

DOCUMENT RESUME

ED 142 992

CS 003 638

AUTHOR Munro, Allen; Rigney, Joseph W.
 TITLE A Schema Theory Account of Some Cognitive Processes in Complex Learning. Technical Report No. 81.
 INSTITUTION University of Southern California, Los Angeles. Behavioral Technology Labs.
 SPONS AGENCY. Office of Naval Research, Arlington, Va. Personnel and Training Research Programs Office.
 PUB DATE Jul 77
 CONTRACT NO0014-75-C-0838
 NOTE 76p.

EDRS PRICE MF-\$0.83 HC-\$4.67 Plus Postage.
 DESCRIPTORS Artificial Intelligence; *Cognitive Processes; Cognitive Style; *Comprehension; Literature Reviews; *Memory; Models; Psychological Studies; *Semantics; Theories; *Thought Processes
 IDENTIFIERS *Inference; *Schemata

ABSTRACT

Procedural semantics models have diminished the distinction between data structures and procedures in computer simulations of human intelligence. This development has theoretical consequences for models of cognition. One type of procedural semantics model, called schema theory, is presented, and a variety of cognitive processes are explained in terms of the theory. Specific cognitive processes discussed include semantic and episodic memory, consciousness, cognitive strategies, orienting tasks and self-direction, inference and depth of processing, and insight in various contexts. Three dimensions for distinguishing or comparing schemata are proposed: function, abstractness, and scope. The contrasts between multistore models of cognition and schema theory are summarized. (Author/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 * Documents acquired by ERIC include many informal unpublished *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. *

Unclassified

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

ED142992

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM	
1 REPORT NUMBER	2 GOVT ACCESSION NO.	3* RECIPIENT'S CATALOG NUMBER	
Technical Report # 81			
4 TITLE (and Subtitle)		5 TYPE OF REPORT & PERIOD COVERED	
A Schema Theory Account of Some Cognitive Processes in Complex Learning		1 Jan. - 30 June - 1977	
		6 PERFORMING ORG. REPORT NUMBER	
7 AUTHOR(s)		8. CONTRACT OR GRANT NUMBER(s)	
Allen Munro and J. W. Rigney		N00014-75-C-0838	
9 PERFORMING ORGANIZATION NAME AND ADDRESS		10 PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS	
Behavioral Technology Laboratories University of Southern California Los Angeles, California 90007		Program Element: 61153N Project: RR042-06 Task Area: RR042-06-01 Work Unit: 154-855	
11 CONTROLLING OFFICE NAME AND ADDRESS		12 REPORT DATE	
Personnel and Training Research Programs Office of Naval Research (Code 458) Arlington, Virginia 22217		June - 1977	
		13. NUMBER OF PAGES	
		66 + iv	
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			
ACKNOWLEDGMENTS			
17. DI			
This research was sponsored by ONR Contract N00014-75-C-0838. The support and encouragement of Marshall Farr and Henry Halff, Personnel and Training Research Programs, Office of Naval Research, and of Harry F. O'Neil, Jr., Program Manager, Cybernetics Technology Office, Defense Advanced Research Projects Agency, is gratefully acknowledged. We thank Nick Bond for discussions on the process of verbal logic puzzle solution and Kathy Lutz for discussions on creativity and conscious control.			
18. SL			
19. KEY WORDS (Continue on reverse side if necessary and identify by block number)			
Schema theory, Procedural Semantics, Conceptually-driven Processing, Episodic Memory, Semantic Memory, Specific Memory, Generic Memory, Consciousness, Cognitive Strategies, Orienting Tasks, Self-direction, Inference, Levels of Processing, Insight			
20. ABSTRACT (Continue on reverse side if necessary and identify by block number)			
Procedural semantics models have diminished the distinction between data structures and procedures in computer simulations of human intelligence. This development has theoretical consequences for models of cognition. One type of procedural semantics model, called schema theory, is presented, and a variety of cognitive processes are explained in terms of the theory. In schema theory, the flow of processing control is			

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

THIS DOCUMENT HAS BEEN REPRODUCED EXACTLY AS RECEIVED FROM THE PERSON OR ORGANIZATION ORIGINATING IT. POINTS OF VIEW OR OPINIONS STATED DO NOT NECESSARILY REPRESENT OFFICIAL NATIONAL INSTITUTE OF EDUCATION POSITION OR POLICY.

85003638

DD FORM 1 JAN 73 1473

EDITION OF 1 NOV 65 IS OBSOLETE
S/N 0102 LF 014-6601

Unclassified

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)



determined not by a central monitor, but by interactions among the conceptual entities (schemata) that make up the model. Intelligence is distributed in this model. Schemata interact by providing activation resources to each other.

Instantiation is the special process whereby a partial copy of a strongly activated schema is created. In this copy, the variables of the schema are filled with particular values. Such copies make up specific or episodic memory. The schemata on which they are based comprise generic or semantic memory.

Many of the phenomena of consciousness and of short-term and long-term memory are explained on the basis of the activation processes of schema theory. Unactivated schemata are equivalent to all the unconscious knowledge in a person's long-term memory. Schemata that are activated, but are below the threshold for instantiation, are in a preconscious or subconscious state. Those schemata that are more fully activated, that are above the instantiation threshold, are the stuff of conscious thought, and may be thought of as roughly equivalent to the contents of short-term memory.

Conscious cognitive strategies are treated as the activations of abstract prescriptive schemata. A treatment of creativity is presented, along with the outlines of an approach to individual differences in creativity. The effects of orienting tasks are explained in schema theory, and the relationship between orienting tasks and self-direction in complex learning and problem solving is discussed.

Inference and depth of processing receive related schema theory treatments. Both concepts are treated in terms of the extent to which activation spreads to include related schemata. In general, the more schemata activated to the level of instantiation by some datum, the more deeply processed that datum is. Inference is seen as a kind of delayed deeper processing.

Types of insight phenomena from several contexts can also be treated in schema theory. Each type of insight involves the instantiation of one or more new schemata that take some pre-existing concept in memory as a parameter.

Three dimensions for distinguishing or comparing schemata are proposed: function, abstractness, and scope. The contrasts between multi-store models of cognition and schema theory are summarized.

The views and conclusions contained in this document are those of the author's and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the Office of Naval Research, the Advanced Research Projects Agency, or the U. S. Government.

Approved for public release: Distribution unlimited.
Reproduction in whole or in part is permitted for any purpose of the United States Government.

ARPA TECHNICAL REPORT

1. ARPA Order Number	: 2284
2. ONR NR Number	: 154-355
3. Program Code Number	: 1 B-729
4. Name of Contractor	: University of Southern California
5. Effective Date of Contract	: January 1, 1977
6. Contract Expiration Date	: 30 September 1977
7. Amount of Contract	: \$150,000
8. Contract Number	: N00014-75-C-0838
9. Principal Investigator	: Joseph W. Rigney (213) 741-7327
10. Scientific Officer	: Marshall Farr
11. Short Title	: A Schema Theory Account of Some Cognitive Processes

This Research Was Supported

by

The Advanced Research Projects Agency

and by

The Office of Naval Research

And Was Monitored by

The Office of Naval Research

TABLE OF CONTENTS

<u>Section</u>		<u>Page</u>
I.	INTRODUCTION	1
	Assumptions of the model	3
	The relation of this work to a general theory of complex learning	4
II.	SCHEMA THEORY EXPLANATIONS OF COGNITIVE PROCESSES . .	8
	Semantic and Episodic Memory	8
	Consciousness	11
	Cognitive Strategies	18
	Creativity and conscious control	20
	Orienting Tasks and Self-Direction	27
	Inference and Depth of Processing	33
	Inference	33
	Depth of Processing	43
	Insight	51
	Insight during complex learning	52
	Insight during problem-solving	54
	Insight in verbal logic puzzle solutions	57
	Common aspects of insight in these three cases	61
III.	SOME REFLECTIONS ON SCHEMA-THEORY	62
	Schema-Theory and Cognitive Component Theories	62
	A Demonology: The Varieties of Schemata	63

LIST OF FIGURES

<u>Figure</u>		<u>Page</u>
1.	Conditions and Processes of Complex Human Learning	5
2A.	Consciousness in a "Memory Components" Model of Cognition	15
2B.	Consciousness in a "Mutually Activating Schemata" Model of Cognition	16
3.	Distributed Control During Problem-Solving in a Schema- System	25
4.	Task-Response Schemata on a Dimension of Abstractness.	32
5.	A Schema-Theory Representation of "Sam Took a Book from John"	37
6.	A Schema Theory Representation of "Sam Took a Book from John" and "Sam Knows the Information Contained in the Book"	38
7.	A Schema-Theory Representation of "Sam Took a Book from John" and "Sam Knows the Information Contained in the Book" After Further Inference has Taken Place	40
8.	A Possible Representation for the List Item "Lion" in the Mind of an Experimental Subject who was Required to Make a "Size" Judgement About "Lion."	49
9.	A Possible Representation for the List Item "Lion" in the Mind of an Experimental Subject who was Required to Make Both a "Size" Judgement and a "Happiness" Judgement About "Lion."	50
10.	Restructuring	53
11.	Solving the Pendulum Problem: A Schema-Theory Representation	56
12.	A Simple "Insight" Necessary to the Solution of the Mr. Scott Problem: A Schema-Theory Representation	58

LIST OF TABLES

<u>Table</u>		<u>Page</u>
1.	An Overview of Some Schemata	65

A SCHEMA THEORY ACCOUNT OF
SOME COGNITIVE PROCESSES
IN COMPLEX LEARNING

I. INTRODUCTION

In the course of the recent evolution of semantic network theories, a qualitative change in the nature of these theories has taken place. Early models of semantic memory (Quillian, 1968; Anderson, 1972; Anderson & Bower, 1972) were characterized by a completely structural approach. Information in memory was represented by data structures. The processes that searched for information, that added new data structures, that compared data structures, that were, in short, responsible for the representation of thought in such systems, were completely distinct and different in nature from the data structures. Friendly, 1977, presents a method for characterizing the features of such structural representations. In more recent models (Minsky, 1975; Norman, Rumelhart & LNR, 1975; Bobrow & Winograd, 1977), the distinctions between data and processes have been blurred. From a computational viewpoint, data structures in these models can function as procedures (and vice versa). In psychological terms, this means that memories can function as thought-processes. These theories constitute a kind of new demonology, in which one kind of validity test is possible through computer simulations. In such simulations, the demons are modeled by procedures that also have data-characteristics. Models that can be so described are procedural semantics models.

