12

KNOWLEDGE REPRESENTATION IN COGNITIVE PSYCHOLOGY

David E. Kieras
University of Arizona

Technical Report No. 7
March 17, 1981

This research was supported by the Personnel and Training Research Programs, Office of Naval Research, under Contract Number N00014-78-C-0509, Contract Authority Identification Number NR 157-423/5-5-78. Reproduction in whole or in part is permitted for any purpose of the United States Government.

UNCLASSIFIED
SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE	READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER Technical Report No. 7	2. GOVT ACCESSION NO. AD-A111 464
	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) Knowledge Representation in Cognitive Psychology	5. TYPE OF REPORT & PERIOD COVERED Technical Report March 17, 1981
	6. PERFORMING ORG. REPORT NUMBER
7. AUTHOR(s) David E. Kieras	8. CONTRACT OR GRANT NUMBER(s) N00014-78-C-0509
9. PERFORMING ORGANIZATION NAME AND ADDRESS Department of Psychology University of Arizona Tucson, AZ 85721	10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS 61153N; RR 042-06 RR 042-06-02; NR 157-423
11. CONTROLLING OFFICE NAME AND ADDRESS Personnel and Training Research Programs Office of Naval Research (Code 458) Arlington, VA 22217	12. REPORT DATE March 17, 1981
	13. NUMBER OF PAGES 46
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office)	15. SECURITY CLASS. (of this report) unclassified
	15a. DECLASSIFICATION DOWNGRADING SCHEDULE

16. DISTRIBUTION STATEMENT (of this Report)

Approved for Public Release; Distribution Unlimited

17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)

18. SUPPLEMENTARY NOTES

Reprinted from: Mathematical Frontiers of the Social and Policy Sciences
Edited by: Loren Cobb and Robert M. Thrall
AAAS Selected Symposium 54.

19. KEY WORDS (Continue on reverse side if necessary and identify by block number)

Knowledge Representation, reading, comprehension

20. ABSTRACT (Continue on reverse side if necessary and identify by block number)

Formalisms for representing human knowledge are currently being developed in cognitive psychology, a branch of human experimental psychology that is concerned with inferring and modelling, on the basis of behavioral data, the processes in complex mental activities such as perception, reading, and problem-solving. These formalisms are based on concepts from graph theory, formal logic, linguistics, computation theory, and artificial intelligence. Models based on such representations

DD FORM 1473 1 JAN 73 EDITION OF 1 NOV 65 IS OBSOLETE
S/N 0102-014-6601

UNCLASSIFIED
SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

usually take the form of large computer simulations in a language such as LISP. Although these representation formalisms have clearly been very fruitful in developing theories and suggesting experiments, cognitive psychology is just now beginning to learn how to deal with the serious problems of the non-identifiability and the non-uniqueness of knowledge representations, a problem similar to the classic "black box" problem in automata theory. But these formalisms, with their emphasis on qualitative, symbolic, complex, discrete structures, as opposed to the earlier approaches using the mathematics of stochastic processes and continuous numeric variables, mark the development of a "new mathematics" in cognitive psychology.

Reprinted from:

Mathematical Frontiers of the Social and Policy Sciences

Published by Westivew Press, Boulder, Colorado 80301

Edited by
Loren Cobb and Robert M. Thrall

AAAS Selected Symposium **54**

David E. Kieras

1. Knowledge Representation in Cognitive Psychology

Abstract

Formalisms for representing human knowledge are currently being developed in cognitive psychology, a branch of human experimental psychology that is concerned with inferring and modelling, on the basis of behavioral data, the processes involved in complex mental activities such as perception, reading, and problem-solving. These formalisms are based on concepts from graph theory, formal logic, linguistics, computation theory, and artificial intelligence. Models based on such representations usually take the form of large computer simulations in a language such as LISP. Although these representation formalisms have clearly been very fruitful in developing theories and suggesting experiments, cognitive psychology is just now beginning to learn how to deal with the serious problems of the non-identifiability and the non-uniqueness of knowledge representations, a problem similar to the classic "black box" problem in automata theory. But these formalisms, with their emphasis on qualitative, symbolic, complex, discrete structures, as

---

This research was supported by the Personnel and Training Research Programs, Office of Naval Research, under Contract Number N00014-78-0-0509, Contract Authority Identification Number NR 157-423. This paper profited from the author's conversations with James Greeno and John Anderson.

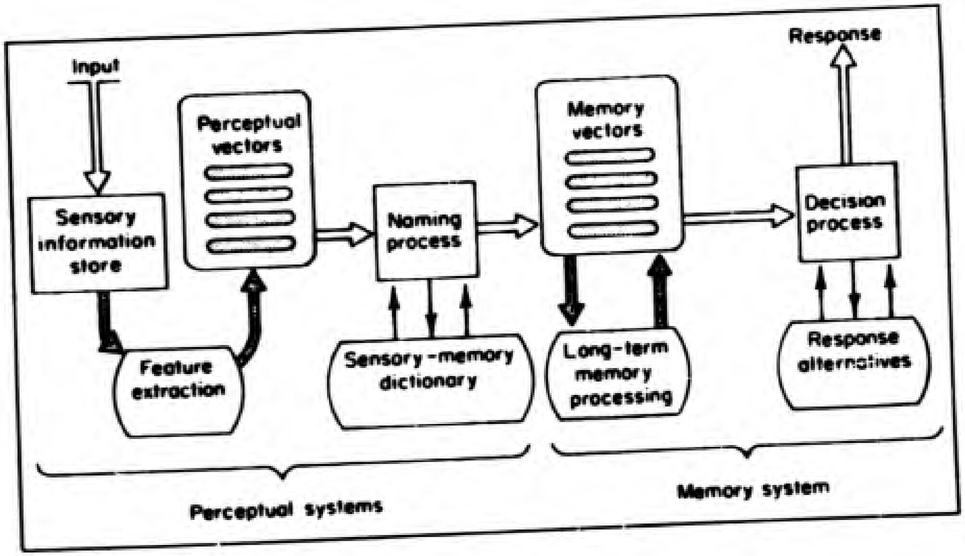

<u>Figure 1.</u> Flowchart for an information-processing model of perception and memory. From D. A. Norman and D. E. Rumelhart, A system for perception and memory. In D. A. Norman (Ed.), <u>Models of Human Memory</u>. Academic Press, 1970. Reprinted by permission.

opposed to the earlier approaches using the mathematics of stochastic processes and continuous numeric variables, mark the development of a "new mathematics" in cognitive psychology.

Introduction

Cognitive psychology is a branch of human experimental psychology that is interested in inferring and modelling, on the basis of behavioral data, the internal mechanisms involved in complex mental activities such as reading and problem-solving. This area has been a distinct discipline within experimental psychology for roughly 15 years. Since this paper is concerned with issues of formal theory, it is important to note that cognitive psychology has always been closely associated with formal approaches to theory, because many of its proponents are also members of the discipline known as "mathematical psychology".

Human Information Processing

The first decade of modern cognitive psychology was largely concerned with models of information processing, many of which were in mathematical form. An example is the model of perception and memory proposed by Norman and Rumelhart in 1970, illustrated in flow-chart form in Figure 1. The human, in performing various perception and memory tasks, was treated as a series of processing stages, with information traveling from one stage to the next, being transformed and manipulated along the way. Hence, information flow diagrams, like that in Figure 1, became a standard way of expressing the overall structure of the human information processing system. These models used probability theory to relate the probability of various types of responses to the stimulus type and the experimental conditions. They did a rather impressive job of predicting and fitting quantitative data from a variety of learning and perception experiments (e.g., Norman & Rumelhart, 1970; Atkinson & Shiffrin, 1968). The primary concern with these models was to represent the processes within the system, with issues of internal structure largely ignored. Consequently, although these theories and models were concerned with information processing, little attention was paid to just what the information being processed consisted of, and how it was represented in the system. Thus, in some of these information-flow models, the information apparently had the consistency of an amorphous fluid to be pumped from one portion of the system to another.

The Representation of Knowledge

The current theories and models have a rather different flavor. There is now a great concern with how information is represented in the human mind, in particular, in memory. An analogy may be drawn to computerized data bases, in which large amounts of information are stored in highly structured ways, with cross-references and retrieval indexes. A similar view arose in cognitive psychology about human long-term memory, which is the permanent store of our general knowledge. The questions were: What does knowledge consist of? What is its format and organization? How can knowledge be represented?

In response to these questions, several theories of knowledge representation are currently being developed. These theories are formal systems, or mathematical systems in the general sense of the term. But they are markedly different from the earlier efforts in mathematical psychology. Instead of being based on probability theory and the mathematics of continuous variables, knowledge representations are qualitative, discrete, highly complex structures of symbols and relations, often represented as graph structures. Usually, computer simulation must be used to explore the properties of these systems, rather than analysis on paper as before.

