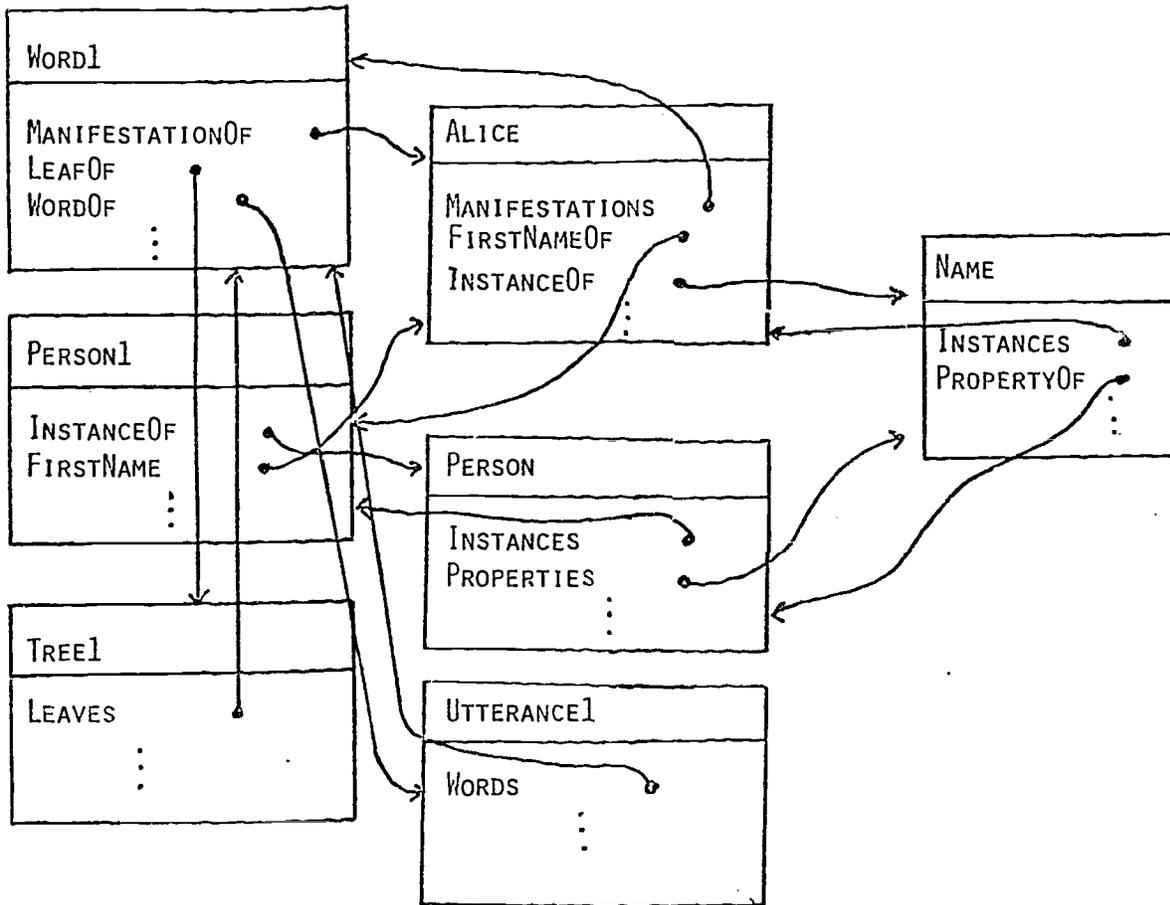


Figure 4

shown in Figure 5. (Reverse pointers are omitted for the sake of clarity). This representation is basically a schema in which an action is assumed to have various slots which need to be filled, the notions of "Mtrans", "Speak", and "Conscious Processor" being taken from Schank, 1975. Note that this representation allows different interpretations of the first sentence. For instance, the voice could be calling to Alice, or merely invoking her name (as in anger at a discovered wrong).