Many telling objections to particular procedural semantics models of language processing and thought have been made. Weisenbaum (1976) and Dresher & Hornstein (1977) have argued convincingly that the models in extant are too concerned with the structure and application of knowledge

in the mature mind. Insufficient attention has been paid to the attributes of human intelligence that control the acquisition of this knowledge. Dreyfus (1972) has argued that artificial intelligence has largely ignored the parallel processing aspects of human cognition. Hayes-Roth & Hayes-Roth (1977) show that the conceptual entities that represent words (as opposed to more abstract semantic "primitives" or "features") have a more important status in memory than they are granted in many current models. As of yet, however, no convincing criticisms of the claim that conceptual units have a dual nature--data and processes-- have been made. It is this claim that is the core of the procedural semantics approach to cognition.

The procedural semantics viewpoint is still incomplete, and, as a result, opportunities exist to remedy the deficiencies that have been pointed out. Rumelhart & Norman (in press) have sketched a procedural semantics approach to the acquisition of generic information. A more detailed version of their theory could answer some of the objections made by Weisenbaum (1976) and by Dresner & Hornstein (1977). Fiskel & Bower (1976) present a formalism for a semantic network composed of finite automata. This is one of a number of possible means of modeling the parallelism Dreyfus (1972) pointed out as an essential part of cognition. Rigney & Munro (1977) present a schema theory model of human text processing in which words are presumed to be represented explicitly by their own special data/process entities in memory. Although their theory also provides for the existence of many non-lexical conceptual entities, the effects due to their lexical-level units can account for the results reported by Hayes-Roth & Hayes-Roth (1977).

The present paper represents an attempt to further extend the explanatory scope of procedural semantics theories to a wider range of cognitive

processes. An account is given in terms of such a theory for the constructs of semantic and episodic memory, consciousness, levels of processing, and some "insight" phenomena. We view this account as an exploratory, even speculative, attempt to extract the maximum possible explanatory power from the theory. We hope that it will stimulate further thinking about the possible applications of this theory and that it will encourage the development of more formal procedural accounts of psychological phenomena.

The procedural semantics model we present here is a lineal descendent of that presented in Norman, Rumelhart, & LNR (1975). The LNR model was implemented in a computer simulation called MEMOD. The current model has not been so implemented, but the data/process structures (hereafter referred to as schemata) have been written in a format compatible with advanced MEMOD data bases. (The format makes use of a predicate calculus-type exposition, which makes procedure-parameter relationships-- scope relationships--explicit). Given appropriate supporting procedures, the schemata discussed below could be expected to function in MEMOD.

Assumptions of the model

Some features of the present schema-theory are different, at least in emphasis, from those described in Norman, Rumelhart, & LNR (1975). One such feature is that of distributed intelligence (discussed in Gentner, in press). The activities of schemata and their interactions account for all cognitive processing, making it unnecessary to postulate some higher level executive function to coordinate this processing. The interactions of schemata occur simultaneously, as parallel processing, in a population of interrelated schemata. These interactions result in interactive data-driven and conceptually-driven processing (Palmer, 1975).

Intelligence is distributed in the sense that intellectual processing is not the responsibility of a single general-purpose device; rather, it is simply the sum of all the activated schemata at the time in question.

Goldstein and Papert (1977) express a similar point of view: "...intelligence is based on the ability to use large amounts of diverse kinds of knowledge in procedural ways, rather than on the possession of a few general and uniform principles."

Another feature of the current model is mutual activation. Following Levin (1976), one of the most important aspects of a schema is its ability to activate other schemata. Each schema is limited in the extent to which it can provide activation to other schemata. If a schema receives only a small amount of activation from other schemata, its own ability to pass on activation will be further limited. When an activation passes a threshold (called the instantiation threshold) it becomes instantiated.

What this means is that a copy of the schema (which is a concept in generic or semantic memory) is created. This "copy" is not an exact duplicate, but contains information about specifics which are represented only as prototypes or selectional restrictions (Chomsky, 1965) in the generic schema. These "copies" or instantiated schemata constitute episodic memory. An activated Schema that is not instantiated will eventually fade to a background level of activation, and no direct evidence of its activation will remain. These processes are more fully discussed below in the sections on semantic and episodic memory and on consciousness and schemata.

The relation of this work to a general theory of complex learning.

The diagram in Figure 1 outlines variables that seem to us to be among the important considerations for understanding more about complex human learning. Most, but not all of the phenomena that we will attempt to explain in terms of schema theory are identified as memory processes in the figure.




Figure 1. Conditions and Processes of Complex Human Learning

According to the view expressed in this diagram, human learning is an elaboration of more primitive biological mechanisms for adapting to change, and retains many of the features to be expected of such mechanisms. Learning is driven by initiating requirements that are formulated as orienting tasks, either implicitly or explicitly. Orienting tasks can originate with an instructor or with an instructional system or with the learner. Orienting tasks contain instructions and goals. These drive cognitive processing during learning and utilization episodes, which tend to be intermingled, at least outside of formal pedagogical environments. It is likely that utilization of knowledge also results in some additional modification of knowledge representations in memory. Conversely, learning new knowledge occurs only in the context of existing knowledge. Learning and utilization episodes may last for a few minutes or for years. In the laboratory, they tend to be limited to a few minutes. Outside the laboratory, learning and utilization of knowledge goes on throughout life; there is a time-shared processing of the many demands for the different kinds of content that all of us are required to learn and use just to survive, or that we set ourselves to learn out of curiosity or in hope of surviving in particularly favored ways. Learning and utilization processes are both conceptually-driven (from the top down) and data-driven (from the bottom up).

In our view, the concept represented by a schema may be in any one of three states: conscious, subconscious, or unconscious. In the diagram, almost all the cognitive processes included in the box labeled "memory processes" ordinarily function only at the pre-conscious or subconscious level. (Some problem-solving processes have conscious components.) Cognitive processes that are available to the learner for conscious use and that are designed to enhance or facilitate

acquisition, retention, and retrieval can be called cognitive strategies. Cognitive strategies are represented in schema-theory by very high level schemata. In the terminology of Norman and Rumelhart (in press), such schemata are not ordinarily highly "tuned." Since they are general purpose, they "fit" no one situation or context perfectly. Their activations are therefore not simple and automatic, as highly tuned schemata are. Components or subschemata of the strategy schema must undergo a good deal of top down activation, which performs a sort of checking function. The person experiencing an activation of such a strategy schema is therefore usually aware of the relevant components of the schema, and the application of the strategy is a conscious process. This viewpoint is further discussed below in the section on Consciousness.

II. SCHEMA THEORY EXPLANATIONS OF COGNITIVE PROCESSES

Semantic and Episodic Memory

The distinction between semantic and episodic memory in schema theory at first appears to be quite a simple and natural one. Semantic memory consists simply of schemata--concepts which represent what one knows about general types of objects or actions or relationships in the world. Episodic memory is simply the collection of all the instantiated "copies" of schemata, with concepts which stand for particular entities filling the argument slots of the schemata as parameters.

This is a good and useful characterization of two types of memories in a schema-system for representing stored knowledge. There are, however, some potential problems with the use of the terms "semantic" and "episodic". The use of the term "semantic" connotes knowledge which is basically lexical in nature. "Semantics" is, in large part, the study of the relationships between symbols and what those symbols refer to. The term "episodic", on the other hand, connotes particular events, to the exclusion of other types of particular concepts, such as particular individuals. Both of these terms are too limited in their connotations.

In a schema-theory representation of knowledge about types, schemata should probably not be thought of as being nearly so closely bound to particular lexical items as was common in early formulations (such as Rumelhart, Lindsay, & Norman, 1972). Of course, there must be schemata associated with particular lexical items; we could call these "lexical-level content-schemata". (Rigney & Munro, 1977, present a partial functional typology of schemata.) In several recent works which are either primarily theoretical papers within the schema-theory model or which report experimental results which can be interpreted as having implications for this model, a very large number of essentially

non-lexical schemata seem to be called for. One type of non-lexical schema is that which is much higher-level or more abstract than lexical schemata. Examples of this type include story- or episode-schemata (Rumelhart 1975; Rumelhart, in press), social-situation-schemata (Schank & Abelson, 1975), and speech act and conversational schemata (Munro, 1977; Levin & Moore, 1977). All of these schemata are designed to account for some of the facts about how people's knowledge about recurring types of situations guides their understanding (and sometimes their actions) in those situations.

A different kind of "non-lexical" schema which seems to be called for in a putatively complete or accurate account of human knowledge is one that captures the fact that people's knowledge about types is often more detailed than can be conveniently expressed in terms of their lexicons. For example, it is perfectly natural to talk about "red hair", "red carpet", and "red eyes", and to mean something very different by each use of "red". Yet each of these uses of "red" should have a reasonably consistent and distinct meaning, at least for individual subjects (Halff, Ortony, & Anderson, 1976). This suggests that we may not want to model people's concepts for red monolithically, but rather to include in our model one concept for the red that can be the color of a person's hair, another concept or schema for the red that can be the color of dried blood, and so on. Other recent work suggests that the problem of polysemy may be more ubiquitous than semantic memory theorists originally thought (Anderson & Ortony, 1975; Anderson Pichert, Goetz, Schallert, Stevens & Trollip, 1976):

The point behind this discussion of non-lexical semantic memory is that it is possible that the term "semantic memory" may be too

restrictive in its connotations to continue to apply to the representation of types in a schema-theory approach to conceptual representations.

As a replacement, consider the hopefully more neutral term "generic memory." (Rumelhart & Ortony, in press, refer to schemata as concepts which represent generic knowledge). A generic concept is one which can fit a variety of situations or objects; it has partially unspecified parameters. In other words, it is a schema.

What of "episodic memory?" Instantiations of generic schemata are the stuff of episodic memory. The term "episodic" seems particularly appropriate when the instantiated schema is one which represents some activity or episode. When the instantiated schema has a more nominal quality, however, the term seems less appropriate. For example, if you hear someone say "A dog bit me yesterday," you would presumably acquire a new instantiation of your DOG schema as part of the process of understanding the sentence. Yet it seems odd, to say the least, to speak of your new concept of this particular dog as an "episode".

A natural replacement for the term "episodic" is suggested by the antonym of "generic", namely, "specific". Let concepts of particular individuals, particular actions, and particular relationships be referred to as "specific concepts". The instantiation of a generic concept is a specific concept.