The Origins of Representation Theories

The basic ideas on the representation of knowledge first appeared during the mid-60's in the field of artificial intelligence, a branch of computer science. The most important early effort was Quillian's (1968) model of semantic memory, the knowledge of word meanings. Figure 2 shows an example of the organization of knowledge as proposed by Quillian. Memory consists of a graph; the nodes represent concepts (type nodes), or references to concepts (token nodes), and the links represent the relationship between the concepts. There are several different types of links, corresponding to logically different types of semantic relationships. Such structures are termed "semantic networks". In Quillian's view, the knowledge of the meaning of a word is the total configuration of connections from that word's type node to all other concepts. Quillian was able to demonstrate how a computer program could retrieve information from such a semantic network and apply it to the task of answering questions in interesting ways.

Around 1970 cognitive psychologists became aware of these developments in computer science (see Norman, 1976,

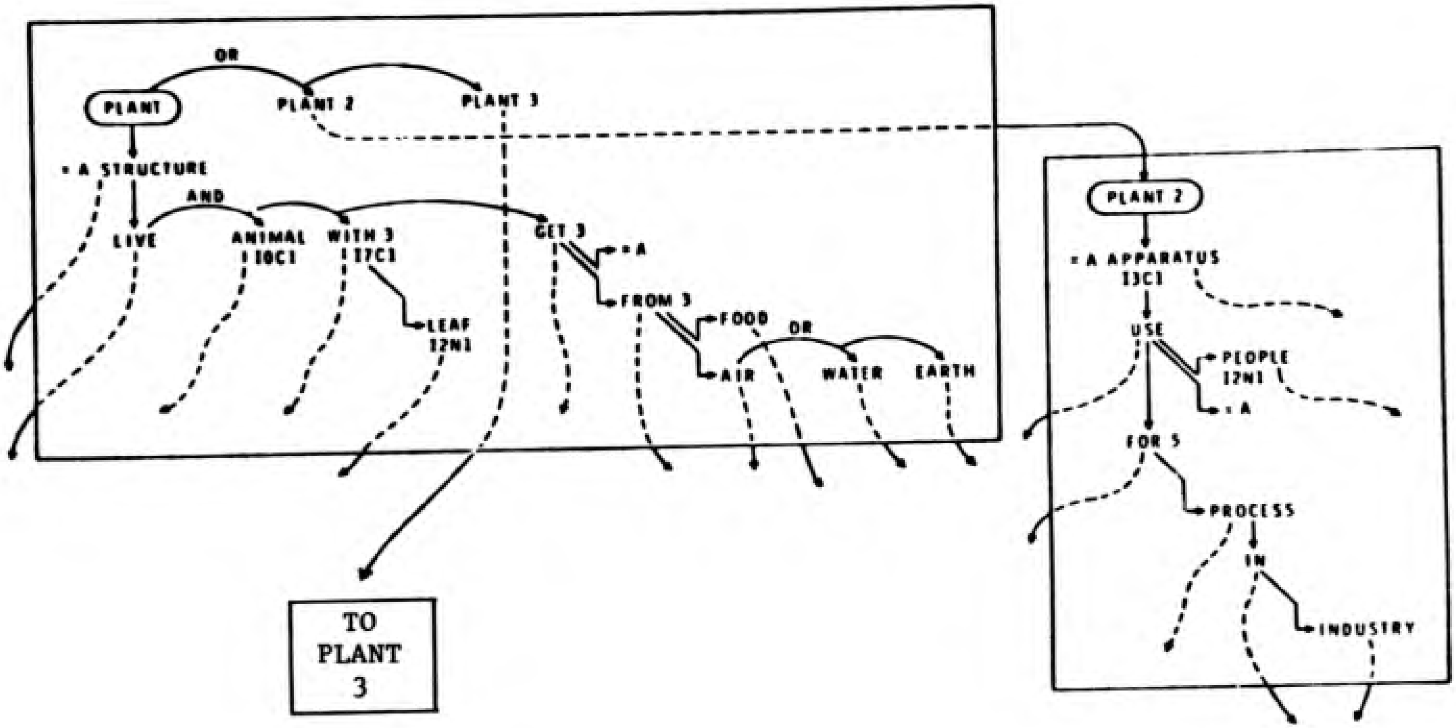

OR
PLANT
PLANT 2
PLANT 3
= A STRUCTURE
AND
LIVE
ANIMAL [0C]
WITH 3 [7C]
GET 3
= A
FROM 3
LEAF [2N]
FOOD
OR
AIR
WATER
EARTH
TO PLANT 3
PLANT 2
= A APPARATUS [3C]
USE
PEOPLE [2N]
= A
FOR 5
PROCESS
IN
INDUSTRY

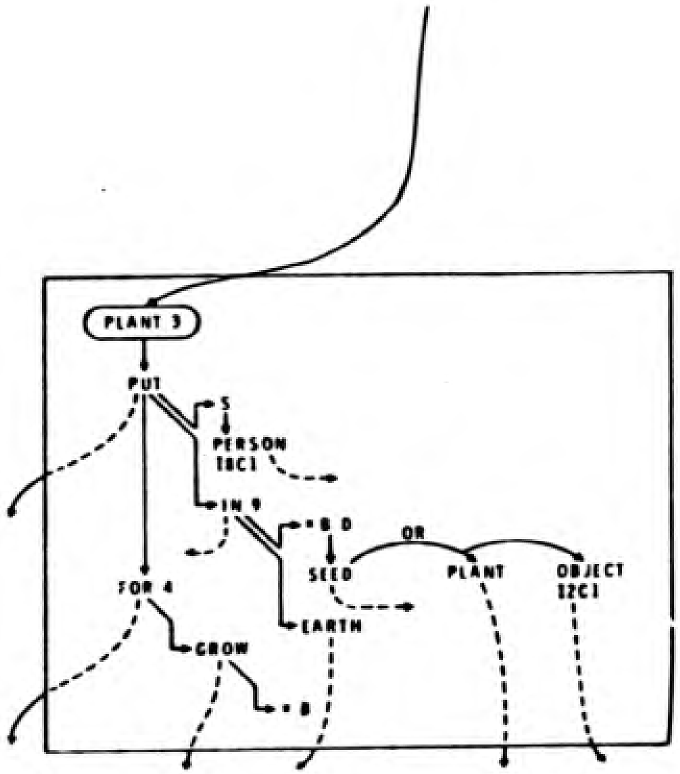

Figure 2. Three definitions of the word "plant" in Quillian's semantic memory system. From M.R. Quillian, Semantic Memory. In M. Minsky (Ed.), Semantic Information Processing, MIT Press, 1968. © by Massachusetts Institute of Technology.

Plant #1: Living structure which is not an animal, frequently with leaves, getting its food from air, water, earth. #2: Apparatus used for any process in industry. #3: Put (seed, plant, etc.) in earth for growth.

for a brief history, and Tulving & Donaldson, 1972, for the first papers). They were interested because semantic networks apparently offered the first approach to the problem of representing our knowledge of word meanings and concepts that was both rigorous and adequate to accounting for our language abilities. That is, the earlier associationistic or stimulus-response approaches either had recognized formal inadequacies (e.g., Bever, Fodor & Garrett, 1968; Fodor, 1965), or seemed to be fundamentally limited in explanatory power (see introduction in Dixon & Horton, 1968). This semantic network idea was easy for information-processing theorists to accept and work with because it had a strong family resemblance to these earlier associationistic and stimulus-response approaches to knowledge and verbal behavior. Thanks to these new ideas of how knowledge could be represented, there has been a very vigorous development in theories of how people understand and use language.

Current Approaches

I will give two examples of some of the systems for the representation of knowledge that were developed by cognitive psychologists in the early 1970's; the descendants of these systems are still being developed. The major theoretical issue in the representation of knowledge has been the detailed rules for the structure of the representation, in particular, the number of the different types of links that are permitted. The first example (Figure 3) is the Rumelhart, Lindsey, and Norman (1972) system. Their system is based on certain ideas from linguistics called case grammar (Fillmore, 1968), in which actions are organized around the verb concept, and there are links for the case categories pointing to the various participants in the action. There are a relatively large number of different link types, due

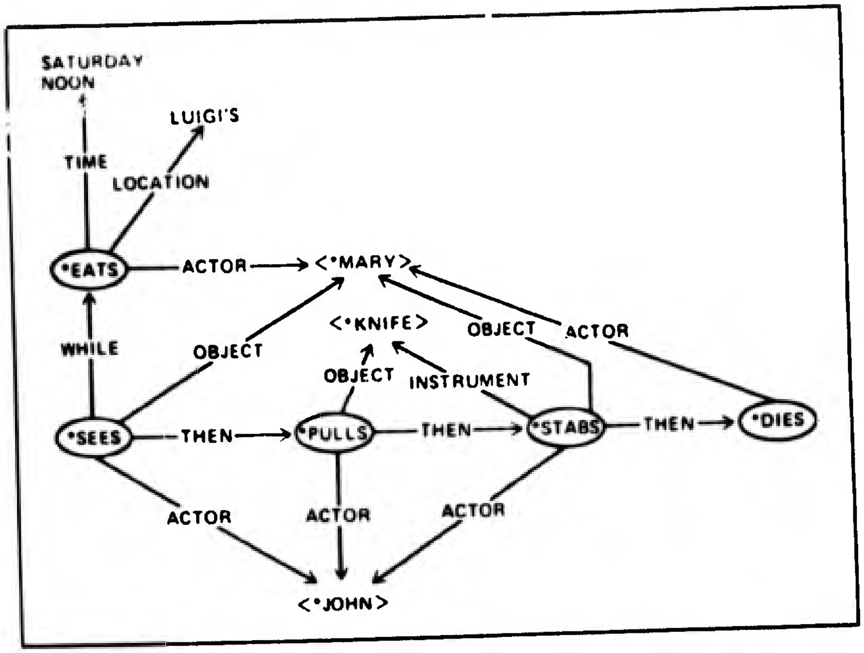

Figure 3. Representation of the event "John murders Mary at Luigi's" in the LNR system. From D.E. Rumelhart, P.H. Lindsay, and D.A. Norman, A process model for long-term memory. In E. Tulving and W. Donaldson (Eds.), Organization of Memory, Academic Press, 1972. Reprinted by permission.