In addition to representations of the text structure (including such immediate inferences as "a voice that can utter 'Alice' probably belongs to a person") there must be representations of relevant world knowledge. For example, the speech act of calling to someone has a number of presuppositions and expectations associated with it which can be used in later structuring of the text. This and similar kinds of knowledge must be readily available for comprehension to occur.

The preceding examples are admittedly sketchy and are intended to show only some of the factors we want to consider in our knowledge representations. Our research will be guided by the demands of actual children's texts and questions such as: 1) Does the model demonstrate how an inference could be made? (2) Can a class of inference failures be described in terms of general features of the

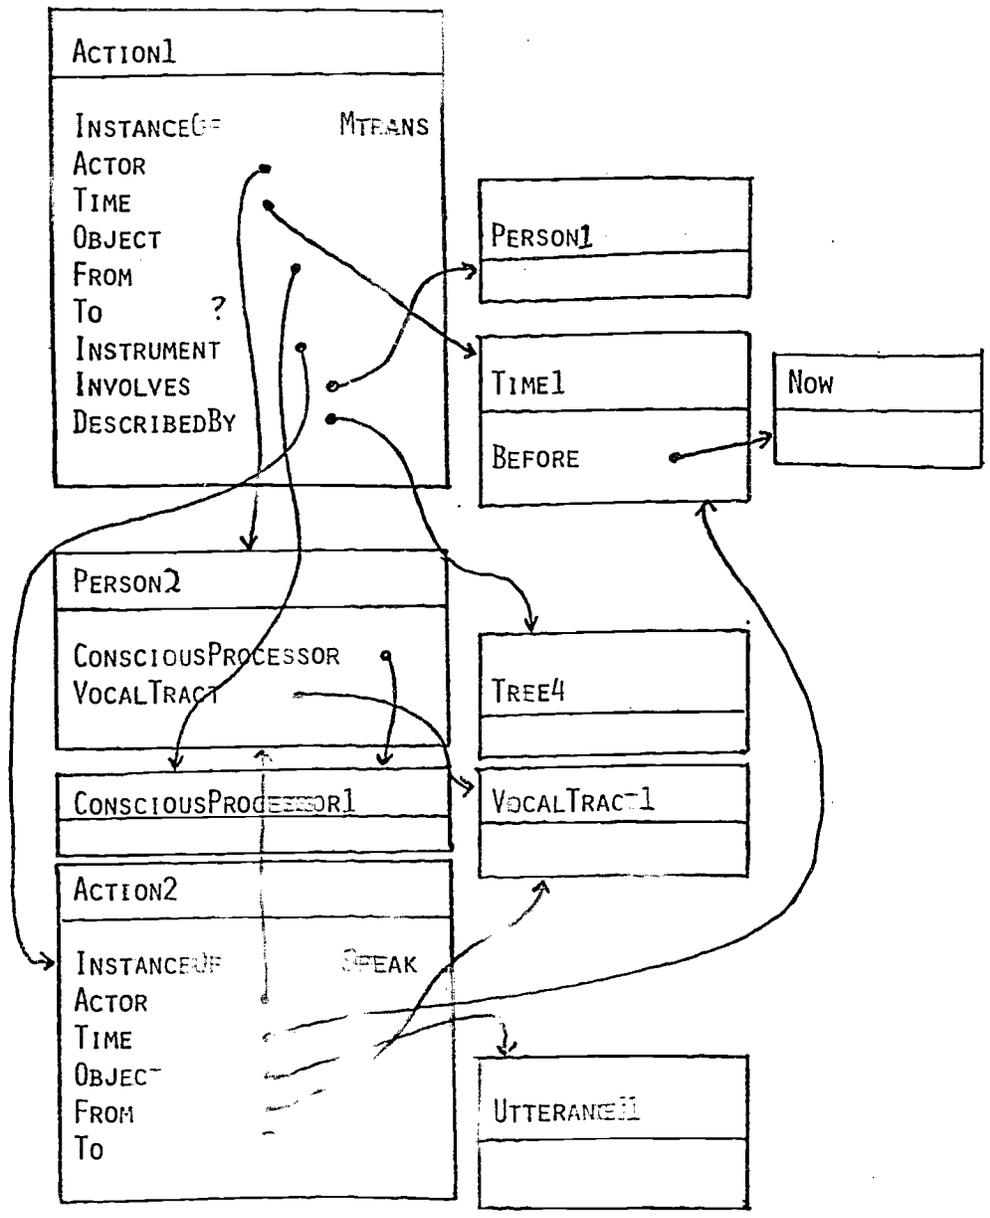


Figure 5

model? (3) Can general features of the model be translated into prescriptions for test and training material design, research procedures, and implications for teaching?

H. Implementation Issues

The first version of our computer model will be used to explore representation and inference control issues which would have a major impact on later versions. We will work with texts selected to share a common body of world knowledge. Thus we will be able to concentrate on general representation issues rather than the specifics of several unrelated texts.

The programs will be written in INTERLISP so that we can quickly incorporate parts of existing programs (e.g., the BBN speech understanding system) which prove useful. A major example in this category is SEMNET, a program which makes it easy to build, change, search through and print out a semantic network.

At first we will use formal representations of the text rather than the raw English. Although both parsing and generation programs are available to us, and could be used at a later date, we feel that the main focus of this programming work ought to be on comprehension problems and not on input/output questions. On the other hand, the formal representation used must allow for expression of

surface syntactic or orthographic information which might interact with comprehension processes.