It should be recognized that, while the use of the terms "generic" and "specific" may constitute an improvement over the less precise terms "semantic" and "episodic", it does nothing to resolve several substantive issues concerning the relationship of generic and specific concepts. For example, there probably should not be an absolute distinction between generic and specific concepts. The primary distinction between a generic and a specific concept is that the first has variables, while

the second has filled parameters (that is, other specific concepts as arguments). This is not always true in a strict sense, however. For example, some "specific" concepts may actually have default values for every parameter--information which is part of the corresponding generic concepts. And some "generic" concepts may, in fact, be extremely specific in nature. If I believe, for example, that "John Peregrine rides his tricycle around our block every morning at 7:30", then this information should probably be represented in generic form, as a schema. But what are the variables in such a schema? Almost every important parameter is already specified in the generic representation. Only a few details about manner and the particular date need to be filled in to create a specific version of this generic schema. (Another inadequately explored relationship between generic and specific concepts is the process by which insight into the similarities among a number of specific schemata results in the creation of a new generic schema. Rumelhart & Norman (in press) sketch some aspects of this process, which they call restructuring). Fuzzy edges aside, we believe that the notions of generic and specific concepts are important, and that the terms "generic memory" and "specific memory" are more precise than the traditional "semantic memory" and "episodic memory".

Consciousness

Schemata, in the mutual activation model, can be thought of as always being in one of three states. They are either (1) quiescent, receiving no activation from any other schema and giving no activation to any other (probably the normal state for the vast majority of a

person's schemata at any single instant);¹ or (2) activated to the level of instantiation, receiving sufficient activation from other schemata that they can become instantiated, so that their copies become part of the person's vast store of specific memories; or (3) activated but not instantiated. Those schemata which are in the latter state may later receive enough activation to become instantiated, or they may subside in their level of activation; in which case there will be no long-term memory for the particular instance of the concept represented by the schema in question. Long-term memory traces for specific concepts are the "copies" of the corresponding generic concepts that were made when instantiations took place. When a generic concept does not receive enough activation to become instantiated, no specific copy is made, and there is therefore no memory (as well as no understanding) in terms of that generic concept.

The identifications of the various states of consciousness in such a model are probably fairly obvious. All those schemata which are activated to the point that they are being instantiated constitute the contents of consciousness. Those schemata which are activated but not instantiated constitute the contents of the subconscious or preconscious at a given moment. All those schemata which are not activated constitute our unconscious knowledge at a given moment.

This sort of cognitive account of states of consciousness has more appeal than some others which could be constructed. In a cognitive model

¹ If the schema-theory presented here made claims to being a physiological model, we would say that schemata in this "quiescent" state were actually functioning at a background level of activation, by analogy to the background firing rate of neurons. However, schema-theory is intended as a model of cognitive function rather than of neurological function. Although we believe that at some time in the future there will be a comprehensive model that accommodates what is known about both cognitive and neurophysiological function (and that such a theory may have much in common with that presented here), we restrict ourselves now to cognition. Hence, we say that schemata can be quiescent.

in which intelligence is not distributed (in a "memory components" system with boxes labeled "STM", "LTM", etc., for example) one might want to account for the subconscious by including in one's diagram a box labeled "Subconscious". In the distributed intelligence system with mutually activating schemata, there is no need to postulate such a special new component. Instead, we can simply identify a portion of the processing already called for by the theory as "subconscious thought." Figure 2 contrasts these two ways of representing the subconscious and the unconscious.

This treatment of "levels of consciousness" seems to account for the basic differences between the conscious, subconscious, and unconscious mind in Freud's theory. However, it does not deal with a number of issues of great concern to him, such as repressed memories. We leave that task to some even more enthusiastic partisan of mutually-activating schemata. Whether this approach has anything to do with another type of "unconsciousness" deserves further consideration. The type of unconsciousness we refer to is that associated with the execution of well-learned motor skills. These kinds of "unconscious" physical activities can be accounted for in schema-theory terms.²

Consider the case of "unconscious" automobile driving. Most experienced drivers have become aware, at one time or another, that they have just driven some number of miles on a familiar route, and yet remember nothing of what they have just done. They say they were not aware of

2

The class of unconscious physical activities that can be accounted for in schema-theory terms is limited to those which must be learned. Innate activities, such as breathing, are not activities we would choose to represent with schemata. However, some very automatic activities, such as the movements made in walking, are, in part, consciously learned and should therefore be representable in schema-theory terms.

their driving. Since the schemata responsible for this driving did not intrude upon consciousness, the theory just outlined should require us to say that those schemata were activated but not instantiated. But why should they not have been instantiated? We are used to schemata not being instantiated because they cannot find confirmation in the data of the context; surely there is enough bottom-up flow in normal driving contexts that this should not be a problem. Perhaps a reasonable explanation for the inadequate (for instantiation) activation of the driving schemata has to do with resource limitations. There is only so much activation for the total system, and a great deal of activation has been allocated to other schemata. Thus, when one becomes aware that one has no memories of having driven the last 15 miles, that does not mean that one has no memories of the last 15 minutes. On the contrary, one is often aware of having had a particularly stimulating conversation or train of thought, many of the details of which are easily remembered. Our driving can become subconscious when unrelated schemata have absorbed most of the total activation resources of the system.

This explanation fails to account for the observation that only very well-tuned motor schemata seem to carry out their functions in reduced-activation situations. Another way of viewing this phenomenon from the schema-theory viewpoint is to say, as was suggested in the Introduction, that the highest-level schema that can account for the data actually was instantiated, but that none of its component subschemata were. The relevant highest-level schema in the "automatic driving" example just given would not be some generalized Driving Schema, but, rather, a very





Figure 2A. Consciousness in a "Memory Components" Model of Cognition



GENERIC MEMORY

Figure 2B. Consciousness in a "Mutually Activating Schemata" Model of Cognition

detailed schema for driving a particular route. An example of such a schema could be a "Route-to-the-University-from-my-house" schema. Such explicit generic schemata have been proposed elsewhere (Munro, 1977).

When such a schema is activated by the driver's knowledge that he or she is on the same old route again, activation resources can be conserved if the schema does not supply activation for its components. It can accept activation from these components, but so long as each component makes its expected contribution to its "parent," there is no need for the parent to drive activations of the subschemata. In effect, the highest-level schema can afford to suspend its top-down processing, because it doesn't need to check for the presence of components--the components have already made their presence known. The absence of this top-down activation to the subschemata may often be enough to keep the activations of these conceptual units below the activation threshold. As a result, the driver ends the trip with no specific memories for any of the details or stages; there have been no instantiations of the subschemata of the "Route-to-the-University" schema. This does not mean that the driver will not know that he has been engaged in "driving to the University," for example. The highest level schema may be instantiated, in which case the driver of the car knows what route he is on, but not, at first, where he is on it.

The basic mechanisms for the activation and instantiation of schemata--mechanisms that are required to represent the processes of understanding, learning, and remembering--can also account for the phenomena of consciousness. We believe that this economy of theoretical primitives constitutes an argument in favor of schema theory in contrast to the "cognitive component" models prevalent in cognitive psychology.

Cognitive Strategies

When people say that someone is applying a cognitive strategy, they seem to mean that he or she has made a conscious decision to allocate processing resources to one kind of cognitive process rather than to another. (See Rigney, in press, for a detailed treatment of the possible varieties of cognitive strategies). From the perspective of a distributed intelligence theory, such as the schema-theory outlined in earlier sections, this meaning of "cognitive strategies" doesn't seem to make much sense. That which is conscious is simply those schemata which are currently being instantiated. It is always true, in the theory, that currently activated concepts (and particularly the highly activated concepts that are being instantiated) play a large role in the determination of which schemata will next receive activation resources. This natural flow of activation doesn't seem to capture the notion of conscious control over attention that most of us believe in. We feel that we are able to make decisions to "pay attention" to something and ignore other things that could compete for attention.

One way to treat this would be to say that we have a number of big, complicated schemata, whose function it is to direct the allocation of processing resources in particular situations. In the terminology of Rigney (1976); these schemata are prescriptions. When we are conscious of their effects (that is, when they are activated to the point of instantiation), we may say that we are "stating rules to ourselves" that apply in this particular situation. Here is an example: Suppose that a computer programmer named Fred is writing a program on the PLATO IV computer system in the TUTOR language. Fred needs to set a variable equal to some value, but has forgotten the name of the function that does this. He "decides" that he will try to find the name of the

function by consulting the on-line TUTOR manual, "aids". There he will look at a table that lists all the TUTOR commands, and see whether one of them "rings a bell". In schema-theory terms, how can we explain his conscious decision to adopt this strategy? One explanation that seems quite natural requires that we posit the existence of a schema which is activated when people don't know the meaning of a term and need to know the meaning of that term. This is a special-purpose schema, first acquired when they originally learned about dictionaries and glossaries. This is a high-level, complex schema which is activated when two of its subschemata (the Don't-Know- and the Need-To-Know-Schemata) are activated. The bulk of the schema is a sequence of instructions on what to do in the various circumstances. (E.g., "Is the meaning of the term part of ordinary English usage? If so, look for a synonym or explanation in the dictionary." Or, "is the term part of some small set of technical terms? If so, is there a list of such terms which could be scanned in a reasonable length of time to search for the term?") The particular attributes of the context one is in at the time (the fact that one is using the TUTOR language, for example) interact with this general schema, filling in the loosely defined parameters of the Look-Up-Schema (such as "a list of such terms") with particular values (such as "the list of TUTOR commands in 'aids'").

If Fred, the PLATO programmer, is very accustomed to looking up particular TUTOR commands in "aids", he may even have a special Look-Up-TUTOR-Command-Schema which is activated in these situations and directs his processing for awhile.

In a schema system, "conscious" allocation of processing resources is simply accounted for by the prior existence of schemata which direct people's information processing activities. When one of these schemata,

which are essentially prescriptive in nature, is activated, one has the impression of "directing oneself" to manage one's attention in particular ways. The example just given involves some fairly overt behaviors, but could equally well be applied to more purely mental activities as well. This would mean positing the existence of special schemata for such information processing activities as scanning a text in search of a particular term, trying to recall the context in which one first learned a concept in order to remember more details of the concept, and so on. In some respects this approach is similar to the "homunculus" of a powerful central processor/monitor, but it does involve much less powerful homunculi, each of which has only one special function and is activated in only one special kind of context.

Creativity and conscious control.

A potential problem for the distributed intelligence system is raised by the nature of creativity in a schema model. Many students of creativity have claimed that there is an important difference between more and less creative people (Guilford, 1968; Dellas & Galer, 1970). They treat this difference as one of cognitive style: more creative people engage in divergent thinking; less creative people typically approach problems with convergent thinking. In schema-theory terms, it seems useful to think of divergent thinking as the distribution of activation resources to a large number of largely unrelated schemata, any one of which does not seem to have a particularly good a priori chance of providing the solution. Convergent thinking in schema theory involves putting almost all activation resources into one schema, perhaps the one that seems to have the best a priori chance of providing the solution to the current problem. (One could think of the difference as a breadth-first approach versus a depth-first approach to schema activation.)