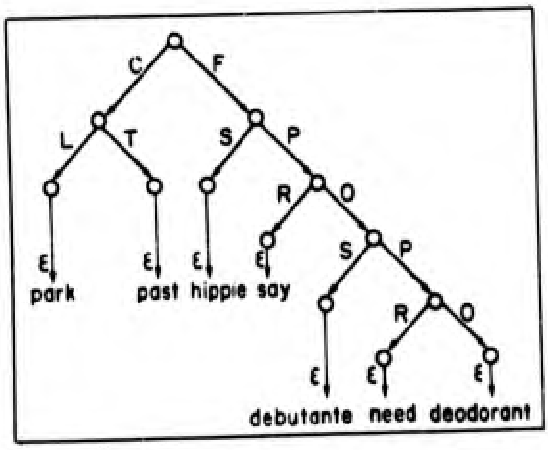

Figure 4. Representation of the event "In the park the hippie said that the debutante needed a deodorant" in the Anderson and Bower HAM system. From J.R. Anderson and G.H. Bower, Human Associative Memory, V.H. Winston (distributed by Wiley), 1973. Reprinted by permission.

to the variety of case categories, but it is also easy to represent fairly complex actions and sequences of actions using these representations. A closely related system was developed by Kintsch (1974), specialized for use as a data-analysis tool in memory experiments. It has been highly successful and of great help in the study of how people understand and remember information presented in the form of prose.

The second example of a knowledge representation system is the Anderson and Bower (1973) system illustrated in Figure 4. Based on formal logic, knowledge is represented as propositions consisting of subject and predicate, with predicates being either simple properties or relation-object pairs. Here there are a relatively small number of connection types; the same information can be represented as in the Rumelhart, Lindsey, and Norman system, but typically more structure would be required. Anderson (1976) has extended this system and developed the formal theory for such networks.

Such graph structures can be represented on a computer as a set of symbols, one for each node. Each symbol has a list of pointers which represent the links. Each pointer consists of a label for the link type and the name of the linked-to node. Programming languages such as LISP make it easy to write simulation programs that can construct, manipulate, or search such networks. A typical simulation model using these knowledge representations, such as that in Kieras (1977), is able to take an English sentence as input, derive its semantic content, locate the relevant portion of the existing memory network, and build new structure at this place to represent the sentence content. Then the system can answer questions about the input information by examining the memory network.

The Impact on Formal Theorization

The use of knowledge representations as part of psychological theory had several effects on cognitive psychology which are of interest to formal theorists. One of the first effects was a new interest in graph theory. Since the knowledge representations were graphs, graph theory should have something to contribute to the theory of memory organization. For example, Kintsch (1972) showed how our knowledge of noun concepts can be considered as a system of set inclusion relations, which form a partial ordering with certain peculiar properties. For example, as shown in Figure 5, dog is a subset of both mammals and pets, which, however, are not subsets of each other. Also, there might be different and mutually exclusive partitions of a noun concept into subsets. An example shown in Figure 6 is that animal may be partitioned into subsets like mammal, bird, insect, reptile, or into a different set of subsets like human, nonhuman, which include the same individual members (Figure 6).

But our initial interest in graph theory was quickly discouraged by the fact that in terms of graph structure, a memory network is essentially an arbitrary graph, with no real constraints on what could be connected ultimately to what. What this lack of constraint means is that most of the important results in graph theory are either inapplicable or uninteresting for the study of memory organization (cf. Anderson, 1976). Consequently, cognitive psychologists are no longer very interested in graph theory.

A second effect of the arrival of knowledge representation theories was that we became interested in other formal tools besides graph theory that dealt with qualitative or symbolic structures. We began to study formal logic, abstract

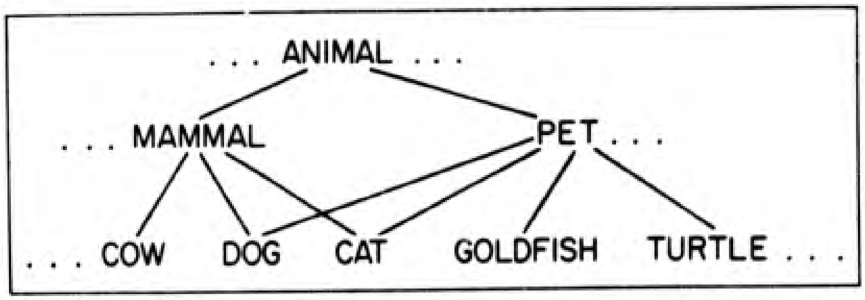

Figure 5. An example of how nodes in the memory graph can be cross-connected, as when concepts are subsets of more than one superordinate. After W. Kintsch, Notes on the structure of semantic memory. In E. Tulving and W. Donaldson (Eds.), Organization of Memory. New York: Academic Press, 1972. Reprinted by permission.

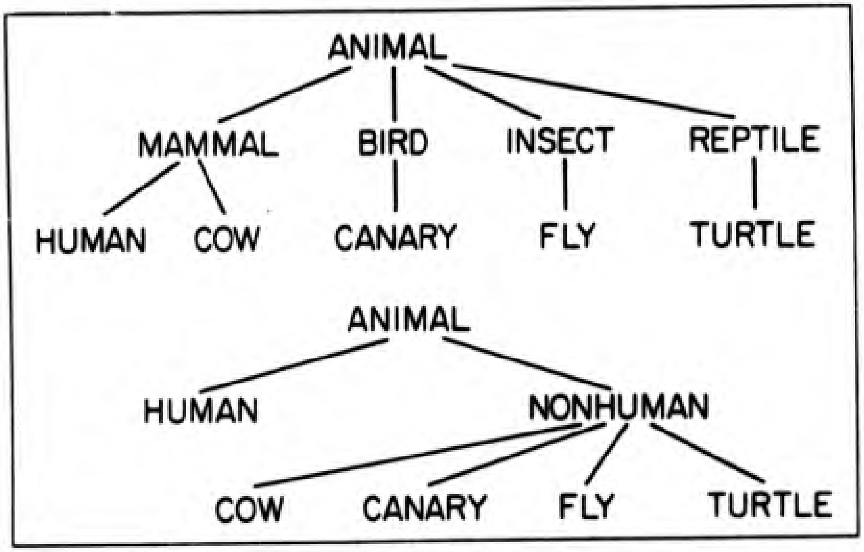

Figure 6. An example of how concepts can be partitioned in more than one way, showing that there may be alternate graph structures for concepts in memory. From E. Tulving and W. Donaldson (Eds.), Organization of Memory. New York: Academic Press, 1972. Reprinted by permission.

algebra, and in particular, computation theory and automata theory (e.g., Minsky, 1967). Since knowledge representations were qualitative, rather than quantitative, structures, these branches of mathematics seemed especially useful. Some examples of early work using these tools to attack problems in theoretical psychology are Suppes (1969), Kieras (1976), and Groen (1974). An example of the recent use of such tools appears in Anderson's (1976) model of cognition, which combines a semantic network for representing knowledge with a production system for representing processes. He provides a formal analysis of the representational power of the model and its computational power as well. The results show that the system is powerful enough to be capable in principle of representing any cognitive process likely to be of interest. However, Anderson's work is not typical of the overall field of cognitive theory. Such powerful analyses are still the exception, rather than the rule.

A third effect of the knowledge representation ideas was a great increase in the use of computer simulation as a theoretical tool. A theory based on complex knowledge representations is usually too complicated to allow characterizations or predictions of behavior to be derived on paper or with common sense. Hence, a computer must be used to simulate the theory. While a computer simulation is a formal system, it differs rather drastically in flavor from the previous conventional mathematical techniques. A major difference concerns evaluating models (see Hilgard & Bower, 1975, p. 484). Mathematical psychologists have been accustomed to rigorous statistical methods of comparing a model's predictions to observed data. However, most computer simulation models of this type operate at a basically qualitative level, and are usually deterministic. Thus many of the earlier comparison techniques are simply inapplicable. Consequently, there is a serious lack of generally accepted methods for testing a simulation against experimental data. Until it is clear that the goodness of fit of a simulation model can be rigorously tested, the full potential of this tool will go unexploited and there will be questions about its value.

The Impact on Experimentation

The most sweeping effect of the theories of knowledge representation was the instigation of a vigorous search for empirical confirmation of the theories. A great deal of experimental work was inspired by the possibility of demonstrating experimental effects of how people's knowledge was organized and structured. The dominant theme of these experiments was attempting to determine what kind of knowledge representation people actually had in their heads.