A sketch of what steps the program should follow do in specifying the difficulty of a test question and what capabilities it is testing is as follows: First a formal representation of a text is read in. Then a structure is built in which some inferences have been made to give coherence to the text. Next the program is asked to answer a question. In the process of answering the question, the program maintains an audit trail which shows just which inferences of each kind were used. This audit trail gives a measure of the difficulty of the question with respect to the text for a given body of world knowledge and inference strategies. Changes in the question, the text, or the stored knowledge can alter the audit trail significantly, and thus show in a precise way the effects of text and task characteristics.

A generalization of the question answering problem is that of text comparison. Given a text and a recalled version of it, the program will apply the same inference rules and knowledge in an attempt to convert the text into its recalled version. Again, the audit trail gives a precise objective measure of the difficulty of the transformation task, and thus, in this case, of the distance between the two versions.

I. Conclusion

Some caution should be exercised in interpreting what we have proposed. Our programming efforts will be directed towards implementing a restricted model which represents selected crucial components. We will incorporate only the knowledge required to handle a few simple texts (e.g., sample test items). For the model to be of more general use would require the incorporation of an enormous amount of world knowledge which is not a realistic undertaking in the foreseeable future. However, once a limited-knowledge version is implemented and working there are several possibilities that could be pursued. For any particular use it could be "primed" with appropriate knowledge as, for example, when one might wish to use it to assist in providing objective scores on recall protocols. It could also be used to handle different texts in the same domain.

Understanding the reading process involves having precise conceptions about the way in which various knowledge sources and critical processes interact. Reading comprehension is a dynamic process; understanding it requires models with dynamic characteristics. The computer is the best way we know of to represent such characteristics, and programs of the kind we propose represent the best way we know of to precisely specify their interactions.