This explanation of the difference in schema-theory between convergent and divergent thinking does not explain, however, why some individuals should appear to consistently follow one mode, and others another. One approach to explaining this would be to claim that the divergent or creative thinkers have a special problem-solving schema that "searches for" the maximum number of potentially applicable schemata for the current problem-solving situation, and then distributes the activation resources available more or less evenly among those schemata. Convergent thinkers could then be characterized as having a problem-solving schema that "searches for" the "best" or most-likely-looking schema and then concentrates all or most of its activation resources to that schema. The only problem with this approach is that these schemata seem unusually powerful and complex. They are more like homunculi than simple-minded demons. An ordinary schema has a structure that specifies precisely which other schemata can be activated by it. There is something disquieting about schemata with unspecified sub-schemata; their functioning is not explicable with the mechanisms of schema theory as they are currently formulated, and they seem to be capable of far too much.

These problems seem to call for further developments either in our understanding of attentional phenomena associated with creativity or in the formulation of schema-theory presented here.

Despite the problems inherent in the requirement that a schema-theory account for differences in problem-solving behavior through the mechanism of very powerful, individualized "problem-solving schemata", there are some aspects of the introspective evidence on creativity that are well-treated in a schema-theory approach.

For example, many creative people (Mozart and Tchaikovsky have been cited as examples) have claimed that entire solutions to very complex problems may spring into their consciousnesses in a relatively complete state. (Mozart is said to have been able to envision an entire new musical piece at once, almost "complete and finished".) The schema-theory solution for the cognitive style of creative people accounts for these reports. Because the creative person's general problem-solving schema distributes activation among a large number of processing schemata, none of them may be activated sufficiently to become instantiated until one of them "fits" the data of the problem so well that the "bottom-up" activation of the schema is almost enough by itself to instantiate the schema. Fitting the data so well means that every aspect of the problem will have a place in the solution to the problem provided by the instantiated schema. This is what accounts for the "full-blown" nature of the solutions to problems that some creative people report "popping into" their consciousnesses.

One consequence to the approach to creativity--and to cognitive strategies in general--suggested here is that a large number of special schemata for attacking problems must be posited. This is not necessarily a bad feature of the schema-theory approach to problem-solving. Professional problem-solving teachers have suggested that effective problem-solvers may have a "library" of problem-solving "routines" that they can consult to try to come up with something useful for particular situations. In Strategy Notebook (discussed in Adams, 1974) a group called Interaction Associates discusses such a group of routines, which they believe a good problem-solver should possess. Here is their list.

Build up	Display	Simulate
Eliminate	Organize	Test
Work Forward	List	Play
Work Backward	Check	Manipulate
Associate	Diagram	Copy
Classify	Chart	Interpret
*Generalize	Verbalize	Transform
Exemplify	Visualize	Translate
Compare	Memorize	Expand
Relate	Recall	Reduce
Commit	Record	Exaggerate
Defer	Retrieve	Understate
Leap In	Search	Adapt
Hold Back	Select	Substitute
Focus	Plan	Combine
Release	Predict	Separate
Force	Assume	Change
Relax	Question	Vary
Dream	Hypothesize	Cycle
Imagine	Guess	Repeat
Purge	Define	Systemize
Incubate	Symbolize	Randomize

Some of these strategies are actually quite complex. To make use of the Eliminate Strategy, for example, one first thinks of all the possible attributes that a solution might have. This is done in a very non-evaluative mode. Only after a very extensive list of attributes has been prepared does one return to the list and begin to eliminate items (attributes of possible solutions) that seem to be undesirable, impractical, or unnecessary. Obviously, a schema for the Eliminate Strategy would also be quite complex, involving a number of specialized subschemata with explicit sequencing controls.

Both creative and non-creative people could be trained to make use of a large collection of such strategies. What would differentiate their use of the strategies would be the manner in which they activated the schemata that represented the strategies. The convergent thinkers would select one of the strategies early in the problem-solving process, and would make use of that strategy in a conscious and deliberate way. The divergent thinkers, on the other hand, would allocate some small amount of activation to each of the Strategy Schemata, and would not be con-

sciously exploring any particular path to a solution until one of the Strategy-Schemata came up with a particularly good fit. At this point, the divergent thinker would present his solution in a completed form, without being aware of any necessary intermediate steps taken to reach the conclusion.

If we accept the claim that problem solvers do have some set of schemata with specific strategies for problem-solving, it becomes an easier task to construct "supervisory schemata" for problem-solving. (For some very practiced or well-trained problem-solvers, the set of specific problem-solving schemata might include those strategies presented above. For most problem-solvers, the set would probably be smaller and the strategies less efficient). The reason that modeling "supervisory schemata" would be easier is that in schema-theory, as we have presented it in preceding sections, schemata must pass activation to other schemata that they already know about. In other words, schemata are not thought of as being so powerful and homunculus-like that they can freely roam through the "library" of schemata, picking and choosing those which are appropriate for a particular purpose. Rather, part of the meaning of a particular schema is the limited set of other schemata which it can activate.

In such a system, the "supervisory schema" (or general problem-solving schema) of a divergent thinker would activate simultaneously all of the problem-solving strategy-schemata. At this point, the supervisor has, in some sense, relinquished control over the ensuing processing. It is not some action of the "supervisor" that now determines which particular strategy is chosen to solve the problem. Instead, the interaction between the data (the facts of the problem) and the activated strategy-schemata determines which strategy will triumph. Figure 3 is an attempt to roughly depict the distributed nature of the control of processing in this sort of system.

The general "supervisory" problem-solving schema

Particular problem-solving strategy-schemata

Schemata activated by the statement of the problem

The problem statement




Figure 3. Distributed Control During Problem-Solving in a Schema-System (arrows represent the flow of activation)

The supervisory Problem-Solver Schema of Figure 3 can be thought of as having the combined structure of two schemata proposed by Rumelhart & Ortony (in press), a Problem-solving schema and a Try schema. Here are the schemata they proposed, expressed in our notation.

PROBLEM-SOLVING (PERSON, EVENT, GOAL)

is when

CAUSE (EVENT, WANT (PERSON, GOAL))

UNTIL (OR (GET (PERSON, GOAL), GIVE-UP (PERSON)), TRY (PERSON, GET (PERSON, GOAL)))

end.

TRY (PERSON, GOAL, ACTION₁, [ACTION₂])³

is when

CHOOSE (PERSON, ACTION₁)

HOPE (PERSON, CAUSE (ACTION₁, GET (PERSON, GOAL)))

WHILE (NOT (SATISFIED (CONDITION (ACTION₁))), TRY (PERSON, ACTION₁, ACTION₂))

DO (PERSON, ACTION)

end.

The link between Problem-Solver and the more explicit problem-solving strategies must lie in the ACTIONS that the TRY schema knows about. The "particular" problem-solving strategy-schemata of Figure 3 are examples of such ACTIONS. In Rumelhart & Ortony's formulation, the problem-solving episode consists of a sequence of attempts at the goal, characterized by a sequence of ACTIONS. (The UNTIL and WHILE subschemata above are responsible for this sequential character). This seems like a good model of the problem-solving behavior of the less creative person, but some less

3

The square brackets denote an optional argument of the schema.

32

temporally constrained activation of ACTION subschemata may be called for to account for "creative" problem-solving. A suitable definition of the CHOOSE subschemata of TRY, one that is capable of returning a set of actions rather than a single ACTION, might suffice to solve this problem.

Orienting Tasks and Self-Direction

Orienting tasks can be of either of two types, instructions or questions. A task of either of these types can be either explicit or implicit. In many cases of interest, the task is implicit. For example, the statement of a problem can often be thought of as an instruction to solve it or to answer the question "What is the answer?" Orienting tasks determine the direction of subsequent cognitive processing after they have been presented. What is the role of orienting tasks in a schema-theory model of cognition?

We think that a productive way to view orienting tasks is as inputs which activate subschemata that, in turn, activate high-level strategy-schemata. Orienting tasks that are instructions will ordinarily activate Prescription-Schemata (Rigney, in press; Rigney & Munro, 1977). "Question" orienting tasks may activate information-retrieval-schemata; or inference-schemata, or problem-solving-schemata. A given orienting task may activate such schemata at several levels of abstractness. To understand what we mean by this, consider the following extended example.

A subject is presented with the following verbal logic problem.

Mr. Scott, his sister, his son and his daughter are all tennis players.
The best player's twin and the worst player are of the same sex.
The best player and the worst player are the same age.
Who is the best player?

Because of the overall setting or context (that he is a subject in a psychology experiment, for example), he takes the presentation of this problem as an implicit instruction to solve it. Because of the explicit query at the end of the problem statement, he also takes it as an explicit question. Now, if the subject is extremely familiar with this kind of verbal logic problem, then it may happen that he already has a very special-purpose problem-solving schema (a prescriptive schema) for processing them. Such a prescriptive schema might be something like "successively fill in the variables in the statements of the problem, using the names or descriptions provided, and look for contradictions in the implications of the resultant statements". Those subjects who do not have enough experience with this kind of problem will, of course, not have any specialized prescription available for solving them. Such subjects are likely to experience the activation of a number of more generalized problem-solving and question-answering schemata. Some subjects may experience an even distribution of activation among these schemata, others may experience sequential instantiations of them. (See the subsection on "Creativity and conscious control" for a schema-theory treatment of individual differences in styles of problem-solving). One such generalized question-answering schema might be called the Information-Retrieval-Schema. This schema does a kind of simple "table look-up" to see if the answer is already stored in specific memory as a result of ~~previous~~ processing. If the subject read the problem-statement in a normal manner, this look up will surely fail, because processing at the normal level of understanding did not result in the deduction of the correct answer to the problem. Another possible activation of a schema is that of an Implications-from-Related-Facts-Schema. (See the section on "inference," below). This schema would drive the further activation

of some of the concepts that were stored during the original comprehension of the text of the problem. For example, if the concept TWIN, mentioned in the problem, had been stored in an unexpanded or surface form, it could now be reactivated. As a result, the meaning of TWIN would be instantiated. Therefore, the subject would understand that the two individuals who were twins were of the same age. More processing of the concept TWIN and of the concepts SON and DAUGHTER would result in an instantiation of the concept that these two twins must be of the opposite sex.

Some orienting tasks activate very simple and very specific response schemata. For example, if an experimental subject is given a piece of paper which has printed on the first line:

NAME _____

most subjects will write their names in the blank. The string "NAME _____" serves as a very simple orienting task that activates a Write-Name-Schema. Other orienting tasks activate very complicated and much less specific response schemata. For most people not familiar with verbal logic puzzles such as the Mr. Scott problem mentioned above, there is no appropriate specific or detailed response schema available. Some problem-solving schema less immediately appropriate is therefore activated (such as the Implications-from-Related-Facts-Schema mentioned above).