The Collins and Quillian Experiment

The seminal experiment that started the search for internal structure was that of Collins and Quillian (1969) on retrieval from semantic memory. They presented simple sentences such as "A canary is a bird" to subjects who judged them as being true or false. The time required to make the judgement was the dependent variable. The hypothesis was that the subject's memory structure was hierarchical and followed a principle of "cognitive economy". That is, facts were stored at the most general possible node in the hierarchy, as shown in Figure 7. Hence at the node for canary is stored properties of canaries, but only those applying to canaries exclusively. The properties of birds are stored at the more general level, at the node for bird. But the properties of birds that are shared by animals in general are stored only at the node for animals. Such a principle eliminates duplicate storage of information, and so is more economical of memory space. However, it requires a time-consuming use of inference to recover the information not explicitly represented. For example, in order to verify that canaries have skin, the system would have to work up the hierarchy until it arrived at the animal level where the property in question is stored. Hence, the time to verify sentences should depend on the distance in the hierarchy that would have to be covered to locate the required information. Collins and Quillian obtained results that verified this prediction; for example, it takes longer to verify that a canary is an animal than that a canary is a bird; their results are shown in Figure 8. It should be noted that there is now much more known about performance in these memory retrieval tasks, which qualifies the simple picture presented here.

Assumptions Underlying the Experimentation

The Collins and Quillian experiment indicated that the structure of memory could be determined by examining experimental data, expecially reaction times. Many experiments were performed as a result; they rested on two major assumptions, which were usually implicit. The first is basically methodological; the second concerns the problem of identifiability, and is subject to formal treatment.

Assumption 1 is that the experimenter has control over the knowledge structures of the subjects. By having a subject learn material with a known organization, a known knowledge structure will be installed in the subject's head.

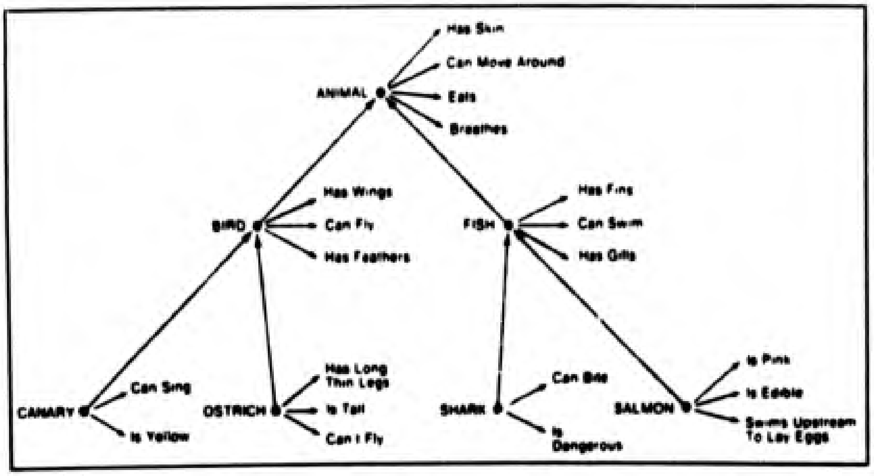

Figure 7. The organization of semantic memory for knowledge about animals, as postulated by Collins and Quillian. From A. Collins and M.R. Quillian, Retrieval time from semantic memory. Journal of Verbal Learning and Verbal Behavior, 1969, 8, 240-247. Reprinted by permission.

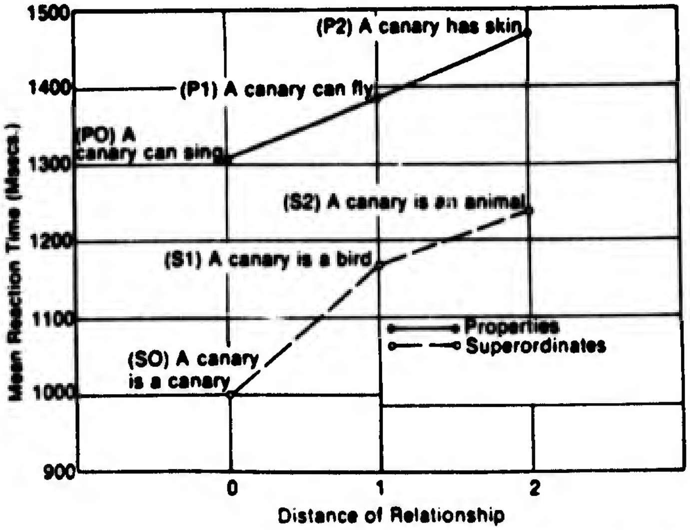

Figure 8. Results of the Collins and Quillian experiment, showing an increase in time to verify sentences as a function of the distance in the hypothesized knowledge structure. From A. Collins and M.R. Quillian, Retrieval time from semantic memory. Journal of Verbal Learning and Verbal Behavior, 1969, 8, 240-247. Reprinted by permission.

Alternatively, the organization of a piece of general knowledge that subjects already had in their heads could be determined by logical analysis on the part of the experimenter, as was done by Collins and Quillian.

This assumption has been recognized as a serious source of error. Subjects in fact often reorganize information in order to learn it, often in idiosyncratic ways. In fact, Hayes-Roth (1977) has described how the organization of knowledge can change radically with practice in retrieving it. Furthermore, the ordinary person can have his or her knowledge organized in ways rather different from the "logical" organization that might be deduced a priori by the experimenter. An example provided by Collins and Quillian (1972) is that it appears that ordinary people tend to use the word animal in a way corresponding to the technical use of mammal, and so the logical organization in which mammals are a proper subset of animals does not in fact appear in the knowledge structures of ordinary people.

Hence, an experimenter can not be confident of what people have for knowledge structures, either beforehand or as a result of experimental treatments. Experimenters currently attempt to plan experiments on the basis of normative information on what the typical subject's knowledge is, or try to avoid having their conclusions contingent on the specifics of the knowledge structures of individual subjects.

Assumption 2 is that the internal knowledge representations of subjects can be determined by examining behavioral data.

There is an immediate problem with this assumption which unfortunately was not readily recognized. As an illustration of this problem, suppose that we observe that a subject's response time is different after two different training or learning conditions. Do we attribute this effect to the subject's having two different knowledge representations, or to the subject's having two different processes which are operating on the same representation? Logically, it could be either structure differences, process differences, or both. But the tendency has been to attribute experimental effects to differences in the representations rather than in the processes. There were several reasons for this: We were more interested in knowledge representations than in internal processes; we had well developed theories of representation in which knowledge was represented as combinations of primitive or simple elements, but we had no corresponding theories of processes; and finally, it was apparently very difficult to control a subject's strategies in an experiment, but according to Assumption 1, we could control the subject's knowledge structure.

What Were the Real Representations?

Under the influence of these assumptions, we believed that an important task was to determine, on the basis of behavioral data, which theory of knowledge representation was the correct, or at least the most accurate, theory. Since it was assumed that the internal knowledge representations of subjects were distinguishable in their behavior, all that was required was to examine the relevant features of reaction time or recall data collected under the appropriate conditions, and a choice between the competing theories of representation could be made.

Some examples of experiments claiming to distinguish knowledge representations on the basis of behavioral data are: (a) Thorndyke and Bower's (1974) attempt to show that the memory representations for simple sentences were more like the Anderson and Bower representations than like the Rumelhart, Lindsey, and Norman representations; (b) Foss and Harwood's (1975) effort to show that certain Gestalt-like qualities of sentence memory could not be accounted for by Anderson and Bower representations; (c) a rather active and continuing debate over whether memory information is always represented in one of the propositional representations such as a semantic network, or can be represented as mental images with some sensory properties (for various sides of the debate see Simon, 1972; Bower, 1972; Pylyshyn, 1973; Shepard, 1975; Paivio, 1976; Anderson, 1978; and Kieras, 1978). Hence the influence of Assumption 2 is rather widespread and is still prevalent.

The Identifiability Problem

While Assumption 2 has clearly led to a lot of research, formal analysis shows that it is completely wrong; behavioral data is simply not strong enough to determine internal structure. Or in other words, internal structure is not identifiable from behavioral data. It should not be necessary to demonstrate this in formal terms, because it follows from the logic of scientific explanation as usually presented in the philosophy of science. As everybody knows, there is any number of theories that predict a given result. Hence an empirical result does not uniquely determine the correct theory.

A related problem is that while one version of a representation theory might be disconfirmed, a different version of the same general theoretical ideas might succeed in predicting the results. Hence a particular set of disconfirming results may leave a generally stated theory intact. This problem is serious because many of the debates over representations have assumed only partially specified theories, or only one of the many possible full specifications of a representation theory. Since only a well-specified theory is testable enough to be falsified by data, when disconfirming results have been obtained, it is not clear just what theory, or what versions of an informal theory of representation, have actually been refuted. Hence many of the debates over representations, such as the imagery-proposition debate, have produced no firm conclusions.