REFERENCES

- Bobrow, R.J. and Brown, J.S. Syntheses, analysis, and contingent knowledge in specialized understanding systems. In D. Bobrow and A. Collins (Eds.), Representation and Understanding: Studies in Cognitive Science. New York: Academic Press, 1975.
- Bormuth, J.R. Comparable cloze and multiple-choice comprehension test scores. Journal of Reading, 1967, 10, 291-299.
- Brown, J.S. and Burton, R.R. Multiple Representations of Knowledge for Tutorial Reasoning. In D. Bobrow and A. Collins (Eds.) Representation and Understanding: Studies in Cognitive Science. New York: Academic Press, 1975.
- Bruce, B.C. Belief Systems and Language Understanding. BBN Report No. 2973, Bolt Beranek and Newman Inc., Cambridge, Mass., 1975a.
- Bruce, B.C. Case Systems for Natural Language. BBN Report No. 3010, Bolt Beranek and Newman Inc., Cambridge, Mass., 1975b.
- Bruce, B.C. A Model for Temporal References and Its Application in a Question Answering Program. Artificial Intelligence, 1972, 3, 1-26.
- Bruce, B.C. and Schmidt, C.F. Episode Understanding and Belief Guided Parsing. NIH Report CBM-TR-32. Computer Science Department, Rutgers University, New Brunswick, N.H. 1974.
- Bruce, B.C. and Harris, G. Procedural Semantics in the Travel System. BBN Report No. 3115, pp. 27-36. Cambridge, Mass: Bolt Beranek and Newman Inc., 1975.
- Carbonell, J.R. and Collins, A.M. Natural Semantics in Artificial Intelligence. Proceedings of Third International Joint Conference on Artificial Intelligence, 1973, pp. 344-351. (Reprinted in the American Journal of Computational Linguistics, 1974, 1, Mfc. 3).
- Collins, A., Warnock, E.H., Aiello, N. and Miller, M.L. Reasoning from Incomplete Knowledge, in D. Bobrow and A. Collins (eds.). Representation and Understanding: Studies in Cognitive Science. New York: Academic Press, 1975.

- Dale, E. and Chall, J. A formula for predicting readability. Educational Research Bulletin, 1949, 27, 11-20, 37-54.
- Fahlman, S.E. A planning system for robot construction tasks. Artificial Intelligence, 1974, 5, 1-50.
- Fodor, J.A., Bever, T.G. and Garrett, M.F. The Psychology of Language: An Introduction to Psycholinguistics and Generative Grammar. New York: McGraw-Hill, 1974.
- Frederiksen, C.H. Effects of context-induced processing operations on semantic information acquired from discourse. Cognitive Psychology, 1975, 7, 139-166.
- Goodman, K.S. Psycholinguistic Universals in the Reading Process, in Smith, F. (Ed.). Psycholinguistics and Reading. New York: Holt, Rinehart and Winston, Inc., 1973.
- Laberge, D., and Samuels, S.J. Toward a theory of automatic information processing in reading. Cognitive Psychology, 1974, 6, 293-323.
- Nash-Webber, B., and Bruce, B. Evolving Uses of Knowledge in a Speech Understanding System. Proceedings of the COLING-76 Conference, Ottawa, Canada, June 1976.
- Newell, A., and Simon, H.A. Human Problem Solving. Englewood Cliffs, NJ: Prentice-Hall Inc. 1972.
- Perfetti, C.A. Language Comprehension and Fast Decoding: Some Psycholinguistic Prerequisites for Skilled Reading Comprehension. Paper presented to the Development of Reading Comprehension, 1975 summer seminar of the International Reading Association, Newark, Delaware, July 1975.
- Rieger, C.J. Conceptual Memory and Inference. In R.C. Schank (Ed.) Conceptual Information Processing. New York: American Elsevier, 1975.
- Rubin, A.D. Hypothesis Formation and Evaluation in Medical Diagnosis. Report AI-TR-316, Artificial Intelligence Laboratory, Massachusetts Institute of Technology, 1975.
- Rumelhart, D.E. Toward an interactive model of reading. In S. Dornic (Ed.) Attention and Performance VI. London: Academic Press, in press.
- Rumelhart, D.E. Notes on a schema for stories. In D. Bobrow & A. Collins (Eds.), Representation and Understanding: Studies in cognitive science. New York: Academic Press, 1975.

Schank, R.C. and the Yale A.I. Project SAM--A Story Understander (Res. Rep. 43). New Haven, Connecticut: Yale University, Department of Computer Science, August 1975.

Schmidt, C. Understanding Human Action: Recognizing the Motives and Plans of Other Persons. Paper for the Carnegie Symposium on Cognition: Cognition and Social Behavior, 1975.

Smith, F. Psycholinguistics and Reading. New York: Holt, Rinehart, and Winston, 1973.