This is where self-direction comes into the discussion. If a subject does not have a highly appropriate concrete response schema for this particular type of task in his generic knowledge repertoire, then he must employ some more complex and more vague task-response schema.

These complex, vague schemata have a prescriptive nature, and when they are instantiated, the subject may experience the sequence of prescriptions as a kind of "talking to himself". (I.e., "If I'm presented a problem with a lot of facts, I should think about each of the facts in some

detail, to see if any of the facts have implications for each other. So what's a fact? Let's see, at least two of these people are twins. Now . . .") When do we say that someone is being self-directed? When that person's behavior (or sequence of processing) is being driven in a top-down fashion by an abstract prescriptive schema. One is self-directed when one encounters an orienting task for which there is no special response-schema in generic memory, and one therefore responds "under the guidance of" a less concrete response-strategy schema.

Presumably, there is some fairly large store of such response-strategy schemata. The list of 66 problem-solving strategy-schemata (in the section on "Cognitive Strategies") is a subset of these schemata. When one of these general-purpose response-schemata guides processing or behavior, the person may report that he is "telling himself" to follow a certain course, and we say that the person is self-directed.

This means that some orienting tasks call for self-directed processing--namely, those tasks which are, in some important way, novel, and therefore not easily accommodated in every respect by some existing concrete strategy-schema. The orienting tasks that do not require self-directed processing are those which can be responded to adequately with a special task-specific schema.

It is possible that another characteristic of self-direction in response to orienting tasks is the activation of a high-level task-response schema. Such a schema would serve as a kind of index to more specific task-response strategy-schemata. It looks at the characteristics of the orienting task and the context in which the orienting task is presented, and compares these features with a sort of "dictionary" that matches particular combinations of task features with potentially useful response-strategy-schemata. The effect of the activation of such

a schema, in terms of subjects' task-response protocols, would be that the subjects would report that they noticed features A, B, and C of the task, which led them to postulate the possible applicability of strategy A.

Figure 4 presents a sample of the kinds of schemata that can be activated in response to orienting tasks. These schemata can vary in abstractness and generality. The most highly "tuned" schemata are those which can usefully guide responses to only a very limited class of task demands. The activation of the higher schemata on the abstractness dimension characterizes self-directed responses to tasks.

What consequences does this theoretical approach have for the training of students to be more effectively self-directed? One is that students should benefit from being taught a variety of response-strategies for different types of situations. It is important that such response-strategies should be formulated so that they are general enough to have a fairly wide range of tasks to which they are applicable. However, they should also be concrete enough that students have some good ideas about what aspects of the task or the context in which the task is presented are especially good discriminators among or diagnostics for the possible response strategies, and hence should be carefully noted. In other words, students should be taught diagnostics for the application of particular strategies.




Figure 4. Task-Response Schemata on a Dimension of Abstractness. The activation of the higher schemata on this dimension is characteristic of self-direction in response to tasks.

Inference and Depth of Processing

These two constructs receive related treatments in schema-theory and so are discussed together here. Craik & Lockhart (1972) used the "levels of processing" notion to discriminate between semantic and non-semantic processing. Klein & Salz (1976) extended the term to apply to different levels of semantic processing as well as to the difference between semantic and non-semantic processing. Both types of levels of processing will be discussed below. The term "inference" has a more varied and less technical range of uses in psychology. We will begin by suggesting a meaning for "inference" in the framework of schema-theory, and will then discuss some of the uses of the term from this point of view.

"Inference"

The approach described here for a treatment of "inference" in a schema-system is influenced in part by ideas about inference expressed to us by Rumelhart (personal communications). Rumelhart & Levin (1975) propose a model for language comprehension in which activated schemata can be "satisfied" without fully activating all of their attendant sub-schemata. (From a linguistic point of view, what this means is that lexical decomposition--as developed by Lakoff, 1970, McCawley, 1968--is not an indefatigable process. The process of understanding a word does not require the activation of the word schema and all of its sub-schemata and even all of their subschemata, as some critics, such as Kintsch, 1975, Fodor, Fodor & Garrett, 1975, have assumed. Instead, the meaning of a word can be only partially lexically decomposed or partially comprehended). What this means is that some of the subschemata of an instantiated schema may not be instantiated themselves; rather the representation of the instantiated schema includes pointers to

uninstantiated subschemata, that is, to generic representations. We could say of such "partially instantiated" calling schemata that they are "frozen" in a state of partial comprehension.

From the viewpoint of our theory, this condition of being only partially instantiated is not an unusual thing for specific concepts in memory--far from it! Rather, we believe that people very rarely process incoming information to the extent that they really come to understand every possible nuance of that information. (Fiction provides some counter-examples to this claim. Sherlock Holmes is a character who seems to process each datum for absolutely every "inference" that can be drawn from it). People instead process incoming information to the extent that some contextually-determined "demand for understanding" is satisfied. In many cases this may simply mean that processing stops when the highest level schema that can "account for the data" is reasonably satisfied that its major arguments can be instantiated.

We just referred to the status of an instantiated but not-fully-comprehended schema as "frozen." By this we mean that it can become active again, or, at least, that its uninstantiated subschemata can become activated and can, in their turn, seek their own subschemata or arguments in order to become instantiated themselves. By this means, people are sometimes able to come to a deeper level of understanding of some information that was at first encoded in a more superficial manner.

The process of "inference" has two varieties in our theory. These are "immediate inference" and "delayed inference." The former is that which occurs because some aspect of the context in which a piece of information is acquired drives (usually in a bottom-up way) more processing than would be usual. The contextually-driven schemata and the normal schema that is the generic representation of the information

being understood, seem to "cooperate." Because these activations reinforce each other, only the contextually appropriate "inferences" are instantiated and, thus, come into consciousness.

"Delayed inference" is the process referred to above, in which a "frozen," partially comprehended schema becomes reactivated. This can happen when first, some aspect of ongoing processing provides activation for the stored specific concept (i.e., something makes the person think about this particular stored episodic knowledge); and, second, some other aspect of ongoing processing (often in the context) drives the activation of one of the uninstantiated subschemata or arguments.

Of course, it is possible for there to be some top-down activation of the inference process. This would ordinarily be provided by some fairly specific problem-solving strategy-schema (such as those discussed in the section on "Conscious allocation of processing resources"). Such a strategy-schema, when activated, would have the effect of reactivating the stored specific concepts that were judged as possibly related to the problem. In effect, when this strategy-schema is activated, the problem-solver says to himself, "Let's see, maybe some of the circumstances of this problem have consequences for the solution. Maybe something can be inferred on the basis of what I already know. Well, I know Fact A. Do facts like Fact A have consequences for problems like this one? How about Fact B?" This kind of introspectively observed "talking to oneself" is simply a byproduct of the activation of the special strategy-schema that could be called the Implications-from-Related-Facts-Schema.

Here is an example (unfortunately crude and simple, but still fairly lengthy to present) of the process of "delayed inference" in schema-theory. In this example, someone is told that "Sam took the book from John"

in some context in which a particular book is being discussed. The person who hears this understands it (by which we mean that some semantic schema or schemata are satisfied by the processing of this string). However, at the time that he hears it, this person doesn't think much about it; he makes no explicit inferences based on the information. Figure 5 represents the hearer's semantic structures for the information in long-term memory. This representation is that of a very simple and superficial understanding of the utterance. Two major schemata account for the bulk of the representation--the Take-
Schema and the Book-Schema. Both of these schemata have many subschemata which happened not to be activated at the time that the person heard the sentence.

Now suppose that the person is told, "Sam knows the information contained in the book." The mention of the information contained in the book probably activates a portion of the schema for Book, a subschema that records the fact that books contain information. Figure 6 represents this simple level of understanding, after the hearer has processed these two sentences in a fairly casual way (i.e., the schemata used in understanding were not driven to produce the maximum possible number of inferences).

If, at some later point, the person who learned these facts is told that it is very important to find all those who know the information in the book, what answers will he be able to come up with? Well, to begin with, some simple information-retrieval routines should be able to automatically come up with the information that Sam knows the information, since it is already directly stored in just this format in long-term memory. Another answer that he should be able to come up with,




Figure 5. A Schema-Theory Representation of "Sam Took a Book from John."




Figure 6. A Schema Theory Representation of "Sam Took a Book from John" and "Sam Knows the Information Contained in the Book."

however, is that it is quite possible that John also knows the information, since he had the book at one time, and possessors of books often read them and thereby learn the information contained within them.

How can the ability to make such inferences be accounted for within our schema-theory notation? The information does not seem to be directly available in the representation in Figure 6. The answer to this question lies in the "reactivation" feature of the schema-model of understanding, memory, and thinking. The Take-Schema and the Book-Schema in Figure 6 can be reactivated; and those reactivated schemata, along with the schemata activated by the structure of the question ("Who might know the information in the book?"), will result in a more complex instantiation for the information about Sam and John and the book (presented in Figure 7).

In order to see how this kind of inferential thinking is done in a schema-model; we must first sketch some kind of representations for the generic meanings of the important terms, like "take" and "book". Here is the generic representation for the meaning of "take".⁴

TAKE (ACTOR, OBJECT, from PATIENT)

is when

CAUSE (DO (ACTOR, ACTION), CHANGE (from POSSESS (PATIENT, OBJECT), to POSSESS (ACTOR, OBJECT))).

end.

⁴

This representation for the meaning of "take" is adapted from that of Gentner (1975). For a related representation, see Miller & Johnson-Laird (1976).

According to this Take-Schema, "taking" is a relationship that holds between three parties, an actor, an object, and a "patient".⁵ "Taking" is when the actor does something that results in a change from a state in which the patient has the object to a state in which the actor has the object. Notice that this representation for "take" makes explicit reference to possession by the actor and the patient.

Here is the generic representation for the meaning of "book":

BOOK (x, [POSSESSOR, WRITER, ...])⁶

is when

HAVE (x, PAGES)

HAVE (x, TEXT)

CONTAIN (x, INFORMATION)

POSSIBLE (READ (POSSESSOR (of x), x))

end .

According to this incomplete Book-Schema, a book has pages and contains information. One who possesses the book may read it.

These schemata alone are not quite enough to ensure that the person who knows about Sam taking the book from John will be able to "infer" that John might know the information contained in the book.

⁵ These terms are used here in a manner very similar to the case relations employed by linguists such as Fillmore (1968), Stockwell, Schacter, & Partee (1973). Part of their function in schemata is to impose "selectional restrictions" (see Chomsky, 1965) on the arguments of the schemata.

⁶ The lower-case x in this line is used to refer to that object which is itself the book. The square brackets are used to enclose arguments of the schema that are, in some sense, optional. They need not be made explicit parts of the representation of an instantiation of the Book-Schema.

We must also specify the meaning of the Read-Schema.