A Formal Approach to the Identifiability Problem

A formal way to show that internal mechanisms are not identifiable from behavioral data is with automata theory, which is the mathematical theory of abstract machines (see Minsky, 1967, or Arbib, 1969, for introductory treatments). An automaton is a system which produces outputs that are a function of the current input and an internal state; the internal state then changes as a function of current input and current state. If the set of possible inputs, outputs, and states are finite, the system is termed a *finite state machine* or *finite automaton*. Since the input-output behavior of the machine may change due to changes in the internal state, the behavior of the system has to be considered in terms of a behavior *history* consisting of sequences of outputs. In other words, a finite-state machine has *memory*, meaning that its response to a particular input may well depend on the previous sequence of inputs.

The finite automaton provides an elegant formal characterization of a behaving organism if we make the altogether reasonable assumption that since the brain is made up of a finite number of particles, it can attain only a finite (although very large) number of states. This characterization makes certain long-standing metatheoretical problems in psychological theory easy to attack, such as the formal status of stimulus-response theory (Suppes, 1969; Kieras, 1976).

But most importantly for the present metatheoretical problem, automata theory is rich in results on identifiability. The main result is that a finite set of behavioral data, that is, input-output sequences of finite length, does not determine which automaton produced the data. For any automaton that one might propose to account for a set of finite input-output sequences, there are an infinite number of other automata that could produce the same behavior. These other automata fall into two classes: those that are equivalent in behavior, and so would produce the same behavior as the proposed automaton over sequences of infinite length, or those that are not in fact behaviorally equivalent, but happen to appear equivalent over the finite sequences of behavior available. Note that in both cases, these other automata might be radically different in structure from the proposed automaton, but yet produce the same behavior.

Hence, all that a set of behavioral data does is determine which infinitely large classes of automata the actual system might belong to. It does not by any means provide a unique characterization of the internal mechanisms of the behaving system. Since the different theories of representation are basically proposals on what kind of finite automaton the human brain is, these results make it clear that we can not arrive at a formally defensible definite conclusion in favor of a single representation theory on the basis of behavioral data.

Anderson's Argument. Anderson (1976, 1978) has presented a more limited form of the nonidentifiability argument, in which the relationship between two competing theories of representation is made more explicit. Since a cognitive theory contains both a theory of internal representations, and a theory of the processes that operate on the representations, differences in representation can be compensated for by differences in process, to yield the same input-output behavior. Given the lack of physiological data that provides an independent set of hard constraints on the possible processes that might go on in the brain, we are left with the situation that it is possible to devise radically different theories of representation that can account for behavioral data equally well.

When Can a System be Identified? One might wonder under what conditions a finite automaton can be identified from behavioral data. This is one of the earliest issues addressed in automata theory (Moore, 1956), and thus has some clear-cut answers (see Arbib, 1969). These are: (a) identification of the class of automata equivalent to the actual automaton producing the behavior is possible if behavior sequences of infinite length are available. This is obviously not a practical solution. (b) If an upper bound on the number of internal states that the unknown automaton might have is known, then behavior sequences of a finite length determined by this bound will contain all the information necessary to specify the equivalence class of the unknown automaton. It seems clear, however, that cognitive psychologists are not in the possession of such an upper bound on the number of states in the brain. In any case, such a bound would be so large that the required behavior sequences would be so many in number, and so long, that knowing the bound would have little practical value. (c) Even if a complete set of behavioral data is available, the identification that can be achieved is not that of the actual automaton producing the behavior, but rather of the equivalence class that contains the actual automaton. Now, for any class of behaviorally equivalent automata, there exists a single unique (except for state labels) minimal automaton that also produces the same behavior, but does so with the fewest number of states. But we have no way to determine from behavioral data whether we are observing this unique minimal automaton or one of its behavioral equivalents. Hence, from a formal point of view, behavioral data is really rather weak as a constraint on the possible internal mechanisms.

A Concrete Example

An example will reinforce the point. Suppose you are given an actual black box that has several input buttons and several output indicators that is built to behave like some finite automaton. Such a box can be built by any electronics hobbyist. Your problem is to "figure out the box". Suppose that what we mean by this goal is to simply infer the rules governing the box's behavior. Then what you are trying to do is to arrive at a description of an abstract machine that is behaviorally equivalent to the box; you are not concerned with whether the structure of this inferred automaton corresponds in any direct way with the internal structure of the box.

From the above discussion, it should be clear that you would not be able to infer this machine with certainty unless you have either an infinite amount of time, or another source of information that provides the upper bound on the number of states. If you lack this information, you would simply follow some form of hypothesis testing heuristic and try to arrive at an adequate characterization of the rules governing the box's behavior. In my informal experiments with just such a box, most people find the task surprisingly difficult for even a very simple machine. Their hypotheses tend to recapitulate the history of psychological theory, assuming first simple stimulus-response connections and then adding internal "thought" or "memory" processes only when definitely required.

Do Reaction Times Help? It could be argued that for real, as opposed to abstract, machines, reaction time data can be used to solve the identifiability problem. For example, two competing theories of representation might predict different times for the system to respond to certain stimulus conditions. Although all cognitive psychologists are agreed that such reaction-time data are extremely valuable and useful in theorizing, it still fails to solve the identifiability problem, despite some arguments to the contrary (Pylyshyn, in press; Hayes-Roth, in press). That is, any assertions about the response time of a finite-state system rest unavoidably on assumptions about the internal "hardware" components of the system and their timing characteristics. Before the overall system timing can be specified, the components must be known, and their timing properties specified. In the lack of such information, any attempt to identify the internal mechanism on the basis of reaction time information may be very useful theoretically, but has in fact no formal basis whatever. Rather, such an exercise is simply another one of our many informal approaches to trying to develop and test theories.

To return to the black box example, suppose your goal was more ambitious than the prior one, being to infer the actual internal mechanisms. That is, you want to determine the circuitry inside the box at some reasonable level of analysis, such as a block diagram of the logic circuits. The reaction time of the box would appear to be very relevant. By using appropriate instruments, you could measure to the nearest nanosecond the delays between input and output as the box responded to inputs under various conditions. Would this help you determine the internal mechanism of the box? In a way it would, because you could entertain much more detailed and interesting hypotheses than you could in the absence of such information. But in fact, the timing data would not tell you much of anything definite about the actual internal mechanism.

If the times fell into certain ranges, you could rule out certain possibilities. For example, if only microsecond timings were observed, it would rule out electromechanical devices as plausible internal components. However this same time range is consistent with many other varieties of hardware components, such as vacuum tubes, discrete solid-state components, or integrated circuits.

But identifying the type of components used in the box is hardly the whole problem. At least as difficult is the question of how the components are interconnected. As any electronics hobbyist would know, there are many different ways a logic circuit can be constructed to produce the same input-output function. For example, the black box could be built so that it first makes the state change, and then generates the output, or it could determine and generate the output first, followed by the state change, or it could generate the output and the state change simultaneously. Since only the output indication is seen, how could you determine when the state change was made? How could you determine even this gross feature of the block diagram for the box?

Does Parsimony Help? It could be argued that the principle of parsimony is a major constraint that eliminates many possibilities. Hence we could assume that the black box contained the simplest circuit that would produce its behavior in the observed times. But what is "simplest" for a logic circuit? But the clincher that completely negates parsimony as a strong constraint is that you have no justification whatever to assume that the builder of the box used the simplest design, or even a very good design. My black box, for example, uses a rather crude design determined by the fact that I had a great many old RTL gates and diodes on hand, but few other components. Hence parsimony also has no formal status as a constraint, but is only another of our informal guides for theory construction.

The Brain as a Black Box

The above argument can be applied to the task of trying to understand the black box that is the human brain. If we use only behavioral data, and attempt to understand only the rules governing behavior, we can not solve the problem with certainty. If we are more ambitious, and attempt to determine what actual mechanisms (in terms of block diagrams) are inside the brain, we could arrive at definite conclusions only with detailed information on the neurological components involved. But even then, in the absence of constraints on the possible neurological designs for the brain, we could not be confident of our conclusions. Parsimony is of no value as a constraint, especially since biological evolution is notorious for using components on hand rather than possibly simpler designs. Finally, reaction time data, in the absence of knowledge of the timing characteristics of the neurological components, may be interesting and useful to our hypothesis testing, but it adds no additional definite constraints on the possible internal mechanisms.

From these arguments, it should be clear that internal representations can not be uniquely identified from behavioral data. Unfortunately, many in cognitive psychology have not understood or recognized this basic limitation on cognitive theory. Apparently their problem lies in a basic confusion on the logical status of various sources of knowledge used in the process of constructing psychological theories. Informal constraints such as parsimony can guide theorizing, but are not necessarily true of the system under study. On the other hand, observed behavior is a definite property of the system under study, but is inherently ambiguous as a constraint on the processes producing it.

Where Do We Go From Here?

Some of those who subscribe to the nonidentifiability argument have despaired, concluding that behavioral approaches to cognitive psychology are essentially worthless. However, others, such as Anderson (1976, 1978), argue that the effort to determine what kind of representation people really use was misguided to start with. Instead of attempting to discriminate between theories that are not behaviorally distinguishable, an intrinsically futile undertaking, we should be simply trying to devise some good models for interesting cognitive processes that can account for a large variety of data. As Anderson points out, we have in fact made very little progress towards this seemingly modest goal. With regard to choosing a knowledge representation, all that is required is that the representation system used in a model meet two criteria: (1) it should be adequately powerful to represent the required information; (2) the representation system should yield an intelligible and intuitively appealing theory. Just how the representation system achieves these criteria is then not important as far as ability to explain behavioral data is concerned.