READ (ACTOR, BOOK)

is-when

Here are subschemata describing some of the motor processes in reading, such as looking at the words, turning the pages, and so on. For the sake of brevity, we refer below to these processes as VIEW (ACTOR, TEXT (BOOK)).

CAUSE (VIEW (ACTOR, TEXT (BOOK)), KNOW (ACTOR, INFORMATION (OF BOOK)))

end

According to this Read-Schema, viewing the text (at least in "reading mode") causes the viewer to come to know the information in the text.

Our model of the mind of the person who hears the sentences about Sam and John can now account for the ability to make the inference that John might know the information from the book. The phrasing of the question, which specifically refers to knowing the information in the book, activates the subschema of the Read-Schema presented above. But this activation is missing a specification for the ACTOR. This drives a search for schemata of the form "POSSIBLE (READ (Unspecified, BOOK))". The search is satisfied by the last subschema of the Book-Schema shown above, in which the Unspecified argument is shown to be the POSSESSOR. This, in turn, activates a search for a schema of the form "POSSESS (Unspecified, BOOK)", where BOOK is not the generic book, but rather the specific book under discussion. If such a search can reactivate the representation for the meaning of "take", then the facts that both John and Sam possessed the book should become an explicit part of the representation.

Figure 7 represents the state of the person's long-term memory representation for the relationships among the book and Sam and John, after the reactivation of the BOOK and TAKE schemata.

In the explanation we have just given for the "inference" that John may have read the book, all the activations were driven in a bottom-up fashion from the phrasing of the question. This is not the only way in which such reactivation can take place. As was mentioned earlier in this subsection, there may be some complex strategy-schemata which are called into action during problem-solving, question-answering, and other cognitive tasks to drive inferential processing. For example, if there were an Implications-from-Related-Facts-Schema, it could have been activated by the orienting task (in this case, a question to which there was not a fully satisfactory answer based on simple retrieval strategies). This schema would then direct the activation of all the related facts known about Sam, John, and the book. This would mean that the "frozen" (in Figures 5 and 6) schema for "take" could be reactivated resulting in the explicit propositions about John and Sam separately POSSESSING the book. The reactivation of the schema for the book would result in an explicit statement about the possibility of the possessor reading the book. If this activated subschema could, in turn, activate the READ SCHEMA, so that the notion of knowing the information in the book (because of reading, because of possessing), then the complete "inference" process would be directed in a conceptually-driven manner, due to the activation of the special Implications-from-Related-Facts-Schema.

It seems likely to us that the best viewpoint on an inference process such as that just described is that the two types of processing, conceptually-driven and data-driven, must both contribute if the "inference" is to be arrived at.

"Depth of Processing".

This term is used in slightly different ways by different researchers.

Craik and Lockhart (1972) used the term "levels of processing" to distinguish between processing at sensory levels (more "superficial" or "shallow" processing) and semantic processing ("deep" processing). Some other psychologists (e.g., Klein and Salz, 1976) have experimented on different levels of semantic processing. Some of the differences in recall of meaningful materials that have been studied by Bransford & Johnson (1973) can be thought of as being due to differences in the level of semantic processing to which the materials were subjected.

The question "How are depth of processing phenomena treated in schema-theory?" is better phrased as two questions: "How are the differences between semantic and non-semantic processing treated in schema-theory?" and "How are the differences among levels of semantic processing treated in schema-theory?". Answering the first question is more difficult, given the current state of development of the schema-theory approach to cognition. Schema-theory evolved primarily as a means of modeling the understanding of meaningful materials--it is primarily a semantic-level model. The suggestions that we sketch in the following paragraph constitute only an outline of an explanation for the differences between semantic and non-semantic, more sensory, processing.

Imagine that, during visual input, a number of "feature-detectors" are activated. Feature detectors are very low-level schemata that become active when figures of certain lengths, curvatures, colors, etc. are present in the visual field. (This activation is due to data-driven processing. Needless to say, these feature-detector-schemata could also be activated in a conceptually-driven fashion, due to their being "called" by a higher schema. See Palmer, 1975). Activations of these feature-schemata can cause the activation of higher schemata,

such as letter-schemata. Letter schemata can provide activation for lower-level feature schemata or for higher-level schemata such as word-schemata. In turn, of course, these word-schemata provide activation for the letter-schemata and for higher schemata--phrase-schemata, etc. (For a more complete discussion of the interactions among these schemata, see Rigney & Munro, 1977, Rumelhart, 1977, or Rumelhart, in press). Now, what is involved in an experiment in which subjects are given instructions that are supposed to arrest their processing of a text at a "sensory" level? Suppose that subjects are instructed to count the number of "s"s in a text. From the point of view of schema theory, these instructions result in the subjects somehow "turning off" their word-level-schemata and other higher schemata. As a result of this, processing essentially stops at the letter level. There may be some activation at the word level due to the stimulus of the activations of individual letter-schemata. However, there is not enough activation of word-schemata to activate any of the more integrative or semantically comprehensive schemata, such as the mid-level content-schemata discussed in Rigney & Munro (1977). Because no higher-level, integrative conceptual structures were instantiated during the initial processing of the text, subjects do not have any "top-level" conceptual structures in memory to represent the text.⁷ This means that there is no single structure, which, when accessed, could guide the entire recall process. Rather, the only record in memory of the text is the assortment of instantiations of a number of letter- and perhaps word- and phrase-schemata, all independent in memory, essentially unrelated to each other.

⁷ In a sense, subjects in such experiments do have an instantiated schema that applies to the whole text as a result of processing it in this manner; this is a Counting-Schema. The final value of the "counter" argument of this schema is the number of "s"s in the text. Needless to say, this instantiated schema will not provide very useful cues for recall of the entire text, despite the fact that its form is determined by the text.

Now consider the second type of "depth of processing", that describes the differences between materials which are subjected to differing degrees of semantic processing. Klein and Saltz (1976) gave different groups of subjects the same lists of words with different orienting tasks. One group of subjects was required to rate each of the words for its location on a semantic dimension (e.g., if the word was "lion", "Is a lion more pleasant or more unpleasant?"). Another group was required to rate each word on two semantic dimensions (e.g., "First, is a lion more pleasant or more unpleasant? Second, is a lion fast or slow?"). Those who rated the words on two dimensions (such as pleasantness and speed) performed significantly better on a later recall test for the presented (and judged) words than did those who rated on only one dimension. Furthermore, those who rated the words on quite different dimensions (such as pleasantness and speed) did better than those who rated them on similar or correlated dimensions (such as pleasantness and happiness).

How can the Klein and Saltz results be accounted for from a schema-theory perspective? Recall the mechanisms for driving "inference" discussed in the previous subsection. In many circumstances, we claimed, the depth and kind of inferences based on a given input made by an understander would depend on some aspect of the context in which the input occurred. To apply that theory to the experiment under discussion, think of each word on the list (e.g., "lion") as the input upon which inferences can be made; the orienting tasks of making judgements about these inputs on one or more semantic dimensions constitute the contexts which determine the direction of the inferences.

For each of the words on the list, subjects have generic representations in memory. In most cases, these representations are quite

complex--they contain a great deal of information. The subject will probably not experience a powerful enough activation of the concept to instantiate all of the subschemata that represent this information. Here is a possible partial representation for someone's generic concept lion:

LION (x)
 is when
ANIMAL (x)
FELINE (x)
SIZE-OF (x, ...)
COLOR-OF (x, ...)
· FIERCE (x)
· DANGEROUS (x, to OTHER-ANIMALS)
POSSIBLE.(LOCATED (x, in ZOO)).

end.

Because activation resources are limited, not all of the above representation will be part of the specific representation for the lexical item "lion" included on the list. If, however, subjects are required to make judgements about the appropriate location of each concept on a "big-little" dimension, then the SIZE-OF subschema above will surely be activated. The more unrelated or uncorrelated dimensions upon which decisions about the lexical items must be based, the more "inferences" about each lexical item must be included in memory. The more different contexts or orienting tasks driving the activation of the generic schema, the more complete will be the specific representation of the list item.

Figure 8 represents an experimental subject's specific representation for the item "lion" on a list for which the subject is required to make a judgement only on the big-little-dimension. Only portions of the generic representation have been instantiated for this particular specific concept in memory, but one of the instantiated subschemata has to do with size. This is because the orienting task of making a size judgement helped drive the activation of the SIZE-OF subschema. Suppose that the subject had also been required to make judgements about the term "lion" on a happy-unhappy dimension. Suppose further that the subject's HAPPY schema includes some information on the implications of dangerous creatures for happiness; and that the ZOO schema includes the subschema that people are often happy at zoos. The activations of these subschemata interact with the appropriate subschemata in the generic representation LION, and, as a result, a more detailed representation for the list item "lion" is part of memory. Figure 9 is an example of such an elaborated representation.

What is the consequence of having a more complex representation in memory for a specific instance, rather than a less complex one? The answer to this question requires a detailed theory of ordinary retrieval of specific concepts from long-term memory. This is not the place to develop such a theory (the reader is encouraged to examine the initial proposals for retrieval in a schema model made by Rumelhart & Levin, 1975), but an essential property of such a system can at least be alluded to here. That property is that the greater the number of schemata that take a specific concept as an argument, the greater is the likelihood that the concept will be retrieved from memory at recall time.




Figure 8. A Possible Representation for the List Item "Lion" in the Mind of an Experimental Subject who was Required to Make a "Size" Judgement About "Lion."




Figure 9. A Possible Representation for the List Item "Lion" in the Mind of an Experimental Subject who was Required to Make Both a "Size" Judgement and a "Happiness" Judgement About "Lion."

We have shown that, in schema theory, "inference" and "semantic depth of processing" phenomena are related. Both are treated as the effect of the more extensive activation of a generic concept or schema, with the effect that more of its subschemata are activated than would otherwise be the case.

Insight

Before saying what "insight" consists of in schema-theory terms, we need to ask what is meant by this word both in ordinary language and when it is used by psychologists as an explanatory construct. Webster's New Collegiate Dictionary tells us that insight means "keen discernment or understanding; penetration; also, intuition; immediate apprehension or cognition." This definition is somewhat vague, but we shall see that various aspects of this ordinary language meaning of "insight" can be applied to some of the more technical uses psychologists have made of the term.

In the following paragraphs, we will ask how certain particular types of "insight" phenomena could be accounted for in schema theory. This is hardly an exhaustive list of the uses of the term "insight" in psychology, but we hope that certain commonalities in the use of the term may become apparent, so that the reader will be able to imagine the kind of schema-theory explanation which could account for some novel "insight" phenomenon.


We begin by treating three uses of the term "insight". The first is the use of the term by Norman (in press) to refer to the experience during complex learning, of realizing that a number of specific items of knowledge that were previously not thought to be especially similar can be thought of as examples of the same previously unknown general type.