Let me give an example of this attitude from my own work. I am attempting to explain how readers identify and make use of the topic of a paragraph while they are reading. While it is possible to simply collect a large variety of behavioral data on this subject, a simulation model using knowledge representations could provide a useful summarizing tool for the data, and more importantly, provide a comprehensive explanatory model for the topic identification processes. The simulation model that I am developing for these results uses memory representations similar to the Anderson and Bower type that I described. A different representation system would probably work just as well; my choice of the Anderson and Bower system is basically one of personal preference. I am more familiar with it, and have an intuition that it will be easier to use than one of the others. Since which representation is actually used by people is not behaviorally decidable, I am free to make the choice on these important, but essentially subjective, grounds. So, rather than attempting to support one theory of representation over the others, the goal is to arrive at a model that explains the interesting and important features of the comprehension process in terms of internal processes and knowledge representations. This in itself can be remarkable difficult.

Conclusion

With the knowledge representation theories, cognitive psychology has embarked on a new mathematical approach in its theories that is based on qualitative formalisms which are best expressed as computer simulation programs written in a general symbol-manipulating language. Good application of formal techniques in these theories is just beginning, but an important result of formal analysis is that there is a basic limitation to cognitive theory: Internal representations are not uniquely identifiable using behavioral data. What we need in the way of mathematics to pursue these new theories is more sophistication in the formal analysis of our models, and good ways to test their empirical adequacy.

References

Anderson, J.R. Language, memory and thought. Washington: Lawrence Erlbaum Associates, 1976.

Anderson, J.R. Arguments concerning representations for mental imagery. Psychological Review, 1978, 85, 249-277.

Anderson, J.R. Further arguments concerning mental imagery: A response to Hayes-Roth and Pylyshyn. Psychological Review, 1979, in press.

Anderson, J.R., & Bower, G.H. Human associative memory, Washington: Winston, 1973.

Arbib, M.A. Theories of abstract automata. Englewood Cliffs, N.J.: Prentice-Hall, 1969.

Atkinson, R.C., & Shiffrin, R.M. Human memory: A proposed system and its control processes. In R.W. Spence & J.T. Spence (Eds.), Advances in the psychology of learning and motivation, Vol. 2. New York: Academic Press, 1968.

Bever, T., Fodor, J., & Garrett, M.A. A formal limitation of associationism. In T.R. Dixon & D.L. Horton (Eds.), Verbal behavior and general behavior theory. Englewood Cliffs, N.J.: Prentice-Hall, 1968.

Bower, G.H. Mental imagery and associative learning. In L.W. Gregg (Ed.) Cognition in learning and memory. New York: Wiley, 1972.

Collins, A.M., & Quillian, M.R. Retrieval time from semantic memory. Journal of Verbal Learning and Verbal Behavior, 1969, 8, 240-247.

Collins, A.M., & Quillian, M.R. How to make a language user. In E. Tulving & W. Donaldson (Eds.), Organization and memory. New York: Academic Press, 1972.

Dixon, T.R., & Horton, D.L. (Eds.), Verbal behavior and general behavior theory. Englewood Cliffs, N.J.: Prentice-Hall, 1968.

Fillmore, C.J. The case for case. In E. Bach & R.T. Harms (Eds.), Universals in linguistic theory. New York: Holt, Rinehart & Winston, 1968.

Fodor, J.A. Could meaning be an r_m? Journal of Verbal Learning and Verbal Behavior, 1965, 4, 73-81.

Foss, D.J., & Harwood, D.A. Memory for sentences: Implications for human associative memory. Journal of Verbal Learning and Verbal Behavior, 1975, 14, 1-16.

Groen, G.J. The representation of sequential concepts as locally testable events. Behavioral Science, 1974, 4, 238-253.

Hayes-Roth, B. Evolution of cognitive structures and processes. Psychological Review, 1977, 84, 260-278.

Hayes-Roth, F. Distinguishing theories of representation: A critique of Anderson's "Arguments concerning mental imagery". Psychological Review, 1979, in press.

Hilgard, E.R., & Bower, G.H. Theories of learning (4th ed.). Englewood Cliffs, N.J.: Prentice-Hall, 1975.

Kieras, D. Finite automata and S-R models. Journal of Mathematical Psychology, 1976, 13, 127-147.

Kieras, D. Problems of reference in text comprehension. In M. Just & P. Carpenter (Eds.), Cognitive processes in comprehension. Hillsdale, New Jersey: Lawrence Erlbaum Associates, 1977.

Kieras, D. Beyond pictures and words: Alternative information-processing models for imagery effects in verbal memory. Psychological Bulletin, 1978, 85, 532-554.

Kintsch, W. The representation of meaning in memory. Washington: Lawrence Erlbaum Associates, 1974.

Minsky, M.L. Computation: Finite and infinite machines. Englewood Cliffs, N.J.: Prentice-Hall, 1967.

Moore, E.F. Gedanken-experiments on sequential machines. In C.E. Shannon & J. McCarthy (Eds.), Automata Studies, Princeton, N.J.: Princeton University Press, 1956.

Norman, D.A. Memory and attention (2nd ed.). New York: Wiley, 1976.

Norman, D.A., & Rumelhart, D.E. A system for perception and memory. In D.A. Norman (Ed.), Models of human memory. New York: Academic Press, 1970.

Paivio, A. Images, propositions, and knowledge. In J.M. Nicholas (Ed.), Images, perception, and knowledge. The Western Ontario Series in the Philosophy of Science. Dordrecht: Reidel, 1976.

Pylyshyn, Z.W. What the mind's eye tells the mind's brain: A critique of mental imagery. Psychological Bulletin, 1973, 80, 1-23.

Pylyshyn, Z.W. Validating computational models: A critique of Anderson's indeterminacy of representation claim. Psychological Review, 1979, in press.

Quillian, M.R. Semantic memory. In M. Minsky (Ed.), Semantic information processing. Cambridge: M.I.T. Press, 1968.

Rumelhart, D.E., Lindsey, P.H., Norman, D.A. A process model for long-term memory. In E. Tulving & W. Donaldson (Eds.), Organization and memory. New York: Academic Press, 1972.

Shepard, R.N. Form, formation, and transformation of internal representations. In R.L. Solso (Ed.), Information processing and cognition. Washington: Lawrence Erlbaum Associates, 1975.

Simon, H.A. What is visual imagery? An information-processing interpretation. In L.W. Gregg (Ed.), Cognition in learning and memory. New York: Wiley, 1972.

Suppes, P. Stimulus-response theory of finite automata. Journal of Mathematical Psychology, 1969, 6, 327-355.

Thorndyke, P.W., & Bower, G.H. Storage and retrieval processes in sentence memory. Cognitive Psychology, 1974, 6, 515-543.

Tulving, E., & Donaldson, W. (Eds.) Organization and memory. New York: Academic Press, 1972.

Navy

1 Dr. Ed Aiken
Navy Personnel R&D Center
San Diego, CA 92152

1 Dr. Robert Breaux
Code N-711
NAVTRAEQUIPCEN
Orlando, FL 32813

1 Dr. Richard Elster
Department of Administrative Sciences
Naval Postgraduate School
Monterey, CA 93940

1 DR. PAT FEDERICO
NAVY PERSONNEL R&D CENTER
SAN DIEGO, CA 92152

1 Dr. John Ford
Navy Personnel R&D Center
San Diego, CA 92152

1 Dr. Henry M. Halff
Department of Psychology,C-009
University of California at San Diego
La Jolla, CA 92093

1 LT Steven D. Harris, MSC, USN
Code 6021
Naval Air Development Center
Warminster, Pennsylvania 18974

1 CDR Charles W. Hutchins
Naval Air Systems Command Hq
AIR-340F
Navy Department
Washington, DC 20361

1 CDR Robert S. Kennedy
Head, Human Performance Sciences
Naval Aerospace Medical Research Lab
Box 29407
New Orleans, LA 70189

1 ONR Branch Office
Millington, TN 38054

Navy

1 Dr. William L. Maloy
Principal Civilian Advisor for
Education and Training
Naval Training Command, Code 00A
Pensacola, FL 32508

1 Dr. Kneale Marshall
Scientific Advisor to DCNO(MPT)
OP01T
Washington DC 20370

1 CAPT Richard L. Martin, USN
Prospective Commanding Officer
USS Carl Vinson (CVN-70)
Newport News Shipbuilding and Drydock Co
Newport News, VA 23607