The second type of "insight" considered here is that which is discussed in one type of problem-solving literature. This is the sort of insight that a chimpanzee achieves when he realizes that two sticks can be put together to reach a banana, or that a human subject achieves when he realizes that a common weighty tool can be tied to the end of a string suspended from the ceiling to make a pendulum. The third kind of "insight" we will discuss is that which people achieve when they solve verbal logic puzzles.

"Insight" during complex learning

Norman points out that at an intermediate stage of complex learning, students often experience a series of "Aha!" reactions to the learning process. Sometimes this happens during a Socratic dialogue, in which a tutor asks a series of questions which make the student aware of a relationship between pieces of knowledge which had hitherto seemed unrelated. Sometimes these insights are experienced when a teacher suggests a metaphor for some poorly understood concept. Sometimes these insights occur when the student simply "happens to be thinking about" several apparently unrelated bits of specific or particular knowledge.

In schema theory terms, what happens in these instances is that a new generic concept is formed in the mind of the student; that is, a new schema is created. Figure 10 is an informal sketch of the relevant mental concepts, in schema-theory form, before and after the flash of insight. In the first state, the student is aware of a number of specific facts, which are not perceived as instances of the same concept. One is a specific version of the generic concept represented by Schema_j, another is a specific instance of Schema_j, another of Schema_k.



LEGEND


-  a specific concept with arguments or parameters
-  a specific concept without arguments
-  a generic concept

Figure 10. Restructuring
 Note: Not all types of normal and expected relationships among concepts are depicted in this figure.

Then a new schema is created (as a result of the simultaneous activation of these three specific concepts, due to the Socratic questions, the metaphor made, or whatever).⁸ In the figure, the new schema is Schema_n. At this point the three specific concepts that played an important role in the attainment of the "insight" (that is, the restructuring) are now thought of by the student as specific cases of the new generic concept or schema (as well as of their old schemata). Norman's use of "insight" refers to a restructuring process, as a result of which previously unrelated concepts are now thought of as being instances of the same type.

We shall return to the question of what makes this restructuring "insightful" after examining the other two types of insight.

"Insight" during problem-solving

Consider the two problem-solving episodes mentioned above. In one, a person realizes that a tool can be used as a weight to make a pendulum, which can be set swinging and caught when it is closest to the other string; then the two strings can be tied together. In the other, a chimpanzee realizes that two sticks can be put together to draw a banana to himself. Part of what seems to be involved in these cases is the realization that some object which had been encoded in terms of one schema could equally well be encoded in terms of some other schema that has more applicability toward the problem solution.

⁸ There is no very detailed theory to account for the construction of new generic concepts. Rumelhart & Norman (in press) suggest that one possible mechanism may be to use an existing schema as a model for the new one, which then constitutes a sort of differentiation or refinement in the characterization of a class of objects, actions, or relationships. Rumelhart & Ortony (in press) refer to this process as one of "schema specialization."

Figure 11 represents some portion of the relevant concepts of the person solving the "pendulum problem" mentioned above. Before "insight" the hammer is thought of as a hammer, a tool with a specialized function, a tool which can be used to join planar objects with nails. After the "insight" in this figure, the person solving the problem is still at least potentially aware of these properties of the hammer, but now this object is no longer encoded solely with the HAMMER schema. Now the PENDULUM schema has also been activated, and the hammer is seen as playing a particular role in this activation of that schema.

It may be worth noting here that not all subjects seem to experience an activation of their PENDULUM schema spontaneously. The experimenter often activated this schema for subjects by "accidentally" brushing against one of the dangling strings while walking by, thus setting the string in an oscillating motion. Subjects ordinarily came to the pendulum solution quite quickly after the experimenter did this, but when they were later questioned, they did not remember that the experimenter had brushed the string. They said that the idea "simply came to them in a flash." This sudden and complete reinterpretation of a problem situation seems to be one of the most important reasons for labeling these solutions "insightful." In schema-theory terms, it could be claimed that the subject's perception that the strings could be put in motion resulted in a weak activation of his PENDULUM schema. At first this activation was not strong enough to result in an instantiation of the schema, but it did permit the schema to search for its arguments or sub-schemata. The identification of the hammer with the pendulum weight then contributed further activation to the schema (bottom-up activation), which, together with the activations





Figure 11: Solving the Pendulum Problem: A Schema-Theory Representation

of whatever schemata represented the subject's idea that the ends of the strings needed to be brought together (a top-down activation), resulted in adequate activation for an instantiation. As a schema is being instantiated, we say it is in consciousness or in working memory. (See the section on "Consciousness"). Therefore, from the schema-theory viewpoint, it is quite natural that there is no awareness of the possible solution to the problem until a relatively complete solution has been obtained, for only then is there adequate activation to allow an instantiation.


"Insight" in verbal logic puzzle solutions

In attacking verbal logic puzzles, subjects may proceed to the solution in a variety of ways. Each path to the solution is marked by a number of "insights" about constraints on the relationships between the individuals mentioned in the problem statement. In the Mr. Scott problem (discussed in the section on "Orienting Tasks and Self-Direction"), for example, a number of intermediate "insights" are required before a subject can achieve the final insight that leads to a solution. One such insight is the realization that there is at least one set of twins in the family: either Mr. Scott and his sister are twins and/or his son and daughter are twins. Another insight is that there are a number of possible same-age relationships in the family (Mr. Scott and his sister, the son and the daughter, the son and the sister, the daughter and the sister), and that some of these possible same-age relationships are contradictory but others are not.

Let's consider what form the first of these "insights" would take in a schema-theory representation of a person's understanding of the problem.



Some of the relevant concepts in the puzzle-solver's mind before "insight"



Some of the relevant concepts in the puzzle-solver's mind after one "insight"

Figure 12: A Simple "Insight" Necessary to the Solution of the Mr. Scott Problem: A Schema-Theory Representation

In the first part of Figure 12, the left portion of the figure represents the problem-solver's understanding of the familial relationships between Mr. Scott, his sister, his son, and his daughter. This structure is part of the representation that results from the processing of the first sentence of the problem. The right portion of the figure is a partial representation of the problem-solver's understanding of the sentence, "The best player's twin and the worst player are of the same sex." The argument of the Best-Player Schema is a dummy. This means that the parameter is really unfilled. Presumably schemata do not like to be instantiated with unfilled arguments, and this state of affairs drives further processing to try to "fill" the dummy argument with a real, specific concept, namely one of the concepts that represent the members of Mr. Scott's family. As we saw in the earlier discussion of inference, this kind of processing will activate some conceptual structures which had been only partially comprehended before. If the TWIN-schema included in Figure 12 is activated, it activates a sub-schema called SIBLING (as well as other subschemata, of course). The SIBLING schema is something like this:

SIBLING (x, y)

is when

AND (OFFSPRING (x, z), OFFSPRING (y, z))

end

The OFFSPRING schema has this form:

OFFSPRING (x, 'z)

is when

OR (SON (x, of z), DAUGHTER (x, of z))

end .

A chain of activations is set off: First TWIN activates SIBLING. Then SIBLING sponsors two activations of OFFSPRING. These OFFSPRING activations, in turn, look for instances of SON or DAUGHTER. In the "before" section of Figure 12 are an instance of each, SON and DAUGHTER, that are able to satisfy the two activations of OFFSPRING. This means that SIBLING and TWIN are also satisfied. As a result, the concepts that represent Mr. Scott's son and daughter are marked as possibly being twins in the "after" portion of Figure 12. (The reason that this instantiation of TWIN is "hedged" by the "Possible" is that there are other important subschemata of TWIN that have not yet been satisfied. For example, there must be a subschema, not represented in the above definition, that expresses the fact that twins must be born at the same time). A second instantiation of TWIN, also marked as "possible" is included in the second part of Figure 12. This instantiation is arrived at more readily than the first, since the SIBLING schema is already present in the first part of Figure 12.

The second part of this figure thus represents the change in the reader's understanding of the Mr. Scott problem after he has realized that there are two possible sets of twins and has assigned the appropriate specific concepts to these-possible Twin Schemata.

There are a number of other "insights" necessary to the solution of this problem. Most of these can be thought of as the kinds of delayed inference discussed in "Inference and depth of processing," above.

Common aspects of insight in these three cases.

In each of these types of insight, new instantiations of schemata occurred (sometimes with old generic schemata, sometimes with new ones). In each case, the new instantiations related two or more pre-existing, specific concepts in a novel way. Perhaps this is exactly what insight consists of: the recognition of a previously unnoticed relationship. The more unexpected such relationships are, the more likely we are to label their discovery as instances of "insight".

III. SOME REFLECTIONS ON SCHEMA-THEORY

A number of explicit suggestions for the treatment of cognitive processes in terms of schemata have been presented in II. Here we intend to step back and present an overview of schema-theory and, hopefully, clear up some questions that remain. The relationship between schema-theory and more traditional information-processing theories is discussed, and the varieties of schemata are explicated.

Schema-Theory and Cognitive Component Theories

The presentations of many information-processing theories are characterized by diagrams such as that in Figure 2A. We refer to theories of this type, which follow the flow of information through a series of discrete cognitive components, as "cognitive component theories". In such theories, concepts in memory that are not part of consciousness are thought of as being in a vast storage depot, called long-term memory. When certain features (provided by a feature extractor) enter a smaller storage area, called short-term memory, they may cause some of the concepts in long-term memory to be transferred (or copied) into short-term memory. The same phenomena that are accounted for in a cognitive components theory in this manner are treated in a slightly different way in schema-theory. In schema-theory there is only one storage area, but the concepts within it are thought of as being in different states of activation. In general, the least activated schemata are similar to the contents of long-term memory in a cognitive components theory. The most activated schemata are those concepts which, in cognitive components terms, are in short-term memory.

The activation construct that is so central to schema-theory provides a natural means for accounting for the resource-limited processing discussed by Norman & Bobrow (1975). The resource limitations they treat can all be thought of as being due to limitations on activation resources.⁹ There is no similar pre-existing feature in cognitive component models to account for the facts of resource limitations.

A demonology: The varieties of schemata

Schemata can be distinguished from each other by three means. The most basic means of distinguishing schemata is in terms of their functions. That is, when a schema acts like a procedure, what does it do? Every schema is unique in terms of function. No two do exactly the same thing in every environment or they would not be two schemata, but one schema. To categorize all the possible functions of schemata would be an arduous task. There are schemata to recognize/identify visual features, to recognize auditory features, to provide the meanings of words, to represent known facts about the world, to guide motor activities, to plan, to solve problems, and to carry out all the other functions of human beings that might be considered even remotely intellectual in nature. We are not prepared to present a classification in terms of function.