1 Dr William Montague
Navy Personnel R&D Center
San Diego, CA 92152

1 Ted M. I. Yellen
Technical Information Office, Code 201
NAVY PERSONNEL R&D CENTER
SAN DIEGO, CA 92152

1 Library, Code P201L
Navy Personnel R&D Center
San Diego, CA 92152

1 Technical Director
Navy Personnel R&D Center
San Diego, CA 92152

6 Commanding Officer
Naval Research Laboratory
Code 2627
Washington, DC 20390

1 Psychologist
ONR Branch Office
Bldg 114, Section D
666 Summer Street
Boston, MA 02210

536 S. Clark Street
Chicago, IL 60605

Navy

1 Office of Naval Research
Code 437
800 N. Quincy SStreet
Arlington, VA 22217

5 Personnel & Training Research Programs
(Code 458)
Office of Naval Research
Arlington, VA 22217

1 Psychologist
ONR Branch Office
1030 East Green Street
Pasadena, CA 91101

1 Special Asst. for Education and
Training (OP-01E)
Rm. 2705 Arlington Annex
Washington, DC 20370

1 Office of the Chief of Naval Operations
Research Development & Studies Branch
(OP-115)
Washington, DC 20350

1 Dr. Donald F. Parker
Graduate School of Business Administrati
University of Michigan
Ann Arbor, MI 48109

1 LT Frank C. Petho, MSC, USN (Ph.D)
Code L51
Naval Aerospace Medical Research Laborat
Pensacola, FL 32508

1 Dr. Gary Poock
Operations Research Department
Code 55PK
Naval Postgraduate School
Monterey, CA 93940

1 Roger W. Remington, Ph.D
Code L52
NAMRL
Pensacola, FL 32508

1 Dr. Bernard Rimland (03B)
Navy Personnel R&D Center
San Diego, CA 92152

1 Dr. Worth Scanland
Chief of Naval Education and Training
Code N-5
NAS, Pensacola, FL 32508

1 Dr. Robert G. Smith
Office of Chief of Naval Operations
OP-987H
Washington, DC 20350

1 Dr. Alfred F. Smode
Training Analysis & Evaluation Group
(TAEG)
Dept. of the Navy
Orlando, FL 32813

1 Dr. Richard Sorensen
Navy Personnel R&D Center
San Diego, CA 92152

1 Roger Weissinger-Baylon
Department of Administrative Sciences
Naval Postgraduate School
Monterey, CA 93940

1 Dr. Robert Wisher
Code 309
Navy Personnel R&D Center
San Diego, CA 92152

1 Mr John H. Wolfe
Code P310
U. S. Navy Personnel Research and
Development Center
San Diego, CA 92152

Army

1 Technical Director
U. S. Army Research Institute for the
Behavioral and Social Sciences
5001 Eisenhower Avenue
Alexandria, VA 22333

1 Dr. Dexter Fletcher
U.S. Army Research Institute
5001 Eisenhower Avenue
Alexandria,VA 22333

1 DR. FRANK J. HARRIS
U.S. ARMY RESEARCH INSTITUTE
5001 EISENHOWER AVENUE
ALEXANDRIA, VA 22333

1 Col Frank Hart
Army Research Institute for the
Behavioral & Social Sciences
5001 Eisenhower Blvd.
Alexandria, VA 22333

1 Dr. Michael Kaplan
U.S. ARMY RESEARCH INSTITUTE
5001 EISENHOWER AVENUE
ALEXANDRIA, VA 22333

1 Dr. Milton S. Katz
Training Technical Area
U.S. Army Research Institute
5001 Eisenhower Avenue
Alexandria, VA 22333

1 Dr. Harold F. O'Neil, Jr.
Attn: PERI-OK
Army Research Institute
5001 Eisenhower Avenue
Alexandria, VA 22333

1 Dr. Robert Sasmor
U. S. Army Research Institute for the
Behavioral and Social Sciences
5001 Eisenhower Avenue
Alexandria, VA 22333

1 Dr. Frederick Steinheiser
U. S. Army Reserch Institute
5001 Eisenhower Avenue
Alexandria, VA 22333

Army

1 Dr. Joseph Ward
U.S. Army Research Institute
5001 Eisenhower Avenue
Alexandria, VA 22333

Air Force

1 Dr. Earl A. Alluisi
HQ, AFHRL (AFSC)
Brooks AFB, TX 78235

1 Dr. Genevieve Haddad
Program Manager
Life Sciences Directorate
AFOSR
Bolling AFB, DC 20332

1 Dr. Marty Rockway
Technical Director
AFHRL(OT)
Williams AFB, AZ 58224

2 3700 TCHTW/TTGH Stop 32
Sheppard AFB, TX 76311

Marines

1 H. William Greenup
Education Advisor (E031)
Education Center, MCDEC
Quantico, VA 22134

1 Special Assistant for Marine
Corps Matters
Code 100M
Office of Naval Research
800 N. Quincy St.
Arlington, VA 22217

1 DR. A.L. SLAFKOSKY
SCIENTIFIC ADVISOR (CODE RD-1)
HQ, U.S. MARINE CORPS
WASHINGTON, DC 20380

Other DoD

12 Defense Technical Information Center
Cameron Station, Bldg 5
Alexandria, VA 22314
Attn: TC

1 Dr. Craig I. Fields
Advanced Research Projects Agency
1400 Wilson Blvd.
Arlington, VA 22209

1 Military Assistant for Training and
Personnel Technology
Office of the Under Secretary of Defense
for Research & Engineering
Room 3D129, The Pentagon
Washington, DC 20301

1 DARPA
1400 Wilson Blvd.
Arlington, VA 22209

Civil Govt

1 Dr. Susan Chipman
Learning and Development
National Institute of Education
1200 19th Street NW
Washington, DC 20208

1 Dr. Joseph I. Lipson
SEDR W-638
National Science Foundation
Washington, DC 20550

1 William J. McLaurin
Rm. 301, Internal Revenue Service
2221 Jefferson Davis Highway
Arlington, VA 22202

1 Dr. Arthur Melmed
National Intitute of Education
1200 19th Street NW
Washington, DC 20208

1 Dr. Andrew R. Molnar
Science Education Dev.
and Research
National Science Foundation
Washington, DC 20550

1 Dr. Frank Withrow
U. S. Office of Education
400 Maryland Ave. SW
Washington, DC 20202

1 Dr. Joseph L. Young, Director
Memory & Cognitive Processes
National Science Foundation
Washington, DC 20550

Non Govt

1 Dr. John R. Anderson
Department of Psychology
Carnegie Mellon University
Pittsburgh, PA 15213

1 Anderson, Thomas H., Ph.D.
Center for the Study of Reading
174 Children's Research Center
51 Gerty Drive
Champiagn, IL 61820

1 Dr. John Annett
Department of Psychology
University of Warwick
Coventry CV4 7AL
ENGLAND

1 DR. MICHAEL ATWOOD
SCIENCE APPLICATIONS INSTITUTE
40 DENVER TECH. CENTER WEST
7935 E. PRENTICE AVENUE
ENGLEWOOD, CO 80110

1 1 psychological research unit
Dept. of Defense (Army Office)
Campbell Park Offices
Canberra ACT 2600, Australia

1 Dr. Alan Baddeley
Medical Research Council
Applied Psychology Unit
15 Chaucer Road
Cambridge CB2 2EF
ENGLAND

1 Dr. Patricia Baggett
Department of Psychology
University of Denver
University Park
Denver, CO 80208

1 Mr Avron Barr
Department of Computer Science
Stanford University
Stanford, CA 94305

Non Govt

1 Dr. Nicholas A. Bond
Dept. of Psychology
Sacramento State College
600 Jay Street
Sacramento, CA 95819

1 Dr. Lyle Bourne
Department of Psychology
University of Colorado
Boulder, CO 80309

1 Dr. John S. Brown
XEROX Palo Alto Research Center
3333 Coyote Road
Palo Alto, CA 94304

1 Dr. Bruce Buchanan
Department of Computer Science
Stanford University
Stanford, CA 94305

1 DR. C. VICTOR BUNDERSON
WICAT INC.
UNIVERSITY PLAZA, SUITE 10
1160 SO. STATE ST.
OREM, UT 84057

1 Dr. Pat Carpenter
Department of Psychology
Carnegie-Mellon University
Pittsburgh, PA 15213

1 Dr. John B. Carroll
Psychometric Lab
Univ. of No. Carolina
Davie Hall 013A
Chapel Hill, NC 27514

1 Charles Myers Library
Livingstone House
Livingstone Road
Stratford
London E15 2LJ
ENGLAND

1 Dr. William Chase
Department of Psychology
Carnegie Mellon University
Pittsburgh, PA 15213

Non Govt

1 Dr. Micheline Chi
Learning R & D Center
University of Pittsburgh
3939 O'Hara Street
Pittsburgh, PA 15213

1 Dr. William Clancey
Department of Computer Science
Stanford University
Stanford, CA 94305

1 Dr. Allan M. Collins
Bolt Beranek & Newman, Inc.
50 Moulton Street
Cambridge, Ma 02138

1 Dr. Lynn A. Cooper
LRDC
University of Pittsburgh
3939 O'Hara Street
Pittsburgh, PA 15213

1 Dr. Meredith P. Crawford
American Psychological Association
1200 17th Street, N.W.
Washington, DC 20036

1 Dr. Kenneth B. Cross
Anacapa Sciences, Inc.
P.O. Drawer Q
Santa Barbara, CA 93102

1 Dr. Ronna Dillon
Department of Guidance and Educational P
Southern Illinois University
Carbondale, IL 62901

1 Dr. Hubert Dreyfus
Department of Philosophy
University of California
Berkely, CA 94720

1 LCOL J. C. Eggenberger
DIRECTORATE OF PERSONNEL APPLIED RESEARC
NATIONAL DEFENCE HQ
101 COLONEL BY DRIVE
OTTAWA, CANADA K1A 0K2