A second means of distinguishing schemata is in terms of abstractness.

⁹

We remain uncommitted on the issue of whether total activation resources available to the system can vary within certain limits or whether they are always fixed at one level. The functions of the reticular activating system argue for a model with variable levels of activation, but, as was explained in footnote 1, schema theory is not intended as a physiological model.

Figure 4, above, presents a sample of the schemata described in this paper, arranged in order of their abstractness. The more abstract a schema is, the more different situations it can apply to. The more concrete or specific a schema is, the more restricted is its range of application. Very highly-tuned schemata tend to be most context-specific or non-abstract. These are schemata that are prepared to respond to one particular type of situation, and not to others. The best examples of these are motor schemata. For example, we would claim that a professional tennis player would have many thousands of specialized schemata for hitting a tennis ball (differing slightly from each other with respect to the speed and position of the ball and the player), while an amateur player would have many fewer, more abstract schemata. (Perhaps the rank beginner would have only three--one for serving, one for backhands, and one for forehands). Another example of extremely concrete schemata is that of the memories for chess positions of master chess players. In general, according to de Groot (1966), the more advanced the player, the larger will be his repertoire of quite detailed, non-abstract memories for types of chess positions.

A third means of distinguishing schemata is in terms of their scope. Table 1 presents a number of schemata, classified according to their scopes. A schema with wide scope is one that can account for a great deal of data. Typically, it has a large number of subschemata, each of which have subschemata of their own. An example of a schema with small scope is a feature-detection schema. An example of a schema with very wide scope is a text-understanding schema such as that proposed by Rumelhart (1975) for narratives.

Examples of Schemata Activated in Three Contexts

	Schemata activated in problem-solving (discussed above)	Schemata activated in text-processing (Rigney & Munro, 1977)	Schemata activated in conversation-understanding (Munro, 1977)
<p style="writing-mode: vertical-rl; transform: rotate(180deg);">Increasing Scope</p> <p>High-level schemata: recognize situation/context</p> <p>Schemata that establish a particular context set</p> <p>Mid-level schemata, more particular</p> <p>Content-schemata: Often represent meanings of lexical items</p>	<p>Problem-Solver (pp. 23-29)</p>	<p>Narrative Explanation Representation Prescription</p>	<p>Conversation</p>
	<p>Build-Up Eliminate Work Forward } (pp. 26-28)</p>	<p>Sperling-Pardigm-Article</p>	<p>Direction-Giving</p>
	<p>Elaborate Implications-From-Related-Facts (pp. 30-31)</p>	<p>Boats-Sail-On-Lakes</p>	<p>Confirmation-Request Landmark-Identify.</p>
	<p>Family Twin Sibling. Same-Age } (pp. 59-60)</p>	<p>Volkswagen</p>	<p>Assertion Information-Request Instruct</p>

Table 1. An Overview of Some Schemata.

Schema-theory is a very powerful system. It is intended to be capable of representing all the varieties of knowledge that people have. In light of the fact that there are at present no adequate theories of human cognitive processes, we believe that a powerful theory such as this is called for. Let Occam's Razor apply when we have two adequate theories to choose between.

REFERENCES

- Adams, J.L. Conceptual blockbusting. Stanford, California: The Stanford Alumni Association, 1974.
- Anderson, J.R. FRAN: A simulation model of free recall. In G.H. Bower (Ed.), The psychology of learning and motivation. New York: Academic Press, 1972. Vol. 5.
- Anderson, J.R. & Bower, G.H. Human associative memory. Washington DC: V.H. Winston & Sons, 1973.
- Anderson, R.C. & Ortony, A. On putting apples into bottles: A problem of polysemy. Cognitive Psychology, 1975, 7, 167-180.
- Anderson, R.C., Pichert, J.W., Goetz, E.T., Schallert, D.L., Stevens, K.V. & Trollip, S.R. Instantiation of general terms. Journal of Verbal Learning and Verbal Behavior, 1976, 15, 667-679.
- Bobrow, D.G. & Winograd, T. An overview of KRL, a knowledge representation language. Cognitive Science, 1977, 1, 3-46.
- Bransford, J.D. & Johnson, M.K. Consideration of some problems of comprehension. In W.G. Chase (Ed.), Visual Information Processing. New York: Academic Press, 1973.
- Chomsky, N. Aspects of the theory of syntax. Cambridge: The M.I.T. Press, 1965.
- Craik, F.I.M. & Lockhart, R.S. Levels of processing: A framework for memory research. Journal of Verbal Learning and Verbal Behavior, 1972, 11, 671-684.
- Dellas, M. & Gaier, E. Identification of creativity. Psychological Bulletin, 1970, 73, 55-73.
- Dresher, B.E. & Hornstein, N. On some supposed contributions of artificial intelligence to the scientific study of language. Cognition, 1976, 4, 321-398.
- Dreyfus, H.L. What computers can't do. New York: Harper and Row, 1972.
- Fillmore, Charles J. The case for case. In E. Bach & R.T. Harms (Eds.), Universals in linguistic theory; New York: Holt, Rinehart, & Winston, 1968.
- Fiskel, J.R. & Bower, G.H. Question-answering by a semantic network of parallel automata. Journal of Mathematical Psychology, 1976, 13, 1-45.

References - continued

- Fodor, J.D., Fodor, J.A. & Garret, M.F. The psychological unreality of semantic representations. Linguistic Inquiry, 1975, 4, 515-531.
- Friendly, M.L. In search of the M-gram: The structure of organization in free recall. Cognitive Psychology, 1977, 9, 188-249.
- Gentner, D. Evidence for the psychological reality of semantic components: The verbs of possession. In D.A. Norman, D.E. Rumelhart & LNR, Explorations in Cognition. New York: Academic Press, 1975.
- Gentner, D.R. The FLOW tutor: A schema-based tutorial system. Paper to be presented at the fifth international joint conference on artificial intelligence, Cambridge, MA, August, 1977.
- Goldstein, I. & Papert, S. Artificial intelligence, language and the study of knowledge. Cognitive Science, 1977, 1, 84-123.
- Guilford, J.P. Intelligence, creativity, and their educational implications. San Diego, CA: Robert-Knapp, 1968.
- Half, H.M., Ortony, A. & Anderson, R.C. A context-sensitive representation of word meanings. Memory and Cognition, 1976, 4, 378-383.
- Hayes-Roth, B. & Hayes-Roth, F. The prominence of lexical information in memory representations of meaning. Journal of Verbal Learning and Verbal Behavior, 1977, 16, 119-136.
- Kintsch, W. Meaning and the representation of knowledge. Hillsdale, New Jersey: Lawrence Erlbaum Associates, 1975.
- Klein, K. & Saltz, E. Specifying the mechanisms in a levels-of-processing approach to memory. Journal of Experimental Psychology: Human Learning and Memory, 1976, 6, 671-679.
- Lakoff, G. Irregularity and syntax. New York: Holt, Rinehart & Winston, 1970.
- Levin, J.A. Proteus: An activation framework for cognitive process models. (Working paper ISI/WP-2). Marina del Rey, California: University of Southern California, Information Sciences Institute, 1976.
- Levin, J.A. & Moore, J.A. Dialogue games: Meta-communication structures for natural language interaction. (Tech. Rep. No. ISI/RR-77-53). Marina del Rey, California: University of Southern California, Information Sciences Institute, January 1977.
- McCawley, J.D. Lexical insertion in a transformational grammar without deep structure. Papers from the Fourth Regional Meeting, Chicago Linguistic Society. Chicago: Chicago Linguistic Society, 1968.

References - continued

- Miller, G.A. & Johnson-Laird, P.N. Language and perception. Cambridge, MA: Belknap Press, 1976.
- Minsky, M. A framework for representing knowledge. In P. Winston (Ed.), The psychology of computer vision. New York: McGraw-Hill, 1975.
- Munro, A. Speech act understanding in context, (Doctoral dissertation, University of California, San Diego, 1977).
- Norman, D.A. Notes toward a theory of complex learning. In press.
- Norman, D.A. & Bobrow, D.G. On data-limited and resource-limited processes. Cognitive Psychology, 1975, 7, 44-64.
- Norman, D.A., Rumelhart, D.E. & The LNR Group, Explorations in cognition. San Francisco: Freeman, 1975.
- Palmer, S.E. Visual perception and world knowledge: Notes on a model of sensory-cognitive interaction. In D.A. Norman, D.E. Rumelhart & LNR, Explorations in cognition. New York: Academic Press, 1975.
- Quillian, M.R. Semantic memory. In M. Minsky (Ed.), Semantic information processing. Cambridge, MA: M.I.T. Press, 1968.
- Rigney, J.W. On cognitive strategies for facilitating acquisition, retention, and retrieval in training and education. In press.
- Rigney, J.W. & Munro, A. On cognitive strategies for processing text. (Tech. Rep. No. 80). Los Angeles: University of Southern California, Behavioral Technology Laboratories, March 1977.
- Rumelhart, D.E. Notes on a schema for stories. In D. G. Bobrow & A. Collins (Eds.), Representation and understanding: Studies in cognitive science. New York: Academic Press, 1975.
- Rumelhart, D.E. Introduction to human information processing. New York: John Wiley & Sons, 1977.
- Rumelhart, D.E. Toward an interactive model of reading. In S. Dornic (Ed.), Attention and performance VI. Hillsdale, New Jersey: Lawrence Erlbaum Associates, in press.
- Rumelhart, D.E. & Levin, J.A. A language comprehension system. In D.A. Norman, D.E. Rumelhart & LNR, Explorations in cognition. New York: Academic Press, 1975.
- Rumelhart, D.E., Lindsay, P.H. & Norman, D.A. A process model for long-term memory. In E. Tulving & W. Donaldson (Eds.), Organization of memory. New York: Academic Press, 1972.

References - continued

Rumelhart, D.E. & Norman, D.A. Accretion, tuning and restructuring: Three modes of learning. In J.W. Cotton & R.L. Klatzky (Eds.), Semantic factors in cognition. Hillsdale, New Jersey: Lawrence Erlbaum Associates, in press.

Rumelhart, D.E. & Ortony, A. The representation of knowledge in memory. In R.C. Anderson, R.J. Spiro & W.E. Montague (Eds.), Schooling and the acquisition of knowledge. Hillsdale, New Jersey: Lawrence Erlbaum Associates, in press.

Schank, R.C. & Abelson, R.P. Scripts, plans and knowledge. In The proceedings of the fourth international joint conference on artificial intelligence, 1975.

Stockwell, R.P., Schacter, P. & Partee, B.H. The major syntactic structures of English. New York: Holt, Rinehart & Winston, 1973.

Weisenbaum, J. Computer power and human reason: From judgement to calculation. San Francisco: W.H. Freeman & Co., 1976.