Non Govt

1 Dr. Ed Feigenbaum
Department of Computer Science
Stanford University
Stanford, CA 94305

1 Dr. Richard L. Ferguson
The American College Testing Program
P.O. Box 168
Iowa City, IA 52240

1 Mr. Wallace Feurzeig
Bolt Beranek & Newman, Inc.
50 Moulton St.
Cambridge, MA 02138

1 Dr. Victor Fields
Dept. of Psychology
Montgomery College
Rockville, MD 20850

1 Dr. John R. Frederiksen
Bolt Beranek & Newman
50 Moulton Street
Cambridge, MA 02138

1 Dr. Alinda Friedman
Department of Psychology
University of Alberta
Edmonton, Alberta
CANADA T6G 2E9

1 Dr. R. Edward Geiselman
Department of Psychology
University of California
Los Angeles, CA 90024

1 DR. ROBERT GLASER
LRDC
UNIVERSITY OF PITTSBURGH
3939 O'HARA STREET
PITTSBURGH, PA 15213

1 Dr. Marvin D. Glock
217 Stone Hall
Cornell University
Ithaca, NY 14853

Non Govt

1 Dr. Daniel Gopher
Industrial & Management Engineering
Technion-Israel Institute of Technology
Haifa
ISRAEL

1 DR. JAMES G. GREENO
LRDC
UNIVERSITY OF PITTSBURGH
3939 O'HARA STREET
PITTSBURGH, PA 15213

1 Dr. Harold Hawkins
Department of Psychology
University of Oregon
Eugene OR 97403

1 Dr. Barbara Hayes-Roth
The Rand Corporation
1700 Main Street
Santa Monica, CA 90406

1 Dr. Frederick Hayes-Roth
The Rand Corporation
1700 Main Street
Santa Monica, CA 90406

1 Dr. James R. Hoffman
Department of Psychology
University of Delaware
Newark, DE 19711

1 Glenda Greenwald, Ed.
"Human Intelligence Newsletter"
P. O. Box 1163
Birmingham, MI 48012

1 Dr. Earl Hunt
Dept. of Psychology
University of Washington
Seattle, WA 98105

1 Dr. Steven W. Keele
Dept. of Psychology
University of Oregon
Eugene, OR 97403

Non Govt

1 Dr. Walter Kintsch
Department of Psychology
University of Colorado
Boulder, CO 80302

1 Dr. Kenneth A. Klivington
Program Officer
Alfred P. Sloan Foundation
630 Fifth Avenue
New York, NY 10111

1 Dr. Stephen Kosslyn
Harvard University
Department of Psychology
33 Kirkland Street
Cambridge, MA 02138

1 Mr. Marlin Kroger
1117 Via Goleta
Palos Verdes Estates, CA 90274

1 Dr. Jill Larkin
Department of Psychology
Carnegie Mellon University
Pittsburgh, PA 15213

1 Dr. Alan Lesgold
Learning R&D Center
University of Pittsburgh
Pittsburgh, PA 15260

1 Dr. Michael Levine
Department of Educational Psychology
210 Education Bldg.
University of Illinois
Champaign, IL 61801

1 Dr. Robert A. Levit
Director, Behavioral Sciences
The BDM Corporation
7915 Jones Branch Drive
McClean, VA 22101

1 Dr. Charles Lewis
Faculteit Sociale Wetenschappen
Rijksuniversiteit Groningen
Oude Boteringestraat
Groningen
NETHERLANDS

Non Govt

1 Dr. Erik McWilliams
Science Education Dev. and Research
National Science Foundation
Washington, DC 20550

1 Dr. Mark Miller
Computer Science Laboratory
Texas Instruments, Inc.
Mail Station 371, P.O. Box 225936
Dallas, TX 75265

1 Dr. Allen Munro
Behavioral Technology Laboratories
1845 Elena Ave., Fourth Floor
Redondo Beach, CA 90277

1 Dr. Donald A Norman
Dept. of Psychology C-009
Univ. of California, San Diego
La Jolla, CA 92093

1 Dr. Jesse Orlansky
Institute for Defense Analyses
400 Army Navy Drive
Arlington, VA 22202

1 Dr. Seymour A. Papert
Massachusetts Institute of Technology
Artificial Intelligence Lab
545 Technology Square
Cambridge, MA 02139

1 Dr. James A. Paulson
Portland State University
P.O. Box 751
Portland, OR 97207

1 MR. LUIGI PETRULLO
2431 N. EDGEWOOD STREET
ARLINGTON, VA 22207

1 Dr. Martha Polson
Department of Psychology
University of Colorado
Boulder, CO 80302

Non Govt

1 DR. PETER POLSON
DEPT. OF PSYCHOLOGY
UNIVERSITY OF COLORADO
BOULDER, CO 80309

1 Dr. Steven E. Poltrock
Department of Psychology
University of Denver
Denver,CO 80208

1 MINRAT M. L. RAUCH
P II 4
BUNDESMINISTERIUM DER VERTEIDIGUNG
POSTFACH 1328
D-53 BONN 1, GERMANY

1 Dr. Fred Reif
SESAME
c/o Physics Department
University of California
Berkely, CA 94720

1 Dr. Lauren Resnick
LRDC
University of Pittsburgh
3939 O'Hara Street
Pittsburgh, PA 15213

1 Dr. Andrew M. Rose
American Institutes for Research
1055 Thomas Jefferson St. NW
Washington, DC 20007

1 Dr. Ernst Z. Rothkopf
Bell Laboratories
600 Mountain Avenue
Murray Hill, NJ 07974

1 Dr. David Rumelhart
Center for Human Information Processing
Univ. of California, San Diego
La Jolla, CA 92093

1 DR. WALTER SCHNEIDER
DEPT. OF PSYCHOLOGY
UNIVERSITY OF ILLINOIS
CHAMPAIGN, IL 61820

Non Govt

1 Dr. Alan Schoenfeld
Department of Mathematics
Hamilton College
Clinton, NY 13323

1 DR. ROBERT J. SEIDEL
INSTRUCTIONAL TECHNOLOGY GROUP
HUMRRO
300 N. WASHINGTON ST.
ALEXANDRIA, VA 22314

1 Committee on Cognitive Research
% Dr. Lonnie R. Sherrod
Social Science Research Council
605 Third Avenue
New York, NY 10016

1 Robert S. Siegler
Associate Professor
Carnegie-Mellon University
Department of Psychology
Schenley Park
Pittsburgh, PA 15213

1 Dr. Edward E. Smith
Bolt Beranek & Newman, Inc.
50 Moulton Street
Cambridge, MA 02138

1 Dr. Robert Smith
Department of Computer Science
Rutgers University
New Brunswick, NJ 08903

1 Dr. Richard Snow
School of Education
Stanford University
Stanford, CA 94305

1 Dr. Robert Sternberg
Dept. of Psychology
Yale University
Box 11A, Yale Station
New Haven, CT 06520

1 DR. ALBERT STEVENS
BOLT BERANEK & NEWMAN, INC.
50 MOULTON STREET
CAMBRIDGE, MA 02138

Non Govt

1 David E. Stone, Ph.D.
Hazeltine Corporation
7680 Old Springhouse Road
McLean, VA 22102

1 DR. PATRICK SUPPES
INSTITUTE FOR MATHEMATICAL STUDIES IN
THE SOCIAL SCIENCES
STANFORD UNIVERSITY
STANFORD, CA 94305

1 Dr. Kikumi Tatsuoka
Computer Based Education Research
Laboratory
252 Engineering Research Laboratory
University of Illinois
Urbana, IL 61801

1 Dr. John Thomas
IBM Thomas J. Watson Research Center
P.O. Box 218
Yorktown Heights, NY 10598

1 DR. PERRY THORNDYKE
THE RAND CORPORATION
1700 MAIN STREET
SANTA MONICA, CA 90406

1 Dr. Douglas Towne
Univ. of So. California
Behavioral Technology Labs
1845 S. Elena Ave.
Redondo Beach, CA 90277

1 Dr. J. Uhlaner
Perceptronics, Inc.
6271 Variel Avenue
Woodland Hills, CA 91364

1 Dr. Benton J. Underwood
Dept. of Psychology
Northwestern University
Evanston, IL 60201

1 Dr. Phyllis Weaver
Graduate School of Education
Harvard University
200 Larsen Hall, Appian Way
Cambridge, MA 02138

Non Govt

1 Dr. David J. Weiss
N660 Elliott Hall
University of Minnesota
75 E. River Road
Minneapolis, MN 55455

1 DR. GERSHON WELTMAN
PERCEPTRONICS INC.
6271 VARIEL AVE.
WOODLAND HILLS, CA 91367

1 Dr. Keith T. Wescourt
Information Sciences Dept.
The Rand Corporation
1700 Main St.